



University of
Nottingham

UK | CHINA | MALAYSIA

Experimental design for discrepant models: selecting mechanistic models of I_{Kr} through informative experiments

Joseph G. Shuttleworth

School of Mathematical Sciences

University of Nottingham

A thesis submitted for the degree of

Doctor of Philosophy

September 2024

Acknowledgements

The completion of this thesis wouldn't have been possible without the help and support of many people, and I would like to take a moment to thank some of them here. First, my thanks go to my supervisors, Gary, Simon, Chon and Dom for their guidance, advice, and patience throughout this process. With their help I have learnt a great deal over the course of my studies regarding biology, statistics and computation (hopefully!), and I am very thankful for all the writing advice, and support with my publications.

I would also like to thank those at the Victor Chang Cardiac Research Institute in Sydney—especially Adam and Monique, without whom we wouldn't have the data presented in this thesis. I'm also very grateful that I got to visit Australia and meet many friendly and insightful researchers.

Also, I have had the pleasure of working alongside a great group of people in Nottingham. Plenty of people in the School have gone out of their way to chat with me, and provide lots of help and career advice. In particular, the postdocs and students that I have worked alongside have made being at the University, and our travel for conferences, great fun.

A special mention goes to my good friend and tech guru Dave for providing support and advice. I also want to thank my amazing partner, Duasha, and her family and friends, for providing plenty of love and support (both emotionally and nutritionally). **බොහොම ස්තූතියි! மிக்க நன்றி!**

Finally, I want to convey my gratitude to my family, especially my mum and dad, Sam, Nanna Pat and my late Grandad Jim for their encouragement – and for buying me lots of tea and coffee in Wollaton Park over the last four years. Without your hard work, guidance and support, this would not have been possible.

Thank you everyone!

Joey

Abstract

In cardiac electrophysiology applications, mathematical models are relied upon to provide quantitatively accurate predictions. The potassium current I_{Kr} is of particular importance because its blockage by drugs is known to cause dangerous changes in heart rhythm. To help quantify the risk presented by potential new drugs, we require mathematical models which produce quantitatively accurate predictions and accurately describe the gating behaviour of a cell's ion channels (particularly those which carry I_{Kr}). Building such models requires experimental data that allows both the selection of appropriate model structures and the inference of model parameters. Recently, short “information-rich” experimental designs have been developed, allowing cell-specific models to be fitted. Here, new experimental designs can be specified, which allow the collection of new data under previously unseen I_{Kr} dynamics. The resulting data promise to improve models of the gating dynamics of I_{Kr} , improve our understanding of cardiac electrophysiology as a whole, and improve the risk-classification of new drugs.

In this thesis, we introduce the methods necessary for the fitting and validation of cell-specific mathematical models of I_{Kr} using a diverse ensemble of information-rich experimental designs. This approach allows us to produce empirical quantifications of predictive uncertainty and permits the comparison of literature models in terms of their predictive accuracy and the variability of parameter estimates across cells. Whilst some models produce more accurate predictions than others, a certain amount of model discrepancy seems unavoidable. Our results suggest that this discrepancy is caused, in part, by the presence of experimental artefact effects, which, when unaccounted for, confound our parameter estimates, contributing appreciably to the apparent cell-to-cell variability of parameters relating to the kinetics of channel gating. Moreover, we demonstrate that our broadly applicable multiprotocol approach allows for thorough validation of our models, a realistic quantification of predictive uncertainty, and the selection of suitable mathematical models.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Selection and validation of mechanistic models	3
1.3	Structure of the thesis	5
2	Literature Review	7
2.1	Electrophysiology	7
2.1.1	Patch-clamp experiments: state of the art	11
2.2	Mathematical models of cellular electrophysiology	12
2.2.1	The Hodgkin-Huxley action-potential model	12
2.2.2	Cardiac action-potential models	15
2.2.3	Isolated ion-channel current models	16
2.3	Fitting models of I_{Kr} to patch-clamp data	24
2.3.1	Maximum likelihood estimation	25
2.3.2	Numerical optimisation	27
2.3.3	Integration of ODE systems	28
2.3.4	Identifiability	31
2.3.5	Local sensitivities for ODE-based models	35
2.3.6	Experimental design	37
2.4	Sources of predictive uncertainty	39
2.4.1	Noise model	39
2.4.2	Parameter uncertainty	40
2.4.3	Structural uncertainty	40
2.5	General approaches for model discrepancy	42
2.6	Summary	42
3	Uncertainty Quantification with an ensemble of experimental designs	43
3.1	Preamble	43
3.2	Introduction	43
3.2.1	Motivating Example	44
3.3	Methods	47

CONTENTS

3.3.1	Fitting models using multiple experimental protocols	48
3.3.2	Application to I_{K_r} models	50
3.4	Results	58
3.4.1	Case I: Misspecified maximal conductance	58
3.4.2	Case II: Misspecified dynamics	60
3.5	Discussion	67
3.5.1	Limitations	68
3.5.2	Concluding remarks	69
4	Fitting an experimental, multiprotocol dataset	70
4.1	Introduction	71
4.2	Experimental methods	72
4.2.1	Design of voltage-clamp protocols	74
4.3	Postprocessing	78
4.3.1	Methods	78
4.3.2	Quality control	85
4.4	Fitting and validation of models	94
4.4.1	Computational methods	103
4.4.2	Model fitting results	105
4.4.3	Predictive accuracy under differing assumptions	106
4.5	Variability of parameter estimates	117
4.6	Discussion	128
5	Accounting for experimental artefacts in fitting improves model accuracy	130
5.1	Preamble	130
5.2	Introduction	130
5.3	Mathematical modelling of experimental artefacts	132
5.3.1	Impact of experimental artefacts on postprocessing	134
5.4	Fitting predictive models of I_{K_r} using the artefact model	143
5.5	Results	143
5.5.1	Model fitting	143
5.5.2	Predictive accuracy	145
5.5.3	Variability in parameter estimates	152
5.5.4	Variability in time constants of decay	160
5.6	Discussion	163
6	Discussion	167
A	Fundamental properties of Markov models	172
A.1	Hodgkin-Huxley style gating variables as Markov models	173

CONTENTS

A.2	Steady states of Markov models	175
A.2.1	Example I: A connected Markov model	179
A.2.2	Example II: A disconnected Markov model	180
A.2.3	Application of the Perron-Frobenius Theorem	182
A.3	Single-channel models and macroscopic currents	185
A.4	Reversible Markov models	187
A.5	Simplifying the governing equation	190
B	Supplementary material for Chapter 3	194
B.1	I_{Kr} model equations	194
B.2	Further Case I results	195
B.3	Further Case II results	197
C	Effects of capacitance spike removal	201
D	Model details	207
D.1	Closed-Open-Inactive model	207
D.2	Kemp model	208
E	Further statistical analysis of variability in parameter estimates	210
F	Dynamics of the artefact model	220
F.1	Dynamics under ramps	221
F.2	Steady states of the artefact model	222
	Bibliography	225

Chapter 1

Introduction

1.1 Motivation

Mathematical models are used across the biological sciences to provide qualitative and quantitative descriptions of the dynamics of biological phenomena, such as the spread of infectious diseases (Zhang et al., 2020), the response of the immune system to antigens (Hu et al., 2007), and the function of organs such as the heart (Noble, 2002). When models are used in safety-critical settings (such as drug safety or clinical decision-making), it is often important that our models produce quantitatively accurate predictions over a range of scenarios—to predict the impact a certain drug will have on the function of the heart under different physiological conditions, for example. In such cases, we rely on the accuracy of our models as quantitative descriptions of the underlying mechanisms. Such models may then be used to produce accurate predictions, even in situations that are dissimilar to experiments used for model fitting. Perhaps more importantly, these models must allow a reliable quantification of predictive uncertainty if they are to be trusted when making safety-critical decisions, and shaping scientific conclusions (Mirams et al., 2016).

Quantifying the predictive accuracy of models is a key aim of the field of Uncertainty Quantification (UQ) (Smith, 2013). Exact models of biological phenomena are generally unavailable, and we resort to using approximate mathematical models instead. When a mathematical model does not fully recapitulate the *data-generating process* (DGP) of a real biological system, we call this *model discrepancy* or model misspecification. This discrepancy between the DGP and our models presents a particular challenge for UQ.

For example, in the field of *electrophysiology*, mathematical models are commonly used to describe the electrical dynamics of many types of cell in the body—especially excitable cells such as muscle cells and neurons. Such models continue to play an important role in our understanding of the electrical signalling which coordinates the pumping of the heart (*cardiac electrophysiology*), and the mechanisms by which *cardiac*

arrhythmia (dangerous, irregular heart rhythms) can arise. Models across a range of scales are used for this purpose—including single-cell models, tissue models, and whole-organ models.

Single-cell models themselves rely on accurate quantitative models of the subcellular dynamics responsible for the production of action potentials. These mathematical models are fitted using data from various types of experiments, and describe various aspects of cellular electrophysiology (with and without the presence of drugs). Such models are, generally speaking, large-scale approximations of the temporal and spatial dynamics dictating the movement of individual molecules, and, as such, are not perfectly accurate.

There are also various sources of uncertainty which affect the predictions of such models, such as statistical uncertainty in parameter estimates, and structural uncertainty regarding the particular mathematical equations included in the models. Nevertheless, these approximate models can offer a good description of the underlying biophysics, and allow an imperfect but very useful risk classification for new compounds. These models are used in other applications, too. For instance, these single-cell models, through their inclusion in patient-specific organ-level models, have been used to guide ablation in ventricular tachycardia patients (Trayanova et al., 2020). Such applications rely on the selection of suitable model structures, and the accurate estimation of model parameters.

With a view to better understand drug action, and improve the classification of pro-arrhythmic risk in novel compounds, we aim to address this structural uncertainty by using previously unseen experimental designs to inform model selection for a particular component of the cellular action potential, I_{Kr} .

The importance of I_{Kr} modelling

The Comprehensive *in-Vitro* Pro-Arrhythmic Assay (CiPA) is a large scale project which aims to better classify the safety of potential drug compounds (Li et al., 2017b), an important aspect of which is the interaction of novel pharmaceutical compounds with a particular ion-channel current called I_{Kr} (Li et al., 2020). This biological current has been the subject of considerable study due to its importance to the cardiac action potential, and the association of its block with arrhythmia (Roy et al., 1996; Li et al., 2017b). In particular, the block of I_{Kr} is associated with Long-QT Syndrome type-2 (LQT2), which, in turn, is associated with the onset of *Torsade de Pointes* (TdP), a potentially fatal type of arrhythmia (Rampe and Brown, 2013).

The risk that drugs may cause unsafe changes to the heart’s function was first highlighted in astemizole and terfenadine (two antihistamines which have since been removed from markets) and their propensity to block I_{Kr} , cause QT prolongation and increase the risk of TdP (Suessbrich et al., 1996). Since then, many other compounds

have been found to interact with I_{Kr} . It has been suggested that the channel's central pore is particularly susceptible to block by small-molecule drug compounds (Maly et al., 2022). Since the danger of I_{Kr} block was first highlighted, the screening of drugs for their interactions with I_{Kr} has become a key part of the drug-development process. The aim of such screening is to quickly and efficiently identify risky compounds (those deemed likely to have unsafe interactions with I_{Kr}) so that they may be removed from consideration at early stages of the drug-development process. Accurate mathematical models of these dynamics are increasingly being used to assess the impact of promising pharmaceutical compounds (Mirams et al., 2011, 2012).

A key focus of the CiPA initiative is the use of mechanistically-informed mathematical models to classify the safety of drugs. As discussed in Chapter 2, such applications require I_{Kr} to accurately describe the “conformational state” of a cell's ion channels, and to produce accurate predictions of I_{Kr} in normal and drug-blocked conditions. These model predictions allow the impact of drugs on the cardiac action potential to be quantified.

1.2 Selection and validation of mechanistic models

In situations where quantitatively accurate models are used to make safety-critical decisions, it is imperative to provide a realistic quantification of uncertainty in the model output. Throughout this thesis, we aim to fit models which provide both quantitatively accurate predictions (under a diverse range of experimental designs), and an accurate description of the underlying biophysical processes.

There are many possible experimental designs which could be used to fit and validate models of I_{Kr} . The designs explored in this thesis consists of different voltage signals which are applied to individual cells (cells which are prepared as to exhibit large I_{Kr} and minimise the presence of other currents). These designs are used to both train and validate models, and the use of novel designs in this thesis promises to drive model improvement, as shown in Figure 1.1.

It would, of course, be desirable for our models to provide near-perfect predictions across each and every possible voltage signal, even when the experimental design used for validation is dissimilar to that used for model fitting. Provided there are sufficient data, we would expect a perfectly accurate model to provide accurate predictions across all possible experimental designs. However, this general predictive accuracy is difficult to test in practice, and such a model is, most likely, unattainable. We instead aim to identify a model which produces accurate predictions (on average) when using a wide range of experimental designs for both fitting and validation—that is, we seek to fit models of I_{Kr} that provide an accurate response to a wide range of different voltage signals.

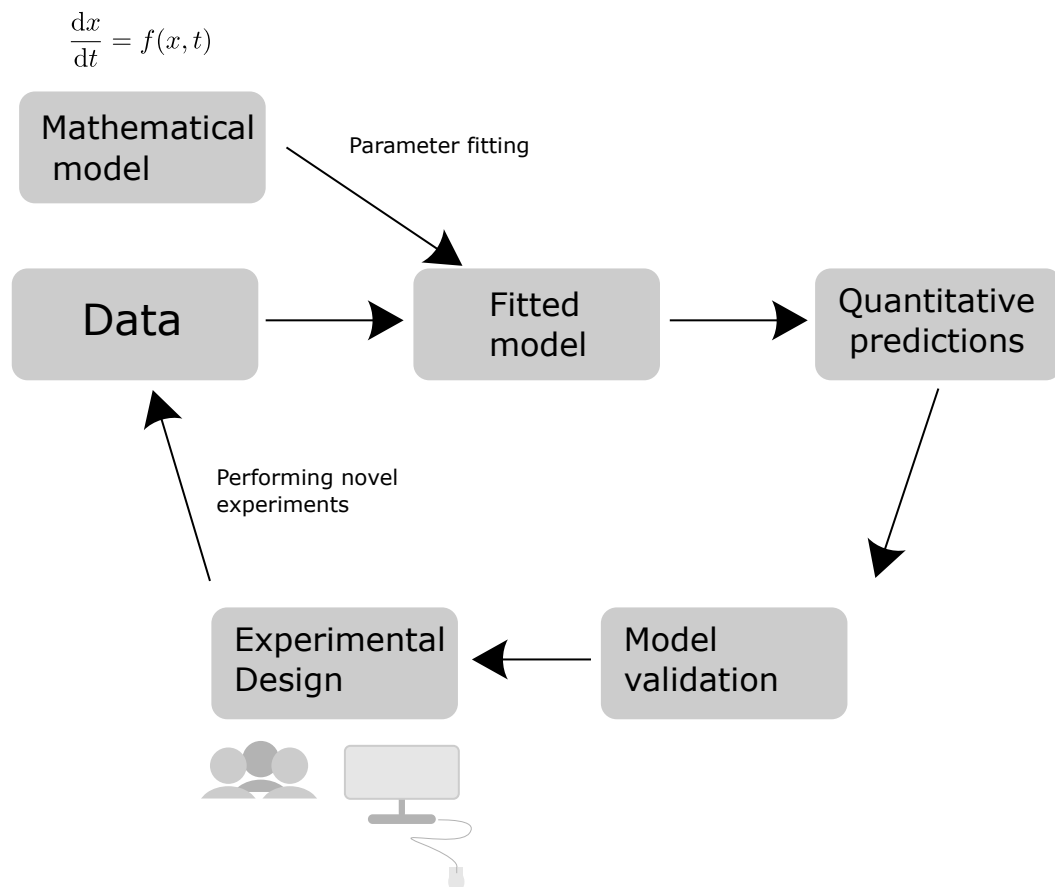


Figure 1.1: The model building process. Mathematical models are proposed and are then fitted (or trained) using experimental data. These fitted models are then used to produce quantitative predictions (sometimes for industry and clinical use). These predictions can also be used to inform the design of further experiments, and may suggest ways in which the original mathematical model can be improved.

In this thesis, we compare a small cohort of I_{Kr} models, seeking to identify the most accurate predictive models by using real, experimental data collected under a diverse range of experimental designs. This approach provides new insights into the dynamics of I_{Kr} gating, and the experimental methods used to record our data. Though this work concerns the specific details of this particular modelling problem, the methods introduced herein are presented as generally as practical with the expectation that they may prove useful for other modelling problems—both in electrophysiology and elsewhere in other biological application areas where mathematical models are used to provide both quantitative and mechanistic insights.

1.3 Structure of the thesis

In Chapter 2 we provide a brief introduction to I_{Kr} , and the mathematical models and experiments which have been used in its study. Particular attention is paid to Markov models—a type of differential-equation-based model commonly used to study ion-channel currents like I_{Kr} . Chapter 2 ends with an introduction to the field of Optimal Experimental Design, and an overview of the limitations of Markov models of I_{Kr} .

In Chapter 3, we investigate the behaviour of the mathematical models discussed in Chapter 2 under a range of *information-rich* experimental designs (our *multiprotocol* approach). Here, we use synthetic data to characterise the effects of model discrepancy. We pay particular attention to the structural uncertainty induced by the variety of conflicting I_{Kr} models in the literature.

Then, in Chapter 4, we use similar methodology to fit and validate four literature models to real patch-clamp recordings of I_{Kr} . We introduce existing and novel methods to handle and process experimental data, and use these methods to provide a thorough analysis of a newly-collected dataset. Here, novel experimental designs are performed on a high-throughput automated patch-clamp machine, allowing the evaluation of our models' predictive accuracy for as-of-yet unseen experimental designs. Our results here suggest promising routes for model improvement.

Building on this work, we modify our models to include *experimental artefacts* in Chapter 5. These effects are omitted in Chapters 3 and 4, and in much of the existing literature, in lieu of simpler models. Following the introduction of these artefact effects and their implications for our mathematical models, we investigate their impact on the postprocessing procedures introduced in Chapter 4. Then, we use similar methodology to that used in Chapter 4 to validate our models under a range of protocols, providing a comparison between those models which include artefact effects, and those which do not. Here, our results highlight the importance of these artefact effects, show that their inclusion leads to more accurate models, and suggest particular I_{Kr} model structures for use in future studies.

In Chapter 6, we conclude by discussing the necessity of including artefact effects in our models, and the importance of thorough model validation (as made possible by our multiprotocol approach). Suggestions for future work then follow, where we suggest further analyses to be carried out using the experimental data introduced in Chapter 4, provide some recommendations for future data collection, and suggest other ion-channel currents for which our methodology may also prove useful.

Chapter 2

Literature Review

This chapter provides an introduction to the experimental and mathematical methodologies used to study I_{Kr} , a voltage-gated ion-channel current that plays a particularly important role in the electrical signalling of the heart. Section 2.1 is a brief introduction to the field of cellular electrophysiology, explaining the importance of ion-channel currents such as I_{Kr} , and the experimental methods used to study them.

Then, a general introduction to literature models of I_{Kr} follows in Section 2.2, where some fundamental properties of these models are summarised. Next, Section 2.3 provides an overview of relevant statistical and computational methods used to fit models and design experiments in this thesis and the wider literature. Finally, in Section 2.4 we discuss the limitations of our models, summarise the different types of uncertainty which can affect the accuracy of our predictions, and discuss some approaches that have been suggested to address these limitations.

2.1 Electrophysiology

Electrophysiology is the study of electrical activity in biological cells and tissues. Like many other types of excitable cells (such as nerve cells and smooth muscle cells), *cardiomyocytes* (heart-muscle cells) produce and respond to electrical signals, coordinating the pumping-action of the heart. This phenomenon is made possible by subcellular mechanisms which selectively allow specific ion species (such as K^+ , Na^+ , Ca^{2+} and Cl^-) to cross the cell membrane (Alberts et al., 2015). In addition to the maintenance of cellular homeostasis, the interplay between these subcellular mechanisms, ionic concentrations and electrical potentials across the cell membrane gives rise to complex electrophysiological dynamics.

Ion channels are one type of subcellular component which play a key role in these electrophysiological dynamics, permitting the transport of ions across the cell membrane. This is important because ions are unable to pass through the cell membrane's




Figure 2.1: The $K_V11.1$ channel *opens*, *closes*, and *inactivates* in response to changes in the transmembrane voltage. Ions like K^+ are unable to pass through the cells' lipid bilayer, but may pass through ion channels (which select a particular species of ion, such as K^+). When the cell is *depolarised*, the electrochemical gradient drives K^+ ions out of the cell, helping the cell return to its resting potential.

impenetrable lipid bi-layer. Ion channels are protein structures that are embedded in the cell membrane, typically consisting of a central pore which opens and closes to allow or block the flow of ions, and a “selectivity filter” which ensures that only a given ion species (such as K^+) may pass through the channel. As a *passive* transport process, ion channels allow the flow of ions along an electrochemical gradient. This is shown in Figure 2.1 for the $K_V11.1$ ion channel, allowing K^+ ions to enter and leave the cell. The resulting current is referred to as the *rapid delayed rectifier current*, and denoted by I_{Kr} . Certain ion channels, such as $K_V11.1$ are *voltage-sensitive*, meaning that their propensity to open and allow the flow of ions changes in response to the *transmembrane potential*—the potential difference between the inside and outside of the cell membrane. In contrast to ion channels, *active* transport mechanisms consume energy in the form of ATP to pump ions against their respective concentration gradients. The Na^+-K^+ -ATPase pump is such an example (Stadt et al., 2022).

The action potential of a cardiomyocyte (and other excitable cells) begins with a large influx of Na^+ ions through the cell's sodium ion channels. This rapid influx of Na^+ can be triggered by neighbouring cells via electrical coupling of their cell membranes, as described by mathematical models of cardiac tissue (Clayton et al., 2011); whereas some other cell types are *self-exciting* and produce action potentials spontaneously—for example, the *pacemaker* cells found in the heart's sinoatrial node.

In either case, the influx of Na^+ rapidly increases the cell's transmembrane potential (*depolarising* the cell membrane). In non-self-exciting cells, after depolarisation, the transmembrane potential *repolarises*, that is, returns to its resting potential. This repolarisation is caused, in part, by the opening of channels which allows K^+ to leave the cell (András et al., 2021), causing the transmembrane potential to gradually decrease.




Figure 2.2: The currents and processes included in O’Hara et al.’s mathematical model of cellular action potentials in human ventricular tissue (O’Hara et al., 2011). This model contains multiple ion-channel currents (including I_{Kr}); active transport process such as the sodium-potassium pump and subcellular calcium dynamics are also included. Reproduced from O’Hara et al. (2011) under the terms of the Creative Commons Attribution License ©2011 O’Hara et al.

A diagram showing processes which transport ions into and out of cardiomyocytes is shown in Figure 2.2. This diagram is a visual representation of the various biophysical subcomponents in O’Hara et al.’s mathematical model of cellular action potentials in human ventricular heart tissue (O’Hara et al., 2011). This model includes multiple potassium ion channel models, such I_{Kr} and, for example, I_{Ks} (the slow delayed rectifier potassium current) and I_{K1} (the inward rectifier current). Other ion species play a key role too. Calcium ions (Ca^{2+}), for example, are released from sub-spaces inside the cell and enable the cardiomyocyte to contract. The contribution of some key currents responsible to the cardiac action potential are shown in Figure 2.3.

To study the dynamics of these constituent currents, individual subcellular mechanisms are studied. This is done by manipulating model cells to express specific genes (Beattie et al., 2018). For instance, in Chapters 4 and 5, we use data collected from *Chinese Hamster Ovary* (CHO) cells which were transfected to stably express high levels of hERG1a, the gene encoding the alpha-subunit of the $K_V11.1$ ion channel. The overexpression of this gene causes the formation of many more $K_V11.1$ channels, (albeit homomeric channels without the presence of hERG1b) resulting in large I_{Kr} currents. This approach provides a way to perform experiments where large I_{Kr} currents are present, but the contribution of other currents is minimised.

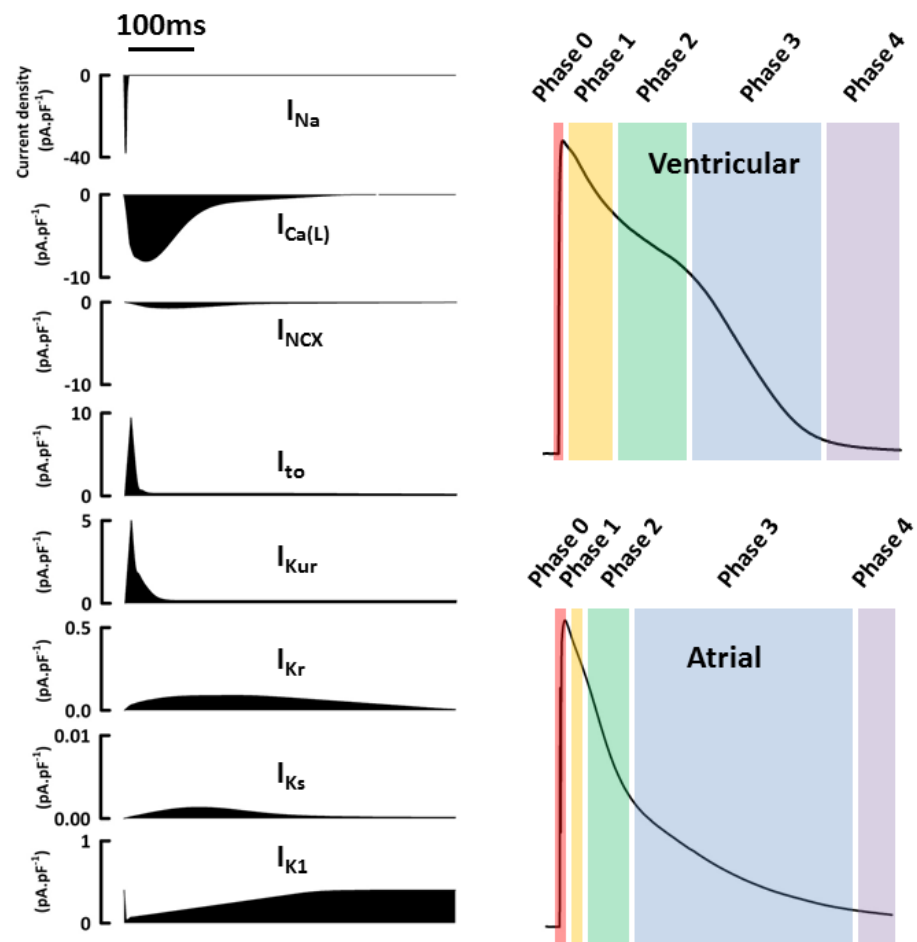


Figure 2.3: The contribution of a selection of currents to the cardiac action potential. The two axes on the right show the transmembrane potential during ventricular (top-right) and atrial (bottom-right) action potentials. On the left, representative currents for a ventricular action potential are shown. Those currents carrying charge into the cell (namely, I_{Na} , $I_{\text{Ca(L)}}$ and I_{NCX}) are shown above those currents which predominantly carry charge out of the cell. PeaBrainC, CC BY-SA 4.0, via Wikimedia Commons.

For individual proteins like $K_V11.1$, the structure of the ion channel can be investigated with imaging techniques such as Cryo-EM which has produced a detailed view of the molecular structure of the channel (Wang and MacKinnon, 2017). Also, gene modulation techniques such as CRISPR are used to mutate the channel, after which changes to its function can be observed (Fajrial et al., 2020). There has been some success using structural information (of relevant channels) to inform macroscopic models of I_{K_S} —a similar voltage-gated potassium ion-channel current (Silva et al., 2009). However, there is limited information regarding, $K_V11.1$, especially with regard to its closed and inactive conformations (Robertson and Morais-Cabral, 2020). This lack of information makes it difficult to apply a similar approach to the resultant macroscopic current, I_{K_T} . Instead, we investigate baseline function of wild-type $K_V11.1$ channels. We do this by comparing mathematical models of I_{K_T} with time-series data from electrophysiology experiments.

2.1.1 Patch-clamp experiments: state of the art

Today, one of the main techniques used to investigate the macroscopic behaviour of ion-channel currents, such as I_{K_T} , is *patch-clamp electrophysiology*. Such experiments may be performed, for example, by using pressure to rupture a small section of the cell membrane such that current can flow from the inside of the cell to the amplifier (see Figure 2.4). The data discussed in this thesis (in Chapters 4 and 5) were collected via a planar patch-clamp setup. Such a configuration is common in automated high-throughput apparatus (Weerakoon et al., 2009). A wide range of cells can be used for such experiments: human cardiomyocytes; stem cells; and commercially available cell lines such as the aforementioned CHO cells.




Figure 2.4: A patch-clamp experiment performed in “whole-cell” mode in a planar patch-clamp setup. Modified from Lei et al. (2020b).

The work in this thesis concerns a specific type of patch-clamp experiment, voltage-clamp experiments, during which an amplifier manipulates the transmembrane potential. In this configuration, the transmembrane potential is “clamped” to some predetermined

voltage trace (the *command voltage*) whilst the current through the cell membrane is recorded.

In single-channel experiments, a minuscule section of the cell membrane is removed, which allows the measurement of the current flowing through an individual channel (Neher and Sakmann, 1976). An analysis of the distribution of times a channel spent in a “conducting” state allowed Colquhoun and Hawkes to suggest a minimum number of distinct channel conformations which could plausibly exist (Colquhoun and Hawkes, 1981).

However, in this thesis, we focus on whole-cell configurations, where an electrode is placed in the intracellular medium, on the inside of the cell membrane, and the cell is kept largely intact. During these experiments, the vast majority of the cell membrane is included in the circuit (as shown in Figure 2.4) The resultant current recordings are typically much larger than those obtained from single-channel experiments because the size of the current is roughly proportional to the number of channels.

Whilst these whole-cell patch-clamp experiments are difficult and time-consuming to perform (even by an experienced experimenter), high-throughput automated patch-clamp machines (such as the SyncroPatch 384 used to collect the data presented in this thesis) allow many such experiments to be performed autonomously in parallel (Di Veroli et al., 2013). The increased throughput of an automated platform is a great advantage for industrial and clinical applications because it allows a large range of compounds (such as drugs) to efficiently be screened for potential dangerous interaction with I_{K_r} (namely, the blockage of current). However, the quality of the data obtained through automated patch-clamp platforms may be of poorer quality of that obtained using manual patch-clamp setups (Lei et al., 2020a).

2.2 Mathematical models of cellular electrophysiology

2.2.1 The Hodgkin-Huxley action-potential model

The first mathematical model of an action potential was fitted to data collected from experiments performed on a squid giant axon—a particularly large nerve cell (Hodgkin and Huxley, 1952). At the time, the precise mechanisms (such as ion channels) by which electrically-excitable cells produce action potentials and transmit electrical signals were unknown. Nevertheless, Hodgkin and Huxley knew that action potentials were produced by the movement of Na^+ and K^+ (as well as other ion species) into and out of the cell. In particular, they proposed that the depolarisation phase at the start of an action potential (as shown in Figure 2.3) is caused by a rapid inflow of Na^+ ; and that a slower, sustained outflow of K^+ is largely responsible for the repolarisation of the cell membrane.

Hodgkin and Huxley reasoned that this was made possible by dynamic changes in the permeability of the cell membrane with respect to K^+ and Na^+ . They also found that these ion-specific permeabilities responded dynamically to changes in voltage. These facts motivated the development of an equivalent circuit model whereby the changing ion-specific permeability of the membrane is modelled as a pair of variable resistors (in parallel). A diagram of this equivalent circuit is shown in Figure 2.5.

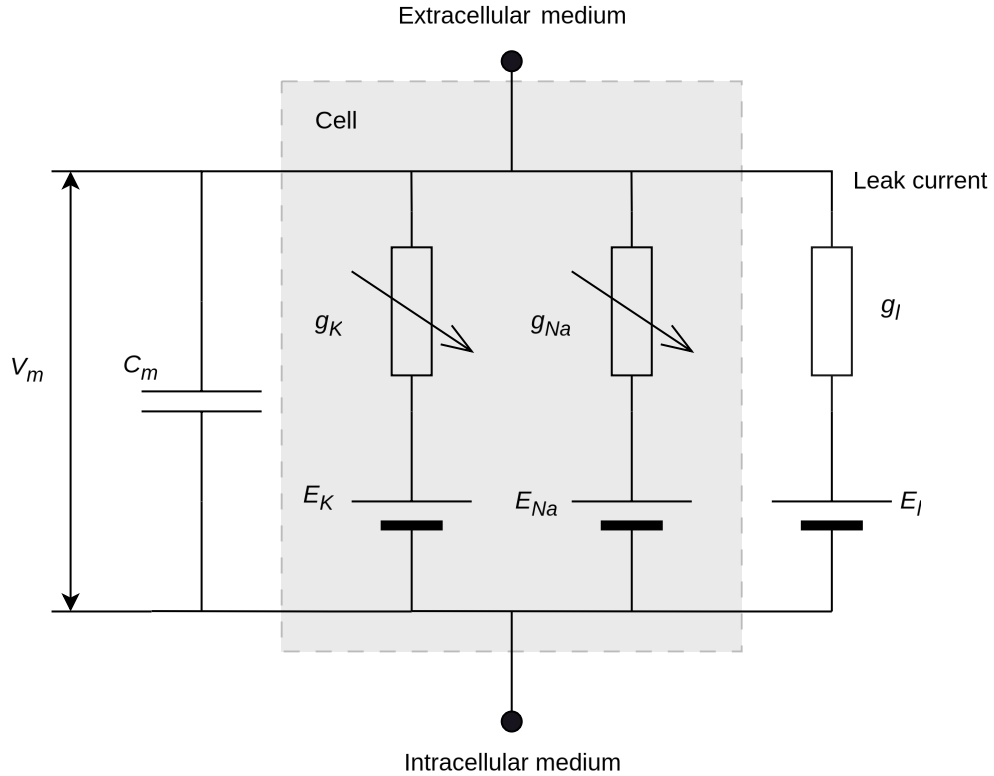


Figure 2.5: A circuit diagram representation of the Hodgkin-Huxley action-potential model. Note that in Hodgkin and Huxley’s original publication, the direction of the transmembrane potential, some currents and reversal potentials are reversed. This diagram displays the reversal potentials as opposing current flowing between the intra-cellular and extracellular mediums, a convention which is used throughout this thesis, and in more modern literature.

The Hodgkin-Huxley model can be seen as an electrical circuit formulation (shown in Figure 2.5), combined with some ordinary differential equations (ODEs) (Chicone, 2006) describing the permeability of the membrane with respect to K^+ and Na^+ (represented in the circuit by variable resistors). The current through these resistors may be

written as,

$$I_K = \bar{g}_K(V_m - E_K), \quad (2.1)$$

$$I_{Na} = \bar{g}_{Na}(V_m - E_{Na}), \quad (2.2)$$

$$\text{and } I_l = g_l(V_m - E_l), \quad (2.3)$$

where \bar{g}_K and \bar{g}_{Na} are conductances which vary dynamically as the membrane allows and restricts the flow of particular ion species and I_l is the *leak* current. Note that this leak current is an amalgamation of various biological currents, unlike the leak currents discussed in later chapters, which are assumed to be non-biological currents arising from the experimental setup. This leak current is modelled using a resistor with constant resistance. These currents have a combined effect on the *membrane potential* governed by,

$$C_m \frac{dV_m}{dt} = I_C - I_K - I_{Na} - I_l, \quad (2.4)$$

where I_C is the capacitive current, V_m is the potential difference between the inside and outside of the cell membrane (such that $V_m > 0$ indicates that the potential inside the cell is greater than the potential outside the cell). Note that in Hodgkin and Huxley (1952), the direction of these currents differ, and we instead use the conventions that are more prevalent in more modern literature (Brown, 2020).

The variable resistors in this model (with conductances \bar{g}_K and \bar{g}_{Na}) are the result of biophysical changes in the membrane which allow certain ions to pass through (carrying a current). Hodgkin and Huxley chose the following equations to describe the dynamics of the membrane's permeability changes,

$$\bar{g}_K = g_K n^4, \quad (2.5)$$

$$\text{and } \bar{g}_{Na} = g_{Na} m^3 h, \quad (2.6)$$

where $m, h, n \in [0, 1]$ are so-called “gating variables”, and g_K and g_{Na} are the maximal values that \bar{g}_K and \bar{g}_{Na} can reach. This only happens when every single channel is open (and so, g_K and g_{Na} are referred to as the *maximal conductance* parameters). These gating variables are each governed by similar ODEs,

$$\frac{dh}{dt} = \alpha_h(V_m)(1 - h) - \beta_h(V_m), \quad (2.7)$$

$$\frac{dm}{dt} = \alpha_m(V_m)(1 - m) - \beta_m(V_m), \quad (2.8)$$

$$\text{and } \frac{dn}{dt} = \alpha_n(V_m)(1 - n) - \beta_n(V_m), \quad (2.9)$$

where $\alpha_h(V_m)$ and $\beta_h(V_m)$ are voltage-dependent *transition rates*. Together, Equations (2.1–2.9) along with a specific choice of transition-rate functions form the Hodgkin-Huxley model.

Since the Hodgkin-Huxley model was first published, it has been adapted and applied to many different types of electrically excitable cells and tissues. Similarly, models of the subcomponent currents, I_K and I_{Na} have been adapted to model various subcellular processes not explicitly present in Hodgkin and Huxley’s original model. As a result, various types of models are referred to as “Hodgkin-Huxley models” (Fink and Noble, 2009). Of particular importance to this thesis are the gating variables in the Hodgkin-Huxley model, Equations (2.7–2.9), which we refer to as *Hodgkin-Huxley style gating variables*.

That is to say, equations similar to Equations (2.2), (2.8) and (2.9) may be used to model I_{Kr} (Fink and Noble, 2009; Clerx et al., 2019a). Whilst the I_{Na} component of Hodgkin and Huxley’s original model describes the dynamics of two independent processes responsible for the conductance of I_{Na} , a generalisation of this general model structure has been applied to I_{Kr} (and other macroscopic ion-channel currents) as described in Section 2.2.3. The role that these currents play in the cardiac action potential was shown through the adaptation of the Hodgkin-Huxley action-potential model to cardiac cells and subsequent iterations of these models.

2.2.2 Cardiac action-potential models

Following the discovery of ion channels, Hodgkin and Huxley’s work was built upon through the development of mathematical action-potential models for many other types of cell. These models often describe the different currents carried by various types of ion channel and other transport mechanisms. In this way, the original model (Hodgkin and Huxley, 1952) has been augmented such that the conductance terms relate to specific biophysical processes, which today are often identified with known protein structures in the cell membrane.

Compared to the Hodgkin-Huxley Model, contemporary action potential models are often more complex, and can be seen as the combination of a number of sub-models, each of which describes a separate current (carried by a specific type of ion channel or some other biophysical other process). Early work on cardiac action-potential models followed Hodgkin and Huxley’s work, including I_K and I_{Na} terms which act as amalgamations of potassium and sodium currents (Noble, 1962; Beeler and Reuter, 1977; McAllister et al., 1975). Whereas, the Hund et al. (2001) model introduced a distinct sub-model for I_{Kr} . Later, ten Tusscher et al. went further, including an I_{Kr} component which was calibrated using data from human embryonic kidney (HEK) cells heterologously expressing hERG (ten Tusscher et al., 2004). Subsequent models

following this approach include the O'Hara et al. (2011) model, the Grandi (2018) model, and more recently in the Tomek et al. (2019) model.

2.2.3 Isolated ion-channel current models

Models of I_{K_r} are also employed in isolation from a whole-cell action potential model (Rudy and Silva, 2006; Beattie et al., 2018; Lei et al., 2019a). For instance, some drug assays aim to characterise the behaviour of I_{K_r} in the presence of pharmaceutical compounds (Li et al., 2020). Typically, these assays are performed on cells which heterologously express hERG and where the presence of currents other than I_{K_r} is minimised. These experiments provide a way of learning about the dynamics of I_{K_r} in isolation. When modelling an isolated I_{K_r} current, an equivalent circuit formulation can be used in which I_{K_r} is the only current passing through the cell membrane. This is similar to the Hodgkin-Huxley formulation, but with the sodium current absent (that is, with $g_{Na} = 0$), and with some modification to the equation governing g_K (Equation (2.1)). Similar models have also been used to describe many other ion channel currents, such as I_{Na} (Mangold et al., 2017, 2021), shaker-related potassium channels (Schoppa and Sigworth, 1998), and $I_{to, f}$ (Mangold et al., 2021; Schoening and Silva, 2024). Mathematical models have also been used to quantify the impact of I_{K_r} block (as well as the block of other currents) on the cardiac action potential (Li et al., 2017b).

There are also models of the interaction between I_{K_r} and drugs. For example, Li et al.'s model of I_{K_r} and its binding to compounds. However, literature models of I_{K_r} are diverse, disagreeing on the number of conformational states included in the model (Beattie et al., 2018). It is, therefore, uncertain how well a given model captures the dynamics of I_{K_r} , even in the absence of any blocking compounds. To alleviate this uncertainty, we seek a model structure (that is, a choice of states and transition rates) which most accurately describes the mass-action dynamics of channel gating.

Though the specific equations used to model I_{K_r} vary between publications, all the I_{K_r} models listed above can be formulated as *Markov models*—a type of compartmental model which describes how channels transition between some number of conformational states. Figure 2.6 shows a how conformational states which correspond to the Beattie et al. (2018) model, and how these states are described by a given Markov model.

For I_{K_r} , the processes corresponding to the *activation/deactivation* process (the upper gates in Figure 2.6) is typically assumed to occur much slower than the process corresponding *inactivation/recovery-from-inactivation* (the lower gates in Figure 2.6). However, there is little information regarding the number of states required to model each gating process, and the independence of the rates governing each process. Typically, there are one or more *open states* such that the conductance of the channel is proportional




Figure 2.6: How the Markov model formulation of the Beattie et al. (2018) model can be seen as a simple description of two independent gating processes (shown in blue and orange). Under the Markov model paradigm, we assume that each gate (each coloured blue or orange) instantly snaps open and closed, which corresponds to a transition between states on the graph. Only when both gates are open does current flow through the channel. The inner arrows show the model's voltage-dependent transition rates, and the outer arrows show how a channel's preference for a given state changes as voltage increases.

to the fraction of channels in these states, though the models considered in this thesis have only a single open state. These models of macroscopic ion-channel currents describe the total current flowing through a large number of channels. Further discussion of the relationship between models of individual channels and these macroscopic current models is provided in Appendix A.

Similar models, using the same theoretical conformations and transitions between conformations, are concordant with recordings of individual channels. These models are typically continuous-time Markov chains, a type of stochastic process that describes how each channel instantly snaps between conformations according to some voltage-dependent transition rates. We focus on the mass-action approximation of a large number of channels, that is, deterministic ODE-based models which may be seen as the limiting behaviour of n independently and identically distributed (IID) copies of an individual channel model as $n \rightarrow \infty$. This is demonstrated in Figure 2.7, and discussed further in Appendix A.

We start by considering models sharing the same form as the I_{Na} component of the Hodgkin-Huxley model (Equations (2.2), (2.8) and (2.9)). These models consist of two gating variables, m and h , which independently and dynamically respond to changes in the transmembrane potential, V_m . Models of this form have been applied to I_{Kr} , where the two gating variables are said to represent two independent processes which control the gating of the channel: the activation and inactivation processes (Clerx et al., 2019a). Activation describes the process by which channels slowly begin to




Figure 2.7: A comparison of stochastic and deterministic models of I_{K_r} , using model structure and kinetic-rate parameters from the Beattie et al. (2018) model. Panel **a**: simulations of a CTMC ion channel model, each independently sampled using the Gillespie (2007) method, with various value of n_{channels} . Panel **b**: the RMSE between the deterministic (mass-action) approximation and our random simulations (averaged across simulations and plus/minus the standard deviation). As n_{channels} increases, the deterministic, ODE-based solution becomes a more accurate approximation of the underlying stochastic process. These simulations were made by setting the transmembrane potential to 0mV, and assuming that an equal number of channels are in each of the model's four states when $t = 0$.

open when the transmembrane voltage is increased from the resting potential. Whereas, the inactivation process occurs much more rapidly, and is most pronounced at higher voltages. Only channels which are activated, but not inactivated, allow the passage of K^+ into or out of the cell (depending on the sign of the driving force term, $V_m - E_{Kr}$).

In the Beattie et al. (2018) model for instance, our activation gating variable, a may therefore subject to voltage-dependent transition rates of the form,

$$\alpha_a(V_m) = p_1 \exp\{p_2 V_m\}, \quad (2.10)$$

$$\beta_a(V_m) = p_3 \exp\{-p_4 V_m\}, \quad (2.11)$$

and is governed by,

$$\frac{da}{dt} = (1 - a)\alpha_a(V_m) - a\beta_a(V_m), \quad (2.12)$$

where p_1, p_2, p_3 and $p_4 > 0$ are model parameters. The recovery-from-inactivation gating variable takes a similar form,

$$\frac{dr}{dt} = (1 - r)\alpha_r(V_m) - r\beta_r(V_m), \quad (2.13)$$

where,

$$\beta_r(V_m) = p_5 \exp\{p_6 V_m\}, \quad (2.14)$$

$$\alpha_r(V_m) = p_7 \exp\{-p_8 V_m\}, \quad (2.15)$$

and p_5, p_6, p_7 and $p_8 > 0$ are model parameters. Note in the Beattie et al. (2018) model, these transition rates are instead labelled such that $k_1 = \alpha_a$, $k_2 = \beta_a$, $k_3 = \beta_r$, and $k_4 = \alpha_r$ (Clerx et al., 2019a).

In the Beattie et al. (2018) model, it is assumed each channel consists of a single copy of two types of gate, each governed by Equation (2.12) and Equation (2.13), respectively. As such, the current may be written as,

$$I_{Kr} = \bar{g}ar(V_m - E_{Kr}). \quad (2.16)$$

However, more generally, we may consider models where the net current through the cell's channels is,

$$I_{Kr} = \bar{g}a^\mu r^\nu (V_m - E_{Kr}), \quad (2.17)$$

for any positive integers, μ and ν . Such models correspond to the assumption that each channel consists of μ copies of one type of gate (corresponding to our a gating variable) and ν copies of a second type of gate (corresponding to our r gating variable). These models, such as the Beattie et al. (2018) Model, may also be represented as Markov models, where the states of the model represent different numbers of gates being open.

Though they were first introduced to model the stochastic behaviour of single-channel currents, Markov models are also used to model whole-cell currents such as I_{Kr} . To do this, Markov models (such as those shown in Figure 2.6) are typically formulated as a system of ODEs of the form,

$$\frac{d\mathbf{x}}{dt} = \mathbf{Q}(V_m)^\top \mathbf{x}, \quad (2.18)$$

where $\mathbf{x} \in \mathbb{R}^N$ is our *state-variables* vector (each element of which describes the proportion of a cell's channels which are in one of $N \in \mathbb{N}$ conformations) and $\mathbf{Q}(V_m)$ is a matrix of voltage-dependent transition rates describing the rate at which channels move between conformations, such that $Q_{i,j}$ is the transition rate between state i and state j . The connection between these deterministic Markov models and the related stochastic models used to model single-channel currents is discussed further in Appendix A.3.

The state-variable vector, \mathbf{x} , is typically mapped to our observables via an equivalent circuit formulation,

$$I_{Kr} = \bar{g}x_O(V_m - E_{Kr}), \quad (2.19)$$

where \bar{g} is the maximal conductance, and x_O is an element of the state vector, \mathbf{x} , describing the proportion of channels in the *open* conformation.

Consider a single Hodgkin-Huxley style gating variable, x , governed by

$$\frac{dx}{dt} = \alpha(V_m)(1 - x) - \beta(V_m)x. \quad (2.20)$$

Though such gating variables describe the mass-action behaviour of a large number of channels, the behaviour of individual channels provides useful intuition. In this way, we may consider a Hodgkin-Huxley style gating variable representing some type of gate of which there are μ independent copies in each ion channel. The probability that a given channel is open (that is, each of its μ independent gates are open), is x^μ . Then, we can construct a Markov model with a state for every possible number of open gates (for example, the Beattie et al. (2018) model as shown in Figure 2.6).

All Markov models considered in this thesis may be represented in the form of Equation (2.18), where the transition-rate matrix, \mathbf{Q}_m has the following properties:

1. the transition rates between states are dependent only on voltage;
2. there is a transition rate from state i to state j if and only if there is a transition rate from state j to state i , that is, $Q(V_m)_{i,j} > 0$ only if $Q(V_m)_{j,i} > 0$;
3. transition rates are always positive, that is, $Q(V_m)_{i,j} > 0$ for some V_m if and only if $Q(V_m)_{i,j} > 0$ for all V_m ;

4. the total occupancy of the system, $\mathbf{1}^\top \mathbf{x}(t)$ remains constant, that is, there are no “source” or “sink” terms in the model.

These properties are sufficient to ensure that a Markov model has a unique, stable steady state whenever V_m is held constant, as discussed further in Appendix A. Note that Markov models for drug-channel interactions may relax Property 2 and allow drug-dependent transition rates to be 0 in the absence of any drug, but such models lie outside of the scope of this thesis.

In either case (whether Hodgkin-Huxley style gating variables or a Markov model is used), our formula for I_{Kr} uses a similar equivalent circuit formulation. Moreover, in Appendix A.1, we follow Keener and Sneyd (1998) and show that any Hodgkin-Huxley style gating variable raised to the power μ may be expressed as an equivalent $(\mu + 1)$ -state Markov model. This equivalence is explained further in Appendix A, and an example is shown in Figure 2.8.

Models consisting of multiple independent Hodgkin-Huxley style gating variables may be similarly represented by considering a product of component Markov models, each representing an independent gating variables (governed by Equation 2.20), as shown in Figure 2.9. The resultant model structure may be seen as a Cartesian product of graphs (Sheridan et al., 2023).

The resulting transition rates (shown in Figure 2.9, for example) are simple to derive from the individual Markov model components. For example, consider the product of two Markov models. Let u_1 and u_2 be states in the first model where the rate for transition from u_1 to u_2 is $k > 0$, and v_1 and v_2 be states in the second model where the transition rate between states v_1 and v_2 is $l > 0$. Then the transition rates between (u_1, v_1) and a second combination of states, (u_2, v_2) are,

$$\begin{cases} k & \text{if } v_1 = v_2, \text{ and } u_1 \neq u_2, \\ l & \text{if } u_1 = u_2, \text{ and } v_1 \neq v_2, \\ -k - l & \text{if } v_1 = v_2, \text{ and } u_1 = u_2, \\ 0 & \text{otherwise,} \end{cases} \quad (2.21)$$

from which the above properties of Markov models naturally follow. Again, such an example is provided by the Beattie et al. (2018) model which may be represented as the product of two Hodgkin-Huxley gating variables, or a single Markov model (Rudy and Silva, 2006; Clerx et al., 2019a).

These facts are useful, because they show that all I_{Kr} models built using the Hodgkin-Huxley gating variables can be expressed as Markov models and characterised by a particular choice of transition-rate matrix, \mathbf{Q} (and a corresponding set of states). This

simplifies the I_{Kr} model selection problem, allowing us to only consider different Markov model structures (that is, different choices of \mathbf{Q}).




Figure 2.8: A Hodgkin-Huxley model represented as a Markov model. Model **a** represents a single Hodgkin-Huxley style gating variable. Model structure **b** is formed by the combination of N independent copies of model structure **a**.

In this way, any Hodgkin-Huxley style model (with two or more gating variables) can be rewritten as a Markov model. Similarly, certain Markov models may be rewritten as the product of two independent Markov models. This is possible for the Beattie et al. (2018) and Kemp et al. (2021) models, which may be seen as representing two independent processes: an activation/deactivation process and an inactivation/recovery-from-inactivation process. This procedure is not possible, however, for models such as the Wang et al. (1997) Model, where these two processes are not independent.

Note that the Markov model ODE system characterised by Equation (2.18), with this particular choice of \mathbf{Q} , has greater dimension (that is, more state variables) than the ODE system governing the equivalent Hodgkin-Huxley model. And, as such, there are some state vectors, $\mathbf{x} = (x_i)_{i=1}^{\mu+1}$, which do not correspond to a state of the Hodgkin-Huxley style model. For $\mu = 2$, there is no $0 \leq y \leq 1$ such that $\mathbf{x} = (0, 1, 0)^\top$, for example. In practice, this is of little consequence because the initial conditions are typically chosen to be at the unique equilibrium point, which is the same for both systems of ODEs, as explained in Appendix A.

Whilst the Markov model of our simple Hodgkin-Huxley style gating variables may be more cumbersome (typically having more state variables) they allow for greater flexibility. This is because where each Hodgkin-Huxley style gating variable, say y , has only 2 transition rates, our equivalent Markov model with open state occupancy y^μ (as constructed in Appendix A.1), has 2μ transition rates. We are free to parameterise 2μ transition rates independently, resulting in a Markov model that cannot be represented in the Hodgkin-Huxley formulation. An example of this equivalence between Hodgkin-Huxley models and Markov models is demonstrated by Beattie et al. (2018) model, which can be expressed as either a four-state Markov model, or a Hodgkin-Huxley style model with two gating variables (Beattie et al., 2018). However, the Kemp et al. (2021) model, as shown in Figure 2.9, cannot be expressed in the Hodgkin-Huxley formulation;

whilst the product of Hodgkin-Huxley style gating variables results in a similar model structure, the two activation rates (a_1 and a_2) and the corresponding deactivation rates, (b_1 and b_2) are independently parameterised in the Kemp et al. (2021) model. This is not possible for a Markov model arising from a Hodgkin-Huxley style gating variable, where we would have $a_2 = 2a_1$ and $b_1 = 2b_2$ instead.

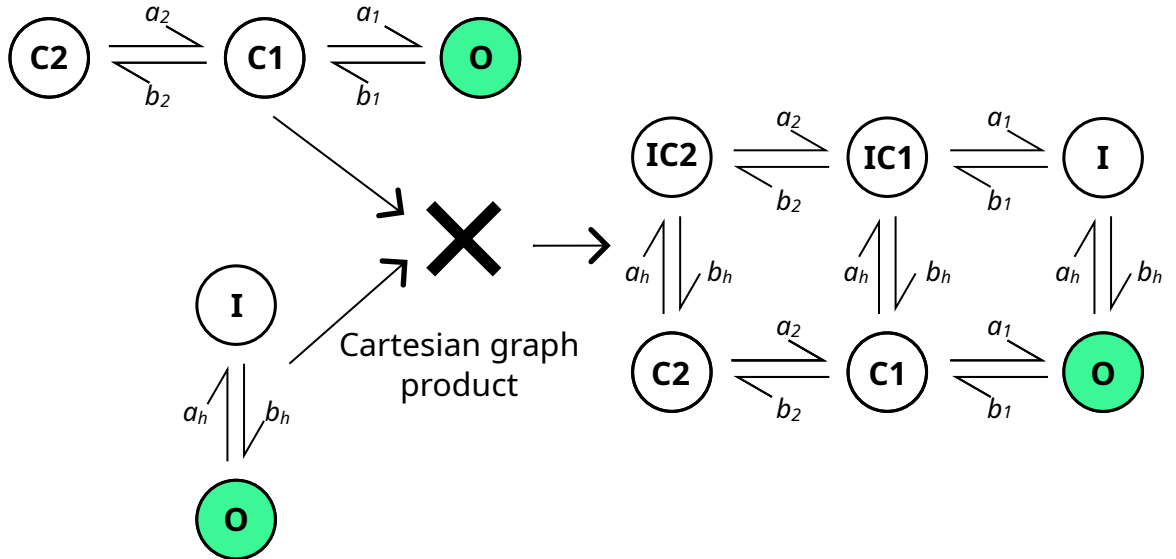


Figure 2.9: The Kemp et al. (2021) model structure represented as the product of two separate networks. Simple chain-like Markov models can be combined into a larger model. Conversely, particular Markov models like the Beattie et al. (2018) model and the Kemp et al. (2021) model can be decomposed into independent activation and inactivation processes.

The Wang et al. (1997) model is another Markov model which cannot be expressed in terms of Hodgkin-Huxley style gating variables. A notable feature of this model is the absence of a path between the *inactivated* state, **I** and the *closed states* **C1** and **C2** which does not pass through the *open state* **O**—this kind of model structure would not be possible with a Hodgkin-Huxley style model. Where the Beattie et al. (2018) and Kemp et al. (2021) models both treat activation and inactivation as independent processes, the Wang et al. (1997) differs and enforces a rather strong form of dependence: an inactivated channel cannot close without first recovering from inactivation, and, conversely, that a closed channel cannot be inactivated without first opening. Both these hypotheses have consequences when modelling the interaction of compounds with the channel, because certain drugs to I_{Kr} have been shown to have greater affinities for open or inactive states (Lee et al., 2016). Therefore, the selection of the most accurate model structure for I_{Kr} may have important ramifications for the modelling of drug-channel interactions.

In Chapters 3, 4 and 5, we fit and validate numerous Markov models of I_{Kr} to patch-clamp data (both synthetically generated data and data from real experiments). All the Markov models considered in these chapters contain transition rates of the

form, $k = A \exp\{bV_m\}$, where $A > 0$ and $b \in \mathbb{R}$ is assumed *a-priori* to be either positive or negative. However, other parameterisations have been suggested (Teed and Silva, 2016). It is possible that such parameterisations improve the predictive accuracy of our models. However, we restrict our attention to a small collection of models (each using the $k = A \exp\{bV_m\}$ parameterisation). This parameterisation of transition rates forms another aspect of structural uncertainty for I_{Kr} models (and macroscopic ion channels more generally). Nevertheless, the methods introduced in this thesis could be easily applied to models utilising alternative parameterisations.

2.3 Fitting models of I_{Kr} to patch-clamp data

In this section, we provide a general introduction to this mathematical and statistical methods used to fit the Markov models introduced in the previous section. No matter what model structure we consider, we may use the same statistical methodology to fit our model to patch-clamp data, that is, compute an estimate, $\hat{\theta}$, of our parameter vector, θ .

Supposing there are n_d observations (possibly depending on the choice of experiment), we consider models of the form,

$$\mathbf{z} = \mathbf{F}(\theta; d) + \boldsymbol{\varepsilon}, \quad (2.22)$$

where: \mathbf{z} is an n_d -vector of observations; F is a function mapping our model parameters, θ and experimental design, d , to n_d -vectors of observables (here \mathbf{F} is called the *mean function*); and $\boldsymbol{\varepsilon}$ is an n_d -vector of observational errors, each with mean, 0. Our mean-function, \mathbf{F} , is a rather general construct which could represent any nonlinear map between our input parameters (the model parameters and choice of experimental design) and the model output. Throughout this thesis, each instance of \mathbf{F} will be computed by integrating a system of ODEs, although the same statistical methodology may be applied to other models (such as those involving partial differential equations). Additionally, we assume that the observational errors (that is, $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n$) are IID Gaussian random variables. Such an assumption is commonplace in literature (Beattie et al., 2018; Clerx et al., 2019a; Lei et al., 2019a), mathematically convenient, and appears reasonable given the experimental data explored later in this thesis. However, an improved noise model (that is, the relaxation of this assumption) is a possible route for future model improvement. The impact of this assumption, and some alternative approaches, are discussed further in Section 2.4.

2.3.1 Maximum likelihood estimation

Given a mean function, \mathbf{F} , some parameter set $\boldsymbol{\theta}$, an experimental design, d and some data $\mathbf{z}(d)$, our observational errors are,

$$\boldsymbol{\varepsilon} = \mathbf{F}(\boldsymbol{\theta}; d) - \mathbf{z}(d), \quad (2.23)$$

using which we may compute the *likelihood* of a parameter set (Bates and Watts, 1988).

Under the assumption of additive Gaussian IID observation errors, $\boldsymbol{\varepsilon}$, our likelihood is,

$$L(\boldsymbol{\theta}) = \prod_{i=1}^{n_d} p(\varepsilon_i) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{\varepsilon_i^2}{2\sigma^2}\right\} \quad (2.24)$$

$$= (2\pi\sigma^2)^{-\frac{n}{2}} \prod_{i=1}^{n_d} \exp\left\{-\frac{\varepsilon_i^2}{2\sigma^2}\right\}. \quad (2.25)$$

where $p(\varepsilon_i)$ is the *probability density function* of the i^{th} observational error, $\boldsymbol{\varepsilon}$, and n_d is the number of observations made during protocol d . Such a noise model (additive, IID Gaussian noise) is often assumed in the literature (Menon et al., 2009; Beattie et al., 2018; Clerx et al., 2019a), and the resulting parameter estimates are, in fact, independent of σ .

This independence may be demonstrated by considering the log-likelihood,

$$l(\boldsymbol{\theta}) = \log(L(\boldsymbol{\theta})), \quad (2.26)$$

$$= -\frac{n}{2} \log(2\pi\sigma^2) - \sum_{i=1}^n \frac{\varepsilon_i^2}{2\sigma^2}, \quad (2.27)$$

which, for any σ , is maximised when the sum of squared residuals, $\sum_{i=1}^n \varepsilon_i^2$, is minimised.

Because $x \mapsto \log\{x\}$ is an increasing function, we can choose to maximise either the likelihood or the log-likelihood to find the *maximum likelihood estimate* (MLE),

$$\hat{\boldsymbol{\theta}} := \operatorname{argmax}_{\boldsymbol{\theta} \in \Theta} \{L(\boldsymbol{\theta})\} \quad (2.28)$$

$$= \operatorname{argmax}_{\boldsymbol{\theta} \in \Theta} \{l(\boldsymbol{\theta})\}. \quad (2.29)$$

Equivalently, we may minimise the root-mean-square error (RMSE) between the mean function, and some data ($\mathbf{z}(d)$),

$$\text{RMSE}(\mathbf{F}(\boldsymbol{\theta}; d), \mathbf{z}(d)) = \sqrt{\frac{1}{n_d} \sum_{i=1}^{n_d} (\mathbf{F}_i(\boldsymbol{\theta}; d) - z_i(d))^2}, \quad (2.30)$$

where n_d is the number of observations in protocol d , and F_i is the mean function for the i^{th} variable (Willmott et al., 1985). Minimising the RMSE is equivalent to minimising the sum-of-square errors (SSE) because,

$$\text{RMSE}(\mathbf{F}(\boldsymbol{\theta}; d), \mathbf{z}(d)) = \frac{1}{n} \sqrt{\text{SSE}(\mathbf{F}(\boldsymbol{\theta}; d), \mathbf{z}(d))}, \quad (2.31)$$

and, under the assumption of IID Gaussian errors, both are equivalent to minimising the log-likelihood (and, therefore, the likelihood itself). This is because, in this case, the log-likelihood consists of a constant term, and a term that is proportional to the SSE (as shown in Equation 2.27).

Throughout this thesis, we quantify the error between our data and models using the RMSE, as this permits easier comparison between protocols with different numbers of observations. This is not the case with the SSE, for example, which almost surely increases as we include more observations. Another advantage is the fact that the RMSE has the same units as the data, which allows for easier interpretation.

Using the RMSE, we define the parameter estimate obtained from protocol d as,

$$\hat{\boldsymbol{\theta}}_d = \text{argmin}_{\boldsymbol{\theta} \in \Theta} \{ \text{RMSE}(\mathbf{F}(\boldsymbol{\theta}; d), \mathbf{z}(d)) \}, \quad (2.32)$$

which is a random variable because it depends on our random data, $\mathbf{z}(d)$ which, in turn, depends on the choice of experimental design, d . Where it is necessary to draw attention to the dependence of the parameter estimate on the experimental design used, we denote the parameter estimate obtained from a particular experimental design, d , as,

$$\begin{aligned} \hat{\boldsymbol{\theta}}_d &= \text{argmin}_{\boldsymbol{\theta}} \{ l(\boldsymbol{\theta}, \mathbf{z}) \} \\ &= \text{argmin}_{\boldsymbol{\theta}} \{ \text{RMSE}(F(\boldsymbol{\theta}; d), \mathbf{z}_d) \}. \end{aligned} \quad (2.33)$$

This is particularly useful when parameter estimates obtained from different experimental designs are discussed.

The log-likelihood, Equation (2.27), can be minimised by first finding the parameter set, $\hat{\boldsymbol{\theta}}$, which minimises the RMSE, and then finding the optimal σ . For the patch-clamp experiments described in Chapters 3, 4 and 5, we estimate σ directly from a stationary section at the start of the trace (where we assume that $I_{Kr} = 0$).

Provided certain regularity conditions hold (such as the existence of this minimum), $\hat{\boldsymbol{\theta}}$ is an *consistent* estimator of our model parameters, meaning (loosely speaking) that our parameter estimate, $\hat{\boldsymbol{\theta}}$, is a good approximation of our model parameters, provided that there is enough data. To be more precise, as we increase the number of observations, n_d , our estimate $\hat{\boldsymbol{\theta}}$ converges *in probability* to the true parameter set, $\boldsymbol{\theta}$ (Seber and Wild, 2005).

2.3.2 Numerical optimisation

Under the assumption of additive, Gaussian IID errors, fitting ODE-based models using maximum likelihood estimation is an example of nonlinear least-squares regression. Bates and Watts (1988) introduce a range of methods for nonlinear regression problems, and pay particular attention to gradient-based numerical optimisation techniques to iteratively find more suitable parameter values (those with greater likelihood than the current proposed value). They suggest the use of the Newton iteration method. However, there are numerous challenges when using numerical methods for optimisation (Bates and Watts, 1988), and alternative methods may prove to be more suitable.

A challenge arises when the *objective function* that we seek to minimise (in this case the negative log-likelihood, $-l(\boldsymbol{\theta})$), is *non-convex*. Local, gradient-based methods such as Newton iteration rely on the initial guess being in the “basin of attraction”—that is, the set of initial guesses for which the method will converge to the global optimum. In the case of a convex objective function, this basin of attraction contains the whole parameter space. For some problems however, basins of attraction may have a non-trivial geometry (Zotos, 2017), and so the parameter estimate obtained from numerical estimation may vary depending on the chosen initial guess—that is, there is no guarantee that such a method will successfully compute the maximum-likelihood estimate. Furthermore, these gradient-based methods rely on the smoothness of the objective function, which may be jeopardised when approximate, numerical methods are used to compute $l(\boldsymbol{\theta})$.

To account for the (likely) non-convexity of our objective functions, we can use a multi-start procedure in which we perform the optimisation from a range of different starting points (Bates and Watts, 1988). The rationale for this approach being that if repeated runs of the optimisation from different starting guesses all yield similar answers, we can be somewhat confident that we have accurately determined the global minimum (in this case, the maximum-likelihood estimate).

Some methods for numerical optimisation incorporate a degree of stochasticity in the iterative procedure itself. For example, the Covariance Matrix Adaptation (CMA) method is often implemented using the Evolutionary Strategy (CMA-ES) in which candidate parameter values are selected according to a constantly updated multivariate Gaussian distribution (Hansen, 2016). The method works by randomly sampling points from this distribution, computing the corresponding likelihoods and updating the mean and covariance of this sampling distribution such that points near those previously found to have high likelihood are favoured. As the method converges, this sampling distribution shrinks such that the sampling distribution is concentrated near to the global optimum. For non-convex functions, a multi-start procedure may prove useful, allowing a somewhat more complete exploration of the parameter space. CMA-ES has been used

to fit models of I_{KR} to information-rich patch-clamp data (Fink and Noble, 2009; Lei et al., 2019a; Clerx et al., 2019a; Whittaker et al., 2020a).

Some methods are more suitable for different problems where modellers may have different priorities regarding the computational expensive of the methods, and their ability to reliably find global optima. In these cases, a careful consideration of the results from multiple optimisation runs is necessary for us to be confident in our parameter estimates. However, for other problems, priorities may vary. For example, local, deterministic, methods such the Newton method and Nelder-Mead algorithm may be more suitable if the speed of computation is a priority and confidence that we have truly identified the global optimum is less important. On the other hand, there are many optimisation methods may be more suitable for “harder” optimisation problems in which there is a substantially larger parameter set, or more complicated optimisation surface (Menon et al., 2009). Nevertheless, throughout this thesis, we use CMA-ES because it has been shown to reliably and quickly identify optimal parameter sets for problems similar to those encountered in this thesis (Clerx et al., 2019a).

2.3.3 Integration of ODE systems

To compute our models’ mean functions, and hence the likelihood, it is necessary to solve systems of ODEs. When our model has a governing equation of the form of Equation (2.18), the transition-rate matrix, \mathbf{Q} is constant for fixed V_m , and so we have a closed-form expression in terms of a matrix exponential, (Moler and Van Loan, 2003),

$$\mathbf{x}(t)^\top = \mathbf{x}(0)^\top \exp\{t\mathbf{Q}(V_m)\}. \quad (2.34)$$

There are numerous methods for computing expressions of this form (Moler and Van Loan, 2003; Teed and Silva, 2016). A number of the methods suggested by Moler and Van Loan (2003) rely on the eigendecomposition of the exponent matrix. For these methods, it is useful to note that \mathbf{Q}^\top (or equivalently, \mathbf{Q}) is guaranteed to have only real eigenvalues for reversible Markov models, as discussed in Appendix A. However, for non-reversible Markov models, the eigenvalues of \mathbf{Q}^\top may be non-real, meaning some care should be taken to account for non-real eigenvalues in the implementation of these methods.

A common trick is to use the property that the number of channels is preserved (as per our definition of Markov models provided in Section 2.2.3), that is, $\mathbf{Q}\mathbf{1}^\top = \mathbf{0}$ to rewrite our governing equation in the form,

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{A}(V_m)\tilde{\mathbf{x}} + \mathbf{b}(V_m). \quad (2.35)$$

There are many such simplified ODE systems, as discussed in Appendix A. One common approach (Colquhoun and Hawkes, 1995) is to take $\tilde{\mathbf{x}}$ to be the first $N - 1$ components of \mathbf{x} .

It is then possible to apply the methods discussed by Moler and Van Loan (2003) to this reduced system, instead. Additionally, the steady-state of the model for when V_m is held constant at V_0 can be computed by simply solving the linear equation,

$$\mathbf{A}(V_0)\mathbf{x}_\infty(V_0) = \mathbf{b}_0. \quad (2.36)$$

The fact that a unique solution to this equation exists follows from our definition of a Markov model as discussed in Appendix A. This is important when applying ODE models to patch-clamp experiments because it is typically assumed that, $\mathbf{x}(0) = \mathbf{x}_\infty(V_0)$, that is, that the system is at equilibrium at the beginning of the experiment, which provides us with the initial conditions necessary to integrate the governing equation.

When V_m is not constant (as it often is during experiments, or in whole-cell simulations) we can no longer use this matrix-exponential formulation (Equation (2.34)). In general, there is no matrix-exponential solution to this system. Instead, we use numerical methods to integrate Equation (2.35). Employing these methods requires some care, however; depending on the system being solved, some methods can be *unstable* and fail to provide a good approximation. Some systems which prove challenging to integrate may be referred to as *stiff*. Sometimes, an ODE system may be called stiff if there are both short and long timescales present in the model, that is, if the ratio of maximum and minimum eigenvalues of $\mathbf{Q}^\top(V_m)$ is large (Fink and Noble, 2009).

Some methods for the numerical solution of ODEs make use of the forward-difference formula, whereby the system can be iterated in discrete steps through time. For example, the forward-Euler method is an iterative scheme defined by,

$$\mathbf{x}^{(i+1)} = \mathbf{f}(\mathbf{x}^{(i)}, t_i) \cdot (t_{i+1} - t_i) + \mathbf{x}^{(i)}, \quad (2.37)$$

where $\mathbf{f}(\mathbf{y}, t)$ is the derivative of the system at time t , and where the state of the system is \mathbf{y} . Note that, in practice, \mathbf{f} may also depend on our model parameters and choice of experimental design. By reducing the step-size, that is, $h := t_{i+1} - t_i$, we can improve the accuracy of this approximation. However, for certain systems, such methods may require very small step sizes to provide a suitable approximation (Gupta et al., 1985).

An alternative is to use ‘backwards’ solution methods, in which a backwards difference formula is used instead. For example, the backwards Euler method uses,

$$\mathbf{x}^{(i+1)} = \mathbf{f}(\mathbf{x}^{(i+1)}, t_i) \cdot (t_{i+1} - t_i) + \mathbf{x}^{(i)}, \quad (2.38)$$

Whilst Equation (2.38) looks similar to Equation (2.37), it is not possible to calculate $y^{(i+1)}$ directly because the right-hand side depends on $\mathbf{x}^{(i+1)}$. Methods such as this are also called “implicit” methods because, unlike forward methods, they do not provide an explicit formula for computing $\mathbf{x}(t + h)$. Instead, iterative methods are used to solve Equation (2.38). In other words, for each step in the solution, where we iterate the state of the system forward, we have to solve a (for our purposes, multivariable) nonlinear equation. The backwards differentiation formulae (BDF) and Adams-Moulton methods are two families of implicit methods for numerical integration of ODEs; both families contain the backwards Euler method, Equation (2.38), as the simplest method. BDF methods are especially suitable for stiff problems (Gupta et al., 1985).

No matter what scheme is used (whether the forward-Euler, backward-Euler, BDF or Adams-Moulton methods), the time-step size may be reduced to minimise error, requiring more steps to solve the system up to the end of the experiment. However, using time steps that are too small results in excessive computation. Adaptive-timestep methods combine an iterative scheme with some way of selecting a suitable timestep, typically by approximating the error between the true solution and the numerical solution, and ensuring that this value is less than some predetermined threshold.

The ODEPACK package is a collection of software containing various methods for the numerical solution of ODEs (Hindmarsh, 1982). One such method is LSODA (Gupta et al., 1985), an adaptive-timestep method which automatically switches between Adams-Moulton and BDF methods depending on the stiffness of the problem at a given point in time, the latter being preferred in the presence of stiffness. The performance of LSODA is compared with other methods for a ODE-based compartmental model in Postawa et al. (2020). We use this method for all numerical integration of ODE systems in this thesis because we require a fast and robust method that will work even when our systems are very stiff. There are other similar methods, such as CVODE, found in the SUNDIALS software package (Balos et al., 2024), which is an adaptive-timestep, backwards integration method developed in the C programming language. CVODE has been shown to perform favourably compared to LSODA across a range of models in systems biology (Städter et al., 2021). Though the majority of models considered by Städter et al. (2021) have many more state variables than the models considered in this thesis, the use of CVODE for future work may provide modest benefits to computational performance.

Besides the various pitfalls of numerical optimisation approaches discussed in this section, other problems can arise when estimating model parameters. As discussed in the following section, some such problems are independent of the particular optimisation methods used, and instead arise from the properties of particular modelling problems—that is, from particular combinations of models and experimental designs.

2.3.4 Identifiability

The identifiability of a model concerns our ability to determine the model parameters. There are multiple ways in which a model may be *unidentifiable*. Perhaps the simplest type of unidentifiability is *a-priori* unidentifiability where the model parameters cannot be uniquely determined—no matter what experiment we perform, or how much data we collect. This can be the case if we include extraneous parameters in our model. Consider, for example, a model with,

$$F_i(\boldsymbol{\theta}; d) = \theta_1^* \theta_2^* d. \quad (2.39)$$

Then, given any $\theta_1, \theta_2 \neq 0$, we have $F_i(\boldsymbol{\theta}; d) = F_i(\boldsymbol{\phi}; d)$ where $\phi_1 = 1$ and $\phi_2 = \theta_1 \theta_2$. This is known as *a-priori* unidentifiability, and can be avoided by changing the way that the model is parameterised. For example, we can simplify F_i such that the product $\theta_1^* \theta_2^*$ is replaced a single model parameter,

$$\tilde{F}_i(\tilde{\boldsymbol{\theta}}; d) = \tilde{\boldsymbol{\theta}} d. \quad (2.40)$$

Here, no matter what experiment is performed, we can only gain information about the product, $\tilde{\boldsymbol{\theta}} = \theta_1^* \theta_2^*$, and not the individual parameters θ_1^* and θ_2^* . Such parameters are therefore said to be *a-priori* unidentifiable.

This is also an example of *structural unidentifiability*, where the output from a given experiment may be insufficient to uniquely identify an optimal parameter set (Bearup et al., 2013). Unlike *a-priori* identifiability, the structural identifiability of a model may depend on what observations are made—for example, which states of a Markov model are observed. Siekmann et al. (2012) provide some examples of Markov models which are structurally unidentifiable from single-channel recordings (at a fixed command voltage). The structural identifiability of linear models may be assessed with well-established techniques. This is especially straightforward when a Gaussian, IID noise model is assumed. The structural identifiability of nonlinear models is more complex, however (Bates and Watts, 1988). Nevertheless, there are a number of analytical approaches. One such approach concerns the local properties of the mapping between model parameters and model outputs (Bearup et al., 2013).

Practical identifiability concerns the combination of model and experimental design. Though a model may be structural identifiable, it may be practically unidentifiable if a suitable experiment is not specified. Fink and Noble (2009) provide a rather broad definition, which encompasses the computational difficulty in finding optimal parameter estimates. This includes even the challenges involved in numerical optimisation, which were discussed in the previous section. However, other definitions are more explicitly based on the statistical properties of models under a particular experimental design. In

particular, a poorly designed experiment for our model could leave us with untenably large confidence regions for our parameter estimates. To illustrate this point, we first consider a linear statistical model (a model in which there is a linear relationship between the model parameters and dependent variables),

$$\mathbf{F}(\boldsymbol{\theta}^*; \mathbf{X}) = \mathbf{X}\boldsymbol{\theta}^*, \quad (2.41)$$

subject to additive IID Gaussian noise, $\boldsymbol{\varepsilon}$. Here we have substituted d (the variable representing our experimental design in Equation (2.22)) with a *design matrix*, \mathbf{X} , which characterises the experiment under consideration. In this well studied case, the MLE satisfies the linear system of equations,

$$\mathbf{X}^\top \mathbf{X} \boldsymbol{\theta} = \mathbf{X}^\top \mathbf{z}. \quad (2.42)$$

If $\mathbf{X}^\top \mathbf{X}$ is singular, this equation has infinitely many solutions stretching out to infinity along some line, plane or hyperplane, each of which has the same likelihood according to the model. In this situation, the MLE is undefined. Such a situation may arise from structural identifiability as discussed previously, in which case the collection of further data will not yield an identifiable model.

Even if $\mathbf{X}^\top \mathbf{X}$ is nonsingular, we may obtain inaccurate parameter estimates. This can happen, for example, when our experiment is designed in such a way that there is great uncertainty in our parameter estimates (arising from the randomness of our observations). This uncertainty can lead to inaccurate parameter estimates and predictions. In the linear case (with IID, additive Gaussian error), such practical (un)identifiability may be characterised using the *Fisher Information Matrix* (FIM).

The FIM is defined as

$$\mathcal{I}_{i,j} = \mathbb{E} \left[\left(\frac{\partial l(\boldsymbol{\theta}; \mathbf{z})}{\partial \theta_i} \right) \left(\frac{\partial l(\boldsymbol{\theta}; \mathbf{z})}{\partial \theta_j} \right) \right], \quad (2.43)$$

where $\boldsymbol{\theta}$ is our model-parameter vector and \mathbf{z} is the data. Alternatively, provided certain regularity conditions hold (Seber and Wild, 2005), we may write this as,

$$\mathcal{I} = \mathbb{E}_{\mathbf{z}} \left[\left(\frac{\partial^2 l(\boldsymbol{\theta}; \mathbf{z})}{\partial \theta_i \partial \theta_j} \right)_{i,j} \right]. \quad (2.44)$$

In the case of our linear model with Gaussian IID noise, direct computation shows,

$$\frac{\partial l}{\partial \theta_i} = \sum_{k=1}^n \frac{\partial l}{\partial F_k} \frac{\partial F_k}{\partial \theta_i}, \quad (2.45)$$

where we denote by F_k the k^{th} component of \mathbf{F} , and n is the number of observations. Then, since our log likelihood is,

$$l = \frac{n}{\sigma\sqrt{2\pi}} + \frac{1}{2\sigma^2} \sum_{i=1}^n (F(\boldsymbol{\theta})_i - z_i)^2, \quad (2.46)$$

we have,

$$2\sigma^2 \frac{\partial l}{\partial F_k} = -2(z_k - F_k), \quad (2.47)$$

for all $1 \leq k \leq n$. Then, by summing over our observations, we have,

$$2\sigma^2 \frac{\partial l}{\partial \theta_i} = \sum_{k=1}^n \frac{dl}{dF_k} \frac{\partial F_k}{\partial \theta_i} = \sum_{k=1}^n 2(F_k - z_k) \frac{\partial F_k}{\partial \theta_i}. \quad (2.48)$$

Hence, using the product rule we obtain,

$$2\sigma^2 \frac{\partial^2 l}{\partial \theta_i \partial \theta_j} = 2 \left(\sum_{k=1}^n \frac{\partial F_k}{\partial \theta_j} \frac{\partial F_k}{\partial \theta_i} + \frac{\partial^2 F_k}{\partial \theta_i \partial \theta_j} (F_k - z_k) \right). \quad (2.49)$$

For all k , we have $\mathbb{E}[F_k - z_k] = 0$. Hence,

$$\mathbb{E} \left[2\sigma^2 \frac{\partial^2 l}{\partial \theta_i \partial \theta_j} \right] = \mathbb{E} \left[2 \left(\sum_{k=1}^n \frac{\partial F_k}{\partial \theta_j} \frac{\partial F_k}{\partial \theta_i} \right) \right], \quad (2.50)$$

and so,

$$\mathcal{I}(\boldsymbol{\theta}; d) = \frac{1}{\sigma^2} \mathbf{X}^\top \mathbf{X}. \quad (2.51)$$

The confidence regions for our model parameter estimates take the form,

$$(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^\top \mathbf{X}^\top \mathbf{X} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) \leq c, \quad (2.52)$$

for some c corresponding to a choice of significance level, and so the size (measure) of this confidence region scales linearly with $\det(\mathbf{X}^\top \mathbf{X})$. Note that under a linear model where $\mathbf{X}^\top \mathbf{X}$ is nonsingular, all confidence regions are finite. However, this is not generally true for nonlinear statistical models (Kreutz et al., 2012).

Provided our log-likelihood function is differentiable (Seber and Wild, 2005), we can substitute \mathbf{X} with the sensitivities matrix in the above derivation. This sensitivities matrix is defined as,

$$\mathbf{S}(\boldsymbol{\theta}) = \left(\frac{\partial F_i}{\partial \theta_j} \right)_{i,j}. \quad (2.53)$$

In this case, the resultant information matrix may be evaluated for a specific parameter vector, $\boldsymbol{\theta}$,

$$\mathcal{I}(\boldsymbol{\theta}; d) = \frac{1}{2\sigma^2} \mathbf{S}(\boldsymbol{\theta}; d)^\top \mathbf{S}(\boldsymbol{\theta}; d), \quad (2.54)$$

where d is our experimental design (replacing the linear case's design matrix \mathbf{X}). This linear approximation of F yields a *quadratic* approximation of the likelihood,

$$2\sigma^2(l(\boldsymbol{\theta}; d) - l(\boldsymbol{\theta}^*; d)) \approx (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^\top \mathbf{S}(\boldsymbol{\theta}^*; d)^\top \mathbf{S}(\boldsymbol{\theta}^*; d)(\boldsymbol{\theta} - \boldsymbol{\theta}^*), \quad (2.55)$$

where $\boldsymbol{\theta}^*$ is the true value of the parameter set. This approximation is discussed in more detail in Bates and Watts (1988); Seber and Wild (2005) where it is used to compute approximate confidence regions.

This approximation is particularly useful when there is a lot of data, in which case asymptotic properties which justify the use of confidence regions of the form of Equation (2.57) (where we substitute \mathbf{X} for \mathbf{S} evaluated at the maximum likelihood estimate). In particular, we have (for sufficiently large n) that the distribution of $\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*$ may be well approximated by the normal distribution,

$$N(0, \sigma^2 \mathbf{S}(\boldsymbol{\theta}^*; d)^\top \mathbf{S}(\boldsymbol{\theta}^*; d)). \quad (2.56)$$

Using this approximation we may compute approximate confidence regions of the form,

$$(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)^\top \mathbf{S}(\hat{\boldsymbol{\theta}}; d)^\top \mathbf{S}(\hat{\boldsymbol{\theta}}; d)(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) \leq c, \quad (2.57)$$

for some chosen $c > 0$ (Bates and Watts, 1988). Note the similarity between these confidence regions and those described by Equation (2.52).

Of the various definitions of practical identifiability in the literature, some are directly related to the size of these confidence regions (which are approximately provided by Equation (2.57) in the nonlinear case). Some authors suggest establishing thresholds for the size of these confidence intervals (Gábor et al., 2017; Tuncer et al., 2018; Wieland et al., 2021), the exceedance of which prompting the model to be labelled “practically unidentifiable”. Whichever definition we adopt, identifiability is an important property of our models which is necessary for us to obtain accurate parameter estimates and produce accurate predictions. For nonlinear models, the computation of the local sensitivities (Equation 2.53) provides a way of evaluating identifiability, though this may fail to characterise the likelihood function outside of a given local neighbourhood. Nevertheless, we introduce a numerical approach for the computation of local sensitivities in the following section.

2.3.5 Local sensitivities for ODE-based models

The models we apply in this thesis are built upon systems of ODEs that describe the kinetics of ion-channel models. In this section, we describe an approach for the efficient computation of the sensitivities matrix, \mathbf{S} and hence, \mathcal{I} , as introduced in the preceding section. Whilst these methods are presented using our equivalent circuit formulation (Equation (2.19)), they may be adapted for use with a wide range of ODE-based models.

As discussed in Section 2.2.3, Markov models of I_{Kr} are a specific type of ODE model in which the model parameters are related to the observables, in part, through the solution of a system of ODEs. In general, we may write,

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t, \boldsymbol{\theta}; d), \quad (2.58)$$

where \mathbf{x} is a state-variables vector, $\boldsymbol{\theta}$ are our model parameters, and d is an experimental design. The solution of this system of ODEs is then related to our observables through an *observation function*. This function maps the states of the system at each observation time, t , to some model output,

$$y_i := h(t_i, \boldsymbol{\theta}; \mathbf{x}(t_i)), \quad (2.59)$$

where h is the observation function. When applied to models of I_{Kr} , the *observable*, y_i , is taken to be the observed current at a given time, t . For example, for I_{Kr} models (Beattie et al., 2018; Clerx et al., 2019a), we typically use,

$$h(t, \boldsymbol{\theta}; \mathbf{x}(t)) = I_{Kr}(t, \boldsymbol{\theta}) \quad (2.60)$$

$$= g x_O (V_m - E_{Kr}), \quad (2.61)$$

where x_O is a particular component of \mathbf{x} .

In this case, our mean function, \mathbf{F} , may be evaluated by first integrating the model's governing Equation (2.18), and then evaluating the observation function at each observation time. This, gives our mean function, \mathbf{F} , which maps each experimental design and possible parameter set to a vector of observations, that is,

$$\mathbf{F}(\boldsymbol{\theta}; d) = (y_i)_{i=1}^n. \quad (2.62)$$

We conclude this section by discussing some specific details regarding the application of the previously introduced statistical theory to ODE-based models (such as the Markov models introduced in Section (2.18), and other models used throughout this thesis). For ODE-models, using the example of the Markov models introduced in Section 3.3.2, the map from our parameters and designs to our observables (\mathbf{F} in Equation (2.62)) is the combination of a system of ODEs and an observation function.

Typically, this observation function is used to map the model's state variables to its observables. For example, for a Markov model (with fixed V_m), we have,

$$I_{Kr} = g x_O (V_m - E_{Kr}). \quad (2.63)$$

By differentiating, we find,

$$\frac{\partial I_{Kr}}{\partial \theta_i} = \frac{\partial x_O}{\partial \theta_i} g(V_m - E_{Kr}) + \frac{\partial g}{\partial \theta_i} x_O (V_m - E_{Kr}). \quad (2.64)$$

We also find,

$$\frac{\partial I_{Kr}}{\partial x_O} = g(V_m - E_{Kr}), \quad (2.65)$$

$$\text{and } \frac{\partial I_{Kr}}{\partial g} = x_O (V_m - E_{Kr}). \quad (2.66)$$

Hence, given $\frac{\partial x_i}{\partial \theta_j}$ (for all i and j) we can compute our sensitivities matrix, \mathbf{S} .

A common approach to finding $\frac{\partial x_i}{\partial \theta_j}$ is to augment the governing equation of our ODE system with additional state variables which describe how $\frac{\partial x_i}{\partial \theta_j}(t)$ evolves over time. For ease of generalisation, we write this governing equation as,

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, \boldsymbol{\theta}), \quad (2.67)$$

(dropping any dependence on the chosen experimental design).

Then we may find second-order partial derivatives,

$$\frac{\partial^2 x_j}{\partial \theta_i \partial t} = \frac{\partial f(\mathbf{x}, \boldsymbol{\theta})}{\partial \theta_i} + \sum_{j=1}^{N_s} \frac{\partial f(\mathbf{x}, \boldsymbol{\theta})}{\partial x_j} \frac{\partial x_j}{\partial \theta_i}, \quad (2.68)$$

where N_s denotes the number of state-variables in our model. To compute these sensitivities, all the terms on the right-hand side of Equation (2.68) may be calculated with pen and paper or automatically using a computer algebra system. It should be noted here, that the number of differential equations added by this method is $N_s N_p$ where N_p is the number of parameters in the model. Although the number of additional ODEs may present a problem for large, complicated models, it is feasible to use these techniques with the models used in this thesis due to the relatively small number of states and parameters in our models. This method may also provide insight into the utility of a given experimental design, as discussed in the following section.

2.3.6 Experimental design

Much work investigating I_{KR} has closely followed Hodgkin and Huxley's original experimental and statistical methodology. Often, simple voltage-clamp protocols (the command voltages applied during the experiment) are repeated with some variation in the voltage during particular steps. Then, “summary curves” are made which describe how, for example, peak currents and decay timescales change with respect to changes in the command voltage, V_{cmd} . In these protocols, a similar three-to-four step voltage-signal is applied to the cell, but the voltage during a single step is manipulated through a range of values under repeats (Fink and Noble, 2009; Clerx et al., 2019a).

More recently, the concept of a *information-rich* protocols has been introduced. These protocols are designed to permit the fitting of models using recordings taken over a shorter duration (relative to the longer experimental designs originally used in Hodgkin and Huxley (1952) and elsewhere). This idea was first proposed (in the context of macroscopic ion-channel currents) by Fink and Noble (2009). Later, more information-rich protocols have been designed and used to fit models (Beattie et al., 2018; Lei et al., 2019a, 2024; Mirams et al., 2024). Such information-rich designs are constructed to permit accurate parameter estimations.

We now provide a brief introduction to techniques used to design optimal experiments for linear models. We start by considering a linear model of the form,

$$\mathbf{F}(\boldsymbol{\theta}; d) = \mathbf{X}\boldsymbol{\theta}, \quad (2.69)$$

where $\boldsymbol{\theta}$ is a vector of parameters, and a *design matrix* \mathbf{X} , which may be subject to experimental design. There are well-established approaches for the design of such experiments to ensure that \mathbf{X} is specified such that the resulting parameter estimates have desirable properties (Bates and Watts, 1988). When this is done using *optimisation* methods, it is known as *optimal experimental design* (OED) (Banga and Balsa-Canto, 2008). In the following sections, we review some OED methods for linear models and show how these methods may be applied to nonlinear models (such as the Markov models described previously in this chapter), also.

Optimal Experimental Design for linear models For our linear model described in Equation (2.69), it is natural to seek a design matrix, \mathbf{X} , which we expect to yield accurate parameter estimates, or an accurate predictive model. The so-called “alphabet criteria” (Ryan et al., 2016) are commonly used ways of quantifying the relative expected accuracy of parameter estimates and model predictions. A selection of relevant “alphabet criteria” are shown in Table 2.1.

Of these criteria, two relate specifically to properties of the design matrix, \mathbf{X} . Under the assumption of additive Gaussian IID noise with 0 mean and standard deviation, σ ,

Name	description
<i>A-optimality</i>	Minimise $\text{trace}(\mathbf{X}^\top \mathbf{X})$
<i>D-optimality</i>	Minimise $\det(\mathbf{X}^\top \mathbf{X})$
<i>E-optimality</i>	Maximise the minimum eigenvalue of $\mathbf{X}^\top \mathbf{X}$
<i>G-optimality</i>	Maximise the maximal entry in $\mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}$

Table 2.1: Commonly used criteria used for design of experiments to find an optimal design. In the case of a nonlinear model, the design matrix, \mathbf{X} is substituted with the local-sensitivities matrix, \mathbf{S} .

the matrix, $\mathbf{X}^\top \mathbf{X}$ is proportional to the FIM, \mathcal{I} and is directly related to the size of the model's confidence regions, as described in Section 2.3.4.

The FIM, \mathcal{I} , may be used to find *D-optimal* designs. Similarly, the size of our marginal confidence intervals scales linearly with,

$$u_A(\mathbf{X}) := \sqrt{\sum_{i=1}^n (\mathbf{X}^\top \mathbf{X})_{i,i}^{-1}}. \quad (2.70)$$

which is used to find *A-optimal* design. That is, the A-optimal design is the design which minimises $\text{trace}(\mathbf{X}^\top \mathbf{X})$, which is proportional to the length of marginal confidence regions, averaged across each model parameter.

Optimal Experimental Design for nonlinear models We can employ a similar approach with nonlinear regression models of the same form (Equation (2.22)) but where F_i is some *nonlinear* map between our model parameters and experimental design and the i^{th} observable (taken from a Markov model, for example). Here, as in the linear-model case, we may apply our alphabet criteria, albeit under the linear approximation (that is, the linear approximation introduced in Section 2.3.4).

The use of this linear approximation simply results in the substitution of the design matrix, \mathbf{X} with the local sensitivities matrix, $\mathbf{S}(\boldsymbol{\theta})$. However, unlike \mathbf{X} in the linear case, the sensitivity matrix of a nonlinear model, $\mathbf{S}(\boldsymbol{\theta})$, may depend on the model parameters $\boldsymbol{\theta}$. This presents a challenge because the alphabet criteria (Table 2.1) depend on parameters which are unknown at the time the experiment is designed. One approach is to assume a fixed, best-guess parameter vector, $\boldsymbol{\theta}^*$, for the purpose of designing the experiment. In doing this, we obtain an optimisation problem similar to the linear model case, where we seek,

$$d_{\text{opt}} := \arg\max_{d \in \mathcal{D}} \{u_{\boldsymbol{\theta}^*}(d)\}, \quad (2.71)$$

where $u_{\boldsymbol{\theta}^*}(\tilde{d})$ is a utility function which depends on a particular value for our parameters, $\boldsymbol{\theta}^*$.

These methods may be used to assess the suitability of a given experimental design for fitting a particular model. For example, in Appendix C, we use the A- and D-optimality utility functions (u_A and u_D) to quantify the effect of data removal on the uncertainty in parameter estimates. Similar utility functions were used to produce the experimental designs used throughout this thesis (Lei et al., 2024; Mirams et al., 2024). Here, in the context of experimental design for voltage-clamp experiments, two main components are considered: the V_{cmd} signal and the sampling frequency at which recordings are made. For the experimental designs used in this thesis, all the V_{cmd} signals consist of piecewise affine functions (that is, a sequence of steps and ramps). However, the protocol proposed in Beattie et al. (2018) is a sum of sine-wave terms, and as such, could not be performed on the particular apparatus used for data collection. The optimal design criteria introduced above may be applied to our models' sensitivity matrices to ensure the suitability of experimental designs in future work—particularly for new model structures and the inclusion of “artefact effects” as discussed in Chapter 5.

2.4 Sources of predictive uncertainty

Throughout this thesis, we aim to develop methods to identify the most accurate mathematical models of I_{K_r} . A key focus is the predictive accuracy of these models—particularly when our models are used to produce predictions for experiments dissimilar to those used for model fitting.

In this section, we briefly summarise some reasons why our mathematical models of I_{K_r} do not produce perfectly accurate predictions. We discuss three sources of predictive inaccuracy, discrepancy in the chosen *noise model* (misspecification of the random distribution of our observation errors), *structural uncertainty* (misspecification of the dynamics in our mathematical models), and uncertainty in our parameter estimates. Whereas the latter may be addressed by improved experimental design (as discussed in the previous section), model misspecification, also known as *model discrepancy*, poses a specific challenge in fitting mathematical models to biological data.

2.4.1 Noise model

Markov models of macroscopic ion-channel currents typically assume that the observations are subject to some form of random error. The most common choice is additive Gaussian noise, which is assumed to be IID. This noise can be estimated using a stationary part of the trace. The observational noise present in time-series traces from patch-clamp experiments arises from a number of sources. Firstly, the measurement of the very small currents present (of the order of nanoamperes) is not perfect and

introduces an amount of observational noise. It is possible that the size of these effects changes over time and, as such, our observational errors may each be distributed with different standard deviation, despite being assumed to be homoscedastic (each distributed with equal variance).

We also anticipate a small amount of noise due to the stochasticity of channel gating (Mirams et al., 2016). However, this is estimated to be negligible due to the number of channels present, as demonstrated by Figure 2.7. Nevertheless, there have been attempts to use the stochastic models describing the behaviour of individual channels to model whole-cell macroscopic currents. This is done by assuming that the current is carried by a population of independent yet identical channels, and, considering the proportion of open channels over time. These techniques amalgamate the behaviour of individual channels and affect the expected probability distribution of observations. Stochastic differential equations provide a way of approximating this behaviour for many channels whilst providing suitable methods for the computation of the likelihood in light of a more complicated error model (Goldwyn et al., 2011).

2.4.2 Parameter uncertainty

When assuming random observational noise, we cannot expect to recover the true data-generating parameter set exactly (even when the data is synthetically generated using the model). This is because the parameter estimates we obtain (such as the maximum-likelihood estimate) depend on the random data, and so, are random variables themselves. However, by using a suitable experimental designs and models, the effect of parameter uncertainty can be minimised. This is demonstrated in Clerx et al. (2019a); Whittaker et al. (2020a), where high-frequency time-series data taken from appropriately designed protocols results in very little variability in parameter estimates. Even so, this parameter variability may increase if more parameters are included in the model.

2.4.3 Structural uncertainty

As discussed in Section 2.2.3, the number of states included in our models and the way they connect together is known as the model structure (and sometimes as the *model topology*) and can be represented mathematically as a graph (Bondy and Murty, 2008). A number of model structures which have been used to model I_{K_r} as discussed in Beattie et al. (2018), and shown in Figure 2.10.

These model structures vary both in terms of the number of states, but also in how the states are connected—for example, as discussed in Section 2.2.3, the Wang et al. (1997) model only allows the transition between open and inactive states via the open conformation; whereas the Beattie et al. (2018) and Kemp et al. (2021) models permit other sequences of transitions.




Figure 2.10: A selection of model structures which have been used to model I_{Kr} . Figure reproduced from Beattie et al. (2018) under the Creative Commons Attribution 4.0 International License .

In the context of I_{Na} models, Mangold et al. (2021) showed that there are an astronomical number of plausible model structures. We similarly have few constraints regarding possible model structures for I_{Kr} , and exhaustively considering each model structure somewhat infeasible. An alternative approach is to perform optimisation over the space of model structures. For example, Menon et al. used a graph-based approach in which the model structure is randomly mutated using a genetic algorithm (Menon et al., 2009).

Additionally, the presence of *experimental artefacts* in patch-clamp experiments (Lei et al., 2020a) means that the (unaltered) Markov models presented in this chapter provide an incomplete description of I_{Kr} under patch-clamp experiments. In particular, the transmembrane potential, V_m is not exactly the command voltage, V_{cmd} due to series-resistance attenuation and capacitive currents. These experimental artefacts impact the observed dynamics of I_{Kr} , and hence, their omission is a source of structural uncertainty. These experimental artefacts, and their combination with mathematical models of I_{Kr} are explored in Chapter 5.

In general, finding a suitable ODE-based model of some real-world phenomena is difficult. One general approach which has recently been popularised is *sparse regression*, in which a weighted combination of library functions is found that matches the data (Klimovskaia et al., 2016). However, we forgo these methods in favour of literature models which have been built using domain-specific knowledge.

2.5 General approaches for model discrepancy

One way of reducing over-confidence in inaccurate parameter estimates in the presence of model discrepancy may be to use approximate Bayesian computation (ABC) (Frazier et al., 2020). With ABC, a likelihood function is not explicitly specified; instead, the model is repeatedly simulated for proposed values of the parameter sampled from a prior distribution. Each proposed value is accepted or rejected according to whether the simulated trajectory is close to the actual data, according to some chosen summary statistics. ABC compares the simulated with the real data using these summary statistics (rather than matching all aspects of the dynamics) and accepts approximate matches (subject to a chosen tolerance). It is suited to inference where there is substantial model discrepancy because this approach can decrease potential over-confidence in the inferred values of parameters. However, it is challenging to select suitable summary statistics, and the computational demands of ABC are much greater than those of maximum-likelihood estimation.

Another approach was first introduced by Kennedy and O’Hagan (2001), who used Gaussian processes to fit approximate, computer models of real-world phenomena. This work has since been applied to electrophysiology models (Lei et al., 2020c). Elsewhere, Sung et al. (2020) introduced an approach to account for heteroscedastic errors using many repeats of the same experiment, although this seems to be less applicable to our I_{Kr} modelling problem because the number of repeats of each experiment (when training individual, cell-specific models) is limited. Alternatively, Lei and Mirams (2021) modelled the discrepancy using a neural network within the differential equations. However, these approaches are less interpretable than a purely mechanistic model.

2.6 Summary

In this chapter, we have introduced the primary aim of this thesis—the identification of suitable model structures for models of I_{Kr} currents. These models were born from the work of Hodgkin and Huxley’s work modelling cellular action potentials. The basic principle of this model—treating the whole cell as an equivalent circuit—allows us to reconcile our mathematical models with our experimental data. Throughout this thesis, we use the statistical and computational techniques introduced in Section 2.3 to fit literature models of I_{Kr} to synthetic and real data from patch-clamp experiments. In the following chapters, we use these methods to explore this structural uncertainty in mathematical models of I_{Kr} , and evaluate the predictive performance obtained by a small collection of plausible model structures.

Chapter 3

Uncertainty Quantification with an ensemble of experimental designs

3.1 Preamble

The work presented in this chapter was published in the Bulletin of Mathematical Biology, Shuttleworth et al. (2023). This work, including all the computational modelling, and statistical analysis, including the software developed for these purposes, was entirely my own.

3.2 Introduction

In mathematical biology, mechanistic models are typically fitted using experimental data from a particular experimental design, and then used to make predictions under different scenarios. We call the set of experimental designs under consideration the *design space*, denoted by \mathcal{D} , and assume the existence of some data-generating process (DGP), which maps each element of $d \in \mathcal{D}$ to some random output. These elements are known as experimental designs, or, as is more common in electrophysiology, *protocols*, and each corresponds to some scenario for which our model can be used to make predictions. For example, in the toy example presented in Section 3.2.1, each protocol, $d \in \mathcal{D}$, is simply a set of observation times, but for our patch-clamp voltage protocols, each d corresponds to some particular choice of voltage steps and durations. By performing an experiment consisting of multiple distinct protocols (a *multiprotocol* experiment), we investigate (and quantify) the *discrepancy* between the DGP and our models in different situations. Furthermore, the resultant parameter estimates provide insight into model discrepancy.

However, in the context of voltage-clamp experiments, these designs consist of a choice of voltage steps and ramps. In this way, each voltage signal (that is, V_{cmd}) is

treated as a forcing function in our ODE-based Markov models. We may also choose the sampling frequency, which dictates our observation times. Note that whilst each voltage-clamp protocol in this chapter uses a 10kHz sampling rate, lower sampling rates (2–5kHz) are used in the following chapters.

When fitting our mathematical models using standard frequentist or Bayesian approaches, it is typically assumed that there is no model discrepancy; in other words, that the data arise from the model (for some unknown, true parameter set). This is a necessary condition for some desirable properties of the parameter estimators which provide some guarantees regarding the accuracy of parameter estimates when there are many observations (as discussed in Section 3.3.1). However, when model discrepancy is not considered, we may find that the ability of a model to make accurate predictions is compromised. In particular, if we try to validate our model with a protocol dissimilar to that used for fitting, there can be a noticeable difference between our predictions and the data—even when the model appears to accurately recapitulate the data used for fitting. This problem is illustrated by a simple, toy problem in the following section.

In Section 3.2.1, we present a simple example, in which each protocol corresponds to a particular choice of observation times. Then, in Section 3.3, we introduce a similar approach using I_{K_r} models and synthetically generated patch-clamp data using a diverse range of voltage-clamp protocols. Here, we demonstrate that, for a discrepant model, the resultant parameter estimates and model predictions depend on the particular protocol used for model fitting, where experimental design relates to the aforementioned V_{cmd} voltage signals. We use this variability in model predictions to evaluate the suitability of our models, and provide empirically-derived, *spread-of-prediction* intervals which provide some quantification of predictive uncertainty due to model discrepancy. We demonstrate the utility of these methods under synthetically generated data by constructing two examples of model discrepancy. These methods are later applied to real, experimental data in Chapters 4 and 5.

3.2.1 Motivating Example

In this section, we construct a simple, toy model example in which a discrepant model is fitted under a range of experimental designs. To do this, we generate synthetic data under a DGP which differs from the model being fitted. This example demonstrates that, when using discrepant models, the experimental design we use for model fitting has an impact on the value of parameter estimates and the resultant model predictions.

First, we construct a DGP formed of the sum of two exponential terms,

$$y^*(t) = \exp\{-t\} + \exp\left\{-\frac{t}{10}\right\} \quad (3.1)$$

$$z^*(t) = y^*(t) + \varepsilon(t), \quad (3.2)$$

for some $t > 0$, where, for each $t > 0$, $\varepsilon(t)$ is an independent Gaussian random variable with zero mean and variance, $\sigma^2 = 10^{-4}$. Here, $z^*(t)$ is a random variable representing an observation of the system at some time, t .

Next, we attempt to fit a model which takes the form of single exponential decay,

$$y(t; \boldsymbol{\theta}) = \theta_1 \exp \left\{ -\frac{t_i}{\theta_2} \right\} \quad (3.3)$$

$$z(t; \boldsymbol{\theta}) = y(t; \boldsymbol{\theta}) + \varepsilon(t_i), \quad (3.4)$$

to these data, denoting the column matrix $[\theta_1, \theta_2]^T$ by $\boldsymbol{\theta}$. We call this a discrepant model because there is no choice of $\boldsymbol{\theta}$ such that $y(t; \boldsymbol{\theta}) = y^*(t)$, for all $t > 0$.

To fit our model, we choose a set of n observation times, $T = \{t_1, t_2, \dots, t_n\}$. We may then find the parameter set, $\hat{\boldsymbol{\theta}}(T)$, which minimises the sum-of-squares error between our discrepant model (Equation (3.4)) and each $z(\boldsymbol{\theta}; t_i)$, that is,

$$\hat{\boldsymbol{\theta}}(T) = \operatorname{argmin}_{\boldsymbol{\theta} \in \Theta} \left\{ \sum_{t_i \in T} (y(t_i; \boldsymbol{\theta}) - z^*(t_i))^2 \right\}, \quad (3.5)$$

where T is some set of observation times.

Then, we consider multiple experimental protocols which we may use to fit this model (Equation ((3.4))). In particular, we consider the following sets of observation times

$$T_1 = \{0, 0.01, 0.02, \dots, 0.01\}, \quad (3.6)$$

$$T_2 = \{0, 0.1, 0.2, 0.3, \dots, 1\}, \quad (3.7)$$

$$T_3 = \{0.2, 0.3, 0.4, 0.5, \dots, 1.2\}, \quad (3.8)$$

$$T_4 = \{0.5, 0.55, 0.6, \dots, 1\}, \text{ and } \quad (3.9)$$

$$T_{\text{all}} = T_1 \cup T_2 \cup T_3 \cup T_4. \quad (3.10)$$

We sample from the DGP 10 times by computing Equation (3.2) for each observation time, t , and adding IID Gaussian noise. Then, for each sample of the DGP, we compute parameter estimates using each set of observation times (T_1, T_2, T_3, T_4 and T_{all}). This process is then repeated with a ten-fold increase in sampling rate, that is, with observation

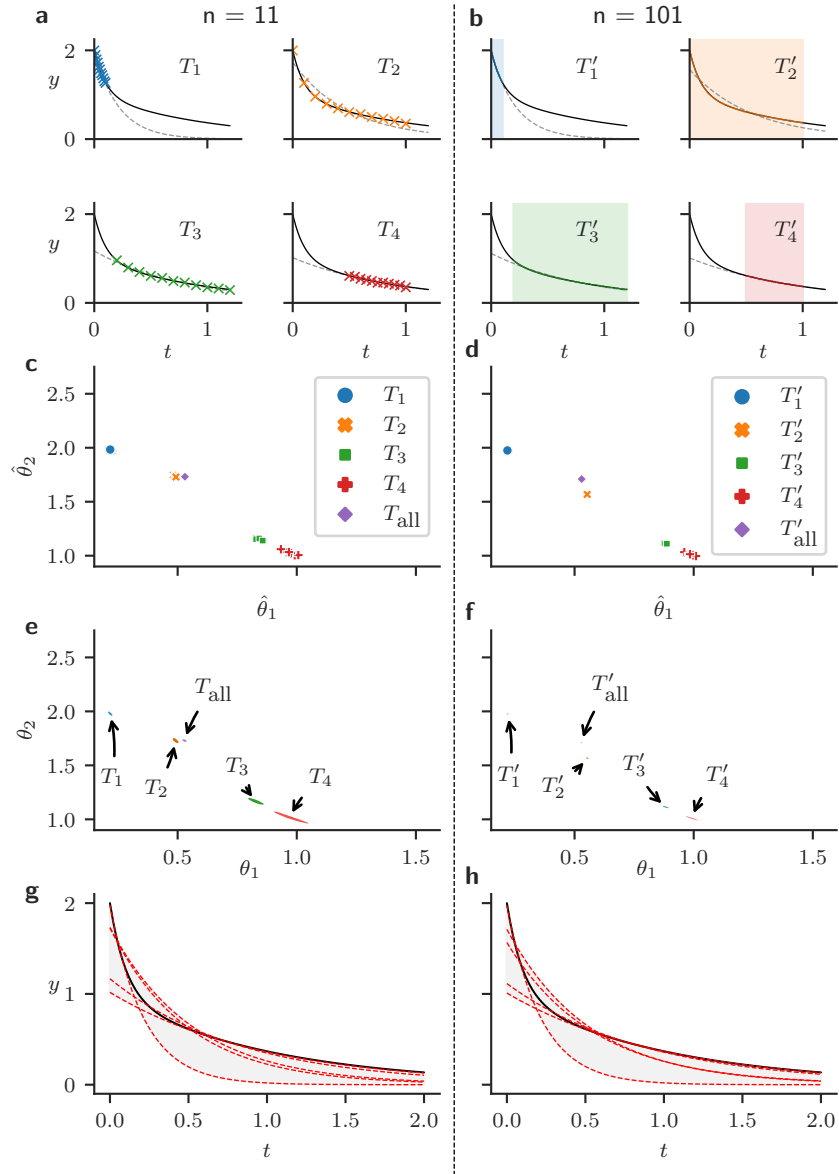


Figure 3.1: Under model discrepancy, parameter estimates depend on the design used for fitting—even when our models are practically identifiable (see Section 2.3.4) and calibrated with a vast amount of information-rich data. Panels to the left of the dotted line correspond to designs containing $n = 11$ observations at times (T_1, \dots, T_4) as shown in Panel **a**). Panels on the right show designs with $n = 101$ observations, (T'_1, \dots, T'_4) as shown in Panel **b**) (**a**) and (**b**): representative datasets generated by the DGP shown with the solid black line (Equation (3.1)) with points indicating observations (sampled using Equation (3.2)) and the fitted discrepant model (Equation (3.4)), with calibrated $\boldsymbol{\theta}$ (grey dashed lines). (**c**) and (**d**): the parameter estimates for each design, each fitted to one of ten repeats of the DGP. (**e**) and (**f**): 99% Bayesian credible regions obtained using MCMC, a uniform prior and a single repeat of the DGP. (**g**) and (**h**): predictions using the discrepant model fitted using a single repeat of each design (using the estimates shown in **e** and **f**), showing the true DGP (black), discrepant model predictions (red), and the difference between predictions (grey).

times,

$$T'_1 = \{0, 0.001, 0.002, \dots, 0.01\}, \quad (3.11)$$

$$T'_2 = \{0, 0.01, 0.02, 0.03, \dots, 1\}, \quad (3.12)$$

$$T'_3 = \{0.2, 0.21, 0.22, 0.23, \dots, 1.2\}, \quad (3.13)$$

$$T'_4 = \{0.5, 0.505, 0.51, \dots, 1\}, \text{ and} \quad (3.14)$$

$$T'_{\text{all}} = T_1 \cup T_2 \cup T_3 \cup T_4. \quad (3.15)$$

As shown in Figure 3.1, for each sample of the DGP, we obtain a parameter estimate from each set of observation times, each with a different distribution. For instance, fitting using T_1 results in a model that approximates the DGP well on short timescales, and fitting using T_4 allows us to recapitulate the behaviour of the system over longer timescales, as can be seen in Panel **a**. From how closely the discrepant model (Equation (3.4)) fits the data in the regions where observations are made (in Figure 3.1, Panels **a** and **b**), we can see that, in either case, a single exponential seems to provide a reasonable approximation to the DGP. However, if we require an accurate model for both the slow and fast behaviour of the system, model discrepancy presents an issue, and this model (namely, Equation (3.4)) may be unsuitable. This is the case for T_2 as shown in Figure 3.1**a**. This variability in behaviour is shown in Figure 3.1 (Panels **g** and **h**) which shows how the model's predictions for $0 \leq t \leq 2$ depend on what protocol was used to fit the model.

The Bayesian posteriors illustrated in Figure 3.1, Panels **e** and **f**, show that we are not able to avoid the problems caused by model discrepancy by simply adopting a Bayesian framework—we will obtain precise parameter estimates that are highly dependent on the chosen fitting protocol, nevertheless. This problem becomes more obvious when we increase the number of observations. In the examples detailed in this chapter, we explore this ‘high-data limit’ where the variability in each parameter estimate (under repeated samples of the DGP) is minuscule compared to the difference between parameter estimates obtained from different protocols.

3.3 Methods

Here, we provide a general overview of our proposed methods before providing two real-world examples of their applications. In Section 3.3.1 we explain our computational methods in detail. In particular, we describe a method for validating our models, in which we evaluate predictive accuracy whilst using many different pairs of protocols for fitting and validation. This motivates our proposed methods for combining parameter estimates obtained from different protocols to empirically quantify the impact of model

discrepancy on predictions of unseen protocols. These methods are expanded upon in later chapters in this thesis, where they are applied to real, experimental data sets instead of synthetically generated data.

We begin this section by providing a general introduction to our multiprotocol, cross-validation approach. Here, we use the same notation as in Section 2.3, where our mean function under protocol d is denoted $\mathbf{F}(\boldsymbol{\theta}; d)$, with components corresponding to our n_d observations under protocol d , $(F_i(\boldsymbol{\theta}; d))_{i=1}^{n_d}$. Likewise, we denote the observations made under protocol d by $\mathbf{z}(d)$, making the dependence on the chosen protocol explicit.

3.3.1 Fitting models using multiple experimental protocols

Evaluation of predictive accuracy and model fitting

We fit our model to a given experimental design, d , by computing a maximum-likelihood estimate, $\hat{\boldsymbol{\theta}}_d$ (using Equation (2.29)). Then, we validate our model by computing predictions for some other protocol, $\tilde{d} \in \mathcal{D}$. To do this, we compute, $\mathbf{F}(\hat{\boldsymbol{\theta}}_d; \tilde{d})$, where \mathbf{F} is the model's mean function as introduced in Chapter 2 such that the i^{th} component, F_i , is our i^{th} observable. This is a simulation of the behaviour of the system (without noise) under protocol \tilde{d} made using parameter estimates that were obtained by fitting the model to protocol d . In this way, our parameter estimates, each obtained from different protocols, result in different out-of-sample predictions (predictions under protocols other than the one used for fitting). Because we aim to fit a model able to produce accurate predictions for all possible $d \in \mathcal{D}$, it is important to validate our model using multiple protocols.

By computing, $\text{RMSE}(\mathbf{F}(\hat{\boldsymbol{\theta}}_d; \tilde{d}), \mathbf{z}(\tilde{d}))$ we are able to perform model validation for each pair of fitting and validation protocols (d and \tilde{d} , respectively). This allows us to ensure our models are robust with regard to the fitting protocol, and allow for the quantification of model discrepancy as demonstrated in Section 3.4.

Consequences of model error/discrepancy

Ideally, we would have a model that is correctly specified. In particular, we assume our vector of observations, $\mathbf{z}(d)$, is such that our i^{th} observation is,

$$z_i = F_i(\boldsymbol{\theta}^*; d) + \varepsilon_i, \quad (3.16)$$

where d is the protocol under consideration and $\boldsymbol{\theta}^*$ is some fixed, unknown parameter vector, and, as in Chapter 2, our observational errors, $\{\varepsilon_i\}_{i=1}^n$ are IID distribution Gaussian random variables with zero-mean and variance σ^2 .

We then consider the distance between an estimate $\hat{\boldsymbol{\theta}}$ and the true value—that is, the error in our parameter estimates. As discussed in Chapter 2, when the model is correctly

specified, and provided certain regularity conditions hold, we can obtain arbitrarily accurate parameter estimates by increasing the number of observations, n . That is, more precisely, that $\hat{\boldsymbol{\theta}}$ converges in probability to $\boldsymbol{\theta}^*$ as $n \rightarrow \infty$ (Seber and Wild, 2005).

However, when fitting discrepant models, we may find that our parameter estimates are heavily dependent on the fitting protocol, as demonstrated for our toy model in Section 3.2.1. For unseen protocols, these discordant sets of parameter-estimate vectors may lead to a range of vastly different predictions, even if each parameter set provides a reasonable fit for its respective protocol. In such a case, further data collection may reduce the variance of these parameter estimates (under repeated data collection under the same protocol), but fail to significantly improve the accuracy of our models when predicting other, unseen protocols. In other words, further data collection does not guarantee a more accurate predictive model.

In Section 3.4, we explore two examples of synthetically constructed model discrepancy. In Section 3.4.1, we have that \mathbf{f} and h (as in Equations (2.58) and (2.23)) are exactly those functions used to generate the data, and the exact probability distribution of the observational errors is known. However, the maximal conductance parameter is fixed to a value that was not used for data generation. Consequently, the true parameter set $\boldsymbol{\theta}^*$ lies outside the parameter space considered for model fitting. Under the assumption of structural identifiability (and a compact parameter space), this is an example of model discrepancy because there is some limit to how well our model can recapitulate the DGP as we will never recover the true data-generating parameter set.

In Section 3.4.2 we explore another example of model discrepancy where our choice of \mathbf{f} (and, in this case, the dimensions of $\boldsymbol{\theta}$ and \mathbf{x}) are misspecified by fitting a model which differs structurally from the one used in the DGP. Again, in this case, the model that we fit differs from that used for data generation, and so, this is an example of model discrepancy.

Ensemble fitting and prediction interval

As outlined in Section 3.3.1, we can obtain parameter estimates from each protocol $d \in \mathcal{D}$ by finding the $\hat{\boldsymbol{\theta}} \in \Theta$ that minimises Equation (2.32). We then obtain an *ensemble* of parameter estimates,

$$\left\{ \hat{\boldsymbol{\theta}}_d : d \in \mathcal{D}_{\text{fit}} \right\}. \quad (3.17)$$

Then, for any validation protocol \tilde{d} , the set,

$$\left\{ \mathbf{y}(\hat{\boldsymbol{\theta}}_d; \tilde{d}) : d \in \mathcal{D}_{\text{fit}} \right\}, \quad (3.18)$$

is an ensemble of predictions where $\mathcal{D}_{\text{fit}} \subseteq \mathcal{D}$ is some set of fitting protocols. Then we may consider the range of our predictions for each observation of interest. For the i^{th}

observation of our validation protocol, \tilde{d} , that is

$$\begin{aligned}\mathcal{B}^{(i)} &= [\mathcal{B}_{\text{lower}}^{(i)}, \mathcal{B}_{\text{upper}}^{(i)}] \\ &= \left[\min_{d \in \mathcal{D}_{\text{fit}}} \left\{ F_i(\hat{\boldsymbol{\theta}}_d; \tilde{d}) \right\}, \max_{d \in \mathcal{D}_{\text{fit}}} \left\{ F_i(\hat{\boldsymbol{\theta}}_d; \tilde{d}) \right\} \right].\end{aligned}\quad (3.19)$$

When all observations are considered at once, Equation (3.19) comprises a band of predictions, giving some indication of predictive uncertainty. We demonstrate below that this band provides a useful indication of predictive error for unseen protocols, and provides a range of plausible predictions. We propose that a wide band of predictions for a given validation protocol suggests that there is model discrepancy and poor prediction accuracy for a particular context of use.

Note that the width of this band at any given point (Equation (3.19)) cannot decrease as more protocols are added, it may only get larger. Therefore, if a large number of protocols are considered, percentiles of our ensemble of predictions may be more suitable. However, in this chapter, we only consider cases where there are a small number of protocols (five protocols are used for fitting in each of the examples discussed in Section 3.4). In later chapters, the number of protocols increases somewhat when 11 distinct fitting protocols are considered.

For the purposes of a point estimate, we may use the midpoint of each interval,

$$\mathcal{B}_{\text{mid}}^{(i)} = \frac{\mathcal{B}_{\text{lower}}^{(i)} + \mathcal{B}_{\text{upper}}^{(i)}}{2}.\quad (3.20)$$

This is used to assess the predictive error of the ensemble in Figure 3.8. There are other ways to gauge the central tendency of the set of predictions (Equation (3.18)). Such a change would have little effect on Section 3.4, but a median or weighted mean may be as (or more) suitable for other problems.

3.3.2 Application to \mathbf{I}_{Kr} models

We now turn our attention to an applied problem in which dynamical systems are used to model \mathbf{I}_{Kr} under the application of various voltage-clamp protocols. The models we employ are Markov models (as introduced in Section 2.2.3) augmented with additive Gaussian observational errors.

\mathbf{I}_{Kr} Markov models under consideration

In this chapter, we use two examples of Markov models for both data generation and model fitting. As discussed in Chapter 2, these two Markov models may be characterised by their respective transition-rate matrices, and, for both models, the evolution of the

proportion of channels in each of N states is governed by,

$$\frac{d\mathbf{x}}{dt} = \mathbf{Q}(V_m)^\top \mathbf{x}, \quad (3.21)$$

where \mathbf{Q} is an $N \times N$ transition-rate matrix, \mathbf{x} is a N -vector where each component describes the proportion of channels in a given state, and V_m is the transmembrane potential. As shown in Appendix A, we may simplify Equation (3.21), and instead compute,

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{A}(V_m)\tilde{\mathbf{x}} + \mathbf{b}(V_m), \quad (3.22)$$

for some voltage-dependent matrix $\mathbf{A}(V_m)$, and voltage-dependent vector $\mathbf{b}(V_m)$, where $\tilde{\mathbf{x}}$ is \mathbf{x} with the bottom element, x_N omitted. This final component, x_N , may be recovered from $\tilde{\mathbf{x}}$ by using the conservation constraint, $\mathbf{x}^\top \mathbf{1} = 1$, from which we see,

$$x_N = 1 - \tilde{\mathbf{x}}^\top \mathbf{1}. \quad (3.23)$$

For both models, the state-variables vector \mathbf{x} is mapped onto our i^{th} observable, F_i , by computing,

$$F_i = g x_O(t_i; \boldsymbol{\theta}, d)(V_m - E_{\text{Kr}}), \quad (3.24)$$

where g is the maximal conductance parameter (an element of $\boldsymbol{\theta}$), x_O denotes the proportion of channels in the open conformation (an element of \mathbf{x}), t_i is the time corresponding to the i^{th} observation, and d is the protocol under consideration. Our i^{th} observable, F_i , is then mapped onto our i^{th} observation, z_i , via the addition of a zero-mean Gaussian random variable (Equation (3.16)).

The first model of I_{Kr} we consider is by Beattie et al. (2018). This is a four state Markov model with nine parameters (8 of which relate to the model's kinetics). For data generation, we use the parameters that Beattie et al. (2018) obtained from one particular cell (namely, Cell #5) by fitting their model to data obtained from an application of the “sinusoidal protocol” with a manual patch-clamp setup. The cells used were Chinese hamster ovary cells, which were made to heterologously over-express *hERG1a*. These experiments were performed at room temperature.

The second model was first presented in Wang et al. (1997). This is a five-state model which has 15 parameters (including the maximal conductance parameter). These parameters were obtained by performing multiple voltage-clamp protocols, all at room temperature, on multiple frog oocytes overexpressing *hERG*. These experiments are used to infer activation and deactivation time constants as well as steady-state current voltage relations, which are, in turn, used to produce estimates of model parameters. Of the two parameter sets provided in Wang et al. (1997), we use the parameter set obtained by using the extracellular solution with a 2 mM concentration of potassium

chloride, as this most closely replicates physiological conditions. The values of the model parameters for both the Beattie et al. (2018) and Wang et al. (1997) models (as used in Sections 3.4.1 and 3.4.2) are provided in Appendix B (Table B.1).

Before and after each protocol, cells are left to equilibrate with the transmembrane potential, V_m , set to the *holding potential*, $V_{\text{hold}} = -80$ mV. Therefore, we require the initial conditions, for at time $t = 0$,

$$\mathbf{x}(0) = \mathbf{x}_{\infty}(V_{\text{hold}}), \quad (3.25)$$

where, $\mathbf{x}_{\infty}(V_{\text{hold}})$ is the unique steady-state solution for the linear system,

$$\frac{d\mathbf{x}}{dt} = \mathbf{Q}^{\top}(V_{\text{hold}}; \boldsymbol{\theta}_f) \mathbf{x}, \quad (3.26)$$

subject to the constraint $\mathbf{x}^{\top} \mathbf{1} = 1$. The existence and uniqueness of such a steady state is demonstrated in Appendix A. Moreover, as described in Appendix A, we rewrite our system of equations in terms of $N - 1$ state variables, where N is the number of states,

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{A}(V_{\text{hold}}) \tilde{\mathbf{x}} + \mathbf{b}(V_{\text{hold}}). \quad (3.27)$$

In particular, we choose \mathbf{A} and \mathbf{b} such that $\tilde{\mathbf{x}} = (x_1, \dots, x_{N-1})^{\top}$. Then the steady state may be found by setting $\frac{d\tilde{\mathbf{x}}}{dt} = 0$ and solving the resulting linear system of equations. We expect that simulating our models using the reduced form (Equation (3.27)) reduces the computation required to fit our models.

As is standard for models of I_{Kr} (Beattie et al., 2018), we take our observation function to be

$$I_{\text{Kr}} = h(\mathbf{x}, t_i; g, d) = g \cdot x_{\text{O}}(t; \boldsymbol{\theta}_f, d) \cdot (V(t; d) - E_{\text{Kr}}), \quad (3.28)$$

where $[O]$ denotes the proportion of channels in an “open” conformation (one of the components of \mathbf{x}); g is the maximal conductance; and the reversal potential, E_{Kr} , is assumed to be the Nernst potential. The Nernst potential is found by computing,

$$E_{\text{Kr}} = \frac{RT}{F} \ln \left\{ \frac{[K_{\text{out}}]}{[K_{\text{in}}]} \right\}, \quad (3.29)$$

where R is the gas constant, F is Faraday’s constant, and T is the temperature and $[K_{\text{in}}]$ and $[K_{\text{out}}]$ are the intracellular and extracellular concentrations of K^+ , respectively. Here, we choose the temperature to be room temperature ($T = 298\text{K}$), and our intracellular and extracellular concentrations to be 120 mM and 5 mM, respectively, which approximately correspond to physiological concentrations. Hence, for all synthetic data generation,

fitting, and validation in Sections 3.4.1 and 3.4.2, we fix $E_{K^r} = -80.24$ mV (using Equation (3.29)).

We assume that our observational errors are additive Gaussian IID random variables with zero mean and variance, σ^2 , that is,

$$z_i = F_i + \varepsilon_i, \quad (3.30)$$

where $\{\varepsilon_i\}_{i=1}^n$ is a set of IID Gaussian random variables, where for each i ,

$$\varepsilon_i \sim N(0, \sigma). \quad (3.31)$$

Accordingly, our models are fitted using MLE by computing the least-squares parameter estimates.




Figure 3.2: The structural differences between the two Markov model structures used in this chapter for synthetic data generation and model fitting. (a): the four-state Beattie et al. (2018) model used in both *Case I* and *Case II*. (b): the five-state Wang et al. (1997) model used only for *Case II*. When a channel is in the open/conducting (O) state (green) current is able to flow. Whereas, when the model is in the other *closed* (C) or *inactivated* (I) states, no current can flow. The arrows adjacent to each model structure indicate the direction in which rates increase as the voltage increases.

Experimental Designs for Voltage Clamp Experiments

A large amount of data can be recorded in voltage-clamp electrophysiology experiments, where the current may be recorded at a several-kHz sampling rate for many minutes. In what follows, we take observations of the current at the same 10 kHz frequency for all protocols. Experimenters have a great deal of flexibility when it comes to designing voltage-clamp experiments (specifying voltage-clamp protocols). A number of studies having been published on the benefits of “information-rich” experimental designs (Beattie et al., 2018; Lei et al., 2019a,b; Clerx et al., 2019a; Kemp et al., 2021). In a real patch-clamp experiment, the amount of data we can obtain from each cell is limited. Hence, it is not feasible to perform many long protocols in sequence on the

same cell. We use five short information-rich protocols for the purposes of model fitting, denoted d_1 to d_5 , as shown in Figure 3.3.

Here, we use simple designs consisting of a combination of sections where the voltage is held constant or “ramps” where the voltage increases linearly with time for compatibility with automated high-throughput patch clamp machines which are restricted to protocols of this form. For the protocols included in this chapter, short identical sequences including ramps are placed at the beginning and end of each protocol. In real experiments, these elements will allow for quality control, leak subtraction, and the estimation of the reversal potential (Lei et al., 2019a,b), as described in Chapter 4. The central portion, consisting of steps during which the voltage is held constant, is what varies between protocols. These constant-voltage steps were chosen as to simplify the optimisation of protocols (Mirams et al., 2024), and to allow the model to be solved using matrix-exponential methods (as discussed in Section 2.2.3).

Not all possible designs are suitable for fitting models. Sometimes we encounter protocols for which the model output is not sensitive to certain changes in parameter values—that is, the model is not *practically identifiable* from the resulting data (as discussed in Chapter 2.3.4). Subsequently, when fitting the model to data generated from such a protocol, many different parameter sets give similar fits. This is, generally speaking, best avoided, particularly when we use our models for predictions dissimilar to those used for training, which may be sensitive to such changes in parameters.

Under both the Beattie et al. (2018) and Wang et al. (1997) models, practical unidentifiability is a problem for the long action-potential protocol (Lei et al., 2024), denoted d_0 . This practical unidentifiability is illustrated for a similar I_{K_r} model and protocol in Figure 3 of Whittaker et al. (2020a)). Nevertheless, d_0 mimics the transmembrane potential of a myocyte in a range of scenarios. The precise waveform of a cardiac action potential varies depending on the precise location of the myocyte (such as in ventricular or atrial tissue), or under different pacing rates, for example (Giles and Imaizumi, 1988; ten Tusscher et al., 2004). Hence, protocol d_0 allows validation of our models’ ability to recapitulate well-studied, physiologically-relevant behaviour. Due to this physiological relevance, and lack of practical identifiability, we use d_0 as a validation protocol, but not for model fitting.

The remaining designs, d_1 – d_5 , were constructed using various criteria under constraints on voltage ranges and the duration of each step as described in Lei et al. (2024). The protocol d_1 was designed algorithmically by sampling from a probability distribution placed over possible parameter sets and maximising the difference in model outputs between all pairs of parameter sets sampled from this distribution; d_5 was the result of the same algorithm applied to the Wang et al. (1997) model. In contrast, d_4 is a manually designed protocol used previously (Lei et al., 2019a) based on a simplification of a sinusoidal design (Beattie et al., 2018). The design, d_2 is a further manual refinement of

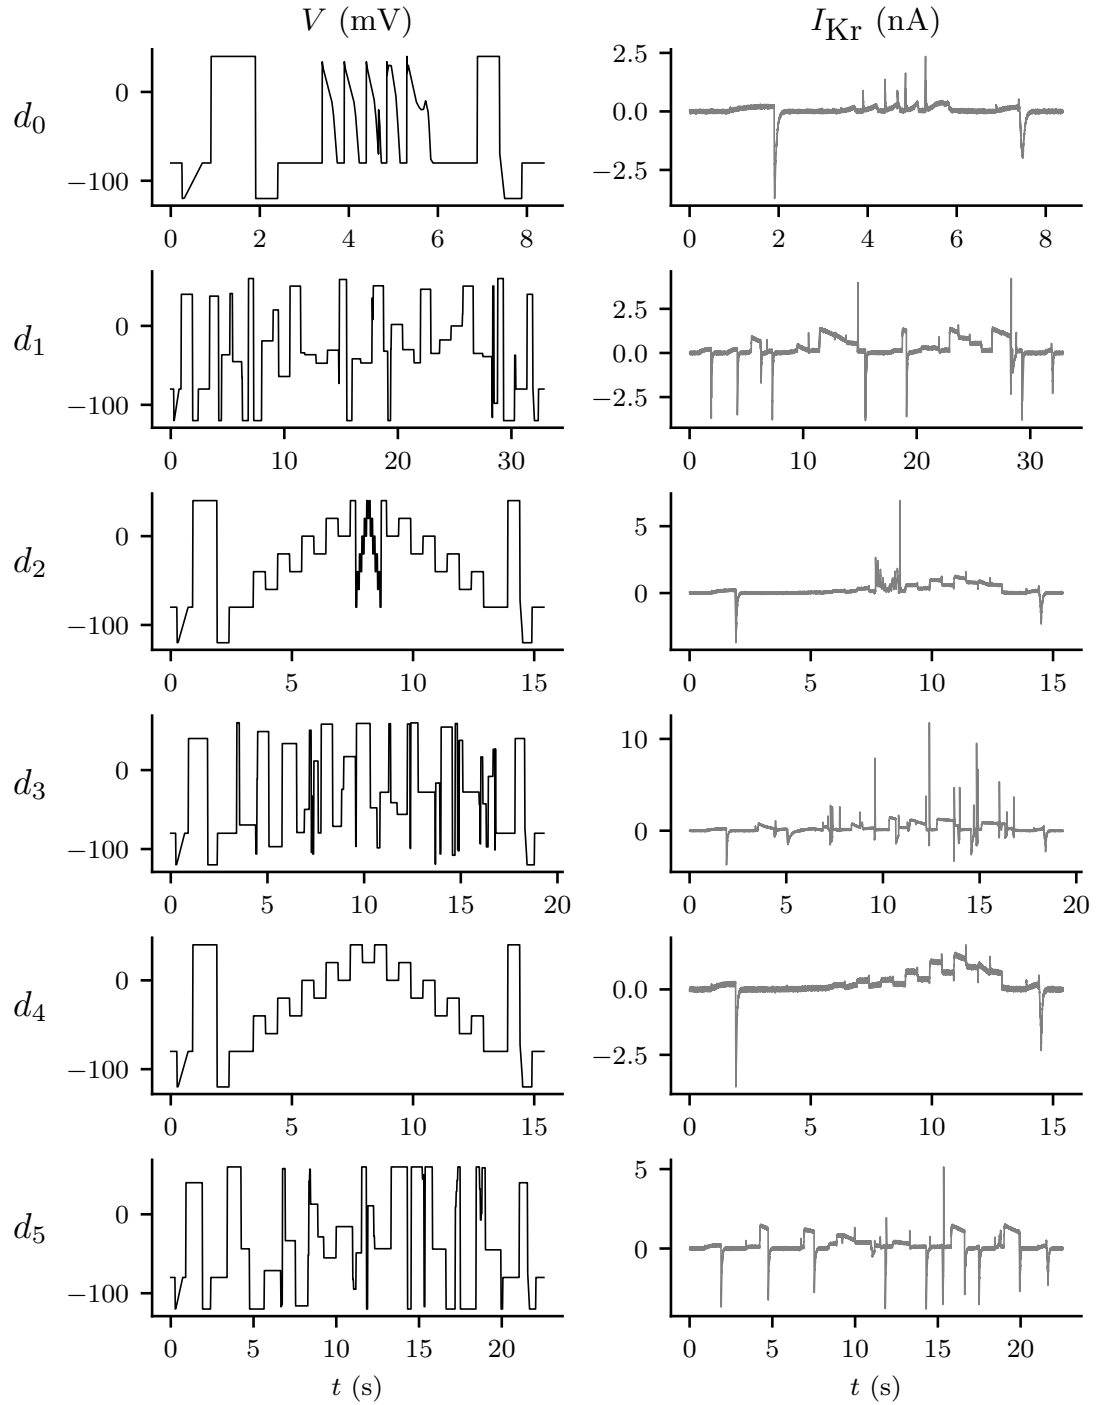


Figure 3.3: **Left:** a range of different input voltage-clamp protocols (forcing functions) used in this study. **Right:** corresponding synthetic output data I_{Kr} simulated using the Beattie et al. (2018) model with noise added as described in Section 3.3.2. Here, we generate and plot data observed at a 10 kHz sampling rate. fitting protocols (all protocols except d_0) were tested for practical identifiability: inverse problems performed on synthetic data with repeatedly sampled random noise yielded parameter estimates with little variability.

d_4 , which aims to explore inactivation processes (rapid loss of current at high voltages) more thoroughly. Finally, d_3 is based on maximising the exploration of the model state space for the Beattie et al. (2018) model, visiting as many combinations of binned model states and voltages as possible (Mirams et al., 2024). The main thing to note for this study, however, is that protocols d_1 – d_5 were designed to result in good practical identifiability (Fink and Noble, 2009) for both models. This is shown in the results presented in Section 3.4 (and in Figure B.2 for $\lambda = 1$ for the Beattie et al. (2018) model, and Figure B.5 for the Wang et al. (1997) model). From these results, we can see that all five fitting protocols yield very low-variance parameter estimates when used to fit model parameters under synthetically generated data (under the assumption of a correctly specified model).

This practical identifiability is a useful property, because it allows us to disregard the (very small) effect of different random noise in the synthetic data on the spread of our predictions (Equation (3.19)). A summary of these protocols is provided in Table 3.1.¹

Protocol	duration (s)	number of segments	code name
d_0	8.4	63	<i>longap</i>
d_1	39.2	57	<i>hhbrute3gstep</i>
d_2	15.0	52	<i>sis</i>
d_3	18.8	64	<i>spacefill19</i>
d_4	22.5	58	<i>staircase</i>
d_5	32.6	58	<i>wangbrute3gstep</i>

Table 3.1: The duration of and number of segments in each protocol that we performed. During these segments, the command voltage is either held constant, or changes according to some constant gradient (a ramp). Here, the number of segments includes segments which are common to all of our protocols—the leak ramp and reversal ramp, and their associated +40 mV pre-conditioning steps. The *staircase* and *longap* protocols (d_1 and d_6 , respectively) were first introduced in Lei et al. (2019a), and the remaining protocols are discussed in Lei et al. (2024).

Computational Methods

Numerical solution of ODEs Any time we calculate $\mathbf{F}(\boldsymbol{\theta}; d)$, we must solve a system of ordinary differential equations. We use a version of the LSODA solver designed to work with the Numba package, and Python to allow for the generation of efficient just-in-time compiled code from symbolic expressions. We partitioned each protocol into sections where the voltage is constant or changing linearly with respect to time

¹Note that the labelling of our protocols in this chapter differs from the labelling in Chapters 4 and 5, where a larger number of protocols are considered. The code names listed in Table 3.1 remain consistent between chapters, however.

because this sped up our computations. We set LSODA’s absolute and relative solver tolerances to 10^{-8} . The fact that the total number of channels is conserved in our models, allows us to reduce the number of ODEs we need to solve from N to $N - 1$ (Fink and Noble, 2009), where N is the number of conformational states in the model. This is discussed further in Appendix A.5.

Synthetic data generation Having computed the state of the system at each observation time, $(\mathbf{x}(t_i, \boldsymbol{\theta}^*, d))_{i=1}^{n_d}$, it is simple to compute F_i by substituting \mathbf{x} into our observation function (Equation (3.16)). Finally, to add noise, we obtain n_d independent samples using Equation (3.16), using *NumPy*’s (Harris et al., 2020) interface to the PCG-64 pseudo-random number generator. Here, because we are using equally spaced observations with a 10 kHz sampling rate, $n_d = 10,000 \times t_{\text{duration}}$ where t_{duration} is the length of the protocol’s voltage trace in seconds.

Optimisation Finding the least-squares estimates, or, equivalently, minimising the RMSE between the mean function and data (as described in Section 2.3) is (in general) a nonlinear optimisation problem for which there exist many numerical methods. We use CMA-ES (Hansen, 2016) as implemented by the PINTS interface (Clerx et al., 2019b). As discussed in Section 2.3.2, CMA-ES is a global, stochastic optimiser that has been applied successfully to many similar problems. We terminate the optimisation once 200 iterations have failed to yield a materially improved parameter set (that is, by decreasing the sum-of-squares error by more than 10^{-7}). We select the population size (that is, the number of parameter vectors sampled for each generation) by computing the integer,

$$n_{\text{pop}} := \lfloor 4 + 3 \ln(n_p) \rfloor, \quad (3.32)$$

where n_p is the number of model parameters. This heuristic is suggested by the PINTS package (Clerx et al., 2019b)

We follow the optimisation advice provided in Clerx et al. (2019a). Here, for parameters ‘ A ’ and ‘ b ’ in state transition rates of the form $k = A \exp\{(bV)\}$, the optimiser works with ‘ $\log A$ ’ and untransformed ‘ b ’ parameters. We enforce fairly lenient constraints on our parameter space, Θ , to prevent a proposed parameter set from forcing transitions to become so fast/slow that the ODE system becomes very stiff and computationally difficult to solve. In particular, we take a similar approach to Clerx et al. (2019a) we require that every parameter is positive, and, for ease of computation,

$$1.67 \times 10^{-5} \text{ ms}^{-1} \leq k_{\text{max}} \leq 10^3 \text{ ms}^{-1}, \quad (3.33)$$

where k_{\max} is the maximum transition rate, $k(V_m)$, for all

$$V_m \in [-120 \text{ mV}, +60 \text{ mV}],$$

which is the range of voltages applied during our protocols (Figure 3.3).

To ensure that we have found the global minimum (that is, that we have accurately computed Equation (2.32)), we repeat every optimisation numerous times (25 repeats for $\lambda = 1$ in Case I, 5 repeats for subsequent λ , and 25 repeats in Case II). Moreover, in Section 3.4.2, when fitting the discrepant model, our initial guesses for the kinetic parameters are randomly sampled using

$$\log_{10}(p) \sim U(-7, -1), \quad (3.34)$$

whereas we set the maximal conductance initial guess (which only affects the observation function) to the value used for data generation (even though these data were generated using a different model structure). We then check that our parameter set satisfies Equation (3.33), and resample if necessary before commencing the optimisation routine.

The examples presented in Sections 3.4.1 and 3.4.2 require the solution of many optimisation problems. For speed, we organised these tasks in such a way that they can be run in parallel on a high-performance computing resource.

3.4 Results

In this section, we use synthetically generated data to explore two cases of model discrepancy in Markov models of I_{Kr} . In this first case, we gradually introduce discrepancy into a model with the correct structure by fixing one of its parameters to values away from the DGP parameter set. Then, in Section 3.4.2, we apply the same methods to another case where the model structure is incorrectly specified. In both cases, we take a literature model of I_{Kr} together with Gaussian IID noise to be the DGP—in the first case, we assume a DGP based on the Beattie et al. (2018) model, whereas in the second, we assume a DGP based on the Wang et al. (1997) model. Again, these conflicting model structures are shown in Figure 3.2.

3.4.1 Case I: Misspecified maximal conductance

In this case, we assume a correctly specified model, but assume increasingly erroneous values of one particular parameter and investigate how this impacts the protocol-dependence of our parameter estimates and the predictive accuracy of our models. Also,

we explore how the spread in our model predictions (Equation (3.19)) increases as the amount of discrepancy increases (in a particular manner).

To do this, we simulate data generation from each fitting protocol, as outlined, ten times using Gaussian IID noise with standard deviation (30 pA). Specifically, we take the true DGP to be the Beattie et al. (2018) model, as shown in Figure 3.2. Then, we fix the maximal conductance, g , to a range of values, and infer the remaining model parameters from the synthetic data, which we generate using the true parameter set, θ^* . We assume that the standard deviation of the Gaussian noise is known because it can be well estimated from the initial portion of each protocol where the current is stationary.

When fitting our models, we use a restriction of the usual parameter space to fit the data by assuming some fixed value, λ , for the maximal conductance, g . In this way, we reformulate the optimisation problem slightly such that Equation (2.32) becomes,

$$\hat{\theta}_\lambda(d) = \operatorname{argmin}_{\theta \in \Theta_\lambda} \{\operatorname{RMSE}(\mathbf{y}(\theta; d), \mathbf{z}(d))\}, \quad (3.35)$$

where Θ_λ is the subset of parameter space where the maximal conductance is fixed to λg . For each repeat of each protocol, we solve this optimisation problem for each scaling factor, $\lambda \in \{\frac{1}{4}, \frac{1}{2}, 1, 2, 4\}$. These calculations are identical to those used in the computation of *profile likelihoods* under the assumption of additive IID Gaussian errors (Bates and Watts, 1988).

Next, we fit these restricted parameter-space models to the same dataset and assess their predictive power under the application of a validation protocol. We do this for each possible pair of fitting and validation protocols. To reduce the time required for computation, we fit our discrepant models sequentially, starting at $\lambda = 1$ and increasing or decreasing λ , using previous parameter estimates as an initial guess. This is done so that, for example, the kinetic parameters found by fixing $\lambda = 2$ are used as our initial guess when we fit the model with $\lambda = 4$, unless the original kinetic parameters (Table B.1) provide a smaller RMSE than the results of the previous optimisation.

The spread in predictions for the validation protocol, d_0 , for, $\lambda \in \{\frac{1}{4}, 1, 4\}$ is shown in Figure 3.4. A more complete summary of these results is provided by Figure 3.5. Here, when $\lambda = 1$ (the central row of Figure 3.5), we can see that no matter what protocol is used to fit the model, the distribution of parameter estimates (Panel **a**) is centred around their true values, and the resultant predictions are all accurate (Figure 3.5, Panels **b** and **c**). In contrast, when the maximal conductance, g is set to an incorrect value our parameter estimates become biased and our predictions become much less accurate overall. This effect on predictive accuracy is also shown in Appendix B.2 (Figure B.1). This effect does not appear to be symmetrical, with $\lambda < 1$ seemingly resulting in more model discrepancy than $\lambda > 1$. Here, we also see that the inaccuracy

in our parameter estimates and our predictions depends on the design used to fit the model.

Further results are provided in Section B.2. Figure B.1 shows the error in our predictions of d_0 for a range of λ . Table B.2 examines the distribution of our parameter estimates for each protocol (under repeated samples of the DGP) for different values of λ and Table B.3 shows the behaviour of our spread-of-predictions interval (Equation (3.19)) and midpoint prediction Equation (3.20) for different values of λ .

3.4.2 Case II: Misspecified dynamics

Next, we apply these methods to an example where we have misspecified the dynamics of the model (the function \mathbf{f} in our governing equation, Equation 3.21). We use two competing Markov models of hERG kinetics, the Beattie et al. (2018) model and the Wang et al. (1997) model. These models have differing structures and differing numbers of states, as shown in Figure 3.2. We generate a synthetic dataset under the assumption of the Wang et al. (1997) model with Gaussian IID noise (with standard deviation 30 pA) and the original parameter set as given in Wang et al. (1997), for all the protocols shown in Figure 3.3. As in Case I, we assume the standard deviation of this noise is known.

Then, we are able to fit both models to this dataset, obtaining an ensemble of parameter estimates and performing cross-validation as described in Section 3.3. In this way, we can assess the impact of choosing the wrong governing equations (our choice of \mathbf{f}), and its impact on the predictive accuracy of the model. We do this to investigate whether the techniques introduced in Section 3.3.1 are able to provide some useful quantification of model discrepancy when the dynamics of I_{K_r} are misspecified.

Our results, presented in Figures 3.6 and 3.7, show how we expect a correctly specified model to behave in comparison to a discrepant model. We can see from the bottom row of Figure 3.7, that when using the correctly specified governing equation, we were able to accurately recover the true maximal conductance using each and every protocol. Moreover, similarly to *Case I*, no matter which protocol the correctly specified model was fitted with, the resultant predictions were very accurate (as can be seen in Figure 3.6).

However, when the discrepant model was used, there was significant protocol dependence in our parameter estimates, and our predictions were much less accurate overall. Moreover, it seems that for the majority of d_0 , the spread in predictions across fitting protocols (Equation (3.19)) was smaller than those seen in Case I, but there are certain portions where the discrepant model and DGP are noticeably different (as highlighted in Figure 3.6). This may be due to the structural differences between the Wang et al. (1997) and Beattie et al. (2018) models. In particular, in the Wang et al. (1997) Model, channels transitioning from the high-voltage inactive state (**I**), must

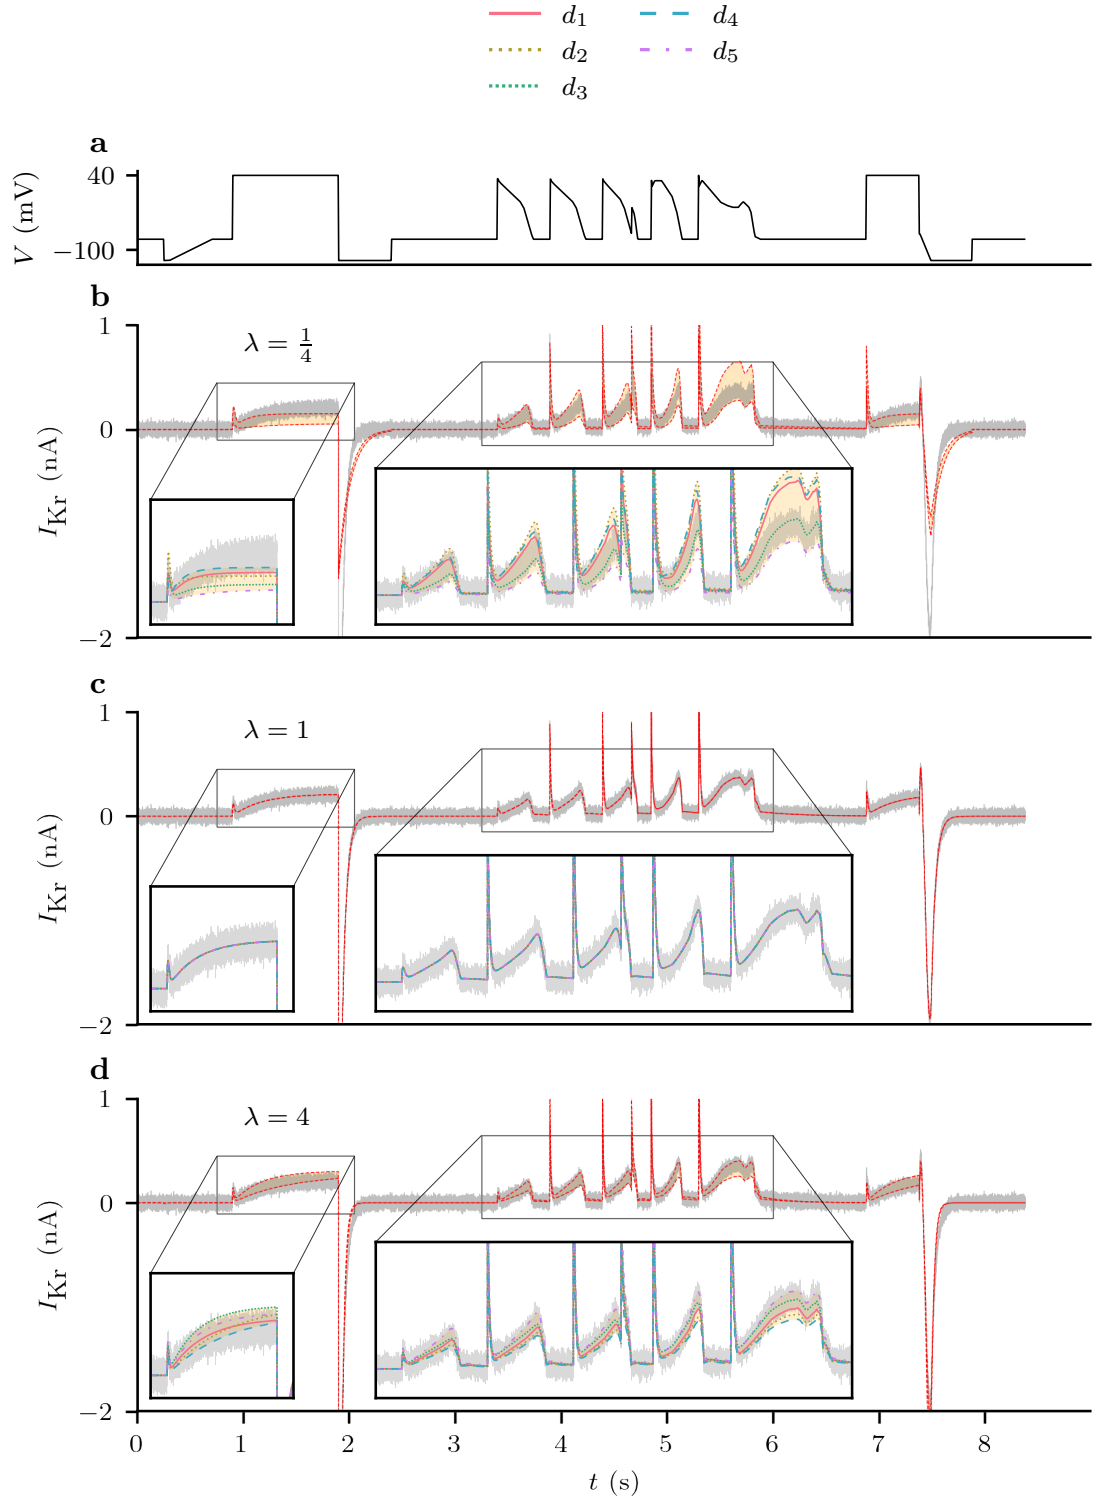


Figure 3.4: The set of predictions (Equation (3.18)) shown for parameter estimates obtained by fitting with different values of λ to synthetic data under d_1, \dots, d_5 (using the Beattie et al. (2018) model). The synthetically generated data used for model validation are shown in grey and the spread of the predictions is highlighted in yellow. Panel **a**: the voltage trace for d_0 . Panel **b**: The set of predictions with $\lambda = \frac{1}{4}$. Panel **c**: the set of predictions with $\lambda = 1$, that is, under the assumption of the correct maximal conductance (g). Panel **d**: the set of predictions with $\lambda = 4$.

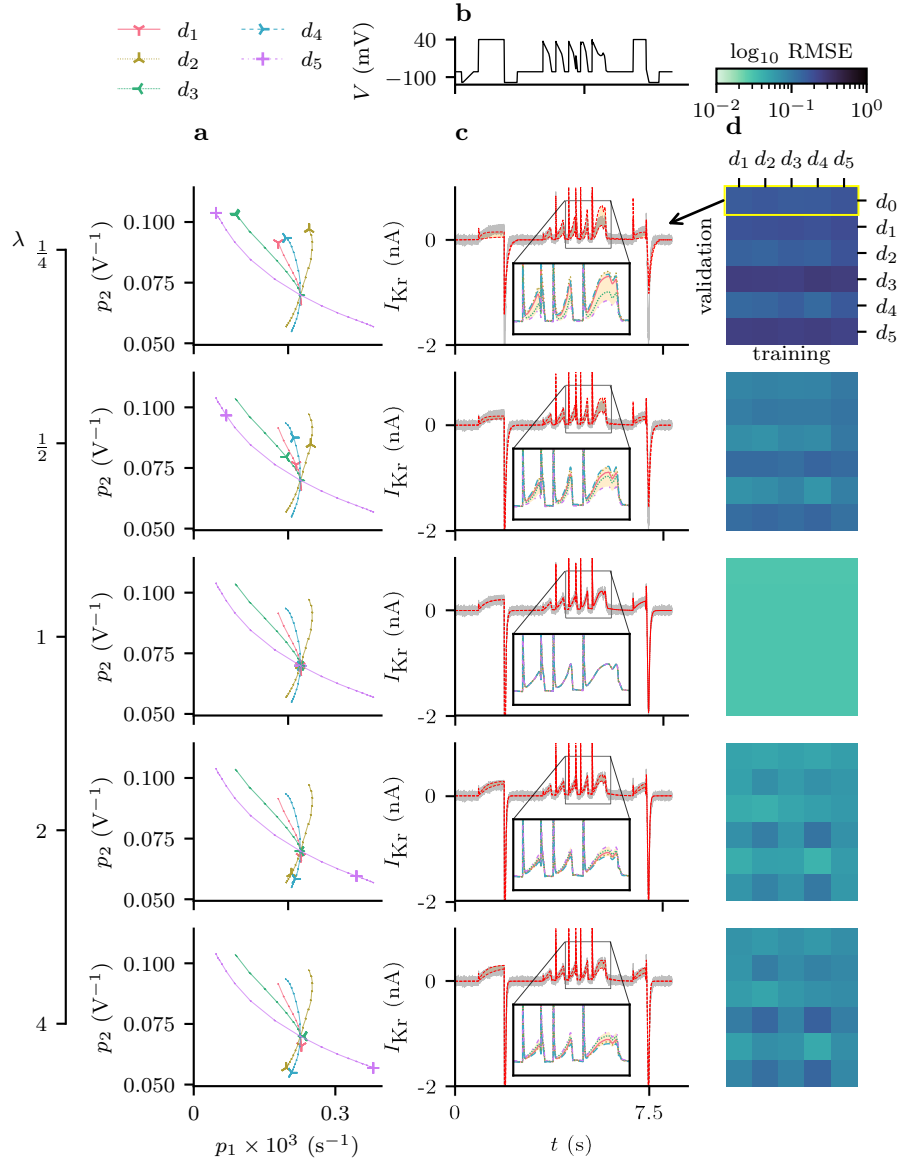


Figure 3.5: Discrepancy in parameter estimates and subsequent currents when a non-discrepant model is fit to synthetic data, with all parameters free except the maximal conductance, g , which is scaled by some factor λ , ($g = \lambda g^*$), where g^* is the true value. Panel **a**: estimates of θ_1 and θ_2 obtained by fitting with different protocols for 10 repeats of the DGP. The lines (linearly interpolated using 17 values for $\lambda \in [\frac{1}{4}, 4]$) show how the estimates from each protocol improve as $\lambda \rightarrow 1$, and how they disperse as λ increases and decreases (that is, when model discrepancy increases). Panel **b**: d_0 voltage protocol. Column **c**: the spread of predictions of I_{Kr} under the d_0 protocol using the parameter estimates in Column **a**. Column **d**: heatmaps showing the predictive error obtained by fitting and validating for each pair of protocols. Here, Column **c** corresponds with the top row of each heatmap, as indicated.

transition through the conducting, open state (**O**) in order to reach low-voltage closed states (**C1**, **C2**, **C3**), causing a spike of current. Instead, channels in the Beattie et al. (2018) Model may transition through the inactive-and-closed state (**IC**) on their way between **O** and **C**, resulting in reduced current during steps from high voltage to low voltage.

Nevertheless, our methods provide a useful indication of this model discrepancy. Figure 3.8, examines the behaviour of our prediction interval (Equation (3.19)) in more detail. Importantly, we can see that our interval shows little uncertainty during sections of the protocol where there is little current (this is also seen in Figure 3.7a and Figure 3.6). This is ideal, because no reasonable model would predict a sizeable current here. On the other hand, we see that our intervals show significant uncertainty around the spikes in current that occur at the start of each action-potential waveform. This is to be expected because it is known that these sections of the current time-series are particularly sensitive to differences in the ‘rapid inactivation’ process in these models (Clerx et al., 2019a).

Further results regarding Case II are provided in Section B.3. In particular, Tables B.4 and B.5 summarise the behaviour of our parameter estimates for each choice of model.

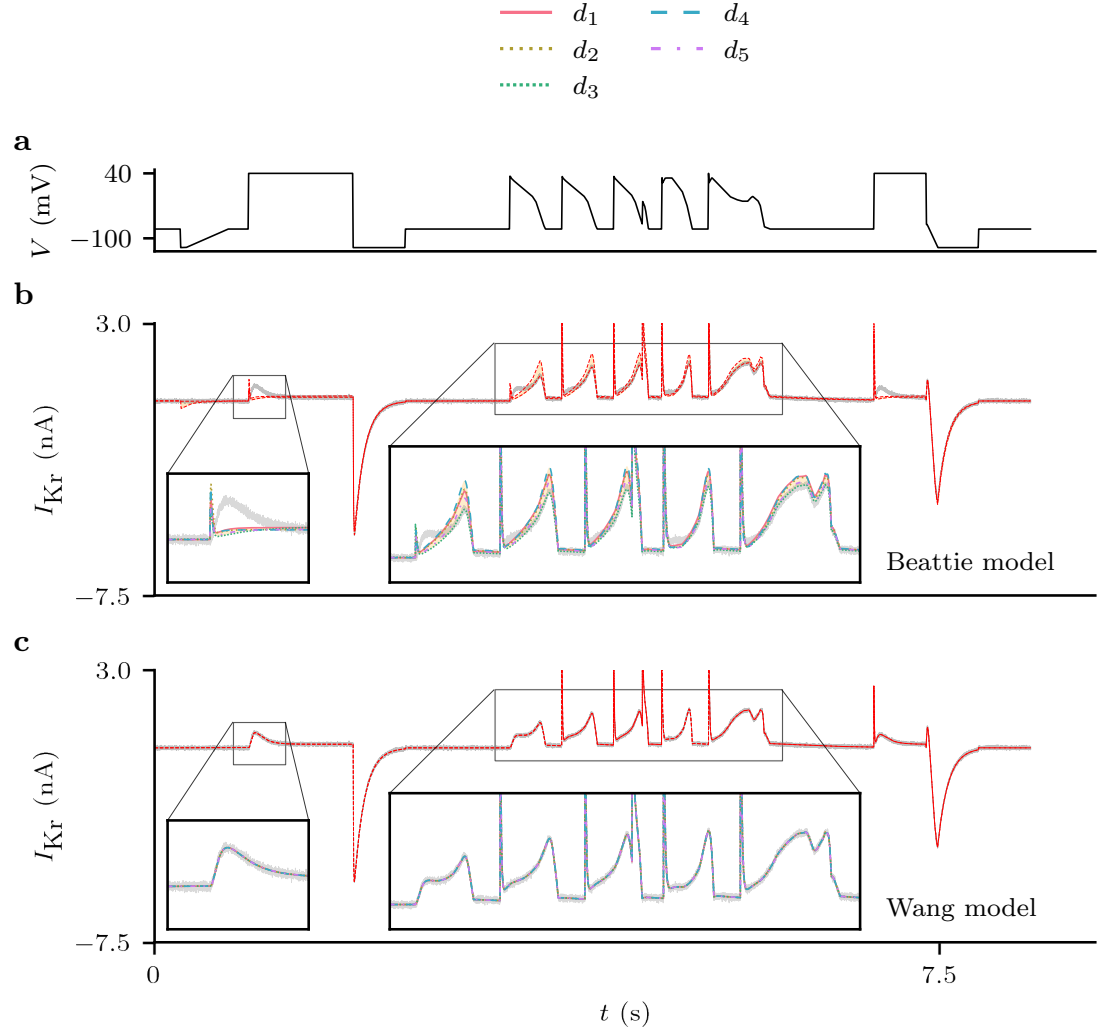


Figure 3.6: Case II: the set of predictions (Equation (3.18)) shown for parameter estimates obtained by fitting Beattie et al. (2018) and Wang et al. (1997) models with data synthetically generated using the Wang et al. (1997) Model. Panel **a**: the d_0 voltage-clamp protocol. Panel **b**: the set of predictions using the Beattie et al. (2018) Model. Panel **c**: the set of predictions with the Wang et al. (1997) Model, that is, with under the assumption of the correct model structure.

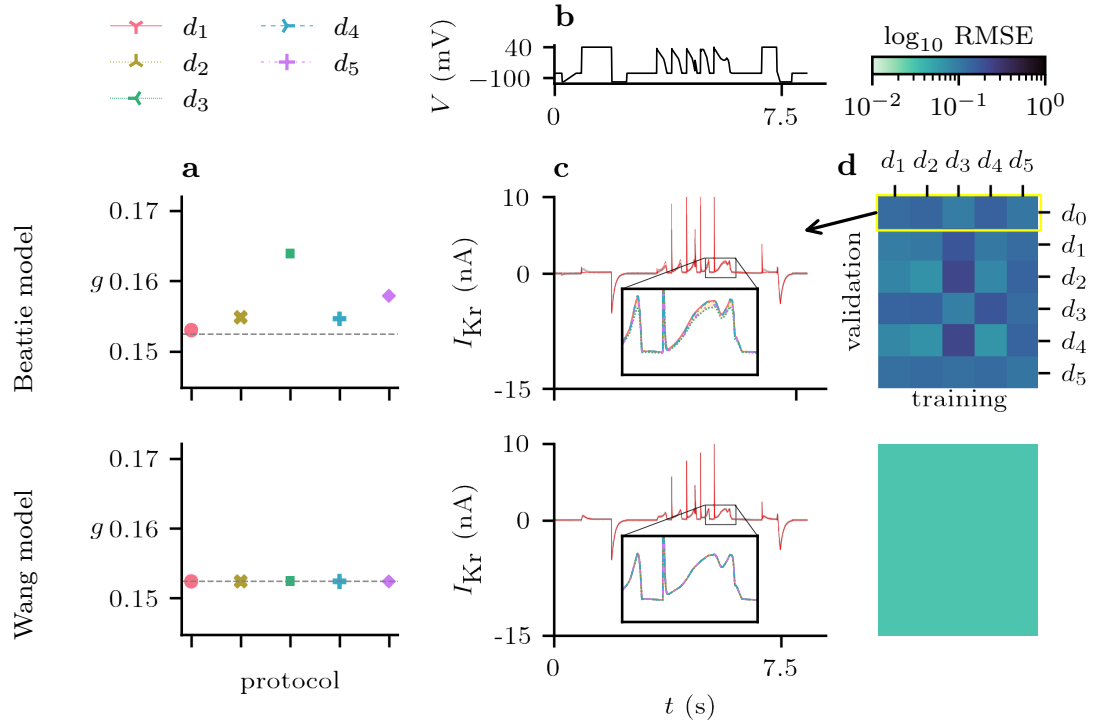


Figure 3.7: Model discrepancy between the Beattie et al. (2018) Model and synthetic data generated using the Wang et al. (1997) Model. Panel **a**: estimates of the maximal conductance obtained by fitting with different protocols for ten repeats of the DGP. There is a noticeable protocol dependence for estimates obtained using the (discrepant) Beattie et al. (2018) Model, but the true underlying parameter (dashed line) can be accurately determined from any protocol when using the (correct) Wang et al. (1997) Model. Panel **b**: d_0 voltage protocol. Panel **c**: the spread of predictions for I_{Kr} under the d_0 protocol for discrepant (Beattie et al., 2018) and correct (Wang et al., 1997) models, which are shown in more detail in Figure 3.6. Panel **d**: cross-validation heatmaps for both the Beattie et al. (2018) and Wang et al. (1997) models fitted to this suite of protocols, averaged over ten repeated samples of the DGP for each protocol.

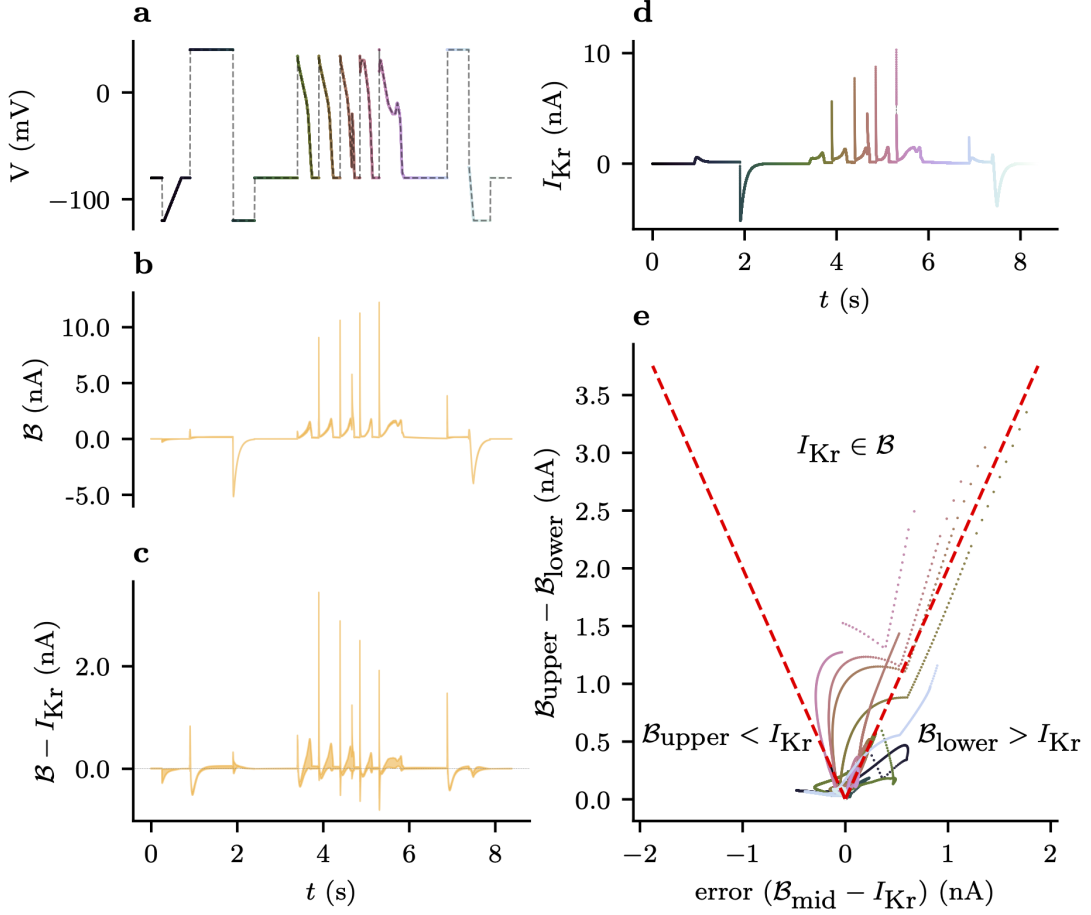


Figure 3.8: The spread in predictions obtained from different protocols provides a useful indicator of model discrepancy for *Case II*. Panel **a**: the validation voltage-clamp protocol, d_0 , with colours corresponding to Panels **d** and **e**. Panel **b**: the spread-of-predictions interval (Equation (3.19)) for d_0 using the Beattie et al. (2018) model fitted with d_1, \dots, d_5 . Panel **c**: the true DGP subtracted from the spread-in-predictions interval. Panel **d**: the true DGP with the colour of each observation corresponding to Panels **a** and **e**. Panel **e**: a scatter plot of the midpoint prediction (Equation (3.20)) and the width of the predictive interval (Equation (3.19)) for every observation in d_0 . Here, the red, dashed lines show the true value of I_{Kr} lies on the extremes of the range of predictions. Accordingly, points above these lines show the observations for which the DGP lies inside this range, and the points below the line correspond to observations for which the true DGP lies outside this range. The colours of these points correspond to those in Panels **a** and **d**.

3.5 Discussion

We have introduced an uncertainty quantification (UQ) approach to highlight when discrepancy is affecting model predictions. We demonstrated the use of this technique by providing insight into the effects of model discrepancy on a model of I_{Kr} in electrically excitable cells. Here, we saw that under synthetically constructed examples of model discrepancy, there was great variability between the parameter estimates obtained using different experimental designs. This variability is a consequence of the different compromises that a discrepant model has to make to fit different regimes of a true DGP's behaviour. Consequently, these parameter estimates produced a wide range of behaviour during validation, despite each individual parameter estimate having little variability under repeated samples of the DGP.

We propose that the variability in the model predictions stemming from this ensemble of parameter estimates is, therefore, an empirical way of characterising the predictive uncertainty due to model discrepancy. Usefully, our spread-of-prediction intervals (Equation (3.19)) correctly indicated little uncertainty when the ion channel model was exhibiting simple dynamics decaying towards a steady state, but more uncertainty during more complex dynamics, which was indeed when the largest discrepancies occurred. For many observations under our validation protocol, the true, underlying DGP lay inside this interval, indicating that Equation (3.19) may provide a useful indication of predictive error under unseen protocols. We expect that the presented methods may be of use for problems where the variability in parameter estimates (from repetitions of each individual protocol) is smaller than the variability between parameter estimates obtained from different protocols—because there is little noise, and lots of observations for example. In such cases, the variability in the extremes of our ensembles ($\mathcal{B}_{\text{upper}}$ and $\mathcal{B}_{\text{lower}}$) is immaterial compared to the width of the interval ($\mathcal{B}_{\text{upper}} - \mathcal{B}_{\text{lower}}$).




Figure 3.9: the Beattie et al. (2018) and Wang et al. (1997) models may be seen as special cases of this more complicated model. The **C3** state is labelled **C** in the Beattie et al. (2018) model and **C3** in the Wang et al. (1997) model. The arrows outside the Markov state diagram indicate the direction in which rates increase with more positive voltage.

At first, Case I may seem like an artificial example—in practice, the maximal conductance is taken to be a model parameter and fitted along with the rest of the model. But Case I and Case II are similar: any two Markov models may be regarded as two special cases of a more general model with some transition rates fixed to 0 (as shown in Figure 3.9 for the models used in this chapter). As in Case I, this means that different model structures can be seen as restrictions of this larger model’s parameter space. Misspecified model structures can then be identified with subsets of parameter space which do not contain the true, data-generating parameter set (provided this larger model is structurally identifiable).

This means there is a setting in which Case II (misspecified governing equations) is an example of the type of discrepancy explored in Case I, where a “true” parameter value exists in the more general model, but is excluded in the parameter space being optimised over when fitting the model. This may prove a valuable perspective for modelling ion channel kinetics, where there are many candidate models (Mangold et al., 2021), and each model may be seen as corresponding to some subset of a general model with a shared higher-dimensional parameter space. Model selection problems have been framed in this way previously (Ball et al., 1989; Akaike, 1998). In Chapter 5, we extend our models to include additional, experimental effects. The choice to set fix some of these additional model parameters to possibly biased estimates is analogous to Cases I and II, and is expected to induce similar model discrepancy, to some extent.

3.5.1 Limitations

Whilst the spread of predictions under some unseen protocol may provide some feasible range of predictions, we can see from Figure 3.8, that our observables (the DGP without noise) often lie outside this range. This is also shown in Figure 3.6 where certain structural differences between the model and DGP (the inclusion or exclusion of inactive-to-closed path which bypasses the open state) mean that no Beattie et al. (2018) Model parameter estimates are able to recapitulate the data. For this reason, Equation (3.19) is best interpreted as a heuristic indication of predictive uncertainty stemming from model discrepancy, rather than providing any guarantees about the output of the DGP. As shown in Figure 3.8, small errors between the median prediction \mathcal{B}_{mid} and the observed data mostly occur where the bound is narrow, and so large values $\mathcal{B}_{\text{max}} - \mathcal{B}_{\text{min}}$ are indicative of predictive inaccuracy (though small values provide no guarantees).

Using more training protocols in the training set may increase the coverage of the DGP by our interval. Whilst the number of protocols that can be performed on a single biological cell is limited by time constraints (Beattie et al., 2018; Lei et al., 2019b), the utilisation of more protocols is likely preferable.

Besides the examples of model discrepancy considered in Section 3.4, there are other ways that the DGP can differ from the fitted models. For example, the DGP may not be accurately described by an ODE system, especially when ion channel numbers are small and the stochasticity of individual channels opening and closing is apparent. In this circumstance, the models can be cast in terms of stochastic differential equations (SDEs), as in Goldwyn et al. (2011), and we can again consider an ensemble of parameter estimates (Equation (3.17)) and an ensemble of model predictions (Equation (3.18)). The assumption of IID Gaussian errors for the observation noise model could also be inaccurate. Here, auto-correlated noise processes (for example, as explored in Creswell et al., 2020; Lambert et al., 2023) could be considered, also. These effects could be included in the DGP, and it remains to be seen how well our method would perform in these cases.

3.5.2 Concluding remarks

The spread of predictions of our ensembles, based on fitting to data from multiple experimental designs, provides a good indication of possible predictive error due to model discrepancy. Ultimately, whilst our ensemble approach is no substitute for a correctly specified model, it is a useful tool for quantifying model discrepancy, predicting the size and direction of its effects, and may guide further experimental design and model selection approaches. The multiprotocol, cross-validation approach and model fitting methodology in this chapter is applied to real, experimental data in the following two chapters.

Chapter 4

Fitting an experimental, multiprotocol dataset

Preamble

In 2019 a selection of voltage-clamp protocols were designed and run in series in a “multiprotocol” experiment. However, the data quality was noticeably improved once Nanion’s “fluoride-free” plates were obtained. Whilst these plates seemed to lead to a decrease in success rate according to Nanion’s quality control procedures, we found them to be much more suitable for our purposes, and decided to re-run our multiprotocol experiments only in a fluoride-free set-up. My analysis of these datasets showed that a fluoride-free experimental approach was feasible, and would provide the high-quality data necessary for our multiprotocol approach. Following this analysis, the data used in this chapter were collected from an experiment by Monique Windley at the Victor Chang Cardiac Research Institute (VCCRI) in Sydney, Australia on the 25th of November 2022 (25/11/2022).

The same experiment (with a fluoride-free setup) was later repeated during my working visit to the VCCRI (November–December 2023), though the 25/11/2022 dataset remained the experiment with the highest success rate (according to the quality-control procedures introduced in this chapter). For this reason, our analysis focuses on this 25/11/2022, fluoride-free dataset, though work to extend this analysis to other datasets is ongoing. The majority of this work is included in a manuscript submitted for publication, which is currently available as a pre-print (Shuttleworth et al., 2024). Some details of the exact experimental methodology and equipment used are deferred to Shuttleworth et al. (2024). What remains in this chapter is my own work.

4.1 Introduction

In the previous chapter, we fitted ODE-based models (that is, Markov models) to synthetically generated data, and investigated the effect of model discrepancy on our parameter estimates and model predictions. In this chapter, we use the same methods to analyse *real experimental data* collected from patch-clamp experiments performed on an automated patch-clamp machine. Applying these methods to real experimental data (in place of synthetically generated data) presents a number of additional challenges. For instance, we must perform quality control (QC) to ensure that data from failed experiments is discarded (so that it does not confound our results), and we *postprocess* the raw data, resulting in the postprocessed traces of time-series data (where non- I_{Kr} currents are removed) used to fit our models. In this chapter, we introduce these QC and postprocessing methods, perform some exploratory analysis on our experimental data, and, finally, evaluate the suitability of some literature models of I_{Kr} .

Ideal-patch assumptions In this chapter, we apply our mathematical models under the assumption of *ideal patch-clamp* conditions (Lei et al., 2020a). The first of our ideal-patch assumptions is that the transmembrane voltage is always equal to the command voltage plus some constant, systematic offset, that is,

$$V_m = V_{cmd} + V_{off}, \quad (4.1)$$

where V_{off} is some (ideally small) systematic voltage offset. Our second ideal-patch assumption is that the current, after postprocessing, is entirely due to I_{Kr} , but subject to random observation errors. That is to say, we assume that the postprocessed trace used for model fitting and validation is,

$$I_{post}(t_i) = I_{Kr}(t_i) + \epsilon_i, \quad (4.2)$$

where t_i is our i^{th} observation time, and $\{\epsilon_1, \epsilon_2, \dots, \epsilon_n\}$, is a set of Gaussian IID random variables with mean zero, and n is the number of observations. The postprocessing methodology used to produce our postprocessed trace, I_{post} , are introduced in Section 4.3.

We can see that these ideal-patch assumptions do not perfectly align with our equivalent-circuit view of the model, which includes the capacitive effect of the cell membrane (as shown in Figure 2.5 in Section 2.1). In particular, the capacitive effect of the cell membrane is absent in models built under the ideal-patch assumptions (including those fitted in the previous chapter). However, the ideal-patch assumptions (Equations (4.1) and (4.2)) are commonly used to provide simple mathematical models of I_{Kr} and other ion-channel currents (Fink and Noble, 2009; Beattie et al.,

2018)—though these assumptions may not be stated explicitly. In the following chapter, Chapter 5, we relax this assumption and adapt our models to account for additional effects, such as the capacitive effect of the cell membrane. The work presented in this chapter forms a basis of comparison for these augmented models.

In the following section, Section 4.2, we briefly summarise the experiment methodology used to collect these data. During these experiments, we obtain data from a number of different *wells*, each of which is designed to apply each of our voltage-clamp protocols to a single cell, resulting in an independently observed current trace (as shown in Figure 4.1). Here, suction is used to catch and attach (ideally) a single cell into the patch-clamp circuit. To select suitable data for the fitting and validation of our models, we introduce a suite of methods for postprocessing and *Quality Control* (QC). These methods, introduced in Section 4.3.1 are necessary to discard poor-quality data from failed experiments.

As discussed in Chapter 2, the selection of an accurate model structure for I_{K_r} , which is suitable for general use, remains an open problem, and so we compare four literature models of I_{K_r} : the Beattie et al. (2018) and Wang et al. (1997) models explored in the previous chapter, as well as a simple three-state model and the Kemp et al. (2021) model (as introduced in Chapter 2). In Section 4.4 we detail the mathematical and statistical methodology used to fit these models. These methods are mostly similar to those presented previously in Chapters 2 and 3. The results of this model fitting allow a detailed comparison of these models’ predictive accuracy, and the variability of their parameter estimates under an ensemble of experiments (following the methodology introduced in Chapter 3). In addition to the predictions of our fitted models, we consider the variability of the corresponding parameter estimates in Section 4.5. This is done by introducing a linear, multivariate regression model which describes the dependence of parameter estimates on the well and protocol from which they were obtained. Standard statistical methods are then used to quantify the importance of the well- and protocol-dependent effects. This type of well-to-well variability has been explored previously in the literature, where it is suggested that artefact effects are a leading cause (Lei et al., 2020a).

4.2 Experimental methods

Our experiments were designed specifically to provide data for some aspects of our postprocessing methodology—the removal of leak currents, for example. These methods are necessary for the evaluation of our *Quality Control* (QC) criteria in Section 4.3. These criteria are used to identify and remove data from failed experiments, and are largely similar to those introduced in Lei et al. (2019a), albeit with some additional criteria for use with our multiprotocol approach.

Apparatus Experiments were performed on a Nanion SyncroPatch 384 at the Victor Chang Cardiac Research Institute, Sydney, Australia. The platform used (a Nanion SyncroPatch 384) uses a wellplate containing 384 separate “wells”, each of which allowing a single cell to be “patched”—that is, attached to the machine as shown in Figure 4.1. Once the cells are patched, multiple voltage-clamp experiments are performed in parallel, by simultaneously applying the same voltage-clamp protocol to each cell.




Figure 4.1: A diagram of a patch-clamp experiment performed on a high-throughput, automated patch-clamp platform. A seal is formed between the plate and the cell. The cell membrane is then punctured using automatically applied pressure. Here, the removal of a miniscule section of cell membrane allows the formation of an electrical circuit, where current flows from the inside of the cell, through the membrane, to the amplifier where it is recorded. Figure modified, with permission, from Lei et al. (2020b).

During an experiment performed on the Nanion SyncroPatch, many cells are deposited into each of the 384 wells. Negative pressure is then used to attach a single cell to a hole in each well. Once a cell is attached to the wellplate, an increase in pressure punctures a small section of the cell membrane and allows the patch-clamp circuit to be completed (this circuit is shown in Figure 4.1). Here, an imperfect seal is formed between the cell membrane and the well plate, allowing some amount of leak current, I_L , to flow.

In our equivalent circuit, this leak current is modelled using an Ohmic resistor (Equation (4.3)) that is, a resistor with constant resistance/conductance. Here, the strength of this seal is characterised by the seal resistance, $R_{\text{seal}} = \frac{1}{g_L}$. In manual patch-clamp experiments, a gigaohm seal, that is, $R_{\text{seal}} \geq 1 \text{ G}\Omega$, is considered desirable (Sakmann and Neher, 1984), but automated patch-clamp experiments typically result in reduced seal quality—that is, decreased R_{seal} and, equivalently, increased g_L (Lei et al., 2019a).

In experiments using the Nanion SyncroPatch, CaF_2 is often used as a *seal enhancer*—a compound designed to increase R_{seal} . However, it has been shown that the presence of CaF_2 complicates the leak current, introducing a nonlinear component (Lei

et al., 2020b). In fluoride-free experiments, a different preparation is used, which may allow suitably good seals to be formed without the need for CaF_2 . We used such an approach to avoid the introduction of additional model discrepancy through the presence of nonlinear and time-dependent leak current dynamics.

Solutions replicating Nanion’s standard External NMDG 60 and Internal K Gluconate 120 were used, as the potassium concentrations roughly correspond to physiological concentrations. However, our experiments were performed at room temperature (taken to be 25°C). The SyncroPatch’s cooling fans were used to actively maintain this temperature throughout the experiment. Further details of the experimental methodology and equipment used are provided in Shuttleworth et al. (2024).

Pharmacological Isolation of I_{Kr} current After applying each voltage protocol to our cells, we add *dofetilide* at $1\ \mu\text{M}$ (a concentration known to almost completely block I_{Kr}) and repeat each protocol in the same order as shown in Figure 4.3. By performing leak correction and subtracting the post-drug, leak-corrected trace from the pre-drug leak-corrected trace, we are able to isolate I_{Kr} with minimal contamination from any endogenous or other unwanted currents. These postprocessing methods are explained in Section 4.3.

4.2.1 Design of voltage-clamp protocols

In this chapter, we employ the methodology of Chapter 3, using an ensemble of voltage protocols to investigate the dynamics of I_{Kr} in CHO cells. To do this, we perform a sequence of protocols on the SyncroPatch 384, such that for each well we obtain data under the application of each protocol. Provided that the experiment is successful, this multiprotocol approach provides us with repeats of all protocols across multiple wells, allowing us to calibrate multiple cell-specific models for each protocol. In particular, for any given model structure, we obtain a vector of parameter estimates for each repetition of each protocol in each well—that is, for each *sweep*.

To ensure the consistency of our I_{Kr} recordings over the course of the experiment, we repeat one of the protocols, Lei et al.’s *staircase* protocol (Lei et al., 2019b) four times—twice at the beginning of the experiment and twice again after all other protocols. All other protocols are performed once only, as to minimise the total duration of the experiment.

Leak ramp As shown in Figure 4.1, the seal between the wellplate and the cell is imperfect, resulting in a small leak current. We assume that this current is of the form,

$$I_{\text{L}} = g_{\text{L}}(V_{\text{m}} - E_{\text{L}}), \quad (4.3)$$

where g_L is a constant *leak conductance* and E_L is the *leak reversal* (that is, the transemembrane potential at which there is no leak current). Our protocols are designed to allow the easy estimation of these leak-model parameters (g_L and E_L), and the subsequent subtraction of this current from our traces. This is achieved by the inclusion of a *leak ramp*, as described below. We refer to estimation and subtraction of the leak current from our raw, experimental traces as *leak correction*.

Before each protocol, and for 250ms at the beginning of each protocol, V_{cmd} is fixed to the *holding potential*, -80 mV. The leak ramp follows, beginning with a 50ms, -120 mV step followed by a gradual increase (over 400ms) back to the -80 mV holding potential. Because I_{K_r} is negligible in this range of voltages, the observed current can be used to fit our leak model. The use of this segment to infer our leak-model parameters (g_L and E_L) is described in Section 4.3.1.

Reversal ramp As in Lei et al. (2019a), we include a *reversal ramp* at the end of each protocol. This is intended to allow the inference of I_{K_r} 's reversal potential (denoted by E_{K_r} in our models). This reversal-ramp segment is preceded by a 500ms long $+40$ mV conditioning step, after which the command voltage is set to -70 mV. Then, over the course of 100ms, the command voltage rapidly decreases to -110 mV. This part of the protocol, the reversal ramp, is useful, because the conditioning step ensures a sizeable current whilst V_{cmd} steadily decreases. During this decrease in V_{cmd} , the sizable current reverses (goes from positive to negative), allowing us to observe the time, t^* , and hence command voltage ($V_{\text{cmd}}(t^*)$), at which the postprocessed trace, I_{post} , is zero. We denote the observed reversal potential by $E_{\text{post}} = V_{\text{cmd}}(t^*)$, and estimate this quantity using polynomial interpolation. This method is described in further detail in Section 4.3.1.

The gradient of this ramp-like section must be steep enough such that the reversal potential is reached before too many channels close. However, must not be so gradual as to allow too many channels to close before the current reverses. This reversal ramp employs the same design as the one used in (Lei et al., 2020a). Further details and examples of this reversal-potential inference are presented in Section 4.3.1.

Design of individual voltage-clamp protocols A range of information-rich voltage protocols were applied sequentially to each well. The differences lie in the specified 'command voltage', V_{cmd} , that is, the voltage the amplifier is instructed to clamp the membrane potential to at each time point during the experiment. These protocols were developed using a range of techniques, as detailed in Lei et al. (2024) and summarised in Table 4.1. Additionally, each protocol is shown in Figure 4.2. We briefly describe the rationales for their designs again here. Note that the labelling of protocol used in this chapter differs from Chapter 3, where we used only a subset of these protocols.

Lei et al. (2024) designed a number of these protocols via the numerical optimisation of various objective functions (Mirams et al., 2024; Lei et al., 2024). In particular: protocols d_3, d_8 and d_9 were designed using the *phase-space-filling* approach described in Mirams et al. (2024); protocols d_{10} and d_{12} designed to maximise Sobol sensitivities (Lei et al., 2024; Sobol, 2001) of the Wang et al. (1997) and Beattie et al. (2018) models, respectively; protocols d_4 and d_5 were found using a brute-force approach to maximise the sensitivity of model output to changes in parameters for the Beattie et al. (2018) and Wang et al. (1997) models, respectively; protocol d_7 was found by considering 3-step blocks, randomising the durations of each step and optimising the voltages; whereas, protocol d_{11} was found by randomising the voltages and optimising the durations.

The remaining three protocols (d_1, d_2 and d_6) were designed manually, without the use of an algorithm. Lei et al. (2019b) *staircase* protocol, d_1 was shown to permit the estimation of transition-rate parameters in models of I_{Kr} . A similar protocol, d_2 is included because it includes a central section in which there are many short-duration segments. We expect that this protocol highlights more of I_{Kr} 's short-timescale behaviour (namely the inactivation/recovery-from-inactivation process, which occurs very rapidly). Finally, d_6 was performed without the intention of providing useful parameter estimates and, in fact, our models are *practically unidentifiable* under d_6 . Nevertheless, this protocol consists of a sequence of action-potential voltage traces and, hence, provides physiologically relevant data for the purpose of model validation.

Protocol	duration (s)	# segments	frequency (kHz)	codename
d_1	15.0	32	5	<i>staircase</i>
d_2	15.0	52	5	<i>staircase-in-staircase</i>
d_3	18.8	64	5	<i>spacefill19</i>
d_4	32.9	58	2	<i>hhbrute3gstep</i>
d_5	22.5	58	2	<i>wangbrute3gstep</i>
d_6	8.4	63	5	<i>longap</i>
d_7	21.4	43	5	<i>rtov</i>
d_8	17.0	64	5	<i>spacefill10</i>
d_9	14.6	64	5	<i>spacefill26</i>
d_{10}	32.6	58	2	<i>wangsobol3step</i>
d_{11}	21.1	43	2	<i>rvot</i>
d_{12}	22.4	40	5	<i>hhsobol3step</i>

Table 4.1: The duration of, number of segments and sampling frequency of each protocol included in our dataset. During these segments, the command voltage is either held constant, or changes according to some constant gradient (which we call a ramp). Here, the number of segments includes segments which are common to all of our protocols—the leak ramp and reversal ramp, and their associated +40 mV preconditioning steps.

The protocols described above were performed sequentially before and after the application of a 1 μ M dose of dofetilide, a drug known to specifically block I_{Kr} at




Figure 4.2: The protocols used in our experiment, shown in the order that they are applied. Common features present at the start and end of each protocol are used for postprocessing. All protocols are repeated exactly once except protocol d_1 , Lei *et al.*'s *staircase* protocol. All protocols are used for model fitting except Lei *et al.*'s *longap* protocol, d_6 , which is used only for validation.

moderate concentrations when compared with other drugs (Mirams et al., 2011). For QC, we perform four repeats of the *staircase* protocol, d_1 , allowing us to discard data from wells that do not remain stable over the course of the experiment. In particular, the *staircase* protocol, d_1 , was repeated four times, twice at the beginning of the experiment and twice after all other protocols. All other protocols were performed exactly once before, and once after, the addition of dofetilide. This procedure is illustrated by the schematic shown in Figure 4.3.

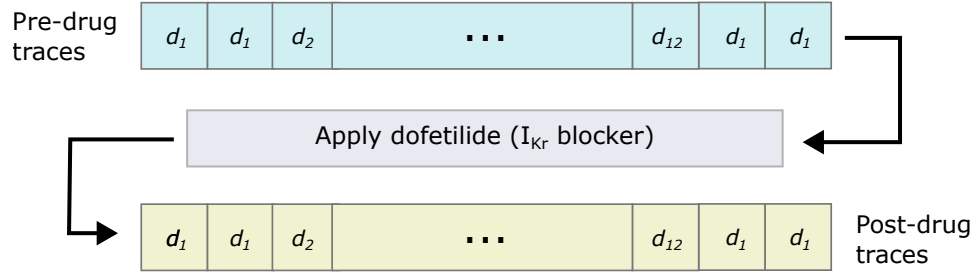


Figure 4.3: A cartoon showing the order in which protocols were performed, our post-drug repeats. Each protocol is run before and after the addition of dofetilide, an I_{Kr} blocker. All protocols are performed once before and once after the addition of drugs except d_1 , which is repeated eight times in total.

4.3 Postprocessing

4.3.1 Methods

Leak correction Leak correction is performed by first inferring the parameters of our leak-current model (Equation (4.3)) and then computing I_L . This predicted leak current is then subtracted from our observations, resulting in traces where the presence of leak current has been removed. Because I_{Kr} is typically small for the range of voltages used in the leak-ramp section, it is possible to use these observations to infer g_L and E_L independently of I_{Kr} . Hence, we assume that our observed current during the leak ramp is entirely leak current, that is,

$$I_{\text{obs}} = I_L + \varepsilon, \quad (4.4)$$

for the duration of the leak ramp where $\varepsilon \sim \mathcal{N}(0, \sigma)$ for some known standard deviation σ . As usual, we assume that the observation errors ε are independent of each other.

The constant conductance of our leak-current model (Equation (4.3)) means that, when $I_{Kr} = 0$, there is a linear relationship between the command voltage V_{cmd} and I_{obs} , the resulting current. Thus, it is natural to use simple linear regression to fit our leak-model parameters, writing,

$$I_L = \beta_0 + V_{\text{cmd}}\beta_1, \quad (4.5)$$

where β_0 and β_1 are our regressors—parameters to be estimated.

As is standard, we estimate these parameters via simple linear regression by computing the unbiased estimators (Jammalamadaka, 2003),

$$\hat{\beta}_0 = \sum_{i=1}^{n_{\text{leak}}} \bar{I}_{\text{obs}} - \beta_1 \bar{V}_{\text{cmd}}, \quad (4.6)$$

$$\text{and } \hat{\beta}_1 = \frac{\sum_i^{n_{\text{leak}}} (I_{\text{obs}} - \bar{I})^2 (V_{\text{cmd}} - \bar{V}_{\text{cmd}})}{\sum_{i=1}^{n_{\text{leak}}} (V_{\text{cmd}} - \bar{V}_{\text{cmd}})^2}, \quad (4.7)$$

where n_{leak} is the number of observations made during the leak ramp, and \bar{V}_{cmd} and \bar{I} are the respective averages of V_{cmd} and I_{obs} taken over the course of the leak-ramp section of the protocol. For ease of interpretation, we map these leak-model estimates (Equations (4.6) and (4.7)) to estimates of our leak parameters by writing $\hat{g}_L = \hat{\beta}_1$ and $\hat{E}_L = -\frac{\hat{\beta}_0}{\hat{\beta}_1}$. We then compute the *leak-corrected* trace,

$$I_{\text{corrected}} = I_{\text{obs}} - I_L, \quad (4.8)$$

$$= I_{\text{obs}} - \hat{\beta}_1 V_{\text{cmd}} - \hat{\beta}_0 \quad (4.9)$$

$$= I_{\text{obs}} - \hat{g}_L (V_{\text{cmd}} - \hat{E}_L), \quad (4.10)$$

Though it would be possible to calibrate this model using only two distinct V_{cmd} values—sometimes referred to as a “leak pulse” (Li et al., 2017a)—using a leak ramp allows our linear leak-current model to be validated across a range of voltages. This approach, referred to as *leak-correction*, is demonstrated in Figure 4.4. Note that when $V_{\text{off}} \neq 0$, our estimate of the leak reversal, \hat{E}_L , is instead an estimate of $E_L + V_{\text{off}}$ in accordance with Equation (4.3). In this way, the resulting, fitted leak current, I_L , is completely independent of V_{off} , and so we simply perform leak correction under the assumption that $V_{\text{off}} = 0$ and adjust our estimate of E_L as necessary.

The +40 mV step immediately following the leak ramp is also useful for validation because during this step, I_{K_r} activates slowly. This means that I_{K_r} remains small for some time, during which the observed current, I_{obs} , mostly consists of I_L . Hence, the observation of a significantly negative current here would indicate that the leak has been “overcorrected”. Checking for such leak overcorrection forms part of our QC, as outlined below.

Drug subtraction Although our leak-correction method allows us to infer and correct for linear leak currents, there may be biological currents (other than I_{K_r}) that are active in our cells. These currents are called *endogenous* currents, and we wish to minimise their presence as to avoid the pollution of our data.

As mentioned in Section 4.2, we apply dofetilide to our cells and repeat each of our protocols after I_{K_r} has been almost completely blocked. We use these post-drug traces to

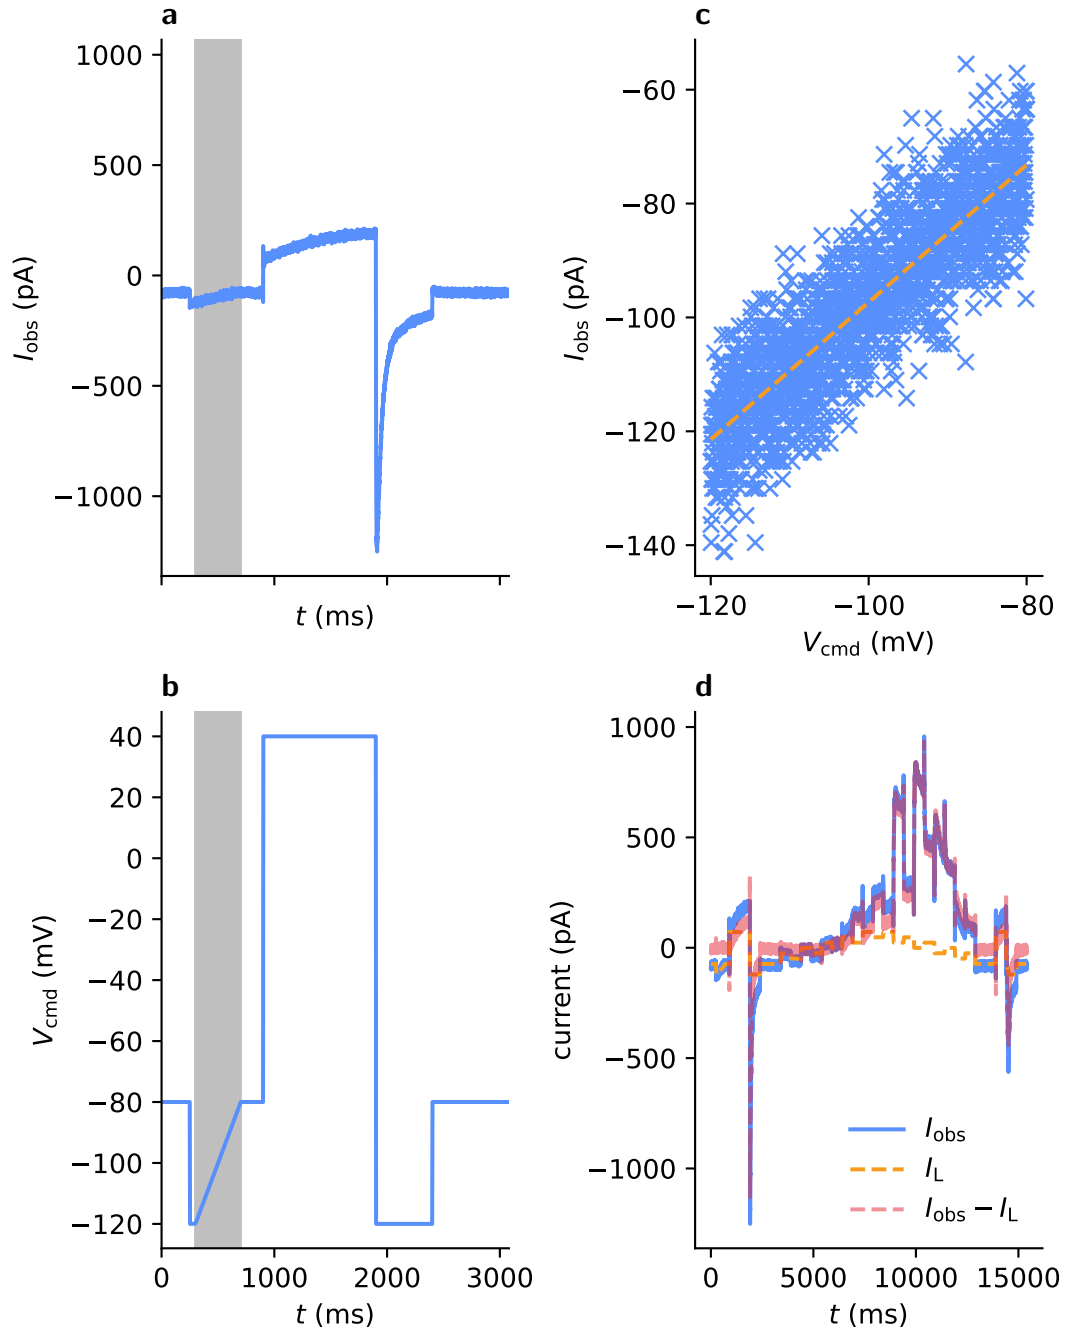


Figure 4.4: Our linear leak-current model is fitted to data from Well B20 using simple linear regression. Panels **a** (showing the observed current) and **b** (showing the command voltage) highlight the leak ramp. Panel **c** shows the simple linear regression used to fit the linear leak model. Panel **d** shows the observed current (I_{obs}), the leak current (I_{L}) and the leak-corrected current ($I_{\text{obs}} - I_{\text{L}}$). This data was taken from the application of the *staircase* protocol to Well B20 (the first repeat). Note how there is little leak-corrected current at the beginning and end of the experiment where I_{K_r} is small.

remove endogenous currents from our data. However, we must be careful to account for leak current, because the pre- and post-drug repeats are performed approximately 10 minutes apart, during which time we may observe some changes in the experimental conditions. For instance, it is plausible that the quality of the seal would degrade over time, causing g_L , and hence, the amount of leak current, I_L , to increase.

Hence, to compute our postprocessed trace I_{post} , we subtract the leak-corrected pre-drug trace, $I_{\text{obs}}^{(\text{after})}$, from the leak-corrected post-drug trace $I_{\text{obs}}^{(\text{before})}$,

$$I_{\text{post}} = \left(I_{\text{obs}}^{(\text{before})} - I_L^{(\text{before})} \right) - \left(I_{\text{obs}}^{(\text{after})} - I_L^{(\text{after})} \right). \quad (4.11)$$

We then use I_{post} to fit and validate our models.

We may quantify the amount of leftover nuisance current by considering the relative size of our leak-corrected traces before and after the addition of dofetilide,

$$R_{\text{leftover}} = \frac{\left\| I_{\text{obs}}^{(\text{after})} - I_L^{(\text{after})} \right\|_2}{\left\| I_{\text{obs}}^{(\text{before})} - I_L^{(\text{before})} \right\|_2}. \quad (4.12)$$

Under our ideal patch-clamp assumptions, we expect that R_{leftover} is small (though non-zero due to the presence of random noise). This statistic provides a useful quantification of the relative size of current that remains after drug subtraction and leak correction. Such leftover current may be due to endogenous currents, nonlinear leak, or miscalibration of the leak model (perhaps due to data pollution, or changes in the quality of the seal over time).

Reversal potential inference Recall from Chapter 2 that our models of I_{K_r} contain an observation function of the form,

$$I_{K_r} = g x_O (V_m - E_{K_r}),$$

where V_m is the transmembrane potential, E_{K_r} is the current's reversal potential, g is the maximal conductance (a model parameter), and x_O is the proportion of $K_v11.1$ channels in the (sole) open conformation—found by solving our models governing ODE (Equation (2.18)).

We typically assume that the reversal potential, E_{K_r} , is some known quantity, and hence, kept constant during model fitting. In Beattie et al. (2018), for example, E_{K_r} is set to the Nernst potential which is calculated as in the previous chapter,

$$E_{\text{Nernst}} = \frac{RT}{F} \ln \left\{ \frac{[K_{\text{out}}]}{[K_{\text{in}}]} \right\}, \quad (4.13)$$

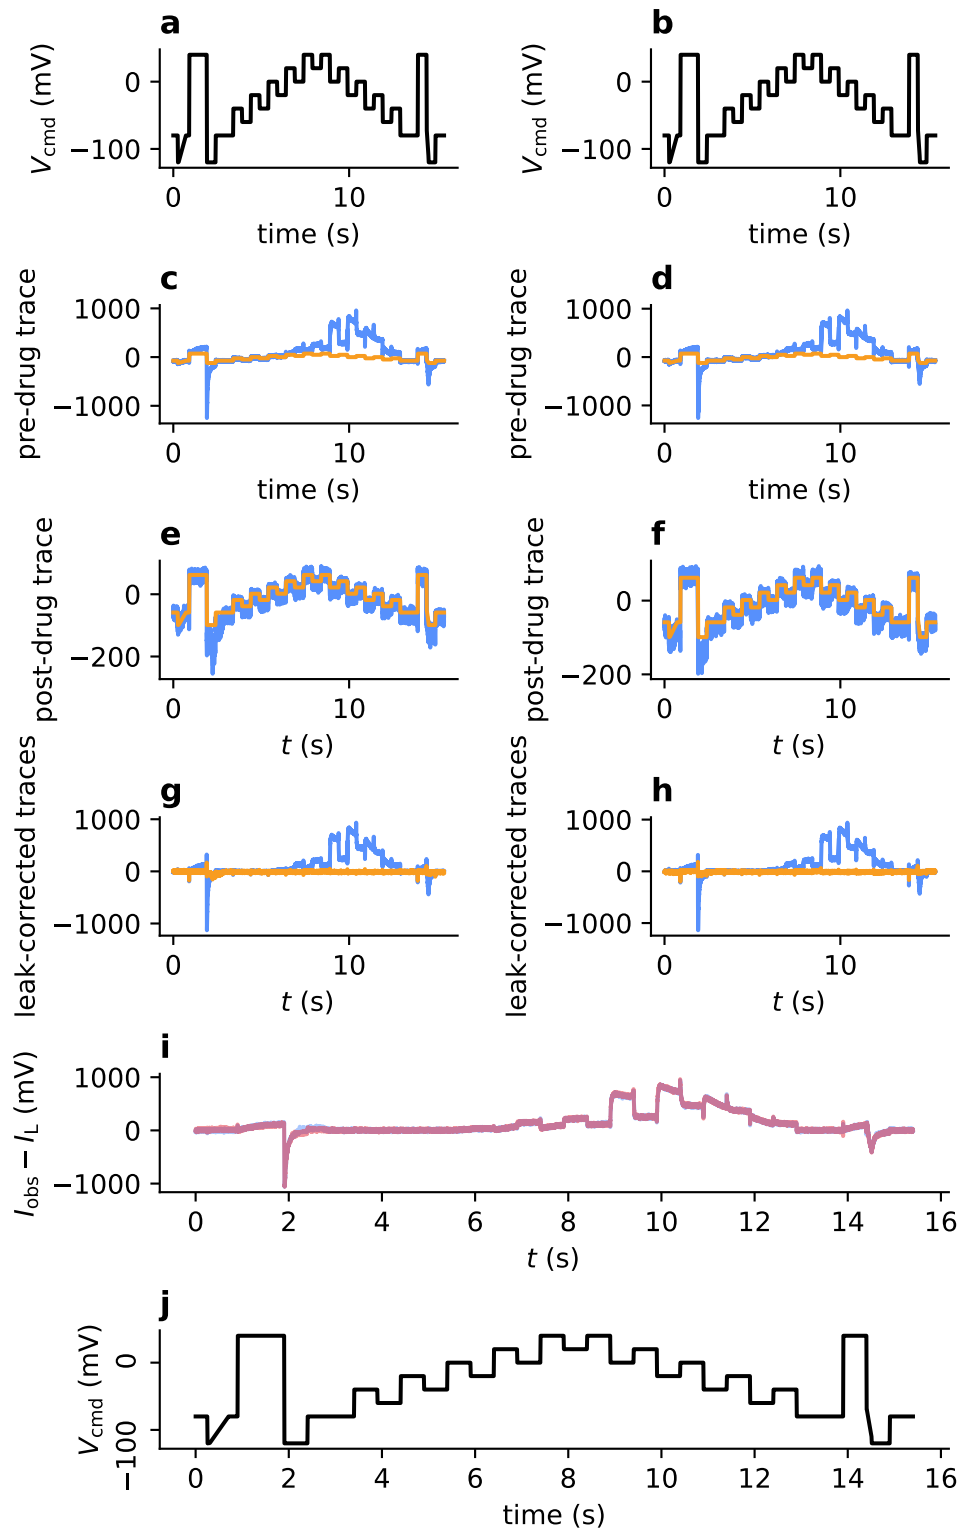


Figure 4.5: Our postprocessing methods, leak correction and post-drug subtraction, shown for data obtained from B20 under the application of the *staircase* protocol (first and second sweeps). The fourth row shows the leak-corrected post-drug trace in orange, and the leak-corrected pre-drug trace in blue. All currents are shown in picoamperes (pA).

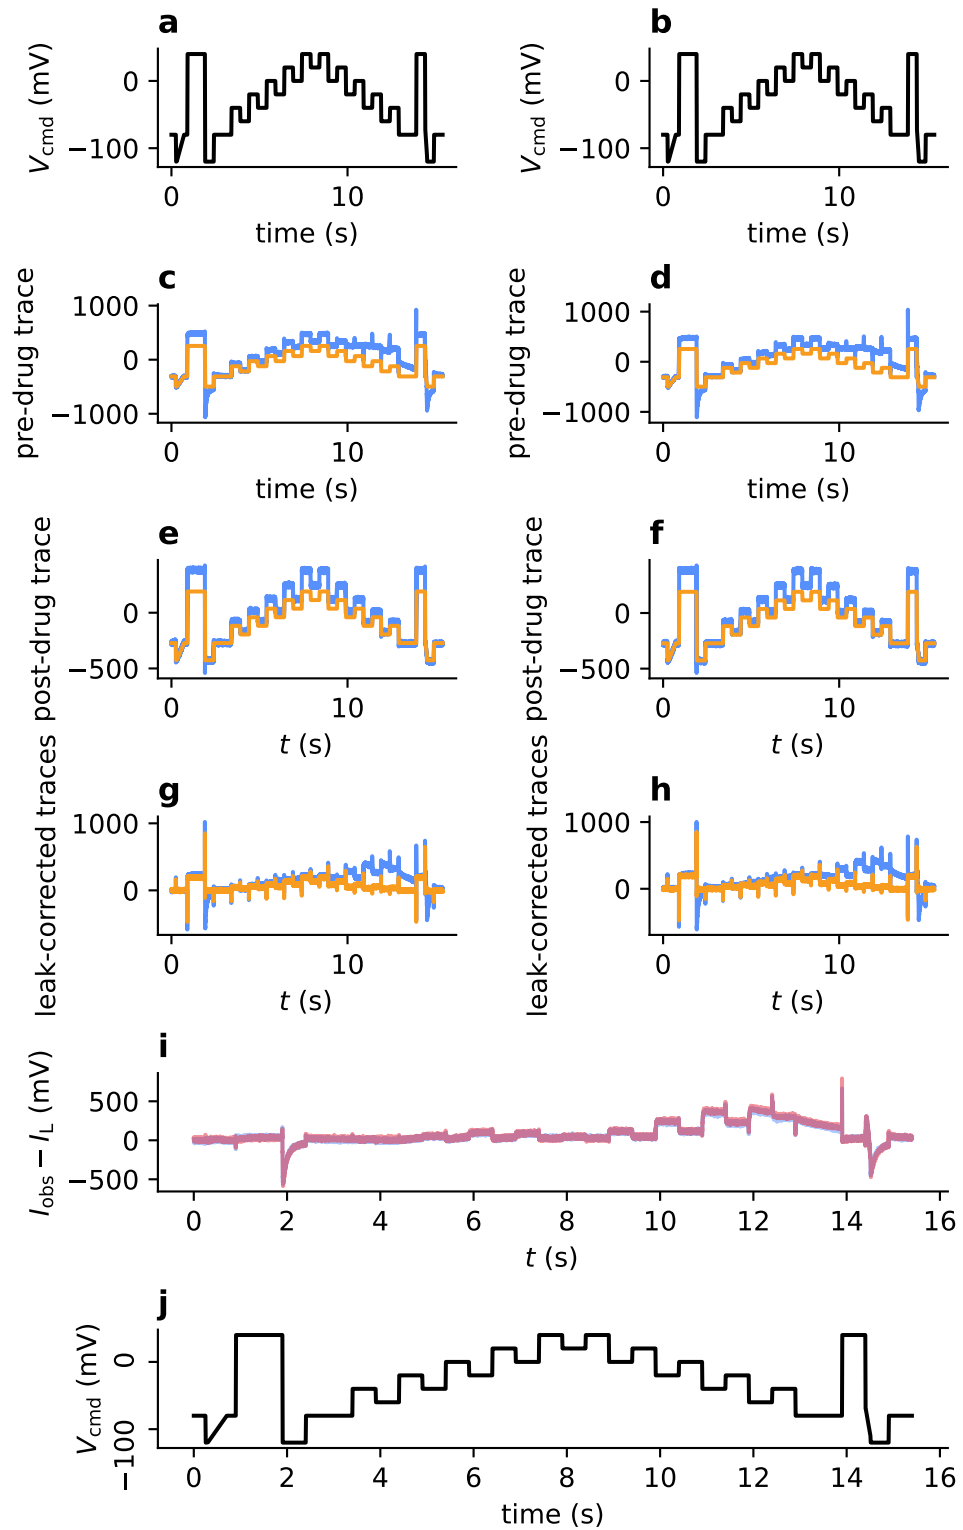


Figure 4.6: Leak correction and drug subtraction, shown for data collected from M06 under the application of the *staircase* protocol (first and second sweeps). This was discarded because of the large leak-corrected post-drug traces, as shown in panels **g** and **h**. All currents are shown in picoamperes (pA).

where $[K_{\text{out}}]$ denotes the extracellular potassium concentration, $[K_{\text{in}}]$ denotes the intracellular potassium concentration, R is the gas constant, F is Faraday's constant, and $T = 298.15 \text{ K} = 25^\circ\text{C}$ is the temperature at which our experiments were performed (Lei et al., 2019a), and E_{Kr} is the transmembrane potential at which, according to the model, K^+ ions do not flow through the channel in either direction (see Equation (4.25)). Using the known potassium concentrations of our intracellular and extracellular solutions, we find $E_{\text{Nernst}} \approx -89.83 \text{ mV}$.

However, it is also possible to estimate E_{Kr} directly from our time-series data. The reversal-ramp segment (Lei et al., 2019a) was included in each protocol to facilitate estimation of E_{Kr} . Following leak correction and drug subtraction, we estimate the reversal potential by fitting a order-4 polynomial to I_{post} , and use this polynomial to identify the time, t^* , at which $I_{\text{Kr}}(t^*) = 0$. We then let $E_{\text{post}} = V_{\text{cmd}}(t^*)$ denote the *observed* reversal potential. Figure 4.8(a) shows how τ (the time-constant in our single exponential model) varies over the course of the experiment, in each well. From this figure, we can see that for all but one well, $E_{\text{post}} \approx -80 \text{ mV}$ is consistently observed, which is markedly different to our computed value of $E_{\text{Nernst}} \approx -89.8 \text{ mV}$. Three separate approaches for fitting our models in light of this discrepancy are discussed in Section 4.4.

Time-course variability To quantify differences between sweeps, we may use the time-course of decay after the first $+40 \text{ mV}$ step in our protocols. Here, after the command voltage drops to -120 mV , typically we see a large negative which slowly decays. By fitting an exponential curve through this section of the sweep (a non-leak-corrected, pre-drug trace), we can quantify the rate of this decay, and compare this value between sweeps.

To do this, we find scalar parameters, α , and $\tau > 0$) such that,

$$\sum_{i=i_{\text{start}}}^{i_{\text{end}}} \left(I_{\text{obs}}^{(i)} - \alpha \exp \left\{ -\frac{1}{\tau} (t_i - t_{\text{peak}}) \right\} \right)^2, \quad (4.14)$$

is minimised, where t_{peak} is the time at which the peak current is observed, and i_{start} and i_{end} are the indices of the first and last observations considered during the fitting of the tail current. Here, the characteristic time of decay is, τ , which describes the time taken for the current to decay. We choose i_{peak} to be the index where the maximal negative current is observed, and i_{end} such that the last datum considered is 50 ms before the end of the -120 mV segment.

Figure 4.7 highlights the relevant section of our protocols, and shows these parameters fitted to an example trace. Figure 4.8(a) shows how this quantity varies over the course of the experiment, in each well. From this figure, we can see that there is a noticeable degree of variability between wells, but the values from each well remain

mostly consistent (with the notable exception of Well C12). This indicates that (broadly speaking), the dynamics of I_{Kr} (for a given well) are consistent throughout the experiment. These figures do, however, indicate variability between wells. Similarly, we show the variability of the peak current during the second -120 mV step in Figure 4.8(b).

It is also possible to fit two decay time constants to this current (say, τ_1 and τ_2). Though a two-exponential model may better fit the data, our single-exponential model results in a single statistic describing the rate of decay, allowing us to more easily summarise our dataset. Moreover, using a single exponential helps avoid practical identifiability in some cases where a single exponential appears would be more suitable. For comparison, a double-exponential fit is shown in Figure 4.9.

Now, consider the size of a leak current, I_L (either the pre-drug leak current, $I_L^{(\text{before})}$, or the post-drug leak current, $I_L^{(\text{after})}$). This depends on: the leak reversal E_L , the leak conductance, g_L , and the command voltage, V_{cmd} . The latter corresponds to the particular voltage-clamp being performed. Hence, the size of the leak current varies between protocols. So, in order to allow a simple, protocol-independent comparison of leak current sizes, we consider the average size of the squared leak current across our voltage range,

$$\bar{I}_L^2 := \frac{1}{V_{\text{max}} - V_{\text{min}}} \int_{V_{\text{min}}}^{V_{\text{max}}} g_L^2 (V - E_L)^2 dV \quad (4.15)$$

$$= \frac{g_L^2}{3(V_{\text{max}} - V_{\text{min}})} ((V_{\text{max}} - E_L)^3 - (V_{\text{min}} - E_L)^3), \quad (4.16)$$

where we choose $V_{\text{max}} = +40$ mV $+ V_{\text{off}}$ and $V_{\text{min}} = -120$ mV $+ V_{\text{off}}$. Figures 4.10 and 4.11 show how \bar{I}_L , when computed for each individual pre-drug and post-drug sweep, changes over the course of the experiment. Here we can see that the size of the leak current in a given well changes little over the course of the experiment, though there seems to be substantial leak-magnitude variability between wells. It is noteworthy, however, that Well B09 shows uncharacteristically large leak currents, and it appears that the size of this well's leak currents steadily increased over the course of our pre-drug recordings. We expect that this may have some effect on the resultant I_{Kr} parameter estimates, perhaps leading to bias if such leak currents are large enough to impact V_m and violate our ideal-patch assumptions.

4.3.2 Quality control

After performing leak-correction, post-drug subtraction and reversal on our data, we use a set of *quality control* (QC) criteria to filter out failed experiments from our dataset. These QC-criteria are detailed in Table 4.2. These criteria are mostly the same as those used in Lei et al. (2019b), but only a subset of these QC criteria may be applied to all




Figure 4.7: A single-exponential model is fitted through a common section of the trace to quantify the variability between sweeps. Panel **a**: an example of a postprocessed trace—the portion of the trace used to fit the exponential is highlighted (grey). Panel **b**: the exponential decay that is fit to these data, along with the characteristic time-constant of decay. Though a single-exponential does not provide a perfect fit to the data (and, perhaps, a double exponential would be better suited), the resultant time-constant is a useful statistic.

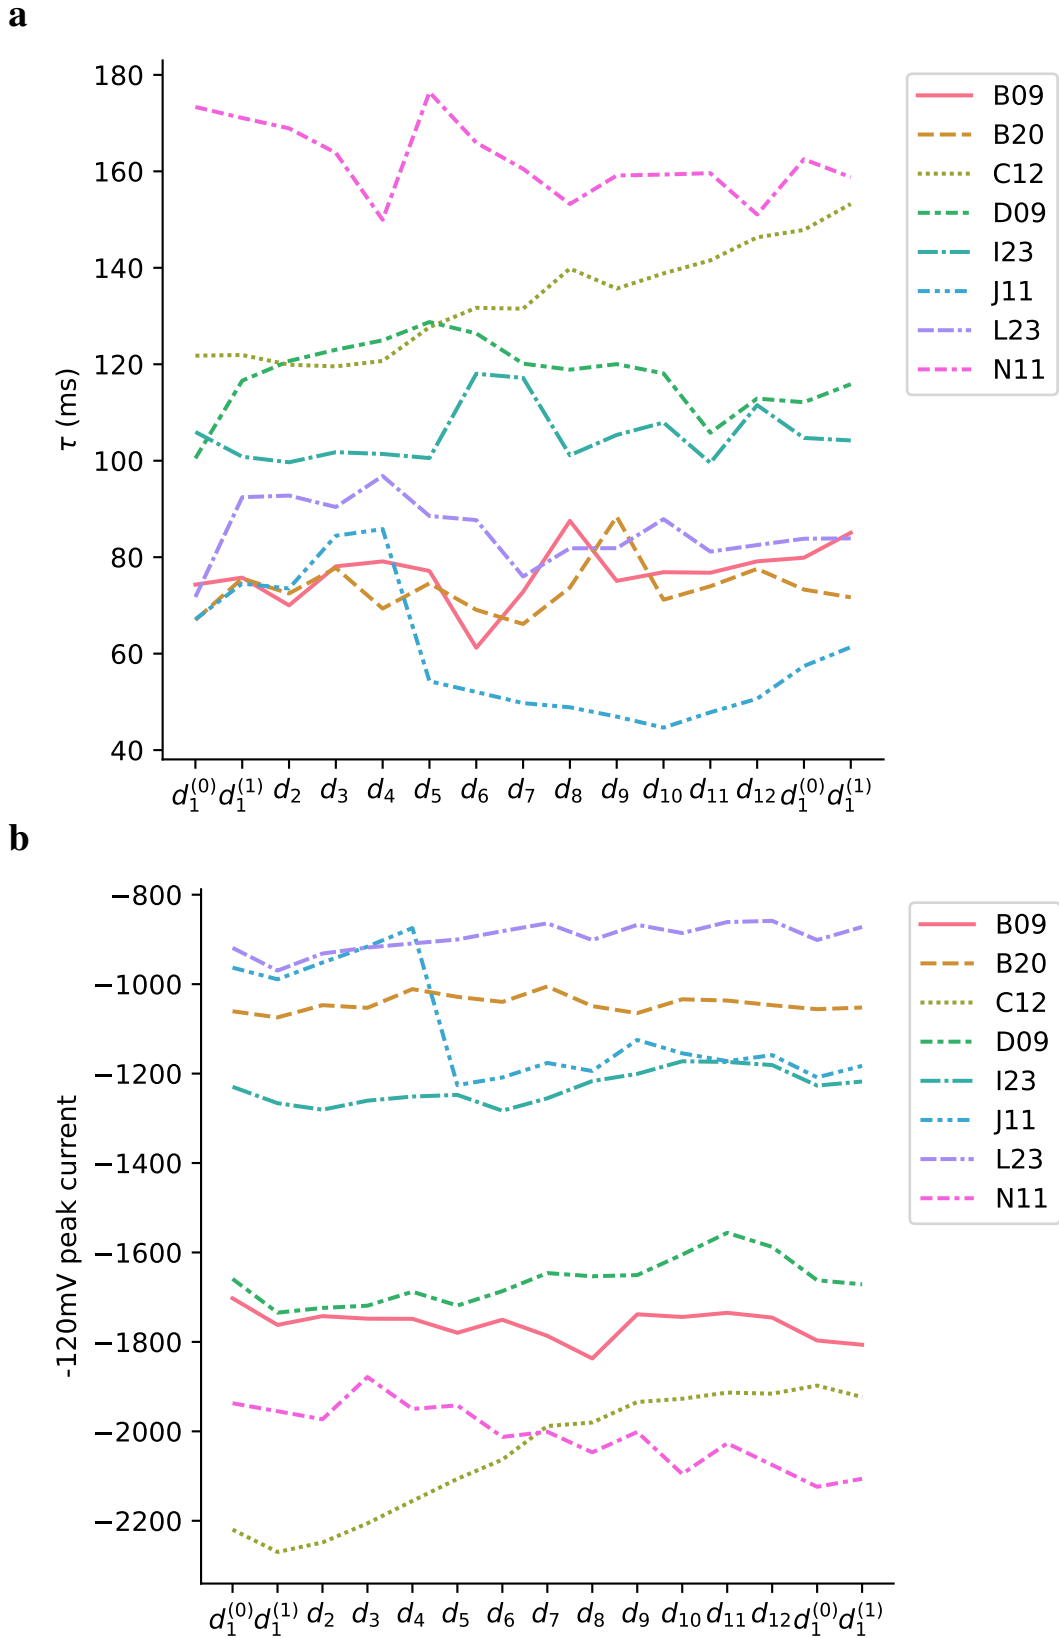


Figure 4.8: Panel **a**: the variability of the time-constant during the first -120 mV step across each protocol chronologically throughout the full experiment, plotted for each well. Panel **b**: The peak (negative) current. In both panels, we see noticeable variability between wells, and some variability in the values obtained from each individual well.

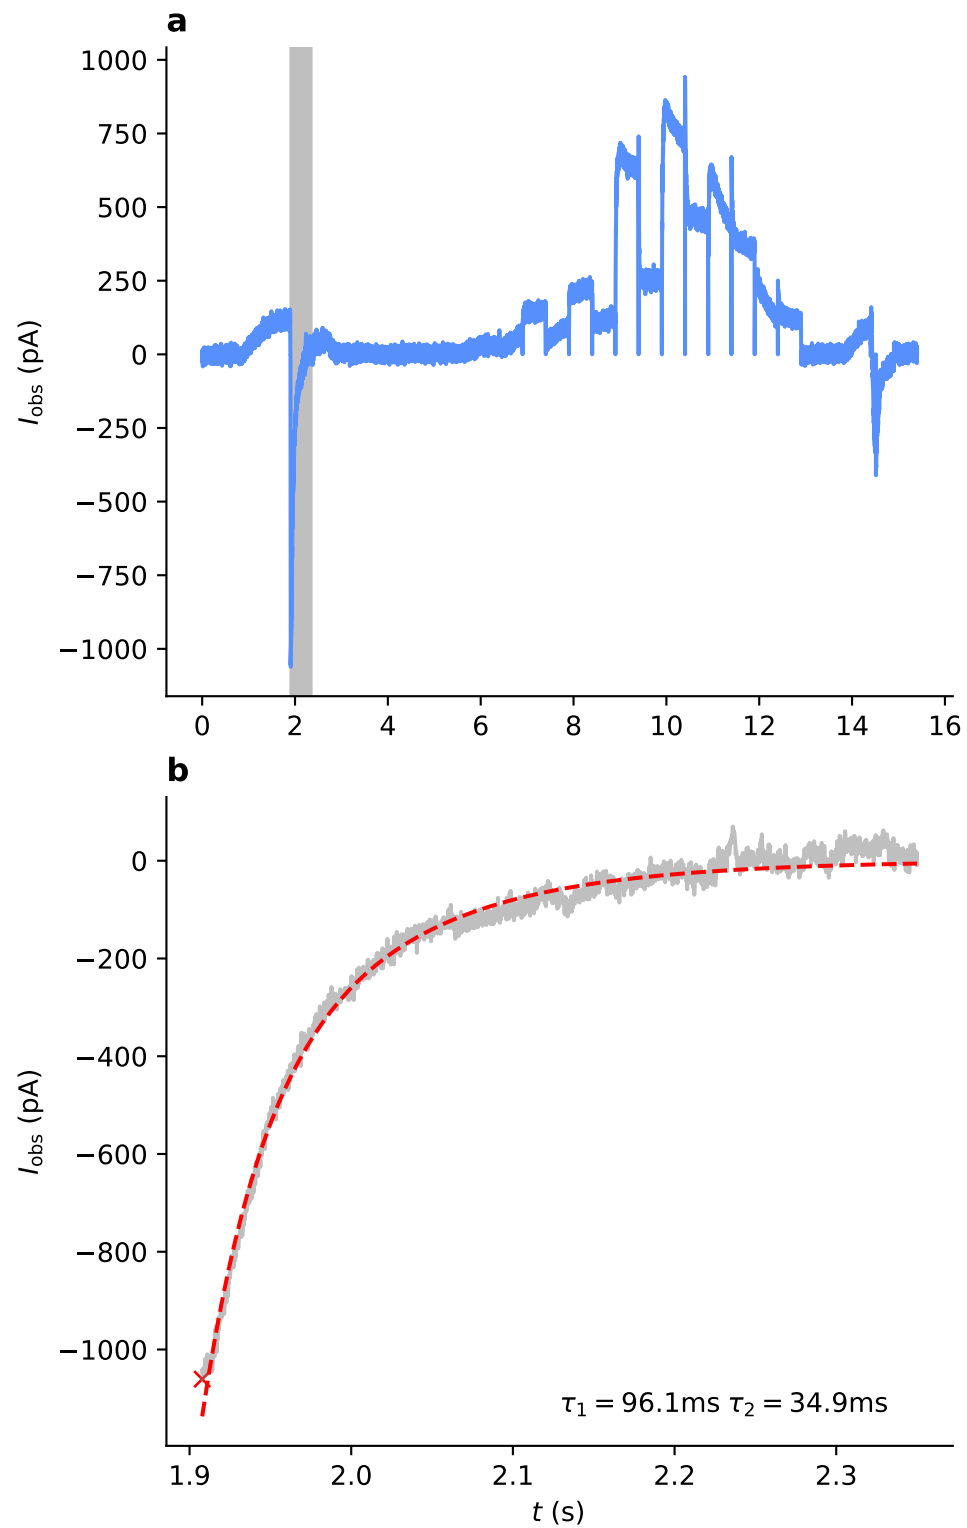


Figure 4.9: A double-exponential model fitted to the first -120mV step found at the beginning of each protocol. Panel **a**: an example of a postprocessed trace—the portion of the trace used to fit the exponential is highlighted (grey). Panel **b**: the exponential decay that is fit to these data, along with the characteristic time-constant of decay. The fitted variables τ_1 and τ_2 describe the timescale of each exponential decay.

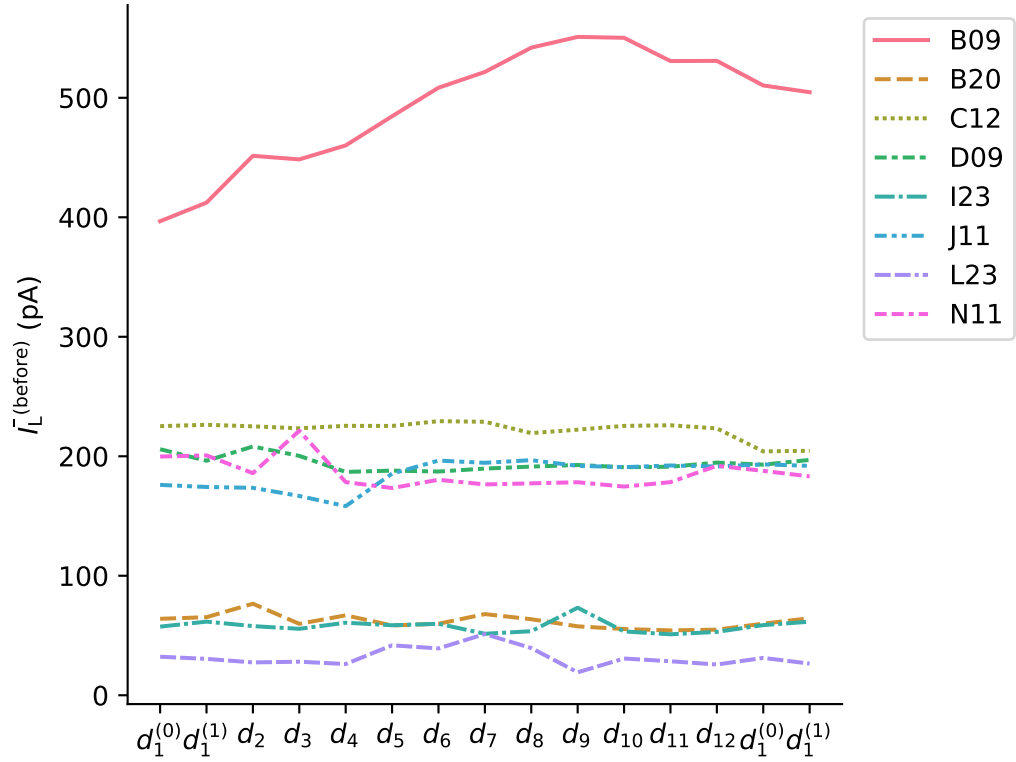


Figure 4.10: The size of the leak current remains largely constant in most wells. We show the pre-drug leak magnitude, $\bar{I}_L^{(before)}$ for each well over the course of the experiment.




Figure 4.11: The size of the leak current remains largely constant in most wells. We show the post-drug leak magnitude, $\bar{I}_L^{(after)}$, for each well over the course of the experiment.

protocols. Furthermore, the fact that we run the *staircase* protocol at the beginning and end of the experiment allows us to include an additional criterion which concerns the stability of the experiments.

Lei et al.’s *staircase* criteria The Nanion SyncroPatch automatically applies a short protocol and infers a number of parameters before each of our user-specified voltage-protocols (d_1 – d_{12}). Namely, these parameters are the seal resistance R_{seal} , the membrane capacitance C_m , and the series resistance, R_{series} . We remove data from each well where any of these values lie outside a predetermined range (**QC1**). For this purpose, we use the values suggested in Lei et al. (2019a) (those listed in Table 4.2. We deem any value outside these ranges as indicative of a failed experiment. For example, such cases arise from wells where: there has been a failure to attach a cell to the hole in the well plate; the cell has detached part-way through the experiment; or we have failed to puncture the cell membrane (that is, we failed to “patch” the cell).

Lei et al. (2019a) introduced and applied a thorough selection of QC criteria to high-throughput, automated patch-clamp data (collecting using a similar Nanion SyncroPatch 384). These quality control criteria use data obtained from Lei et al.’s *staircase* protocol alongside parameter estimates provided by the machine (namely R_{seal} , C_m and R_{leak}). The first selection of criteria, **QC1**, ensures that the machine estimates of these values within some reasonable bounds for each sweep of each protocol—these are the same bounds used in Lei et al. (2019b).

In **QC2**, we check the signal-to-noise ratio in our data. The noise, σ , is estimated from the first 200 observations (where the current is assumed to be stationary) by taking the sample standard deviation. We then compute the standard deviation of the entire postprocessed trace (I_{post}) and ensure that this is at least five times larger than our estimated noise. This condition is checked for both the raw trace, and the drug-subtracted, leak-corrected trace. These criteria are not satisfied due to, for example, an unexpectedly small current, or excessive noise in the trace. A small current may be due to insufficient expression of hERG1a in the cell, whereas excessive noise may indicate a failure of the experimental equipment (faulty electrodes, for example).

In **QC3**, we check if two traces of post-processed, time-series data differ appreciably. A large difference between the raw traces for the first and second sweep of the staircase protocol, for instance, would indicate that something has unexpectedly changed through time whilst the experiment progressed—perhaps indicating an unstable leak current, for example. This instability undermines our modelling assumptions, and so, we discard such data. Lei et al.’s criteria ensure that,

$$\frac{1}{\sqrt{n}} \|\mathbf{I}_1 - \mathbf{I}_2\| < 0.2 \times \frac{1}{2n} (\text{std}(\mathbf{I}_1) + \text{std}(\mathbf{I}_2)), \quad (4.17)$$

where \mathbf{I}_1 and \mathbf{I}_2 are the two time-series current traces that are being compared (Lei et al., 2019b), and $\text{std}(\mathbf{I})$ denotes the standard deviation of the observations from a trace, \mathbf{I} . This check is performed for raw traces (**QC3.raw**), post-drug traces (**QC3.drug**) and leak-corrected, drug-subtracted traces (**QC3.subtracted**).

In **QC.4** we test that our machine estimates, (C_m , R_{seal} and R_{series}) have not drastically changed between back-to-back sweeps of the staircase protocol. To do this, consider the ratio of the mean and standard deviation of the two values, testing the condition,

$$\frac{\text{std}\left(C_m^{(2)}, C_m^{(1)}\right)}{\text{mean}\left(C_m^{(2)}, C_m^{(1)}\right)} < 0.25. \quad (4.18)$$

where $C_m^{(i)}$ is machine estimate of C_m obtained during the i^{th} sweep of the *staircase* protocol (**QC4.Cm**). Similarly, we check consecutive estimates of R_{series} (**QC4.Rseal**) and R_{seal} (**QC4.Rseries**).

In **QC5**, we ensure that the drug subtraction has resulted in a significant reduction in current. This is done by first checking that the maximal staircase current (observed during the second half of the protocol) decreases after the addition of dofetilide (**QC5.staircase**), and then checking whether $2\|I_{\text{after}}\|_2 < \|I_{\text{before}}\|_2$ (**QC5.1.staircase**). In **QC6**, we ensure that a positive leak-corrected, drug-subtracted current is observed during +40 mV steps. This is done by checking the first, second and fourth +40 mV steps in the *staircase* protocol (d_1)—these criteria are denoted **QC6.subtracted**, **QC6.1.subtracted** and **QC6.2.subtracted**, respectively.

Originally, Lei et al. applied these criteria to two sweeps (repeats) of the staircase protocol. However, we perform four staircase sweeps in total: two sweeps before, and two after all other protocols, as shown in Figure 4.3. This means we can apply **QC 1–6** separately to the first two sweeps and the second two sweeps of the *staircase* protocol, before evaluating the remaining QC criteria (concerning the other protocols). Since we discard wells failing any single criteria, this reduces the computational resources (namely execution time and storage) necessary for postprocessing our data.

We include one additional criterion to sweeps obtained using the *staircase* protocol, **QC.R_leftover**. Here, we ensure that each sweep exhibits a small amount of nuisance current which is not explained by our leak model, and retain only those wells for which each of our four *staircase* protocol (d_1) sweeps satisfy,

$$R_{\text{leftover}} \leq 0.25, \quad (4.19)$$

where R_{leftover} is calculated as in Equation (4.12). Recall that $R_{\text{leftover}} \approx 0$ indicates that the post-drug trace consists almost entirely of linear leak.

Application of criteria to other protocols After applying Lei et al.’s QC to our repeats of the *staircase* protocol, we apply a subset of these criteria to the other protocols. In particular, protocols **QC1** and **QC4** rely only on the machine estimates of C_m , R_{seal} and R_{series} , and **QC6.subtracted** concerns only the first +40 mV step in the protocol, until which point all protocols are identical. We denote the application of these criteria to the remaining protocols by **QC1.all**, **QC4.all** and **QC6.all**, respectively.

Reversal potential criteria We also include some additional checks concerning the apparent reversal potential of the current, E_{post} , (inferred as described in Section 4.3). We ensure that E_{post} lies in the range covered by our reversal ramp,

$$-110 \text{ mV} < E_{\text{post}} < -70 \text{ mV}, \quad (4.20)$$

removing any wells for which there is a sweep where E_{post} lies outside this range.

Regarding the variability in our inferred E_{post} values, we ensure that the spread, in values obtained from a given well, is no greater than 5 mV, that is,

$$\max_i \{E_{\text{post}}^{(i)}\} - \min_i \{E_{\text{post}}^{(i)}\} < 5 \text{ mV}, \quad (4.21)$$

where $E_{\text{post}}^{(i)}$ is the reversal potential estimated from a given sweep and i indexes the sweeps performed on the well under consideration. We denote this criterion **QC.Erev.spread**

Ensuring consistency across the experiment We apply **QC3.subtracted** to compare the very first and very last repeats of the *staircase* protocol. This, again, is done to ensure that the experimental conditions have not changed significantly over the course of the experiment.

All of our QC criteria are evaluated after the removal of “capacitive spikes” from the data. In particular, we zero the data observed no more than 2ms after each discontinuity in V_{cmd} —as is done in Lei et al. (2019b). We zero this data because are large current spikes after each voltage “step” resulting from the capacitive effect of the cell membrane, which violates our *ideal patch-clamp* assumptions during a short period of time after each voltage step. It is also possible to discard this data entirely, which we expect to have little impact on the evaluation of our QC criteria.

QC for leak correction and drug subtraction Figure 4.12 shows how R_{leftover} changes over time, between sweeps. Because the size of I_{K_r} varies between protocol, we should expect that R_{leftover} is somewhat protocol dependent. This is shown in the values taken from d_4 and d_5 , where R_{leftover} increases for all but one well, as shown in Figure 4.12.




Figure 4.12: The relative size of the current leftover after drug subtraction and leak correction, as quantified by our R_{leftover} statistic. Only wells passing each of our QC criteria are shown. One such criterion is that $R_{\text{leftover}} < 0.25$ for each *staircase* protocol (d_1) repeat. The value of R_{leftover} is protocol dependent, which is apparent from the values obtained from protocols d_5 , d_9 and d_{12} , where we see increased R_{leftover} .

Ideally, following leak-correction and drug-subtraction, we should be left with a trace that is consisting entirely of I_{K_r} and observational noise, resulting in a small R_{leftover} value. Conversely, if R_{leftover} is large, this may indicate the presence of a nonlinear leak (violating our modelling assumption that the leak current is due to a simple, Ohmic resistor in accordance with Equation (4.3)). It may also indicate the presence of other, endogenous currents. The presence of either nonlinear leak, or endogenous currents is undesirable, and so a small R_{leftover} is indicative high-quality data. Data from Well M06, which passed all QC criteria except **QC.R_leftover**, is shown in Figure 4.6.

Order of application of criteria In practice, we may apply these criteria sequentially—starting with only the criteria that require only data from the staircase protocol, and eliminating the wells that fail **QC1** for example, without the need for processing the data from other protocols in these wells. This sequential approach can save unnecessary computation because a significant portion of wells tend to fail criteria which relate only to the staircase protocol and onboard parameter estimates. Nevertheless, we apply these

criteria in their entirety to every well, allowing a more informative summary of the quality of our data.

Success rate The number of wells passing all of our QC criteria (for every sweep of every protocol) is 8, representing a 2.1% success rate. This is a lower success rate than that of Lei et al. (2019b), which is perhaps not surprising given the additional protocols and QC criteria we have included. The total duration of the experiment (over 10 minutes) and the use of fluoroide-free plates may have also contributed to lower success rates.

The location of each passing/failing well on the wellplate is shown in Figure 4.14. This figure shows no obvious spatial clustering, which, if present, may indicate an equipment failure affecting part of the plate. A more comprehensive summary is shown in Table 4.3, where the number of wells passing each criterion is shown. These results are also visualised in Figure 4.13, which shows which wells are removed by which criteria.

Reversal potential variability There was a noticeable difference between the observed reversal potential, E_{post} and its theoretical value, E_{Nernst} , as shown in Figure 4.15. Many of the wells that passed QC have a mean reversal potential somewhat greater than the calculated Nernst potential ($E_{\text{Nernst}} \approx -89.83$).

4.4 Fitting and validation of models

Modelling assumptions Throughout this chapter, we assume that our data arises from an *ideal* patch-clamp setup described in Section 4.1 (Equations (4.1) and (4.2)). Here, we fit our models to the postprocessed trace, I_{post} (Equation (4.11)) under the assumption of IID additive, gaussian errors, with mean 0 and known standard deviation σ . The standard deviation, σ , may be estimated from the beginning, stationary portion of the trace but, our parameter estimates are in fact independent of σ as explained in Chapter 2.

Modelling the reversal potential We saw in Section 4.3 that there is variability in E_{post} , the value we infer from the reversal ramp. Under ideal conditions, we would expect that this value is identical to the Nernst potential, that is,

$$E_{\text{Nernst}} = E_{\text{post}}. \quad (4.22)$$

However, we can clearly see that this is not the case in Figures 4.15 and 4.16. This raises the question: which value should we use for the reversal potential in our mathematical

Label	Protocols	Criterion
QC1.Rseal ¹	all	Check $R_{\text{Seal}} \in [0.1, 1000] \text{ G}\Omega$
QC1.Cm ¹	all	Check $C_m \in [1, 100] \text{ pF}$
QC1.Rseries ¹	all	Check $R_{\text{Series}} \in [1, 25] \text{ M}\Omega$
QC2.raw ¹	<i>staircase</i>	Check $\frac{\text{std}(I_{\text{raw}})}{\sigma} > 5$
QC2.subtracted ¹	<i>staircase</i>	Check $\frac{\text{std}(I_{\text{subtracted}})}{\sigma} > 5$
QC3.raw ¹	<i>staircase</i>	Check Condition (4.17) with $I_{\text{obs}}^{(1)}, I_{\text{obs}}^{(2)}$
QC3.drug ¹	<i>staircase</i>	Check Condition (4.17) $I_{\text{drug}}^{(1)}, I_{\text{drug}}^{(2)}$
QC3.subtracted	<i>staircase</i>	Check Condition (4.17) $I_{\text{subtracted}}^{(1)}, I_{\text{subtracted}}^{(2)}$
QC3.bookend ¹	<i>staircase</i>	Check Condition (4.17) with $I_{\text{subtracted}}^{(1)}, I_{\text{subtracted}}^{(4)}$
QC4.Cm ¹	all	Check $C_m^{(\text{before})}$ and $C_m^{(\text{after})}$ satisfy Condition (4.18)
QC4.Rseal ¹	all	Check $R_{\text{seal}}^{(\text{before})}$ and $R_{\text{seal}}^{(\text{after})}$ satisfy Condition (4.18)
QC4.Rseries ¹	all	Check $R_{\text{series}}^{(\text{before})}$ and $R_{\text{series}}^{(\text{after})}$ satisfy Condition (4.18)
QC5.staircase ¹	<i>staircase</i>	Check $0.25 \times \max I_{\text{raw}} > \max I_{\text{drug}}$
QC5.1.staircase ¹	<i>staircase</i>	Check that $2\ I_{\text{after}}\ _2 < \ I_{\text{before}}\ _2$
QC6.subtracted ¹	all	$\max\{I_{\text{subtracted}}^{(+40\text{mV})}\} > -2\sigma$ where $I_{\text{subtracted}}^{(+40\text{mV})}$ is the current during the first +40mV step
QC6.1.subtracted ¹	<i>staircase</i>	$\max\{I_{\text{subtracted}}^{(+40\text{mV})}\} > -2\sigma$ where $I_{\text{subtracted}}^{(+40\text{mV})}$ is the current during the second +40mV step
QC6.2.subtracted ¹	<i>staircase</i>	$\max\{I_{\text{subtracted}}^{(+40\text{mV})}\} > -2\sigma$ where $I_{\text{subtracted}}^{(+40\text{mV})}$ is the current during the third +40mV step
QC.Erev	all	Check that $-110\text{mV} < E_{\text{obs}} < -70\text{mV}$
QC.Erev.spread	all	Check that E_{obs} varies by less than 5mV across all repeats and all protocols

Table 4.2: Quality-control criteria used to process our data. The central column shows which protocols are applied only to the *staircase* protocol, and which criteria are applied to all protocols. Note that **QC3.subtracted** is applied to the first and second sweeps of the *staircase* protocol, as well the third and fourth sweeps. Whereas, **QC3.bookend** is identical but applied to the first and fourth sweeps. However, **QC3.raw**, **QC3.drug**, **QC6.1.subtracted** and **QC6.2.subtracted** are applied only to the first two sweeps and the last two sweeps of the *staircase* protocol.

Criterion	Protocols used		all
	first staircase	second staircase	
QC1.cm	86	87	97
QC1.rseal	181	178	194
QC1.rseries	229	229	242
QC2.raw	77	90	94
QC2.subtracted	92	115	121
QC3.E4031	86	93	104
QC3.raw	119	108	140
QC3.subtracted	123	105	144
QC4.cm	86	86	99
QC4.rseal	92	103	117
QC4.rseries	110	90	125
QC5.1.staircase	167	211	224
QC5.staircase	150	152	172
QC6.1.subtracted	117	179	188
QC6.2.subtracted	118	177	190
QC6.subtracted	118	181	191
QC.Erev.all_protocols			121
QC.Erev.spread			229
QC1.all_protocols			309
QC3.bookend			229
QC4.all_protocols			183
QC6.all_protocols			192
QC6.R_leftover			326
All criteria	335	344	376

Table 4.3: The number of wells failing each individual QC criterion. For Lei et al. (2019b)’s original QC criteria, we list the amount of wells failing the first applications of the *staircase* protocol (the first and second sweeps) and those failing the second (third and fourth sweeps).




Figure 4.13: How many wells are discarded using each stage of our quality-control (QC) pipeline. The displayed stages of this pipeline are: Stage 1, Lei et al.’s **QC1** criteria applied to each repeat of the *staircase* protocol; Stage 2, Lei et al.’s QC criteria applied to both the first two, and final two, repeats of the *staircase* protocol; Stage 3, Lei et al.’s QC applied to all protocols (where applicable); Stage 4, check that that observed reversal potential is observed during the reversal ramp, that is, $-11 \text{ mV} < E_{\text{post}} < 70 \text{ mV}$ for all protocols (**QC.Erev.all_protocols**); Stage 5, check that E_{post} remains consistent throughout the experiment (**QC.Erev.spread**); finally, we ensure that the post-drug trace is dominated by leak (**QC.R_leftover**), discarding a single well, Well M06. The remaining wells are used for the fitting and validation of our models.




Figure 4.14: A plan view of the wellplate, with 384 wells arranged in 24 columns. Wells that passed all QC criteria are highlighted in orange. There is no obvious clustering of cells which passed QC. Such clustering may suggest a partial equipment failure as each quadrant of the wellplate shares certain components.

models? For every model in this chapter, we use an observation function of the form,

$$I_{Kr}(t) = gx_O(t)(V_m - E_{Kr}), \quad (4.23)$$

where g is a model parameter representing a cell's maximum conductance with respect to I_{Kr} , x_O denotes the proportion of channels in an open conformation at time t , and E_{Kr} is the reversal potential.

Some possible explanations for this discrepancy between E_{Nernst} and E_{post} are outlined in Table 4.4. Each case corresponds to slight changes in our model with regard to how it is fit to the data, and how it is used to make predictions. The first case, Case I, is the simplest and corresponds to the assumption that $E_{Kr} = E_{Nernst}$, Cases II and III correspond to two possible hypotheses regarding the discrepancy between E_{Nernst} and E_{post} : that the reversal potential is truly (and materially) different from E_{Nernst} ; and that there is a simple, constant offset between V_{cmd} and V_m .

These are not the only possible sources of error, nor are they necessarily mutually exclusive. However, comparing the parameter estimations and predictive models obtained from these different assumptions may allow the dominant cause for the discrepancy between E_{post} and E_{Kr} to be determined, leading to improved predictive models. In this section, we compare each of these subtly different ways of configuring our models, as well as our different model structures.




Figure 4.15: The reversal potential for each well averaged over our 15 repeats of the reversal ramp. Wells passing QC are shown in orange, and wells passing Lei et al. (2019b)’s original QC criteria (for all sweeps of the staircase protocol), but not our full criteria are shown in blue (stacked above segments showing the wells passing QC).




Figure 4.16: The reversal potential for each well plotted (inferred from our drug-subtracted traces) for each protocol in our experiment (listed chronologically). The grey dashed line shows E_{Nernst} .

Case	Parameterisation	Physical Interpretation
I	$E_{\text{Kr}} = E_{\text{Nernst}}$ and $V_{\text{m}} = V_{\text{cmd}}$	I_{Kr} reverses at the Nernst potential. We fit the model treating E_{Nernst} as a known constant calculated from the Nernst equation.
II	$E_{\text{Kr}} = E_{\text{post}}$ and $V_{\text{m}} = V_{\text{cmd}}$	$E_{\text{Kr}} \neq E_{\text{Nernst}}$ (e.g. due to temperature/solution changes). We set $E_{\text{Kr}} = E_{\text{post}}$ using the reversal-ramp estimate and treat this as a known constant when fitting the maximal conductance and transition-rate parameters.
III	$E_{\text{Kr}} = E_{\text{Nernst}}$ and $V_{\text{m}} = V_{\text{cmd}} + V_{\text{off}}$	$E_{\text{Nernst}} \neq E_{\text{post}}$ due to a systematic voltage offset, V_{off} . We fit the model with the known constant $V_{\text{off}} = E_{\text{Nernst}} - E_{\text{post}}$ (E_{post} inferred from the reversal ramp as in Case II).

Table 4.4: Different assumptions regarding the models' reversal potential, E_{Kr} , and the observed reversal potential, E_{post} , and how this informs model fitting.

Capacitive spike removal As mentioned in Section 4.3, we observe “capacitive spikes” whenever there is a discontinuity in V_{cmd} . These capacitive spikes violate our modelling assumption that $V_{\text{cmd}} = V_{\text{cmd}} + V_{\text{off}}$, because during this time, the cell membrane charges and discharges in response to the changing transmembrane potential. To address this problem, we simply discard 5ms of data after each voltage jump when fitting our models. This approach to model fitting is exactly the approach taken in Beattie et al. (2018). In Appendix C, we explore how removing different amounts of data affects the statistical properties of our parameter estimates under various values. Note that we explicitly include the capacitive effect of the cell membrane in the models presented in the following chapter.

Markov models We apply Markov models of the same form introduced in Chapters 2 and 3. In particular, these are ODE-based models with a governing equation of the form,

$$\frac{d}{dt}\mathbf{x} = \mathbf{Q}(V_m)^\top \mathbf{x}, \quad (4.24)$$

together with an observation function of the form,

$$I_{\text{Kr}}(t) = g x_{\text{O}}(t)(V_m(t) - E_{\text{Kr}}), \quad (4.25)$$

where $\mathbf{Q}(V_m)$ is a transition-rate matrix.

As described in Section 3.3.2, various Markov models, characterised by a particular choice of \mathbf{Q} , may be chosen here. We consider four models—those shown in Figure 4.17. These models differ in the number of states and parameters, and in the existence of a path between the inactivated state (**I**) and the closed states (**C/C1/C2/C3**) that avoids the open state (**O**). These models are chosen such that a comparison in the models’ predictive accuracy may provide insight into the gating dynamics of I_{Kr} .

Model fitting Similarly to the previous chapter, we use maximum likelihood estimation under the assumption of additive, IID Gaussian noise. This is mathematically equivalent to nonlinear least-squares regression, as explained in Section 3.3.1.

To constrain our kinetic rate parameters, we use the same kinetic-rate constraints used in Chapter 3. In particular, for rates of the form $k(V) = A \exp\{bV\}$, we require that,

$$10^{-7} \text{ ms}^{-1} < a < 10^5 \text{ ms}^{-1}, \quad (4.26)$$

$$10^{-7} \text{ mV}^{-1} < b < 10^5 \text{ mV}^{-1}, \quad (4.27)$$

$$k(V) < 10^3 \text{ ms}^{-1} \text{ for all } V \in [-120 \text{ mV}, +60 \text{ mV}], \quad (4.28)$$

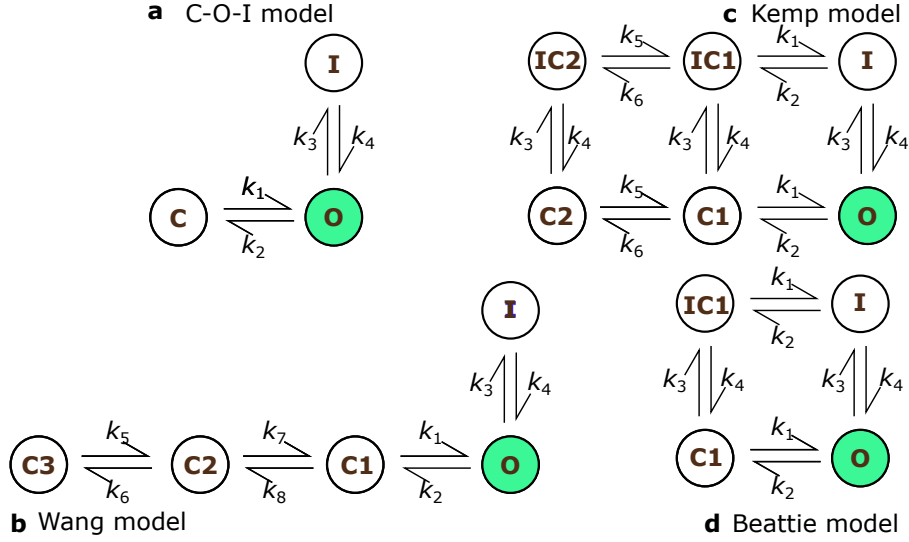


Figure 4.17: The four model structures used in this chapter: the C-O-I model (Panel **a**), the Wang et al. (1997) model (Panel **b**), the Kemp et al. (2021) (Panel **c**), and the Beattie et al. (2018) (Panel **d**). Note that the **C1** to **C2** transition in the Wang et al. (1997) model is constant, that is, determined by a single model parameter. Panel **b**: A table summarising the models

and

$$k(V) > 1.67 \times 10^{-5} \text{ ms}^{-1} \text{ for some } V \in [-120 \text{ mV}, +60 \text{ mV}]. \quad (4.29)$$

When calibrating our models with real data, we expect the maximal conductance, g_{Kr} to vary between cells. This is because we expect the maximal conductance to be roughly proportional to the number of channels, which varies between cells. Likewise, we may expect our maximal conductance to vary if we were to use different cells for our experiments. For this reason, we use maximal conductance constraints that are based on the size of the observed current.

In particular, we consider the minimum observation from the first -120 mV step of our leak-corrected and drug-subtracted trace (discarding the first 5 ms, as we do when fitting our models), and denote this by I_{max} . Then, according to our model, and assuming $V_{\text{off}} = 0$,

$$I_{\text{max}} = \bar{g}x_{\text{O}}(V_{\text{cmd}} - E_{\text{Kr}}), \quad (4.30)$$

where the proportion of states in the open conformation satisfy $0 < x_{\text{O}} < 1$ (see Appendix A) and so, by computing,

$$\bar{g}x_{\text{O}} = \frac{I_{\text{max}}}{V_{\text{cmd}} - E_{\text{Kr}}}, \quad (4.31)$$

we can obtain an upper bound for \bar{g} . This is possible because we know that $x_{\text{O}} \leq 1$ and $V_{\text{m}} = -120 \text{ mV}$. Then, under the rather lenient assumption that $x_{\text{O}} \geq 0.01$ when I_{max} is

observed, we have the inequality,

$$\frac{I_{\max}}{0.01 \times (V_m - E_{\text{Kr}})} \geq \bar{g}. \quad (4.32)$$

Similarly, we ensure that,

$$\frac{I_{\max}}{100 \times (V_m - E_{\text{Kr}})} \leq \bar{g}, \quad (4.33)$$

which corresponds to the similarly weak assumption that $x_O \leq 100$.

4.4.1 Computational methods

Fitting As discussed in Section 2.3 and 3.3.1, we fit our models to time-series data by computing *maximum-likelihood estimates* (MLEs) under the assumption that our observations are subject to independently and identically distributed, additive Gaussian noise. That is, for a given protocol, d , we compute,

$$\hat{\boldsymbol{\theta}}_d = \operatorname{argmin}_{\boldsymbol{\theta}} \left\{ \sum_{i=1}^{n_d} (y_i(\boldsymbol{\theta}; d) - z_i)^2 \right\}, \quad (4.34)$$

where n_d is the number of observations in protocol d , our i^{th} observation is denoted by z_i , and the model output (for the i^{th} observation) for a given parameter vector $\boldsymbol{\theta}$ is denoted by $y_i(\boldsymbol{\theta}; d)$. To fit our models, we seek a solution to this optimisation problem (Equation (4.34)). A general closed-form solution is not available, so we resort to numerical optimisation methods instead.

To fit our models, we use CMA-ES (Hansen, 2016), as introduced in Section 2.3.2. Moreover, we repeat our optimisation at least 30 times from different initial sampling distributions, which allows us to explore more of the model's parameter space and, provided we reliably recover the same parameter estimate, demonstrates that we are able to identify the true global optimum (that is, we can reliably compute Equation 4.34). However, for the Wang et al. (1997) model, we repeat the optimisation from 120 initial guesses, owing to the model's increased complexity and greater number of parameters. For all models except the Wang et al. (1997) model, we select the population size (that is, the number of parameter vectors sampled for each generation) using the heuristic suggested by the PINTS package (Clerx et al., 2019b),

$$n_{\text{pop}} := \lfloor 4 + 3 \ln(n_p) \rfloor, \quad (4.35)$$

where n_p is the number of model parameters, In the case of the Wang et al. (1997) model, we increase the population size to 50 because of the model's increased complexity. We terminate the optimisation once 200 iterations have failed to yield a materially improved parameter set (that is, by decreasing the sum-of-squares error by more than 10^{-7}).

A cell-specific parameter set is fitted to every repeat of every protocol for each well. In this way, each repeat of each protocol is a separate fitting task which may be treated independently. As in Chapter 3, this allows us to fit many traces in parallel using a high-performance computing resource. Again, as in Chapter 3, these methods are implemented using Numba such that efficient code can be generated and just-in-time (JIT) compiled for any Markov model structure. Our ODE systems (that is, Equation 4.24) are integrated using the LSODA method (as described in Section 2.3) with absolute and relative tolerances set to 10^{-8} to ensure convergence of the objective with respect to parameters (Clerx et al., 2019a; Creswell et al., 2024).

As in Chapter 3 we perform each optimisation multiple times using different initial guesses. Here, in this chapter, we use 30 initial guesses for each sweep of each protocol. Our initial guesses for parameters relating to our models' transition rates are sampled from our parameter space using log-uniform and uniform distributions, here we choose,

$$\log_{10}\{A_i\} \sim U(-7, 5), \quad (4.36)$$

$$b_i \sim U(10^{-7}, 10^5), \quad (4.37)$$

where A_i and b_i are our initial guesses of parameters characterising some transition rate, $k_i = A \exp\{\pm b_i V_m\}$. Choosing these probability distributions ensures that we initiate our optimisation starting from various locations in parameter space, lessening the chance that we do not identify the global optimum.

However, unlike Chapter 3, we also sample an initial guess for our maximal conductance parameter, g . Our initial guesses for the maximal conductance, g , are sampled according to a log-uniform distribution which has support over our parameter space,

$$\log\{g\} \sim U(\log\{g_{\min}\}, \log\{g_{\max}\}). \quad (4.38)$$

where g_{\min} and g_{\max} are the upper and lower bounds, respectively, from Inequalities (4.32) and (4.33).

We fit our models under the assumption of Case I and Case II independently. However, It is not necessary to re-fit our models under Case III because there is a simple mapping between parameters obtained under Case II and those under Case III, that is,

$$\begin{aligned} k(V_{\text{cmd}} + V_{\text{off}}) &= A \exp\{\pm b(V_{\text{cmd}} + V_{\text{off}})\} \\ &= \underbrace{A \exp\{\pm b V_{\text{off}}\}}_{A^\dagger} \exp\{\pm b V_{\text{cmd}}\}. \end{aligned} \quad (4.39)$$

Hence, Case II's 'A' parameter is mapped to $A^\dagger = A \exp\{\pm b V_{\text{off}}\}$, but our 'b' parameter remains unchanged. It should be noted, however, that this does result in a small change in the effective initial guess distributions and the constraints of our optimisation.

Predictions As when fitting, we compute our model predictions by integrating the governing equation (Equation (4.24)) and computing the observation function, Equation (4.25). In this way, we obtain a prediction for a validation protocol \tilde{d} from a parameter estimate obtained from a training protocol, d , that is, $\mathbf{y}(\hat{\boldsymbol{\theta}}(d); \tilde{d})$, where $\hat{\boldsymbol{\theta}}(d)$ denotes the parameter estimate obtained by fitting to protocol d , and \tilde{d} is the validation protocol (that is, the protocol which we are making a prediction for). In Case I and Case III, E_{Kr} (which appears in the driving term of our observation function) is set to E_{Nernst} . In Case II, we assume the value we inferred for the validation trace during postprocessing, that is, $E_{\text{Kr}} = E_{\text{post}}$.

Case III differs from Case I and Case II, because we assume a non-zero voltage offset, $V_{\text{off}} = E_{\text{Nernst}} - E_{\text{post}}$, which affects not only the driving term in Equation (4.25), but also the channel kinetics as described by Equation (4.24). This is achieved through the assumption that $V_{\text{m}} = V_{\text{cmd}} + V_{\text{off}}$. As explained above, our ‘fits’ (the model output of our fitted model under the fitting protocol) are identical under Cases II and III. Hence, to produce model predictions under Case III, we use the parameters obtained from fitting under the assumption of Case II, and, for rates of the form $k = A \exp\{\pm b V_{\text{m}}\}$, we compute,

$$\begin{aligned} A_{\text{predict}} &= A_{\text{fit}} \exp \left\{ \pm b \left(V_{\text{off}}^{(\text{predict})} - V_{\text{off}}^{(\text{train})} \right) \right\} \\ &= A_{\text{fit}} \exp \left\{ \pm b \left(E_{\text{post}}^{(\text{train})} - E_{\text{post}}^{(\text{predict})} \right) \right\}. \end{aligned} \quad (4.40)$$

Under the assumption of Case III, this allows us to account for any change in voltage offset (V_{off}) between protocols. Note that constant transition rates (which appear only in the Wang et al. (1997) Model) are independent of the transmembrane potential and, as such, are left unaffected by changes in V_{off} .

4.4.2 Model fitting results

To be confident that we have successfully identified the optimal parameter set, we should obtain similar results from repeated runs of our optimisation procedure (which both starts from a randomised initial guess, and is inherently stochastic). Figure 4.18 shows the results of one particular optimisation task with 30 repeats. Here, we can see that amongst our best runs, the resulting parameter set varies only slightly. Those results which correspond to a less than a 1% increase in RMSE when compared to the best parameter set are highlighted. Figure 4.18e shows a profile along the likelihood surface, plotted along the straight line from the optimal parameter set ($\lambda = 0$) to our reference parameter set, the Cell #5 transition-rate parameters in the original publication, Beattie et al. (2018) (at $\lambda = 1$). This profile is smooth and shows no local optima on the line directly between these parameter sets (though, our high-dimensional likelihood surface

is far more complicated than this simple one-dimensional projection). This also shows that the (transition-rate) parameters found by fitting this particular sweep of data are not a radical departure from those found previously in Beattie et al. (2018).

4.4.3 Predictive accuracy under differing assumptions

Following the methodology introduced in Chapter 3, we consider ensembles of predictions for each model. We first consider differences arising from the different modelling assumptions listed in Table 4.4. Figure 4.19 highlights the difference between those models resting on the assumption that $E_{\text{Kr}} = E_{\text{Nernst}}$ (Case I), and those obtained via the assumption that $E_{\text{Kr}} = E_{\text{post}}$ and $V_{\text{off}} = E_{\text{Nernst}} - E_{\text{post}}$ (Cases III). There is a noticeable difference in the models around during the penultimate -80 mV step around 13 s into the protocol. During this step, the models fit under the assumption $E_{\text{Kr}} = E_{\text{post}} \approx -80$ mV, like our recordings, exhibit little current; those models with $E_{\text{Kr}} = E_{\text{Nernst}} \approx -89.8$ mV fail to recapitulate this small current.

We can quantitatively compare the predictive accuracy of our models by computing the RMSE between our model predictions and the data. To do this, we first compute the RMSE in our predictions for each parameter estimate. As in Chapter 3, we may visualise the resulting values with a heatmap. Figure 4.20 shows a selection of such cross-validation heatmaps, each showing the models' predictive accuracy when fitted to and making predictions for Well B09 (and validated using data from the same well).

To summarise of the models' predictive accuracy across wells, in terms of RMSE, we normalise the RMSE by the size of the trace. This accounts for variability in conductance between wells. Here, we compute

$$\text{NRMSE}(\mathbf{y}, \mathbf{z}) := \frac{\sqrt{\frac{1}{n} \sum_{i=1}^n (z_i - y_i)^2}}{\sqrt{\frac{1}{n} \sum_{i=1}^n z_i^2}} = \frac{\|\mathbf{z} - \mathbf{y}\|_2}{\|\mathbf{z}\|_2} \quad (4.41)$$

where \mathbf{z} is our vector of data, and \mathbf{y} is the model output.

We quantify the average predictive accuracy of our models (under cross-validation) by computing,

$$\mathcal{E}_{\text{predict}} = \frac{1}{N_d(N_d - 1)} \sum_{d \in \mathcal{D}} \left\{ \sum_{\tilde{d} \in \mathcal{D} \setminus \{d\}} \text{NRMSE}(I_{\text{Kr}}(\boldsymbol{\theta}_{\tilde{d}}; \tilde{d}), \mathbf{z}_{\tilde{d}}) \right\}, \quad (4.42)$$

Here, we exclude terms where the fitting and validation protocol are identical (that is, those terms for which $d = \tilde{d}$). As such, because of the diverse nature of our set of protocols, this sum describes the performance accuracy of the model when used in

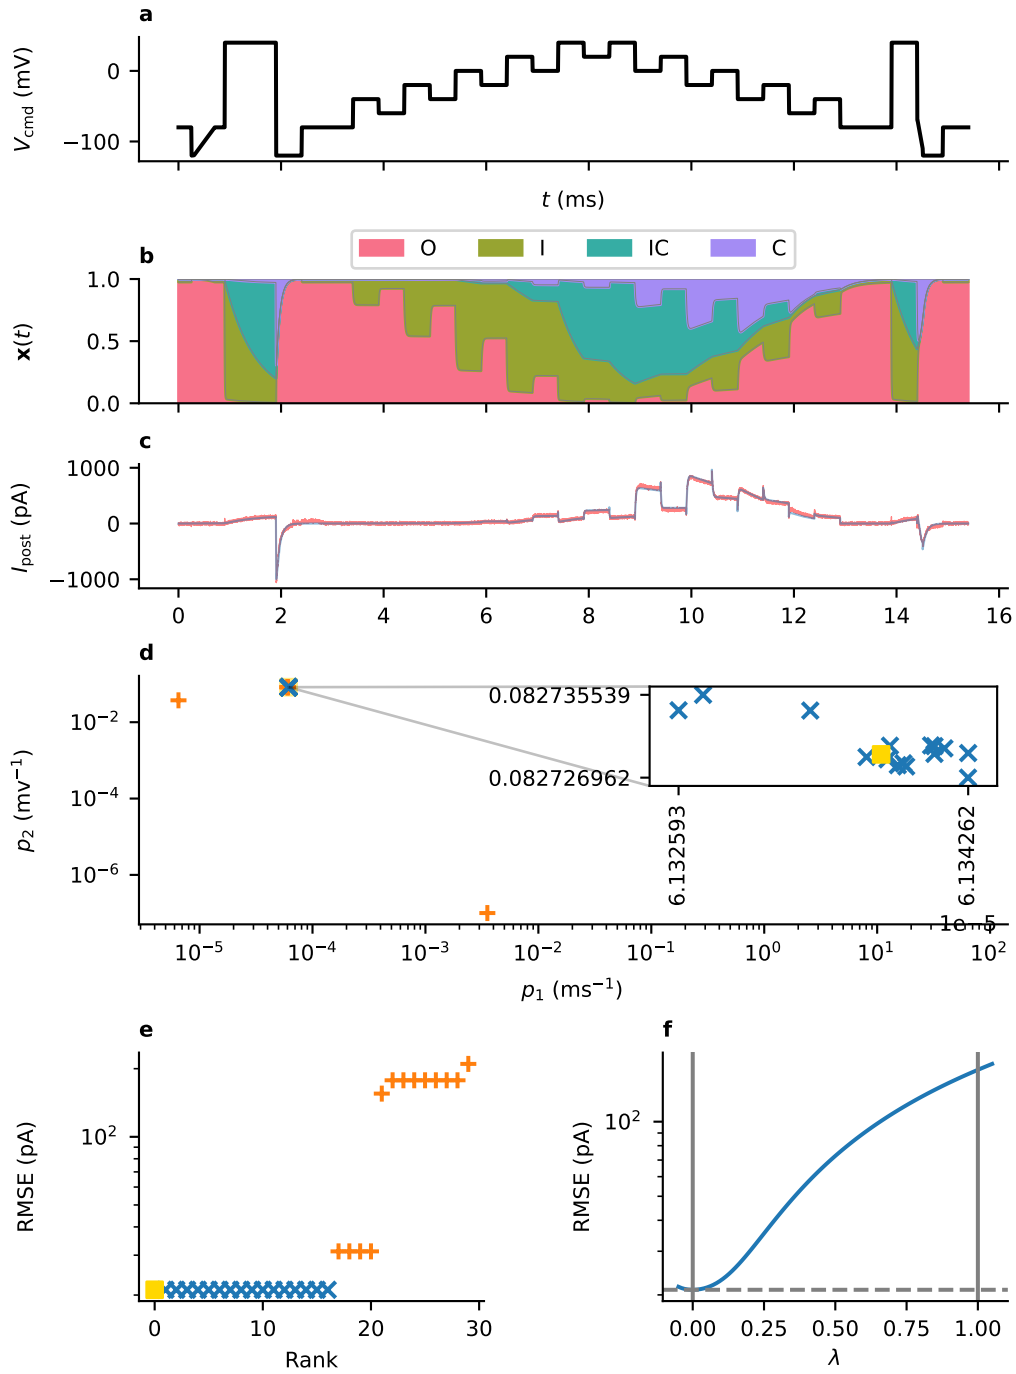


Figure 4.18: Results from fitting the Beattie et al. (2018) model under Case II/III with data collected from Well B20 under the *staircase* protocol (first repeat). Panel **a**: *staircase* protocol (d_1). Panel **b**: the state occupancies according to our fitted model. Panel **c**: the data and model fit. Panel **d**: estimates of p_1 and p_2 obtained from thirty repeats of the optimisation routine— blue markers show optimisation results with RMSE at most 101% of the best-found value. Panel **e**: the amount of RMSE error in each fit. Panel **f**: a profile through the likelihood surface, starting at our best estimate of the parameters ($\lambda = 0$, yellow square) and finishing (when $\lambda = 1$) at the model's original (Cell # 5) parameter set (Beattie et al., 2018) (with fitted conductance).




Figure 4.19: A comparison of the spread of predictions obtained from fitting our models to recordings of each protocol in Well B09. Panel **a**: the *staircase* protocol (d_1). Panel **b**: the spread-of-predictions obtained by fitting the C-O-I model under Cases I and II. Panel **c**: the spread-of-predictions obtained by fitting the Beattie et al. (2018) model to under Cases I and II.

as-of-yet unseen situations. Additionally, we compute,

$$\mathcal{E}_{\text{fit}} = \frac{1}{N_d} \sum_{d \in \mathcal{D}} \text{NRMSE} \left(I_{\text{Kr}}(\boldsymbol{\theta}_d; \tilde{d}), \mathbf{z}_{\tilde{d}} \right). \quad (4.43)$$

Equation (4.43) is an average over the RMSE between our fits and the data (down-weighted by the size of the recorded trace). Where there are multiple repeats of a protocol (as for our four repeats of d_1), we treat these values as separate protocols in Equations (4.42) and (4.43).

These values (Equation (4.42) and Equation (4.43)) are shown for data taken from Well B20 in the cross-validation heatmap in Figure 4.24. This heatmap shows the NRMSE error for each pair of prediction and validation protocols, and as such, shows which protocols the Beattie et al. (2018) model is able to recapitulate. In particular, we can see that parameter estimates obtained from any of our four sweeps of the *staircase* protocol ($d_1^{(1)}$, $d_1^{(2)}$, $d_1^{(3)}$, and $d_1^{(4)}$) result in little error when compared with data from other sweeps, indicating little variability between sweeps. Moreover, we can see that d_1 and d_2 (the latter being the similar *staircase-in-staircase* protocol) result in relatively poor predictions across many of our protocols.




Figure 4.20: Cross-validation heatmap for the Beattie et al. (2018) model under the assumption that $E_{\text{Kr}} = E_{\text{Nernst}}$ (Case I), showing the accuracy of model fits and predictions when compared to data from Well B20. The square in the i^{th} row and j^{th} column shows the normalised RMSE when the parameter set obtained from the i^{th} trace is used to predict the data taken from the j^{th} trace. The *staircase* protocol (d_1) was repeated four times, twice before all other protocols, and twice after all other protocols. The rows and columns are ordered chronologically, except the row showing d_6 predictions (highlighted), as this protocol is used only for validation and not fitting.

Next, we may consider the *mean* predictive accuracy across each well. Figure 4.21 shows the best and worst wells in terms of predictive accuracy of our Case III Beattie

et al. (2018) model, as quantified by NRMSE averaged over each pair of fitting and validation protocols. Additionally, panels **b** and **d** show best and worst fits/predictions from these models for these particular wells. A similar plot is shown for the Wang et al. (1997) model (also under Case III) in Figure 4.22. At their best, our models produce highly accurate fits to the data. At their worst, they produce predictions which are quantitatively inaccurate, but are still reasonable qualitative descriptions of I_{Kr} . Figure 4.20 shows the NRMSE error for each pair of fitting and validation protocols, averaged across all wells (those that passed QC). Figure 4.25 summarises the predictive accuracy under the assumption of each model, and each fitting case (as described in Table 4.4).

Whilst Figures 4.23, 4.24, and 4.25 provide a summary of predictive accuracy, they do not provide an indication of when the discrepancy between the model and our measurements of I_{Kr} is most prevalent. To highlight periods of time when this discrepancy seems significant, we consider the distribution of predictions at each observation time, t_i . Regarding the accuracy of our model fits, we may consider the summary statistic,

$$\frac{y(t_i) - z_i}{\sigma}, \quad (4.44)$$

where $y(t_i)$ is our model fit at time t_i , and z_i is the datum observed at time t_i . Large positive values of this statistic indicate that our model consistently overestimates the true, underlying DGP at time t_i , whereas large negative values indicate a consistent underestimation, highlighting the discrepancy between our ensemble of predictions and the true underlying DGP. We average this test statistic across wells and show where our models are unable to recapitulate the data in Figure 4.26.

Regarding the accuracy of our model prediction ensembles, we consider,

$$T(t_i) = \frac{\text{mean}_j\{y_j(t_i) - z_i\}}{\frac{1}{\sqrt{N_d}}\text{std}_j\{y_j(t_i)\} + \sigma}, \quad (4.45)$$

where z_i , again, denotes the datum observed at time t_i , $y_i(t_i)$ denotes the prediction of the observation at time, and where mean and std denote the sample mean and sample standard deviation, respectively. Large values of T , whether positive or negative, indicate that the mean error in our predictions is much larger than either the spread in predictions, or the apparent noise in the data set. In this way, T , summarises the ensemble of predictions at a given point in time, and indicates whether they consistently overestimate or underestimate the DGP. The σ term in the denominator accounts for the presence of random noise. This problem would be particularly noticeable at a low holding potential (such as the start of each protocol, when $V_{\text{cmd}} = -80 \text{ mV}$). In this

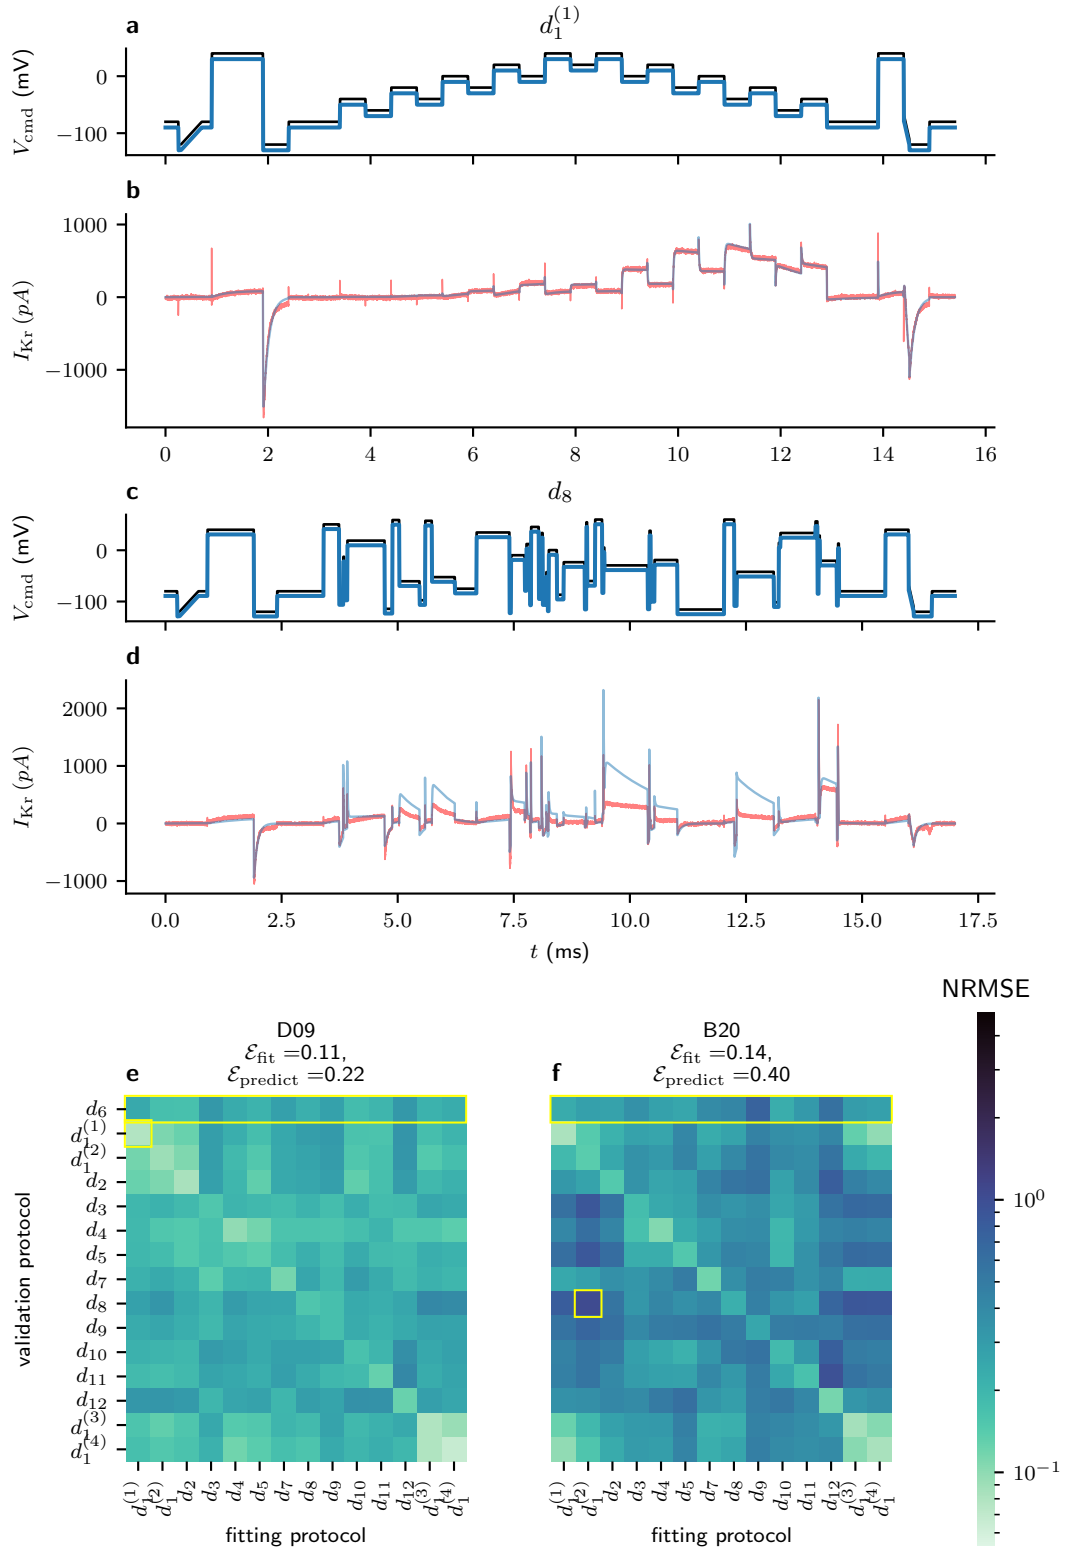


Figure 4.21: A cross-validation heatmap showing the predictive accuracy of the Beattie et al. (2018) model (Case III). The two wells shown, Well D09 (Panel e) and Well B20 (Panel f), are the wells with the least and greatest average NRMSE across each pair of fitting and training sweeps. Panel b shows the best fit (in terms of NRMSE) from Well D09, the highlighted cell in Panel e. The worst model prediction (in terms of NRMSE) for Well B20 is shown in d. Panels a and c show corresponding voltage traces: V_{cmd} (black) and V_{m} (blue).

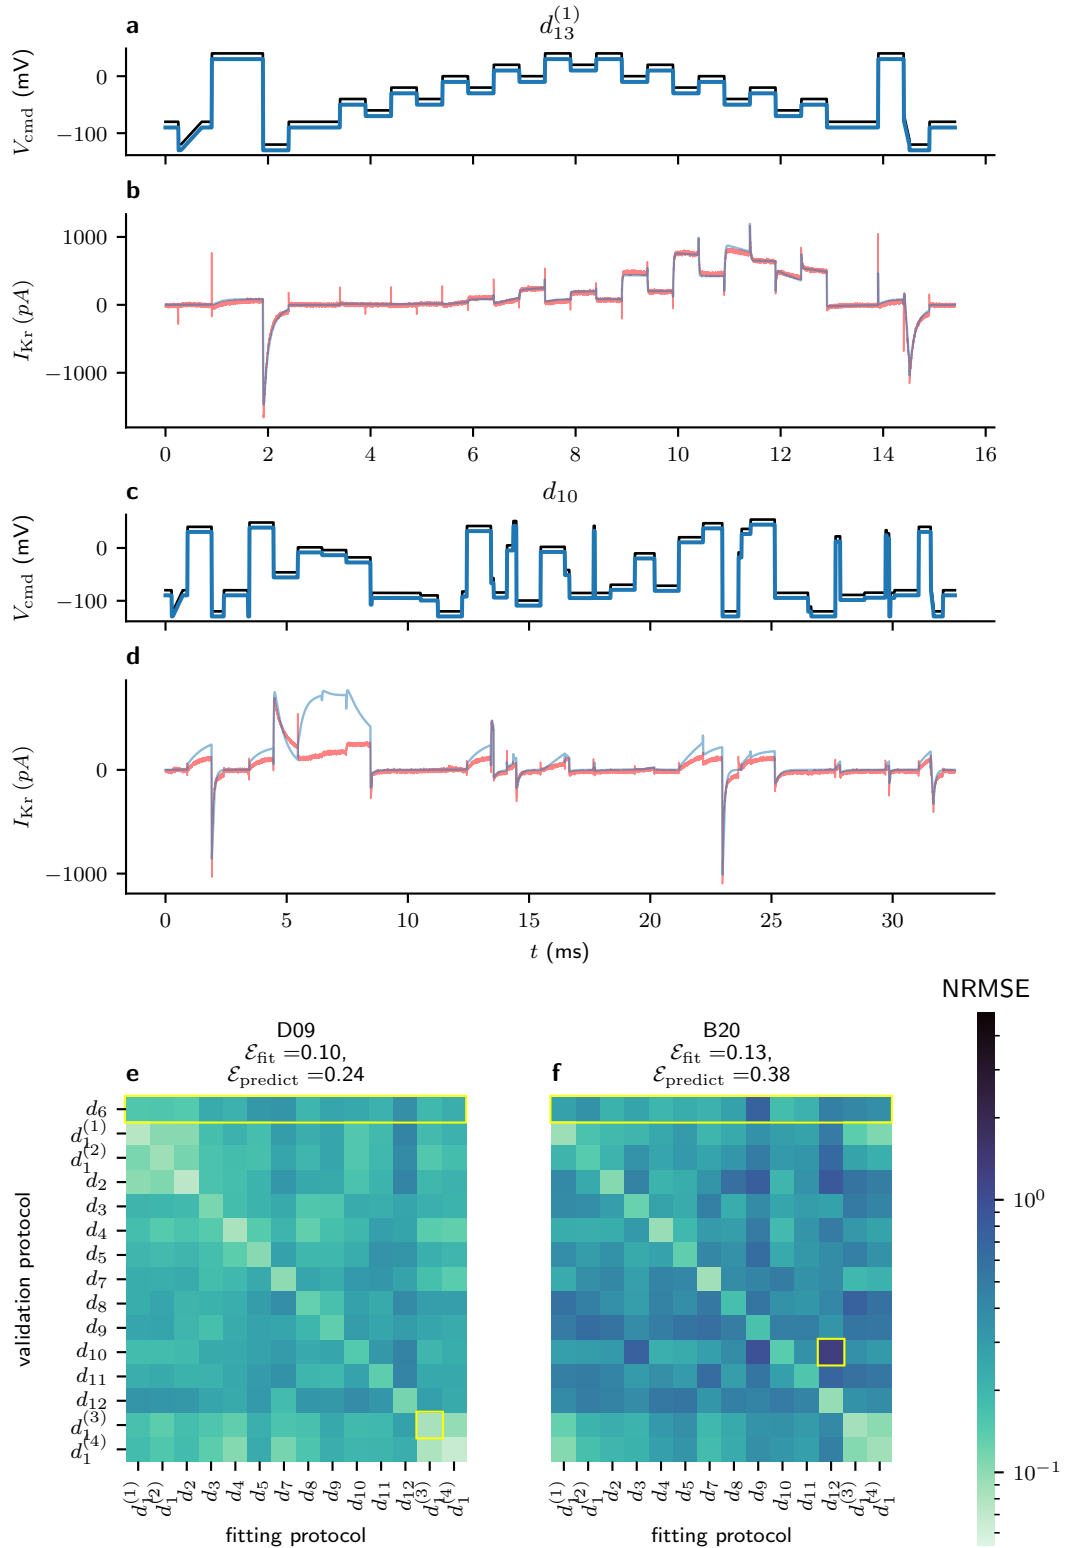


Figure 4.22: A cross-validation heatmap showing the predictive accuracy of the Wang model under Case III. The two wells shown, Well D09 (Panel e) and Well B20 (Panel f), are the wells with the greatest and least average NRMSE across each pair of fitting and training sweeps. Panel b shows the worst model prediction (in terms of the NRMSE) from Well D09, corresponding to the highlighted cell in panel e. Panel c shows the lowest scoring (NRMSE) cell in panel f. Panels a and c show the corresponding voltage traces, V_{cmd} (black) and V_m (blue).

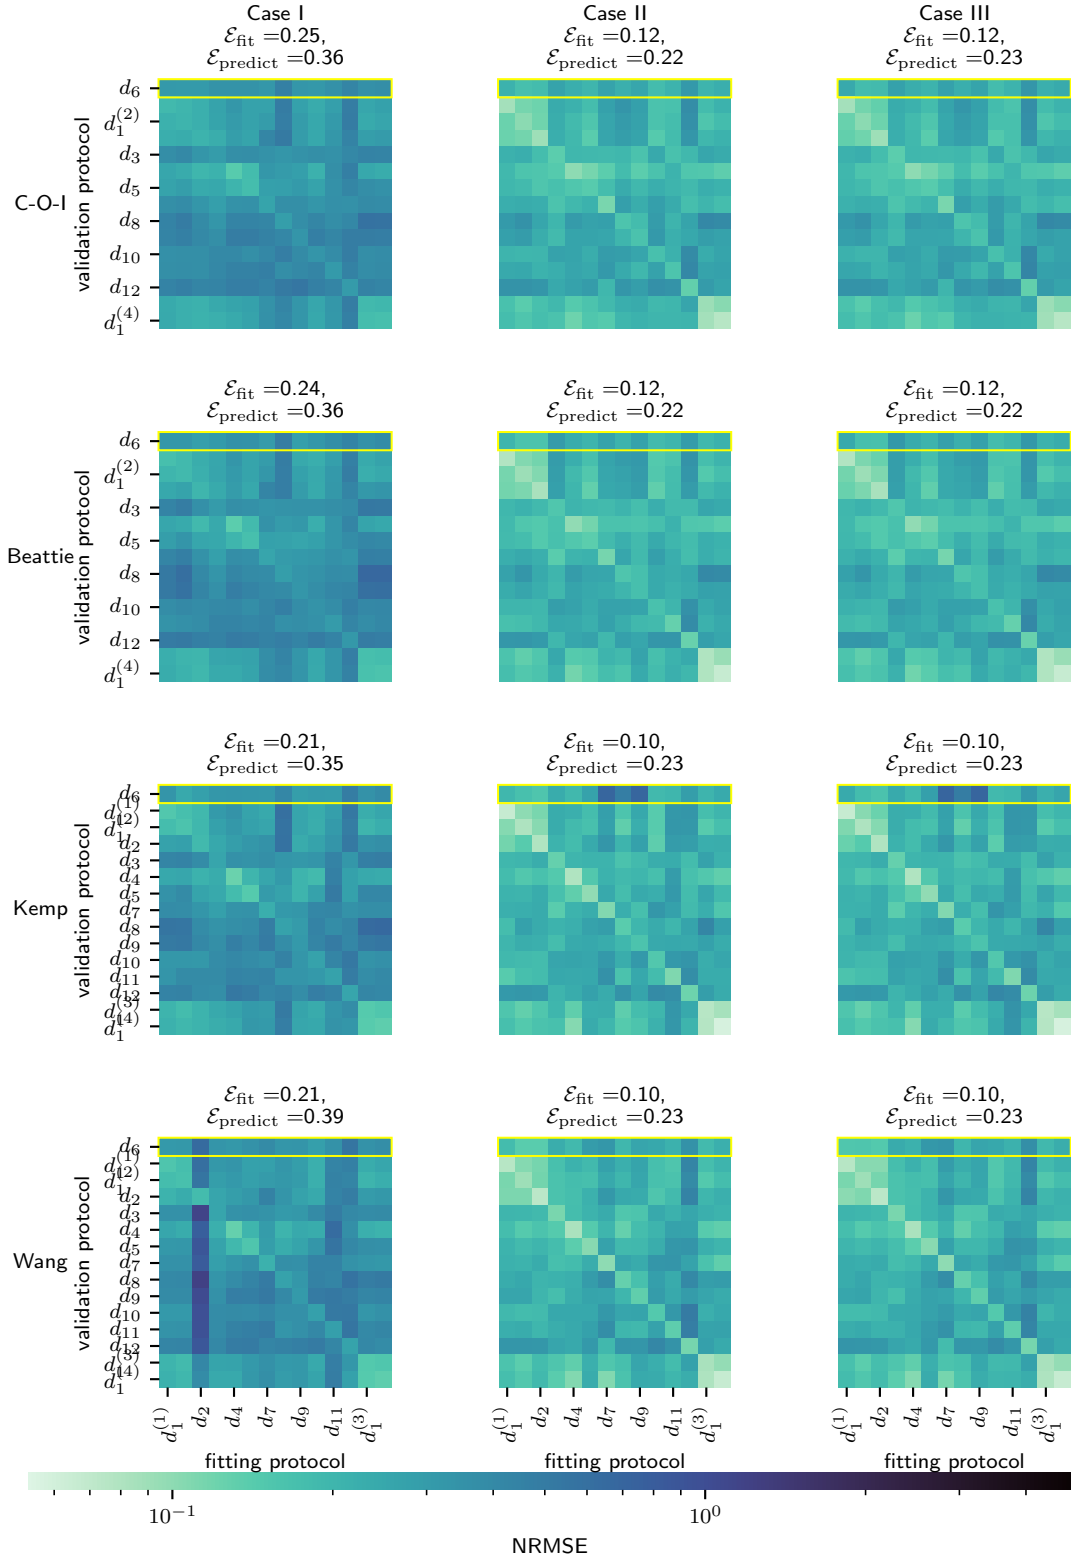


Figure 4.23: A comparison of different combinations of model structures and assumptions, as well as a quantification of the fitting and prediction errors, for data taken from Well D09 (for which our model predictions produce the lowest average NRMSE). Each element of each heatmap shows the prediction error when the model is fitted to data from a particular protocol and used to predict data from another protocol. The traces are ordered chronologically except d_6 (highlighted) which is used only for validation.

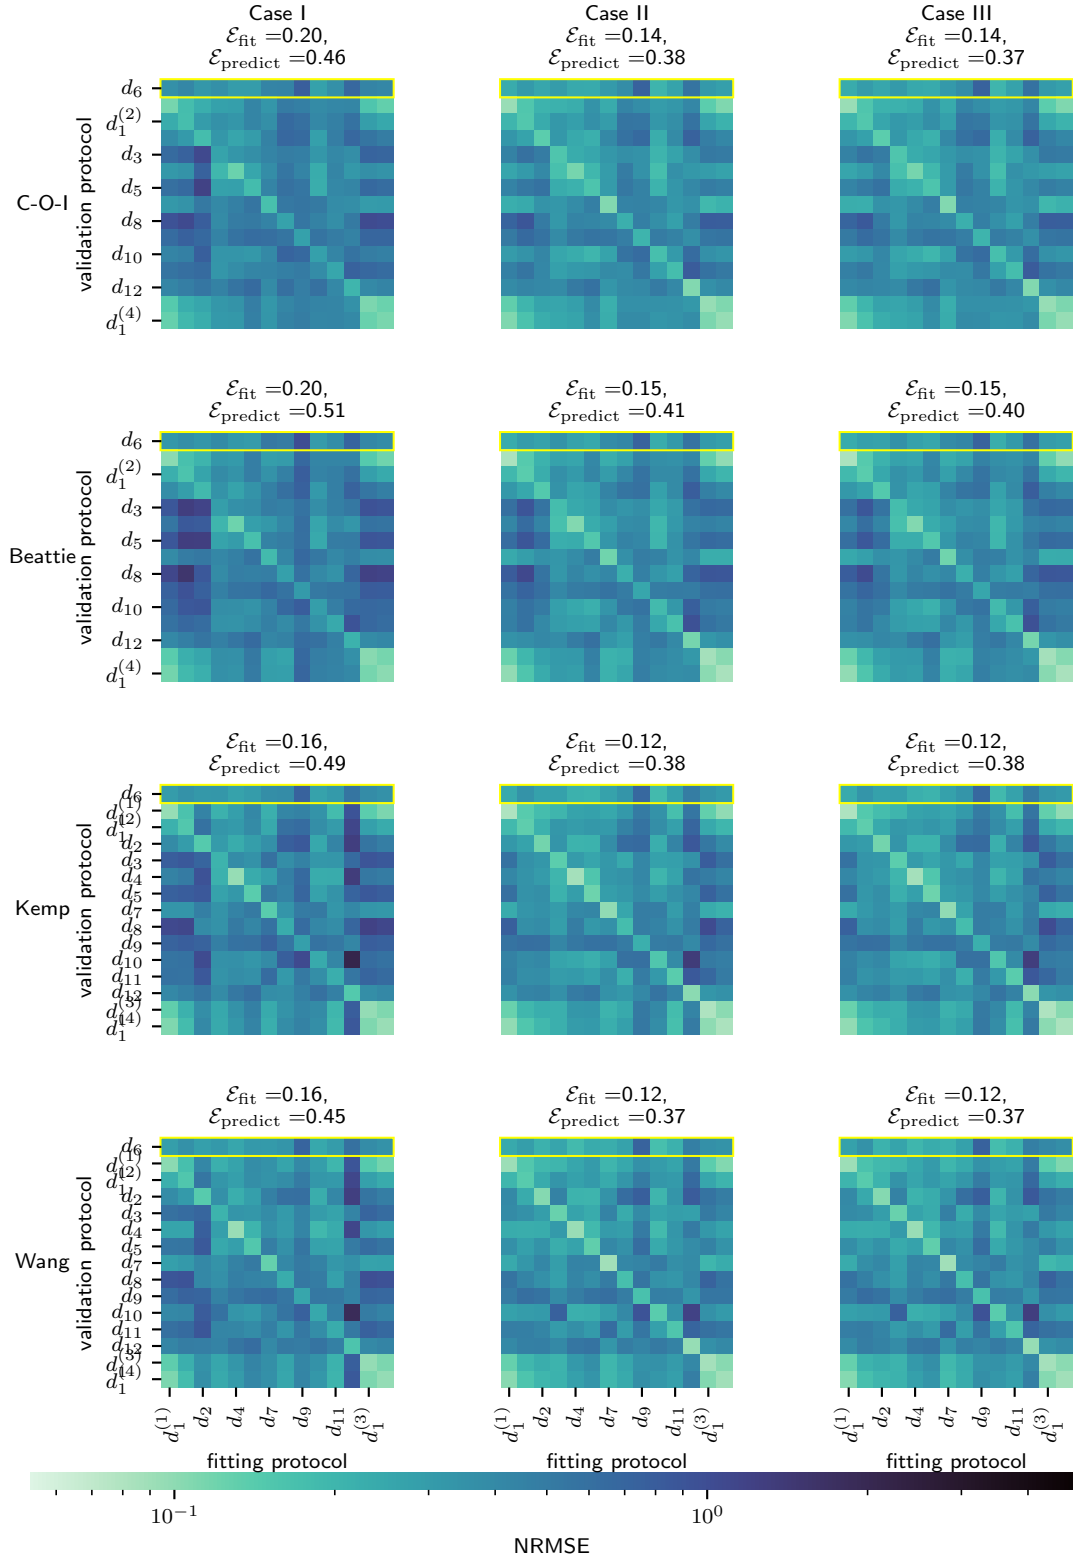


Figure 4.24: A comparison of different combinations of model structures and assumptions, as well as a quantification of the fitting and prediction errors, for data taken from Well B20 (for which our models produce the highest average NRMSE). Each element of each heatmap shows the prediction error when the model is fitted to data from a particular protocol and used to predict data from another protocol. The traces are ordered chronologically except d_6 (highlighted) which is used only for validation.

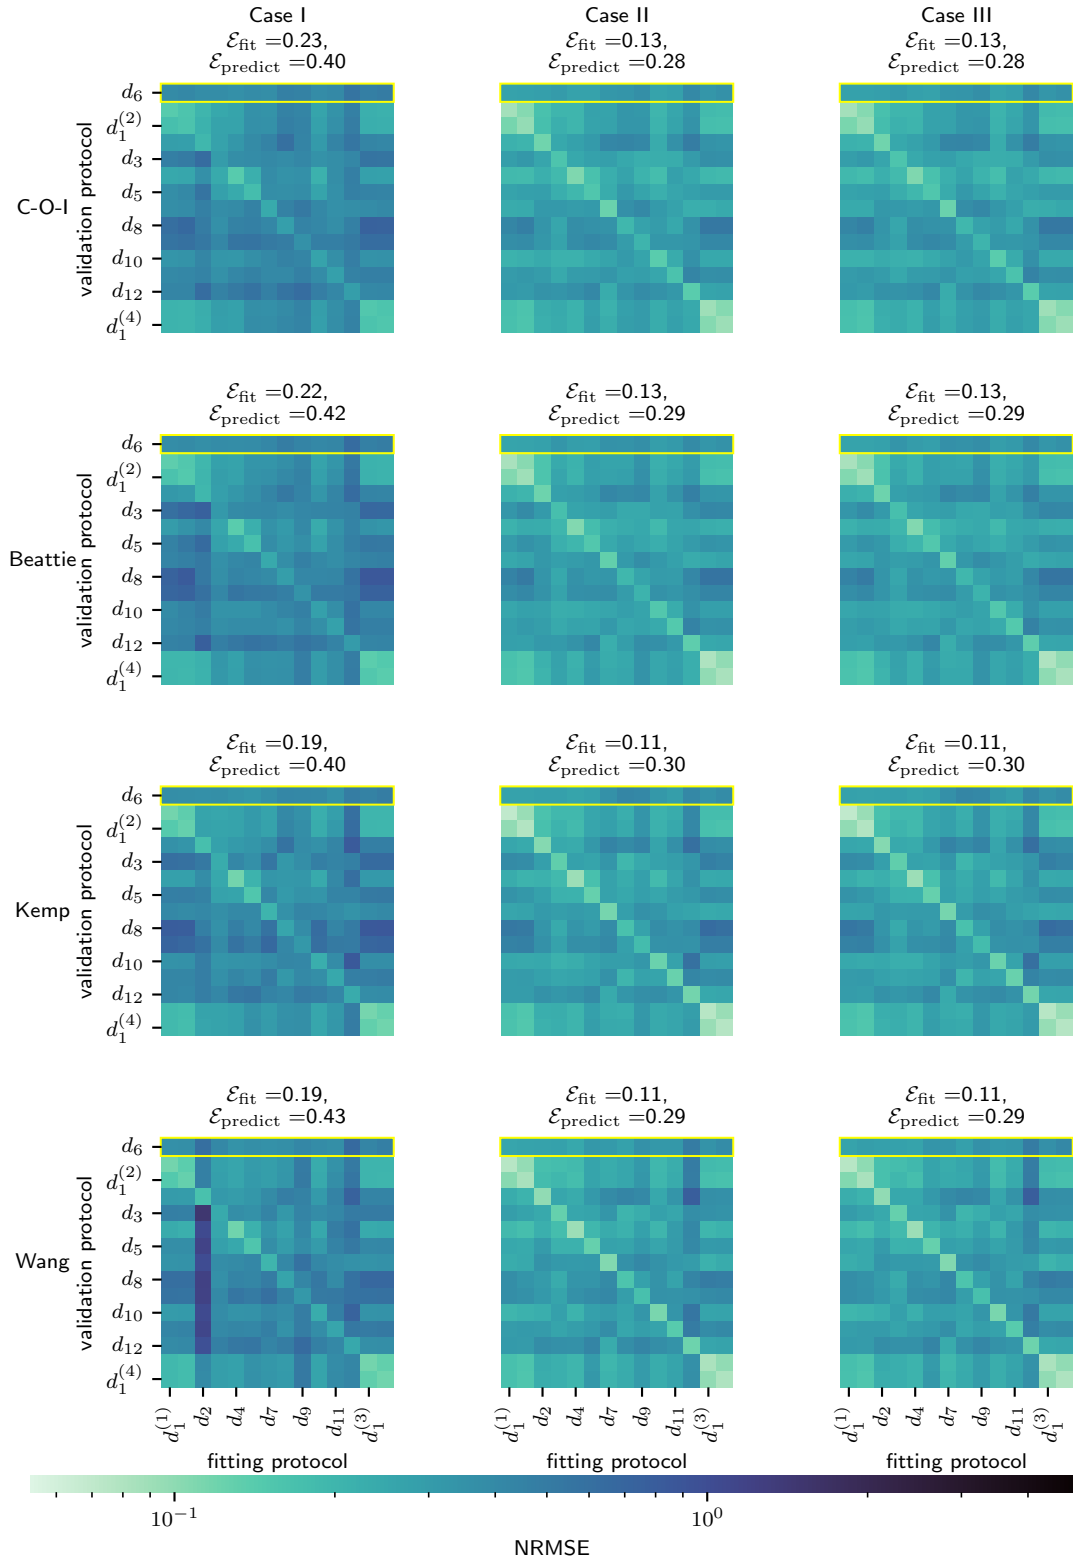


Figure 4.25: A comparison of different combinations of model structures and assumptions, as well as a quantification of the fitting and prediction errors, averaged over each well. Each cell in each heatmap shows the prediction error when the model is fitted to data from a particular protocol and used to predict data from another protocol. The traces are ordered chronologically except d_6 (highlighted) used only for validation.

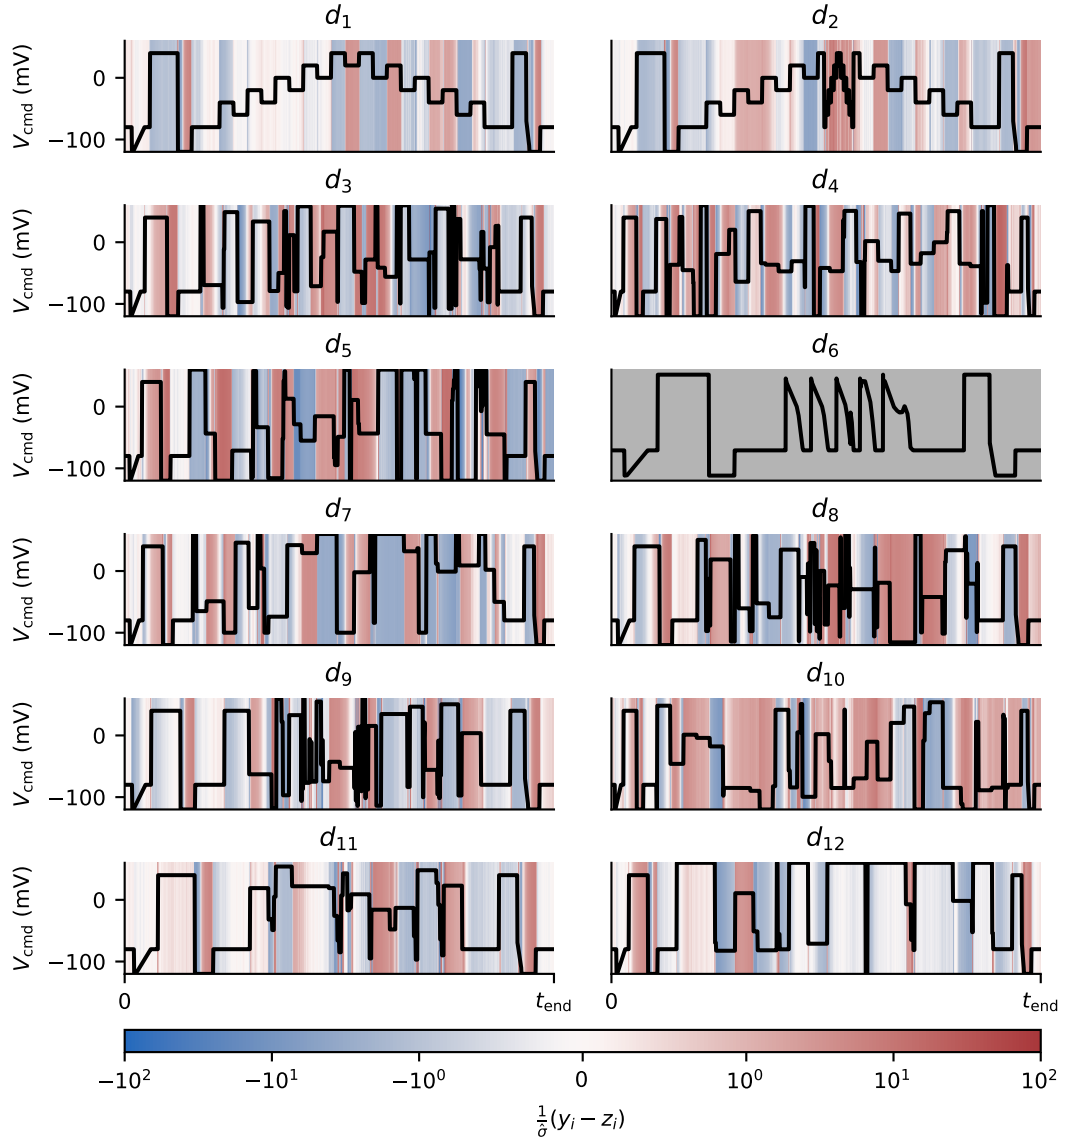


Figure 4.26: The accuracy of our Beattie model fits (under Case II) for each training protocol, as quantified by the normalised residuals averaged across wells. No model was fit to protocol d_6 .

case, I_{Kr} and all model predictions are approximately 0 and we have,

$$\frac{\text{mean}_j\{y_j(t_i) - z_i\}}{\frac{1}{\sqrt{N_d}}\text{std}_j\{y_j(t_i)\}} \approx \frac{\text{mean}_j\{\varepsilon_i\}}{\frac{1}{\sqrt{N_d}}\text{std}_j\{y_j(t_i)\}},$$

where ε_i is the observational error for the i^{th} observation. This value would be highly sensitive to random variability. Moreover, since T is non-dimensional (and is invariant under scaling), it provides a way of averaging this behaviour over multiple wells for which there is variability in maximal conductance, g .

In Figure 4.27, we show T for the Beattie et al. (2018) model (under Case II) for each prediction protocol using data from Well B09. Then, for the same model, we show the same statistic averaged across all wells in Figure 4.28. The fact that the model does not accurately predict the first +40 mV step (which is identical across all protocols), sometimes overestimating it, and sometimes underestimating it, indicates that there are some time-dependent effects which change over the course of the experiment. However, we see larger T values in the central portions of our protocols, indicating dynamics that the model is unable to accurately recapitulate.

4.5 Variability of parameter estimates

By fitting our data to each sweep in our dataset, we obtain a collection of parameter estimates for each model (and fitting case as outlined in Table 4.4). Each of these parameter estimates pertains to a different sweep in our dataset. There are 12 protocols and 8 wells (after the vast majority of wells are removed during QC). All the protocols are performed exactly once (both before and after the addition of dofetilide) except the *staircase* protocol (d_1) of which there are 4 repeats, each of which provides a separate vector of parameter estimates. In this section, we discuss the variability of these parameter estimates and use a simple, linear statistical model to characterise the well-to-well and protocol-to-protocol variability thereof.

By plotting our parameter estimates (Figure 4.29), we can see that there is some relationship between the well the data was obtained from and the resulting parameter-estimate vector (Figure 4.29b). Likewise, we can see that there is some relationship between the protocol under which the data was obtained and the resulting parameter estimate (Figure 4.29c). Some protocols yield consistently high estimates of p_1 when compared to other protocols, for example. However, the relationship is not clear from these plots. Furthermore, due to the high-dimensionality of our parameter space, it is difficult to provide a complete picture of our model. Note that throughout this section, we consider log-transformed ‘A’ parameters, such that for any transition rate of the form $k = A \exp\{bV_m\}$, our model-parameter vector contains $\log(A)$ in place of

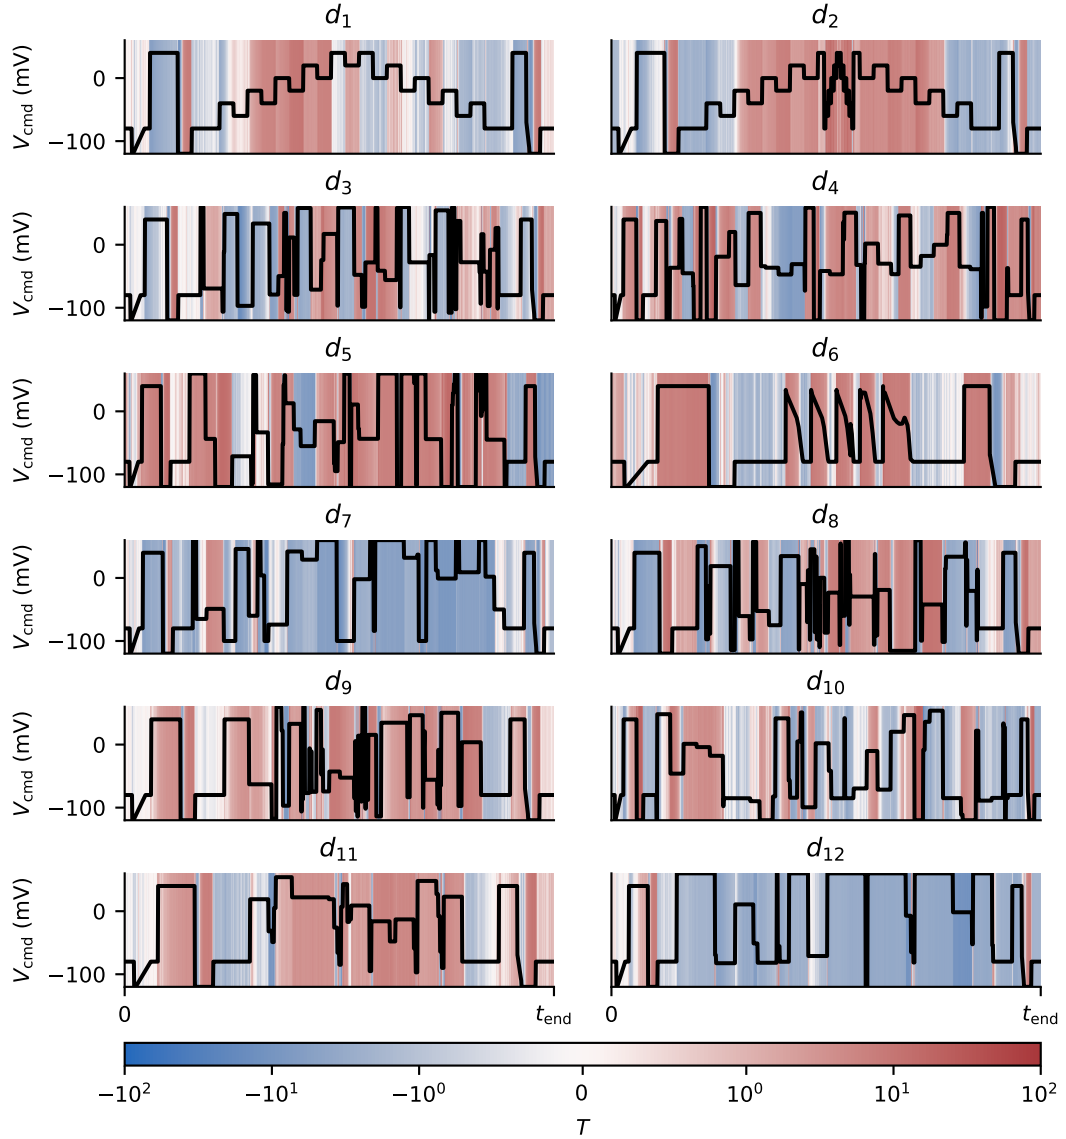


Figure 4.27: Segments of our protocols where the Beattie et al. (2018) model (under Case III) produces inaccurate (yet consistent across parameter sets) prediction, for data obtained from Well B09, as quantified by our T statistic. The areas of our segments where there is significant discrepancy between our ensembles of predictions and the data are shown in dark-red (overestimation) and dark-blue (underestimation).

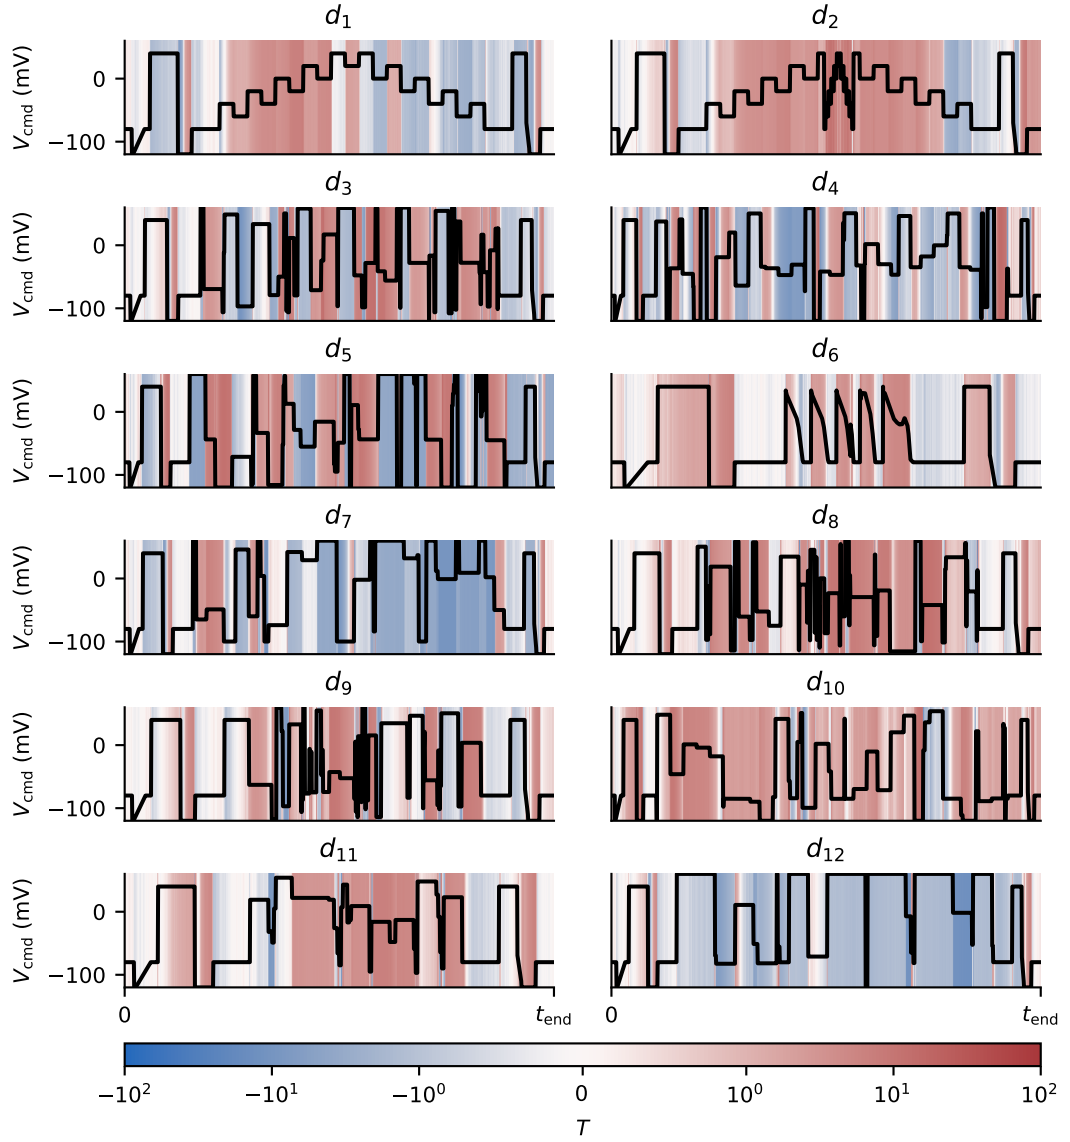


Figure 4.28: Segments of our protocols where the Beattie et al. (2018) model under Case III) produces inaccurate (yet consistent across parameter sets) prediction, as quantified by our T statistic averaged across wells.

A. Nevertheless, all plots shown in this section are plotted in terms of original model parameterisations, that is, those used throughout this thesis.

This fact motivates the introduction of a simple statistical model, which we may use to quantify the variability of our parameter estimates. To do this, we assume that a parameter-estimate vector obtained from well i and protocol j is vector of random variables,

$$\boldsymbol{\theta}_{i,j} = \mathbf{w}_i + \mathbf{d}_j + \boldsymbol{\varepsilon}, \quad (4.46)$$

where: $\boldsymbol{\varepsilon}$ is some random variable with an uncorrelated, zero-mean, Gaussian distribution, $\text{MVN}(\mathbf{0}, \text{diag}\{\sigma_k\}_k)$, where each σ_k being the standard deviation corresponding to the random variation in each parameter; and each \mathbf{w}_i and \mathbf{d}_j is an effect related to said well and protocol respectively.

By performing linear regression, we infer the sizes of these effects for each well and each protocol fitted in the dataset. We denote the matrix of well and protocol effects by,

$$\boldsymbol{\beta}_w = \begin{pmatrix} \boldsymbol{\beta}_w^{(1)} \\ \vdots \\ \boldsymbol{\beta}_w^{(N)} \end{pmatrix}, \quad \boldsymbol{\beta}_d = \begin{pmatrix} \boldsymbol{\beta}_d^{(1)} \\ \vdots \\ \boldsymbol{\beta}_d^{(N)} \end{pmatrix}.$$

These are matrices representing the well and protocol effects, respectively. Each column represents the well or design effects on each parameter, and each row corresponds to the particular well and protocol from which a given parameter estimate was obtained.

We may also write this model in matrix form,

$$\mathbf{Y} = \boldsymbol{\mu} + \mathbf{X}_w \boldsymbol{\beta}_d + \mathbf{X}_d \boldsymbol{\beta}_w + \mathbf{E}, \quad (4.47)$$

where: \mathbf{Y} is a matrix where each row is a parameter estimate obtained from model fitting; \mathbf{X}_w and \mathbf{X}_d are well- and protocol-effect *design matrices* where each row encodes the well/protocol used for a given parameter estimate; $\boldsymbol{\beta}_w$ and $\boldsymbol{\beta}_d$ are our (column) parameter vectors, with each element representing a well and protocol effects (respectively); \mathbf{E} is a matrix of Gaussian-distributed random variables such that the k^{th} column of \mathbf{E} is a vector of IID Gaussian variables with zero mean and standard deviation, σ_k .

In this model, we consider $N_d = 11$ designs (of which d_1 is repeated four times) to $N_w = 8$ wells. Hence, we have $N = (3 + N_d)N_w$ data (each a separate vector of parameter estimates, $\hat{\boldsymbol{\theta}}$). To test whether there is significant well- and/or protocol-dependence in our parameter estimates, we use *log-likelihood differences* (LLDs). We do this by considering a family of nested models where the full model is,

$$\mathcal{M}_{w,d} : \mathbf{Y} = \boldsymbol{\mu} + \mathbf{X}_d \boldsymbol{\beta}_d + \mathbf{X}_w \boldsymbol{\beta}_w + \boldsymbol{\varepsilon}, \quad (4.48)$$

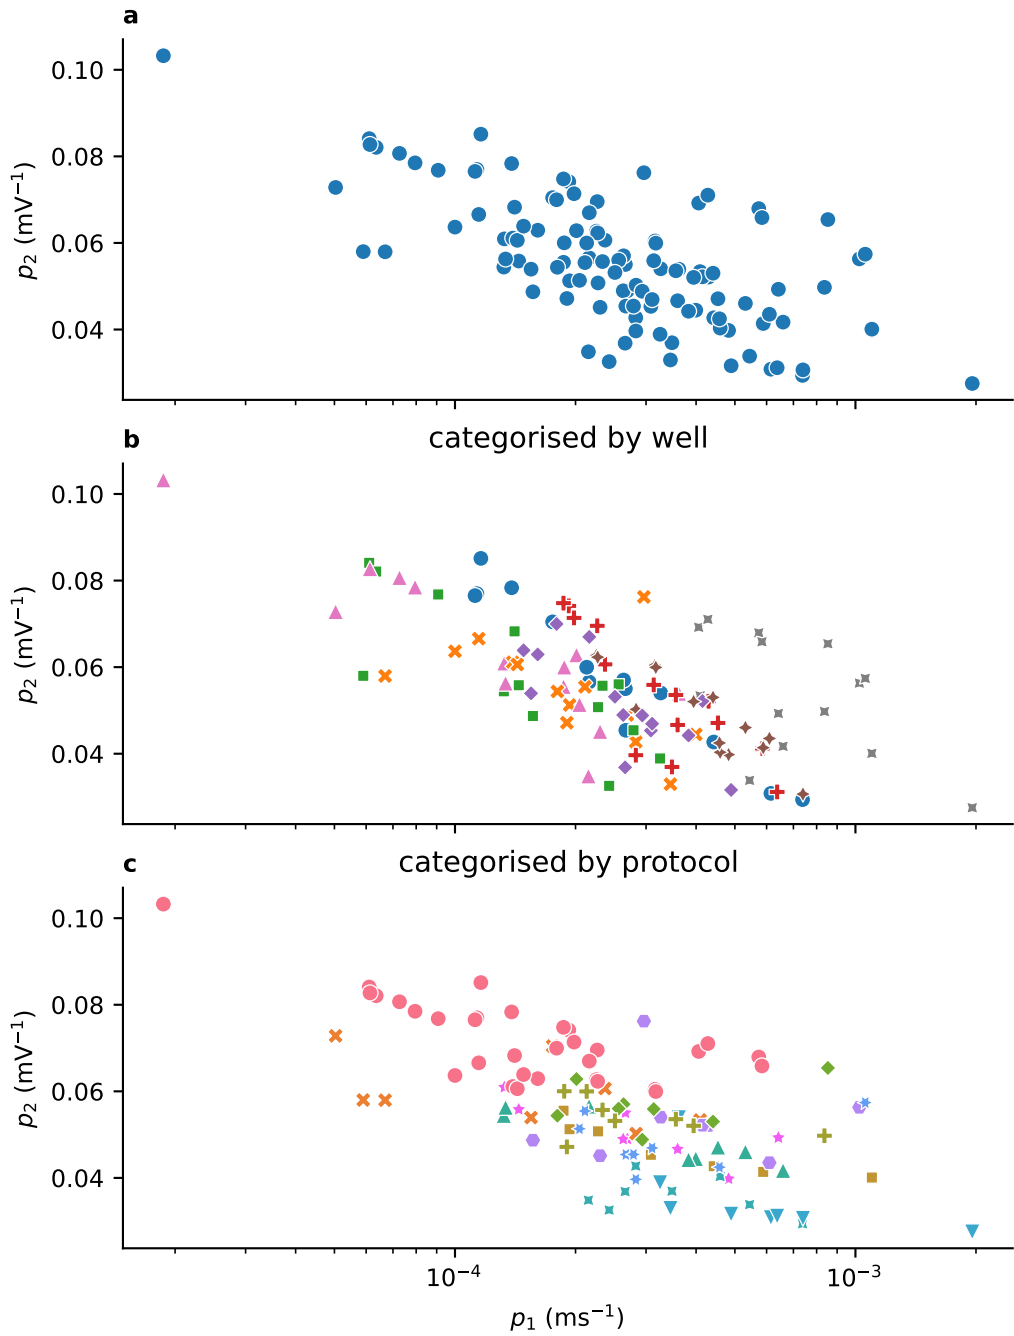


Figure 4.29: Panel **a**: every estimate the p_1 and p_2 Beattie et al. (2018) model parameters obtained under Case III. Panel **b**: the different colours and markers show which estimates correspond to the same well. Panel **c**: the different colours and markers show which parameter estimates were obtained from the same protocol. The protocol and well dependence of these parameter estimates is explored in more detail in Section 4.5.

and the protocol and well effects sum to zero that is,

$$\boldsymbol{\beta}_w^\top \mathbf{1} = \boldsymbol{\beta}_d^\top \mathbf{1} = \mathbf{0}. \quad (4.49)$$

Consequently, μ is an average of our parameter-estimate vectors across all wells and protocols (discounting repeats).

These constraints, together with sufficient observations, guarantee identifiability. Here, our linear model has $N_\theta(N_w + N_d + 1)$ independent parameters (of which N_d correspond to the standard deviations of our random errors, σ_i). Hence, we are able to identify each of the parameters (elements of $\boldsymbol{\beta}_w$ and $\boldsymbol{\beta}_d$) only if there are more data than independent parameters, that is if,

$$N_\theta N \geq N_\theta(N_w + N_d + 1), \quad (4.50)$$

where N is the number of parameter estimate vectors we have (that is, the number of traces used for fitting). Since all protocols are performed exactly once except the *staircase* protocol which is performed 4 times, this inequality holds if,

$$N_\theta N_w(N_d + 3) \geq N_\theta(N_w + N_d + 1) \quad (4.51)$$

$$N_w(N_d + 3) - \geq N_w + N_d + 1 \quad (4.52)$$

$$N_w(N_d + 2) - N_d - 1 \geq 0 \quad (4.53)$$

$$2N_w + N_d(N_w - 1) - 1 \geq 0, \quad (4.54)$$

which is always true when $N_w, N_d \geq 1$, that is, when we have parameter estimates obtained from at least one well and one protocol. Note that this is a family of nested models: \mathcal{M}_0 results from any of the other models subject to $\boldsymbol{\beta}_w = \boldsymbol{\beta}_d = \mathbf{0}$. Similarly, models \mathcal{M}_w and \mathcal{M}_d may be seen as model $\mathcal{M}_{w,d}$ subject to the constraints that $\boldsymbol{\beta}_w = \mathbf{0}$ and $\boldsymbol{\beta}_d = \mathbf{0}$, respectively. This hierarchy of models is shown in Figure 4.30.

Each model may be fit by rewriting it as a standard linear regression model,

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad (4.55)$$

from which we obtain the maximum-likelihood estimate, $\hat{\boldsymbol{\beta}}$, which satisfies,

$$\mathbf{X}^\top \mathbf{X} \hat{\boldsymbol{\beta}} = \mathbf{X}^\top \mathbf{Y}. \quad (4.56)$$

After solving this equation, we may obtain $\boldsymbol{\beta}_w$ and $\boldsymbol{\beta}_d$ from $\hat{\boldsymbol{\beta}}$.

We consider the maximal likelihood under the assumption of each model, \mathcal{M}_i , by $L_i(\hat{\boldsymbol{\beta}}_i)$ where $\hat{\boldsymbol{\beta}}_i$ is the MLE of the model parameters for model i . We list each of these values, for each set of parameter estimates (arising from our candidate I_{Kr} models)

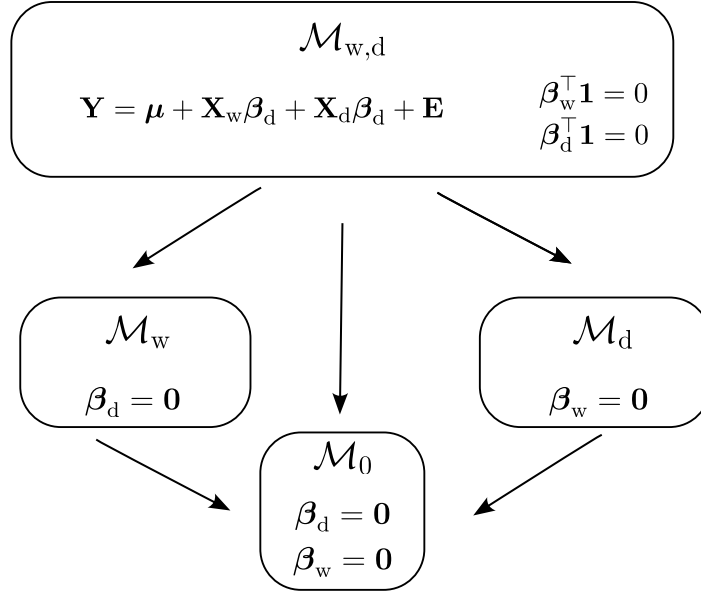


Figure 4.30: The relationship between the linear models applied to our collection of parameter estimates. All of these models are sub-models of $\mathcal{M}_{w,d}$ under various constraints. These constraints correspond to the removal of the well- and/or protocol-dependent effects. In any case, we require that our effects vectors, $\boldsymbol{\beta}_w$ and $\boldsymbol{\beta}_d$. The null model (\mathcal{M}_0) is the simplest model in this family, having no well- or protocol-dependent effects.

in Table 4.5, from which we can see that including the well- and protocol-dependent effects leads to a large increase in likelihood. This indicates that there is a large degree of protocol and well dependence in our parameter estimates.

In Figure 4.31, we see the variability in these estimates for the (Beattie et al., 2018) model (Case I) across each sweep of each protocol for each well. The repeated protocol (the *staircase* protocol) is highlighted and compared with the linear model. Though this plot shows only a pair of transition-rate parameters, we can see that the linear model captures (to some extent) the protocol and well dependence of all transition-rate parameter estimates in a given Markov model. In other words, it appears that well-to-well and protocol-to-protocol differences in parameter estimates can mostly be explained as simple, independent translations in parameter space. The suitability of these models is discussed further in Appendix E, where we show that the assumption of IID normally distributed errors (that is, Equation (4.55)) is reasonable.

The parameter estimates that we obtained are shown in Figures 4.31. From these results, we can see that there is noticeable variability between parameters obtained from different wells under the same protocol, and also variability between wells in those parameter estimates obtained from the same protocol. It is noteworthy that parameter estimates obtained under repeats of d_1 are largely consistent, except in the case of Well B20, where the p_1 parameter estimate obtained from the second repeat is noticeably different, as shown in Figures 4.31 and 4.32. We suggest that this is due to the over-




Figure 4.31: The full linear model, $\mathcal{M}_{w,d}$, containing both well- and protocol-dependent effects describes the variability in our parameter estimates. Each panel shows the parameter estimates obtained from a given well, and highlights the four parameter estimates obtained from the *staircase* protocol (red crosses), the only protocol which was repeated. The blue square shows the well-dependent effect according to \mathcal{M}_w , our linear model with no protocol-dependent effects, and the blue star shows the sum of the well and protocol effects according to $\mathcal{M}_{w,d}$. Note that p_1 , on the x-axis, has been log-scaled (as it is in each linear model).

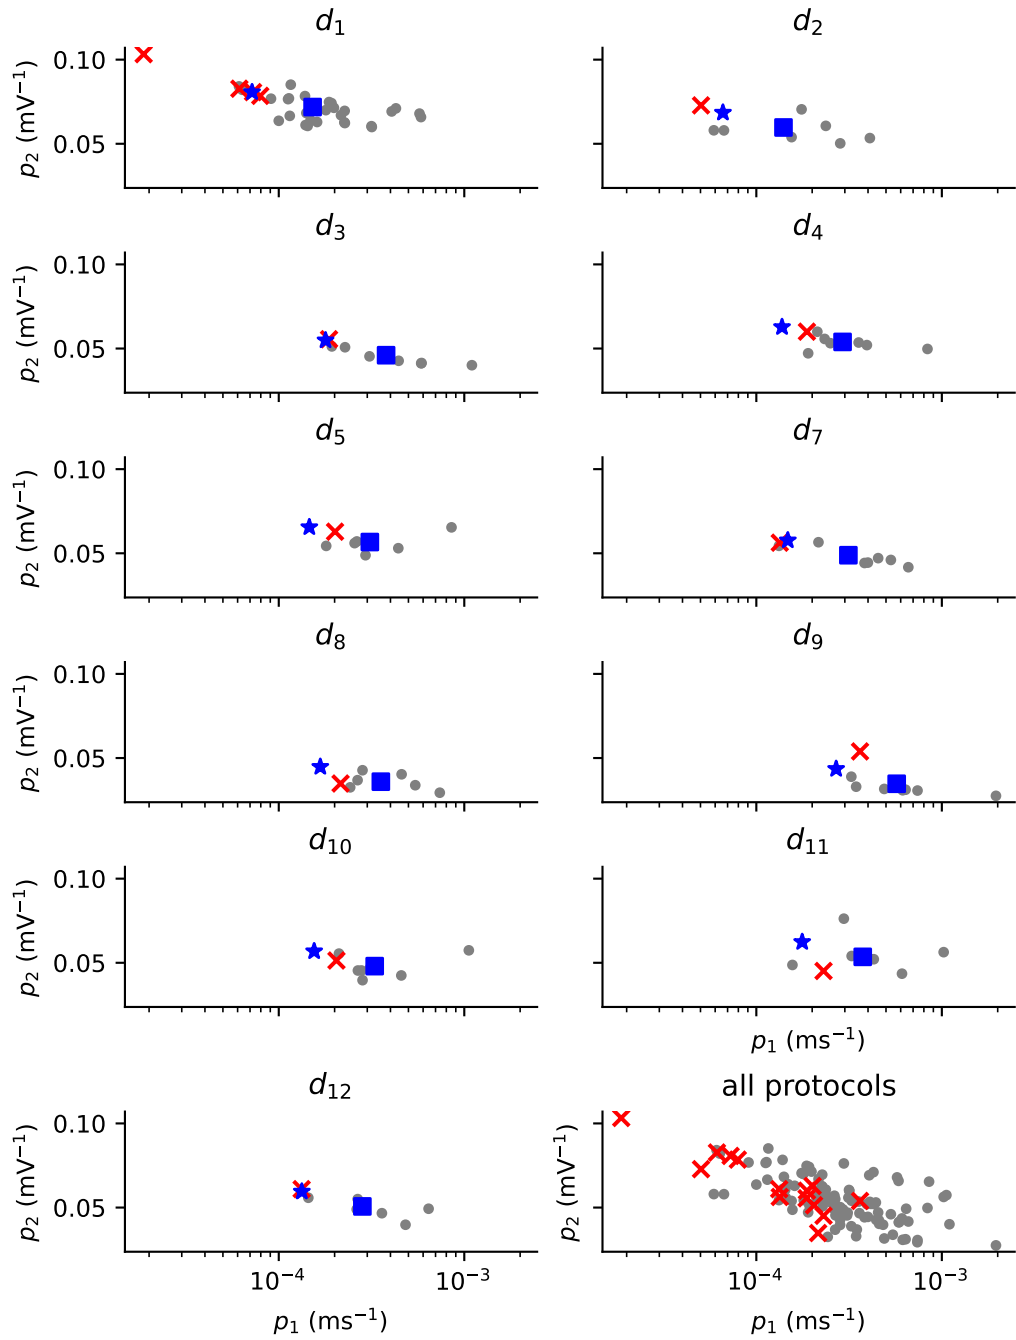


Figure 4.32: Each panel shows the parameter estimates obtained from a given protocol, and highlights the parameter estimates obtained by fitting with data from Well B20 (red crosses), which exhibits noticeable variability under repeats of d_1 . The blue square shows the well-dependent effect according to \mathcal{M}_w , our linear model with no well-dependent effects, and the blue star shows the sum of the well and protocol effects according to $\mathcal{M}_{w,d}$ —our linear model with both well- and protocol-dependent effects. Note that p_1 , on the x-axis, has been logged-scaled (as it is in our each linear model).

correction of the leak current, which is perhaps most noticeable during the first +40 mV step (as shown in Figure 4.33).

Due to the high-dimensional nature of our parameter spaces, it is difficult to provide a complete picture of the variability in our parameter estimates. Nevertheless, we may summarise this information by using some log-likelihood differences (LLDs)—this is made possible by the fact that \mathcal{M}_i are nested models, as illustrated in Figure 4.30. In particular, this provides a way of quantifying the relative significance of the well-dependent and protocol-dependent effects.

The first LLD we calculate, **LLD-W**, shows whether there is a significant well-effect where we find the difference in the maximum log likelihoods of models \mathcal{M}_d and $\mathcal{M}_{w,d}$. We then compute **LLD-D** where we take the difference in log likelihoods between \mathcal{M}_w and $\mathcal{M}_{w,d}$. In either case, the log-likelihood difference is,

$$\log \left\{ \frac{L_1(\hat{\boldsymbol{\theta}}_1)}{L_2(\hat{\boldsymbol{\theta}}_2)} \right\} = \log \left\{ L_1(\hat{\boldsymbol{\theta}}_1) \right\} - \log \left\{ L_2(\hat{\boldsymbol{\theta}}_1) \right\}, \quad (4.57)$$

where: L_1 and L_2 are the maximum likelihoods found under the first and second models, respectively. These statistics, presented in Table 4.5, quantify the importance of well and protocol dependence in our ensemble of parameter estimates.

Model	Case	\mathcal{M}_0	\mathcal{M}_w	\mathcal{M}_d	$\mathcal{M}_{w,d}$	LLD-W	LLD-D
C-O-I	Case I	1081.9	1172.0	1301.0	1467.0	166.0	295.0
	Case II	1451.5	1657.0	1709.7	2104.7	395.1	447.7
	Case III	1465.6	1668.5	1724.4	2111.1	386.7	442.6
Beattie	Case I	1231.2	1361.5	1469.3	1723.4	254.1	362.0
	Case II	1547.4	1804.0	1757.0	2219.3	462.3	415.4
	Case III	1563.3	1817.0	1784.2	2247.2	462.9	430.2
Kemp	Case I	118.9	1220.1	1486.5	1671.3	184.8	451.2
	Case II	914.7	1081.8	1096.0	1337.2	241.2	255.4
	Case III	931.8	1099.8	1113.7	1356.3	242.6	256.5
Wang	Case I	−1919.8	−1851.5	−1578.2	−1474.3	103.9	377.1
	Case II	−506.4	−379.0	−216.6	14.6	202.0	364.4
	Case III	−2082.1	−1976.1	−1836.7	−1673.5	163.2	302.6

Table 4.5: Log-likelihoods and log-likelihood differences (LLD-W and LLD-D) for each of the linear models, applied to all of our models (that is, the collection of parameter estimates obtained using each model). Here we see that in each case, and for each model, that both well- and protocol-effects are sizeable. Whilst large values of **LLD-D** suggests discrepancy between the recordings of I_{Kr} (taken from each well) and the dynamics of our models, the size of **LLD-W** may suggest the presence of latent, well-dependent effects.

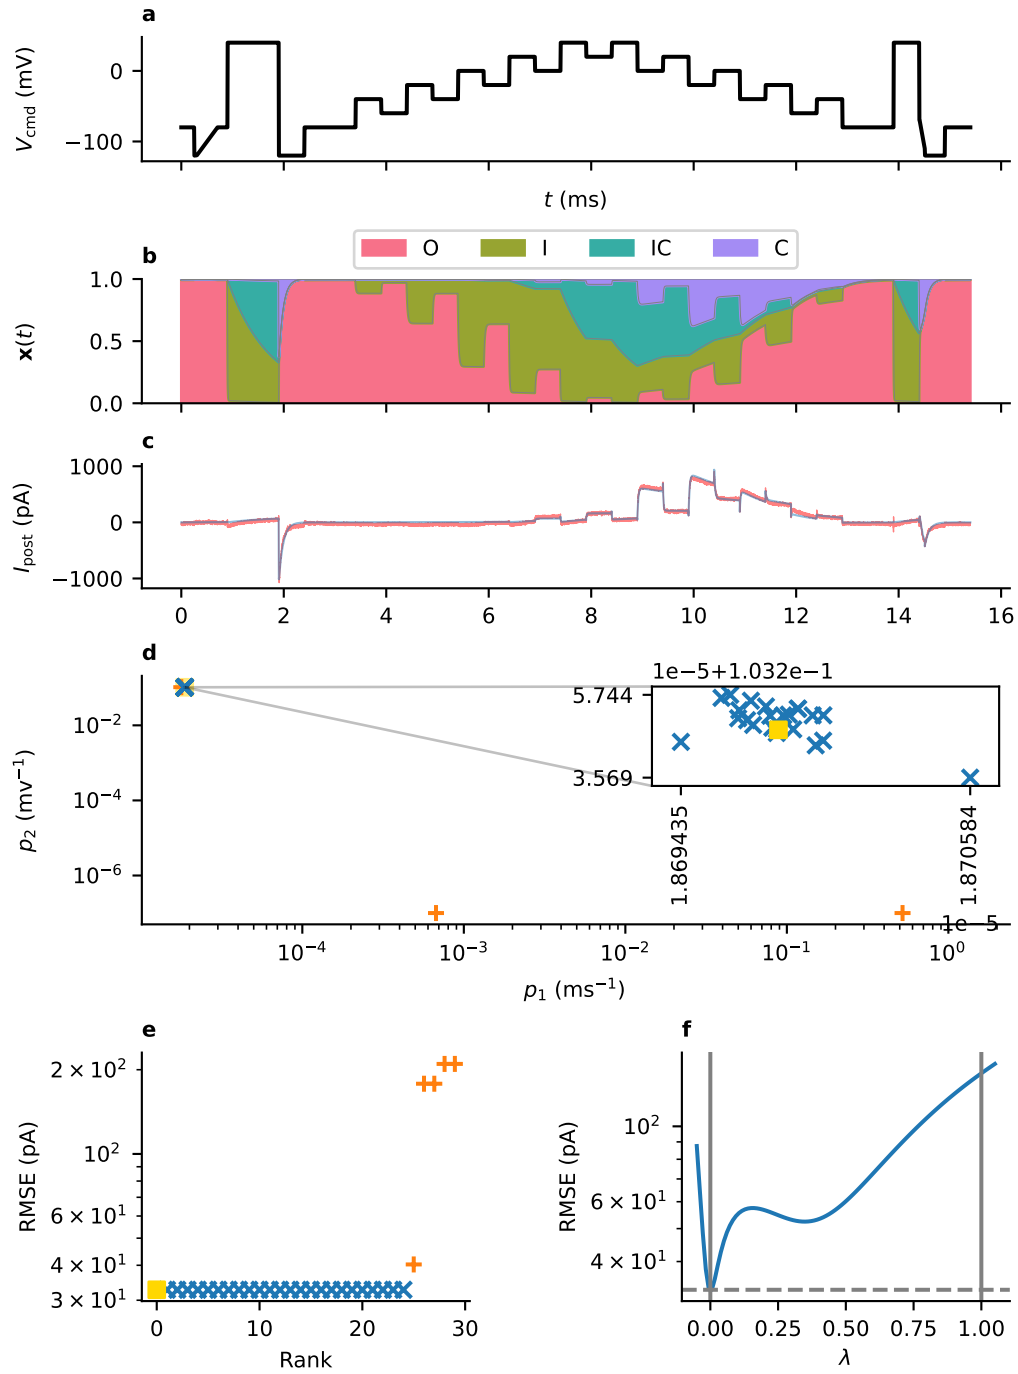


Figure 4.33: Results from fitting the Beattie et al. (2018) model under Case II/III with data collected from Well B20 under the *staircase* protocol, d_1 (second repeat). Panel **a**: the data and the model fit. Panel **b**: the command voltage during this protocol. Panel **c**: estimates of p_1 and p_2 obtained from thirty repeats of the optimisation routine—blue markers show optimisation results with RMSE at most 101% of the best-found value. Panel **d**: the amount of RMSE error in each fit. Panel **e**: a profile through the likelihood surface, starting at our best estimate of the parameters ($\lambda = 0$, yellow square) and finishing (when $\lambda = 1$) at the model's original (Cell # 5) parameter set (Beattie et al., 2018), with its conductance parameter set to that of the fitted model.

4.6 Discussion

In the previous chapter, we demonstrated methods for the application of ODE-based models of I_{Kr} to real time-series data from information-rich experimental designs. However, the previous chapter was focused on synthetically generated data, avoiding many of the challenges that arise when fitting mathematical models to real experimental data. In this chapter we adapted and applied the methodology to work with real experimental data. Our multiprotocol approach necessitates methods which ensure consistent conditions throughout the experiment. Through the diligent selection of QC criteria, we were able to select a subset of the data consisting of high-quality recordings (where the linear leak model seems to work exceptionally well), which exhibit consistency over the course of the experiment (as shown by the minimal variation in leak currents and reversal potentials over time). This aspect of this chapter lays the foundation for further analysis of multiprotocol datasets of ion-channel current recordings, particular those taken from high-throughput automated patch-clamp platforms.

Following the quality control and postprocessing work, in Section 4.4 we used our information-rich multiprotocol dataset to perform thorough cross-validation of the predictive accuracy of a selection of literature I_{Kr} models. Though these models were applied under the assumption of an ideal patch-clamp, these models produced reasonably accurate predictions across a range of protocols. In terms of general predictive performance (when predictive error is averaged across a range of fitting and validation protocols), it appears that the simpler models (with fewer parameters) perform better than models with larger parameter spaces. Generally speaking, models with more parameters (the Kemp et al. (2021) and Wang et al. (1997) models) produced more accurate fits to our data, as quantified by \mathcal{E}_{fit} , whereas the simpler models (the C-O-I and Beattie et al. (2018) models) produced slightly more accurate model predictions for unseen protocols (see Figure 4.25). However, the difference between our competing model structures regarding \mathcal{E}_{fit} and $\mathcal{E}_{\text{pred}}$ are rather subtle.

These results show that whilst our model structures are able to mostly recapitulate the individual voltage-protocol that they are being fitted to, their predictions for other protocols (which they were not fitted to) is markedly worse. This behaviour, combined with the protocol-dependence of our parameter estimates, indicates model discrepancy. We suggest that it may be possible to improve the accuracy of our predictions by selecting other model structures or reconsidering our modelling assumptions (in the next chapter, namely the ideal patch-clamp assumption). Because our data were collected from a diverse range of information-rich experiments, we are able to thoroughly validate the predictive accuracy of a small selection of I_{Kr} models under previously unseen stimuli. Overall, we have shown that each model is able to mostly accurately recapitulate our I_{Kr} recordings, producing quantitatively accurate fits and predictions. This is

particularly true in certain wells, such as Well D09, where our models not only provide very accurate fits to the data, but are able to predict the current during unseen protocols to a high degree of accuracy (see Figure 4.25).

There was some variability between sweeps, for example, in the parameters obtained from repeats of the *staircase* protocol in Well B20. We suggest that this variability was due to the overcorrection of leak current. This variability had a large effect on the resultant p_1 parameter estimate, which shows that small instances of data pollution have the potential to impart large bias onto parameter estimates. In future work, it may be possible to avoid such overcorrection of leak by enforcing stricter QC criteria, or adjusting our inference of the leak-model parameters such that any negative I_{post} during the +40 mV step following the leak ramp is heavily discouraged. However, in this chapter, only 8 of 384 wells were considered, providing a balance between the quality and volume of data—further QC criteria should be carefully selected as to not discard more data. Nevertheless, this particular example of leak over-correction is avoided in the following chapter, where we fit our models directly to the pre-drug trace.

Our statistical analysis of our ensembles of parameter estimates (as discussed in Section 4.5) suggests that there are strong well- and protocol-dependent effects acting upon our parameter estimates. As discussed in the previous chapter, we expect these protocol-dependent effects to be present when there is discrepancy between our mathematical models and the underlying biophysical mechanisms. The strong well-dependent effects, however, suggest substantial experimental variability, unaccounted for by the models presented here. As argued by Lei et al. (2020a), it is possible that *experimental artefacts* are a leading cause of the well-to-well variability in our parameter estimates. Perhaps the variability in our kinetic parameters is caused by the overfitting to account for these experimental artefact effects, which are not present in our models. Moreover, the inclusion of these effects reduced protocol dependence in parameter estimates and led to more accurate model predictions. In the following chapter, we consider models that include these artefact effects (as presented in Lei et al. (2020a)) and investigate whether they provide more accurate predictions and more uniform parameter estimates across wells and protocols.

Chapter 5

Accounting for experimental artefacts in fitting improves model accuracy

5.1 Preamble

The work presented in this chapter is based on the same data as the previous chapter, collected at the Victor Chang Cardiac Research Institute (VCCRI), Sydney, Australia. This work was done before, during and after my working visit to The VCCRI in November/December 2023, where I collected further data and gained insights into experimental methods for patch-clamp electrophysiology. This experience shaped my treatment of artefact effects in our models, leading to the development of the methods presented in this chapter. The work presented in this chapter is my own.

5.2 Introduction

In the previous chapter, we fitted Markov models of I_{Kr} to experimental data. This was done under the assumption of an ideal patch-clamp setup, whereby certain known artefact effects are omitted, and we assumed,

$$V_m(t) = V_{cmd}(t) + V_{off}, \quad (5.1)$$

where V_m is the transmembrane potential, V_{cmd} is the command voltage specified by the experimental design and V_{off} is some constant voltage offset. We also assumed that the current remaining after leak-correction and drug-subtraction is exactly I_{Kr} that is,

$$I_{post} = \left(I_{obs}^{(before)} - I_L^{(before)} \right) - \left(I_{obs}^{(after)} - I_L^{(after)} \right), \quad (5.2)$$

where $I_{\text{obs}}^{(\text{before})}$ and $I_{\text{obs}}^{(\text{after})}$ are the raw current traces observed before and after (respectively) the addition of dofetilide, and $I_{\text{L}}^{(\text{before})}$ and $I_{\text{L}}^{(\text{after})}$ are the corresponding leak currents.

Whilst these models resulted in accurate fits to individual sweeps of data, they were noticeably less accurate when used to predict I_{Kr} under unseen voltage protocols. Furthermore, the well-to-well variability found in our collections of parameter estimates suggested the presence of latent variables. This may indicate that the ideal-patch assumptions themselves, and not a particular choice of model structure, are a dominant source of model discrepancy. This is also supported by the apparent well-dependence of our parameter estimates, because, under the assumption that our models' kinetic parameters are constant across cells, there must be some hidden well-dependent effects causing the apparent variability in our recordings. Experimental artefacts, which vary randomly between wells, may provide an explanation (Lei et al., 2020a). In this chapter, we investigate whether the inclusion of such experimental artefacts in our models can improve the predictive accuracy of our models, and decrease the well-to-well variability in parameter estimates.

As discussed at the end of the previous chapter, and in Lei et al. (2020a), there are known “artefact effects” which affect the output of patch-clamp experiments, and are expected to vary between wells. These effects, if large enough, invalidate the assumption of an ideal patch-clamp setup (the assumptions described in Section 4.1). One example of such an effect is *series resistance* (Lei et al., 2020a), as shown in Figure 5.1. Typically, in manual patch-clamp setups, *series-resistance compensation* is performed, mitigating the impact of this effect, though our automated experiments were performed on a machine where this feature is not present. However, even platforms capable of series-resistance compensation do not completely remove the effect, in which case series resistance, and the circuitry responsible for its compensation, may be included in mathematical models (Lei et al., 2020a).

Another important effect, which is omitted by the ideal-patch model, is the capacitance across the cell membrane. Because charges are stored (in the form of ions such as K^+ , Na^+ and Ca^{2+}) on either side of the cell membrane, any current flowing through the circuit has an effect on the transmembrane potential, V_{m} . This is shown in the circuit diagram in Figure 5.1 where we can see that the membrane charges and discharges in response to changes in V_{cmd} , the command voltage. When there is an abrupt change in V_{cmd} , this results in a smooth change in V_{m} , resulting in a discrepancy in V_{m} and V_{cmd} that is most pronounced after discontinuities in our voltage-clamp protocols (that is, at the beginning of each voltage step). Some patch-clamp equipment is able to perform “supercharging” to mitigate the smoothing effect of the cell membrane’s capacitance, though, again, our experiments were performed on an automated machine with no such feature.

In the following section, Section 5.3 we introduce some methods for integrating the artefact model into the postprocessing and fitting methodology of the previous chapter. Then, we fit each of the I_{Kr} Markov models discussed in the previous chapter, with and without the inclusion of artefact effects. This allows a comparison of the predictive accuracy and distribution of parameter estimates obtained between the artefact-modelling approach, and the ideal-patch approach in terms of their predictive accuracy (Section 5.5.2) and the variability of parameter estimates (Section 5.5.3). Finally, in Section 5.5.4, we revisit the variability in decay time-constants discussed in Section 4.3, and investigate whether this variability can be explained solely by artefact effects under the assumption of identical channel-gating kinetics.

5.3 Mathematical modelling of experimental artefacts

ODE models of patch-clamp experiments have been used to describe the dynamics of patch-clamp experiments, even in the absence of biological currents (Weerakoon et al., 2009; Lei et al., 2020a). Lei et al. (2020a) used ODE-based models to describe the changes in transmembrane potential over time, modelling the interaction between cellular currents and the patch-clamp apparatus. This includes the effect of capacitive currents and voltage attenuation through the series resistance.

Lei et al. (2020a) present two artefact models. Some of the terms included in Lei et al.’s complete model related to the parasitic capacitance, C_p . Typically, the effect of C_p is negligible when considering millisecond-and-above timescales, and values lie in the range 3–5 pF (Lei et al., 2020a). As shown in Figure 5.2, our estimates of C_m lie in the range 12.5–30 pF. Perhaps more importantly, the timescale over which the parasitic capacitance acts is very short—approximately $0.8 \mu s$ (Lei et al., 2020a). Hence, this effect has little impact on the vast majority of our observations when using a 1–10 kHz sampling rate, especially when the presence of capacitive effects is minimised through the removal of 5 ms of data after each discontinuity in V_{cmd} as described in the previous chapter and Appendix C.

Other short-timescale effects are due to various delays within the patch-clamp amplifier itself, which are the result of the solid-state electrical circuitry used to control the injected current. We use Lei et al. (2020a)’s simplified model (henceforth referred to as “the artefact model”) which omits these effects and models the membrane voltage simply as,

$$\frac{dV_m}{dt} = \frac{V_{cmd} + V_{off} - V_m}{R_{series}C_m} - \frac{I_{out}}{C_m}, \quad (5.3)$$

where: $I_{out} = I_{Kr} + I_L$; C_m is the membrane capacitance; R_{series} is the series resistance, and V_{off} is some systematic voltage (similar to the voltage offset discussed in the previous chapter).




Figure 5.1: An equivalent circuit formulation of the artefact model. Here, the subject of our modelling effort, I_{Kr} is just one component of the circuit. We aim to improve our models of I_{Kr} by including other components of the circuit (such as R_{series} and C_m). The grey dashed box highlights the components modelled by our choice of I_{Kr} model.

In this way, the inputs into and outputs from our I_{Kr} model are not exactly V_{cmd} and $I_{Kr} + I_L$ (as we assumed in the previous chapter). Moreover, as I_{out} appears in the right-hand side of Equation (5.3), there is feedback between the ion-channel current (I_{Kr}) and the membrane potential that the amplifier achieves (V_m). The converse is also true, since I_{Kr} is dependent on the transmembrane voltage, V_m .

Estimation of artefact-model parameters The Nanion SyncroPatch 384 provides its own protocols and procedures for the estimation of R_{series} and C_m . These short, simple protocols are not shown to the user, but are performed before each sweep to provide estimates of C_m and R_{series} . Figures 5.2 and 5.3 show how, in each well, these estimates stay largely consistent throughout the duration of the experiment. There is, however, noticeable well-to-well variability in these estimates. We suggest this variability may be related to the well-dependence of parameter estimates discussed in the previous chapter. Because these estimates are readily available, we fix our artefact-model parameters to the respective machine estimates obtained from a particular well immediately prior to the particular sweep under consideration.

Whilst the modelling and prediction of experimental artefacts is not our primary aim, the inclusion of artefact effects may explain the apparent latent effects implied by the well-dependence of our ideal-clamp parameter estimates, as described in the previous chapter. In this way, the inclusion of the artefact model may be necessary to build better predictive models of I_{Kr} and, ultimately, to improve our understanding of the gating dynamics of I_{Kr} . To fit models including these artefact effects, we use methods largely similar to those introduced in the previous chapter. However, the estimation of our leak-model parameters (g_L and E_L), and V_{off} is somewhat complicated by the presence of these artefacts, as explained in the following section.




Figure 5.2: Onboard estimates of the cell membrane capacitance, C_m , obtained over the course of the experiment. These estimates are obtained by the Nanion SyncroPatch before each sweep of data is recorded. Only those wells passing our QC criteria (as introduced in the previous chapter) are shown. We expect that much of the well-to-well variability shown here arises from the varying sizes of each cell. Note that Well J11 was discarded from this dataset due to a failure to infer the reversal potential from its pre-drug traces.

5.3.1 Impact of experimental artefacts on postprocessing

The postprocessing methods introduced in the previous section were built under the assumption of an ideal patch-clamp set-up. Therefore, there are some adjustments necessary to allow the inclusion of experimental artefacts. The most important difference is the use of the pre-drug traces for fitting our models, as opposed to the drug-subtracted traces used in the previous chapter. This is possible due to the fact that our post-drug traces are almost entirely dominated by leak current, as shown by the very small R_{leftover} values obtained during postprocessing (as discussed in Section 4.3).

The artefact model (characterised by Equation (5.3)) describes the feedback between the transmembrane potential and $I_{\text{out}} = I_L + I_{K_r}$, and so we choose to fit the outputted current, I_{out} , to the pre-drug traces. In the previous chapter, we computed E_{post} by applying polynomial interpolation to our postprocessed trace, I_{post} . In this chapter, we instead use the pre-drug traces, from which we estimate E_{before} using the same method. Here, we are unable to compute E_{before} for each pre-drug, leak-corrected sweep taken




Figure 5.3: Onboard estimates of the series resistance, R_{series} , obtained other the course of the experiment. These estimates are obtained by the Nanion SyncroPatch before each sweep of data is recorded. Only those wells passing our QC criteria (as introduced in the previous chapter) are shown. Note that well J11 was also discarded because we were unable to infer E_{before} .

from Well J11 and so, it is removed from consideration for this chapter. The values obtained for E_{before} are shown in Figure 5.4.

Impact of artefact effects on the reversal potential We should expect experimental artefacts to have an effect on E_{post} , the reversal potential observed according to our polynomial interpolation method (Section 4.3). We may expect that this sharp decrease causes V_m to lag behind V_{cmd} . However, our calculations in Appendix F indicate that this effect is small (no more than 1 mV)—and so, would not explain the much larger discrepancy we see between the observed reversal potential, E_{before} and the Nernst potential, E_{Nernst} . Moreover, the effect of this delay would be to reduce E_{before} , but, as shown in Figure 5.4, we have $E_{\text{post}} > E_{\text{Nernst}} \approx -89.8$ for all wells.

In fact, the effect of any delay would lead to a decrease in E_{post} —the opposite of what we see in the dataset (Figure 5.4). Moreover, we show in Appendix F that, under the assumption of the artefact model with $V_{\text{off}} = 0$, such a large difference between E_{Nernst} and E_{post} is not possible unless R_{series} or the recorded currents are unrealistically large. For this reason, we include a voltage offset, and approximate it prior to model fitting.




Figure 5.4: Values of E_{before} computed using each well's pre-drug leak-corrected traces, using the polynomial interpolation method outlined in the previous chapter. It was not possible to compute E_{before} for every protocol from Well J11 due to the pre-drug leak subtracted current not reversing during the reversal ramp.

Leak-model fitting with artefacts We also expect the artefact model to have some effect on the inference of our leak-model parameters because the assumption that $V_{\text{cmd}} \approx V_{\text{m}}$ may not hold during the leak ramp. In the previous chapter, ideal-patch assumptions (Equations (4.1) and (4.2)) allowed us to fit the leak current parameters, g_{L} and E_{L} using simple linear regression. However, the artefact model breaks these assumptions, necessitating a different approach.

Using the artefact model, we keep the assumption that $I_{\text{Kr}} = 0$ during the leak ramp, but fit g_{L} and E_{L} by minimising the difference between I_{out} and I_{before} (the pre-drug trace) using numerical optimisation. If we assume $g = 0$ during this section of the protocol, Equation (5.3) gives rise to a simple, linear ODE, from which we approximate $V_{\text{m}}(t)$ and $I_{\text{out}} = I_{\text{L}}$ (as shown in Appendix F).

Then, we fit g_{L} and E_{L} by finding \hat{g}_{L} and \hat{E}_{L} which minimise the sum-of-squares criterion,

$$\sum_{i=1}^{n_{\text{leak}}} (I_{\text{L}}(t_i; g_{\text{L}}, E_{\text{L}}) - z_i)^2, \quad (5.4)$$

where n_{leak} is the number of observations considered, and $I_{\text{L}}(t_i; g_{\text{L}}, E_{\text{L}})$ is the leak current at time t_i according to the artefact model, n is the number of observations taken during the leak ramp and each z_i is the datum observed at time t_i .

Although this is a nonlinear regression problem unlike the linear regression method presented in Section 4.3, the previous, simple linear-regression method is used to provide an initial guess for g_L and E_L . Additionally, we constrain our parameter space such that, $0 \leq g_L \leq 10^3$ and $-500 \leq E_L \leq 500$. Of the remaining artefact-model parameters, we set C_m and R_{series} to their respective onboard estimates, and V_{off} to 0. If we instead assume a non-zero value of V_{off} , our estimate of g_L is unaffected, but we have,

$$\hat{E}_L = E_L^\dagger + V_{\text{off}}, \quad (5.5)$$

where E_L^\dagger is the estimate of E_L obtained under the assumption that $V_{\text{off}} = 0$. This fact allows us to fit our leak-model parameters under the assumption that $V_{\text{off}} = 0$, and easily adjust these parameters for other values of V_{off} without performing additional optimisation.

Direct estimation of voltage offset It is desirable to estimate V_{off} before we fit our model and compute our parameter estimates. This effectively allows us to simplify our model by reducing the number of parameters—similarly to how we fit our leak-model parameters, g_L and E_L . To estimate the voltage offset (before we begin fitting our I_{K_r} model), we use the artefact model under the assumption that $E_{K_r} = E_{\text{Nernst}}$ together with a reference model of I_{K_r} (the Beattie et al. (2018) Model with Cell #5 parameters). Next, we find the value of V_{off} , which, in conjunction with the other artefact effects, best explains the discrepancy between E_{Nernst} and E_{before} . To do this, we fix the remaining artefact-model parameters, R_{series} and C_m , to their respective machine estimates. We then estimate the voltage offset by computing,

$$\hat{V}_{\text{off}} = \operatorname{argmin}_{\tilde{V}_{\text{off}}} \left\{ (E_{\text{Nernst}} - E_{\text{before}}(\tilde{V}_{\text{off}}, \hat{g}(V_{\text{off}}), \hat{g}_L, \hat{E}_L))^2 \right\}, \quad (5.6)$$

where $E_{\text{before}}(\tilde{V}_{\text{off}}, \hat{g}(V_{\text{off}}), \hat{g}_L, \hat{E}_L)$ denotes the reversal potential that we expect to observe according to the model (using the polynomial interpolation procedure introduced in Section 4.3.2) where: $\hat{g}(V_{\text{off}})$ is the estimate of the channel conductance, g , which is contingent on the proposed V_{off} ; the variables, \hat{g}_L and $\hat{E}_L(\tilde{V}_{\text{off}})$ denote our estimates of the leak-current conductance, the leak-current reversal, respectively. Here, we must find new estimates for $\tilde{g}(\tilde{V}_{\text{off}})$ and $\tilde{E}_L(\tilde{V}_{\text{off}})$ for each proposed value of \tilde{V}_{off} .

As explained above, our estimates of \hat{g}_L are independent of V_{off} and after estimating $E_L(0)$ (assuming $V_{\text{off}} = 0$), it is simple to compute $\tilde{E}_L(\tilde{V}_{\text{off}}) = E_L(0) + \tilde{V}_{\text{off}}$. However, we must use numerical optimisation to find our conductance estimate,

$$\hat{g}(V_{\text{off}}, \hat{g}_L(V_{\text{off}}), \hat{E}_L(V_{\text{off}})) = \operatorname{argmin}_{\tilde{g}} \{ \|I_{\text{obs}} - I_{\text{out}}(V_{\text{off}}, \tilde{g}, \hat{g}_L(V_{\text{off}}), \hat{E}_L(V_{\text{off}}))\|_2 \}, \quad (5.7)$$

where I_{obs} is the data recording during the reversal ramp of the pre-drug sweep and $I_{\text{out}}(V_{\text{off}}, \tilde{g})$ denotes the predicted current according to our reference model with voltage offset, V_{off} and maximal conductance \tilde{g} , and where R_{series} and C_{m} are set to their respective machine estimates. This estimate of the conductance is found using only the central 70% of observations taken during the reversal ramp—that is, those observations taken between 30ms and 70ms after the start of the reversal ramp. This is done such that the model's I_{out} matches our observations, I_{obs} , as closely as possible during the reversal ramp (as quantified by the RMSE). To solve Equation (5.6), we must evaluate Equation (5.7) for every proposed value of V_{off} . As shown in Figure 5.5, this comprises a nested optimisation problem. Nested optimisation problems can be computationally expensive, but the relative simplicity of this particular problem (which has few parameters) renders this approach feasible.

In this way, the relative simplicity of the reversal ramp section of the protocol allows us to accurately recapitulate I_{out} (by fitting the maximal conductance, g) and hence, accurately approximate V_{m} during this time.

Recapitulating artefact-model parameters from synthetic data To validate the leak-fitting and V_{off} estimation methods described above, we use a reference parameter set and randomly sample artefact-model parameters to generate synthetic data, then we attempt to re-infer these artefact-model parameters.

We sample random parameter values according to,

$$\log_{10}(g_{\text{Kr}}) \sim \mathcal{N}(\log_{10}(300), 0.1), \quad (5.8)$$

$$V_{\text{off}} \sim \mathcal{N}(-7.5, 5), \quad (5.9)$$

$$\log_{10}(R_{\text{series}}) \sim \mathcal{N}(0, \log_{10}(2)), \quad (5.10)$$

$$\log_{10}(C_{\text{m}}) \sim \mathcal{N}(\log_{10}(20), \log_{10}(5)), \quad (5.11)$$

$$E_{\text{L}} \sim \mathcal{U}(-50 \text{ mV}, 25 \text{ mV}), \quad (5.12)$$

$$\text{and } \log_{10}(g_{\text{L}}) \sim \mathcal{N}(\log_{10}(5), 0.1), \quad (5.13)$$

where voltages (V_{off} and E_{L}) are expressed in millivolts (mV), conductances (g_{Kr} and g_{L}) are expressed in nanosiemens (nS), C_{m} is expressed in nanofarads (nF), and the series resistance, R_{series} , is expressed in megaohms (M Ω), and where $\mathcal{U}(l, u)$ denotes a uniform distribution between l and u , and $\mathcal{N}(\mu, \sigma)$ denotes a normal distribution with mean, μ and standard deviation, σ . By taking 100 samples from these distributions, we obtain samples consisting of a wide range of plausible combinations of artefact-model parameters. In contrast, our I_{Kr} model parameters (except the maximal conductance) are fixed—set to be the Cell #5 parameters from Beattie et al. (2018) (listed in Table B.2). The synthetically generated datasets are generated from our sampled artefact-model

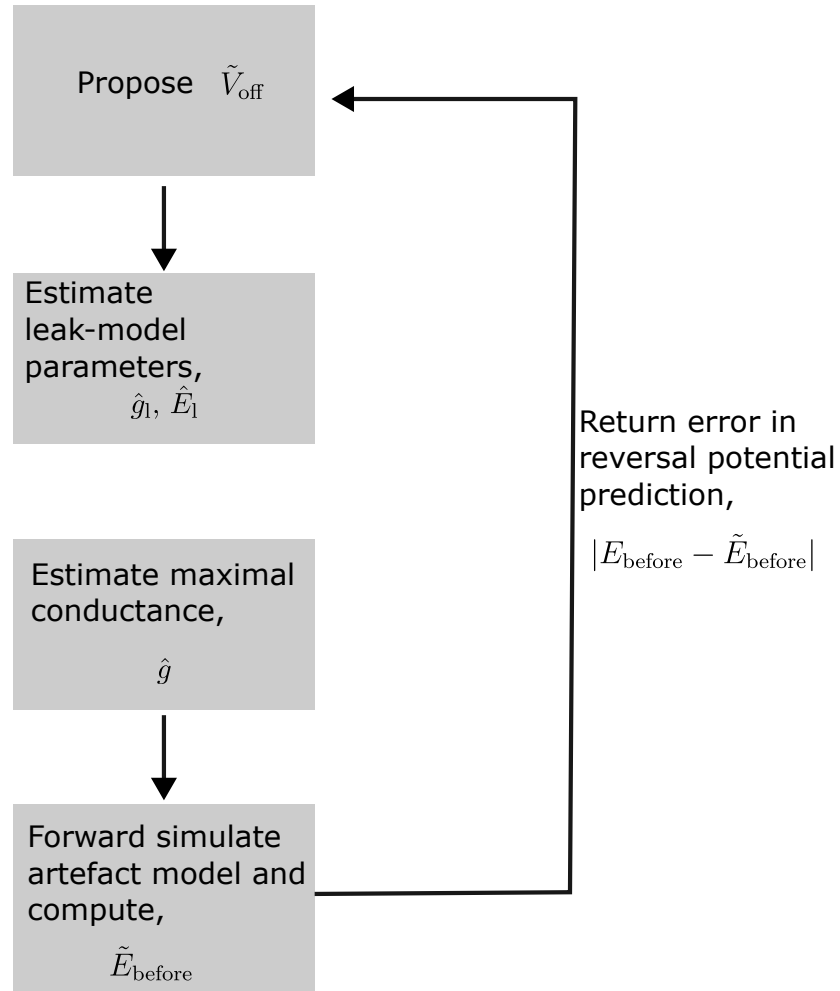


Figure 5.5: The nested-optimisation procedure we use to estimate V_{off} , g_L and E_L , prior to fitting our kinetic parameters, under Case V. The estimation of the leak parameters and maximal conductance are performed by minimising the RMSE between the data and the model during a relevant portion of the protocol.

parameters by first forward-simulating the Beattie model (embedded within the artefact model), before adding IID Gaussian noise which is sampled as $\varepsilon_i \sim \mathcal{N}(0, 10 \text{ pA})$.

We denote our estimates of g_L , E_L and V_{off} by \hat{g}_L , \hat{E}_L and \hat{V}_{off} , respectively. In Figure 5.5 we show the accuracy of these estimates across all of our sampled parameter sets. These results show that we can reliably recover these parameters (for those datasets passing our QC, at least). For our experimental dataset, our estimates of V_{off} obtained

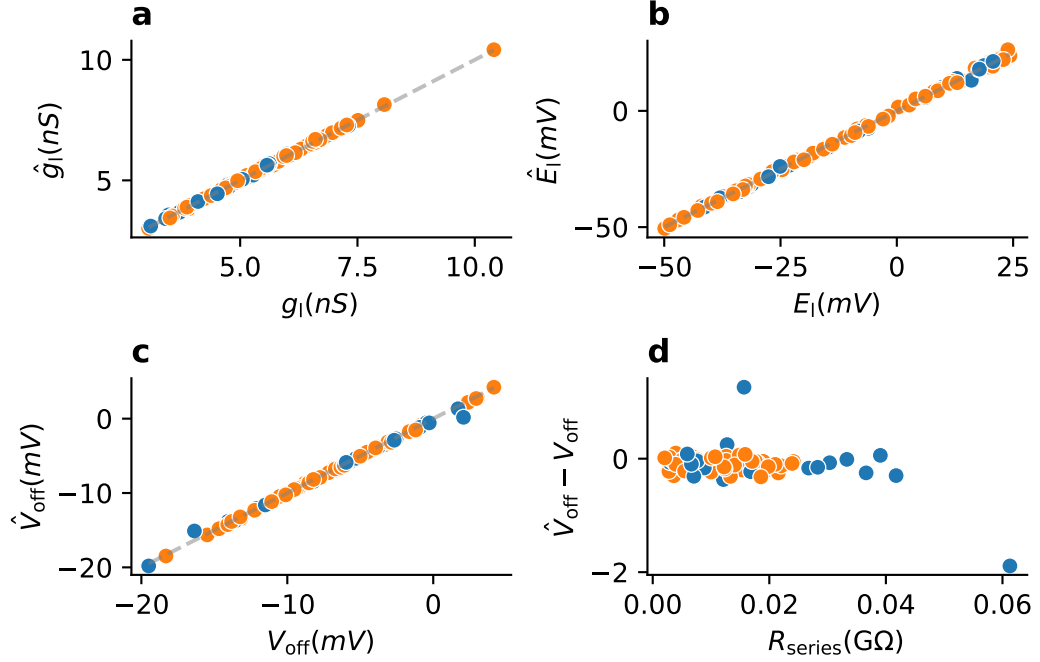


Figure 5.6: Using our nested optimisation approach, we were able to recover artefact-model parameters from synthetically generated data. Data were generated using the Beattie et al. (2018) Model combined with the artefact model. Parameter estimates arising from traces that did not pass QC are shown in blue, whereas those from traces passing QC are highlighted in orange. Panels **a** and **b** show our estimates of the leak parameters, g_L and E_L , respectively. Panels **c** and **d** shows estimates of the voltage-offset parameter. In Panel **d**, we can see that the traces which lead to the largest error in our voltage-offset estimates ($\hat{V}_{\text{off}} - V_{\text{off}}$) came from traces which failed QC (the largest error being from a sample where R_{series} is uncharacteristically large).

by this approach, alongside our inferred values of E_{before} are shown in Figure 5.7. Here we see that, within each well, our estimates of V_{off} vary little (relatively speaking) over the course of the experiment. However, there is noticeable inter-well variability.

Additionally, Figure 5.8 compares a sweep of our data with the Beattie et al. (2018) Model (with its original Cell #5 parameter set) under our ideal-patch assumptions (with $V_{\text{off}} = 0$), and including artefact effects. Here, the conductance and artefact-model parameters are fitted using the methods outlined above. This figure shows that we expect the artefact model has a significant impact on the dynamics of I_{K_r} , and the output of our models. As a result, we might expect that our parameter estimates and the resultant model predictions to be similarly affected by the inclusion of artefact effects.

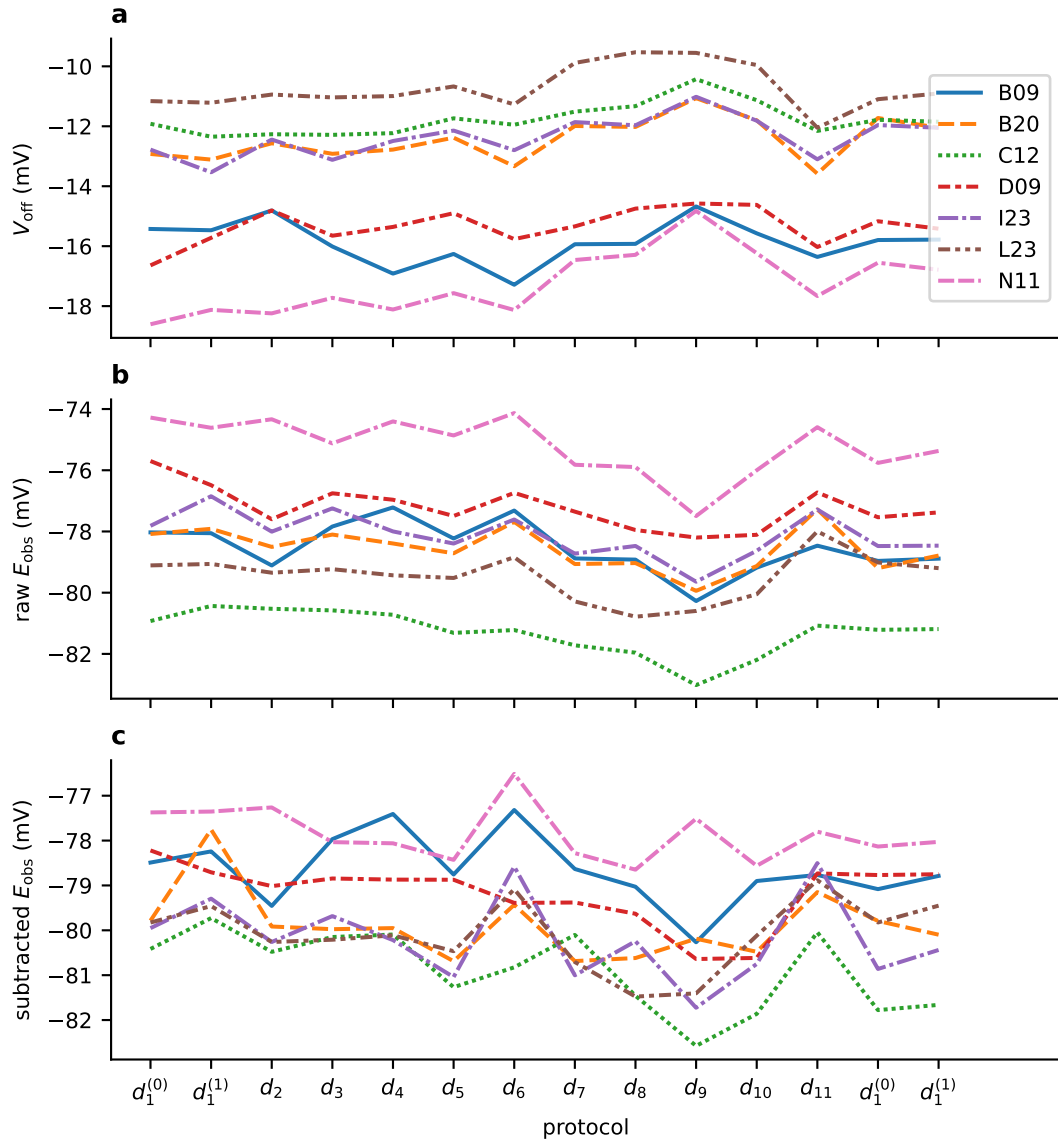


Figure 5.7: Our estimates of V_{off} stay consistent over the duration of the experiment, as shown in Panel **a**. For comparison, the inferred values of E_{before} for our pre-drug, leak-corrected traces are shown in Panel **b**, and those obtained from the corresponding leak-corrected, drug-subtracted traces (that is, E_{post}) are shown in Panel **c**.

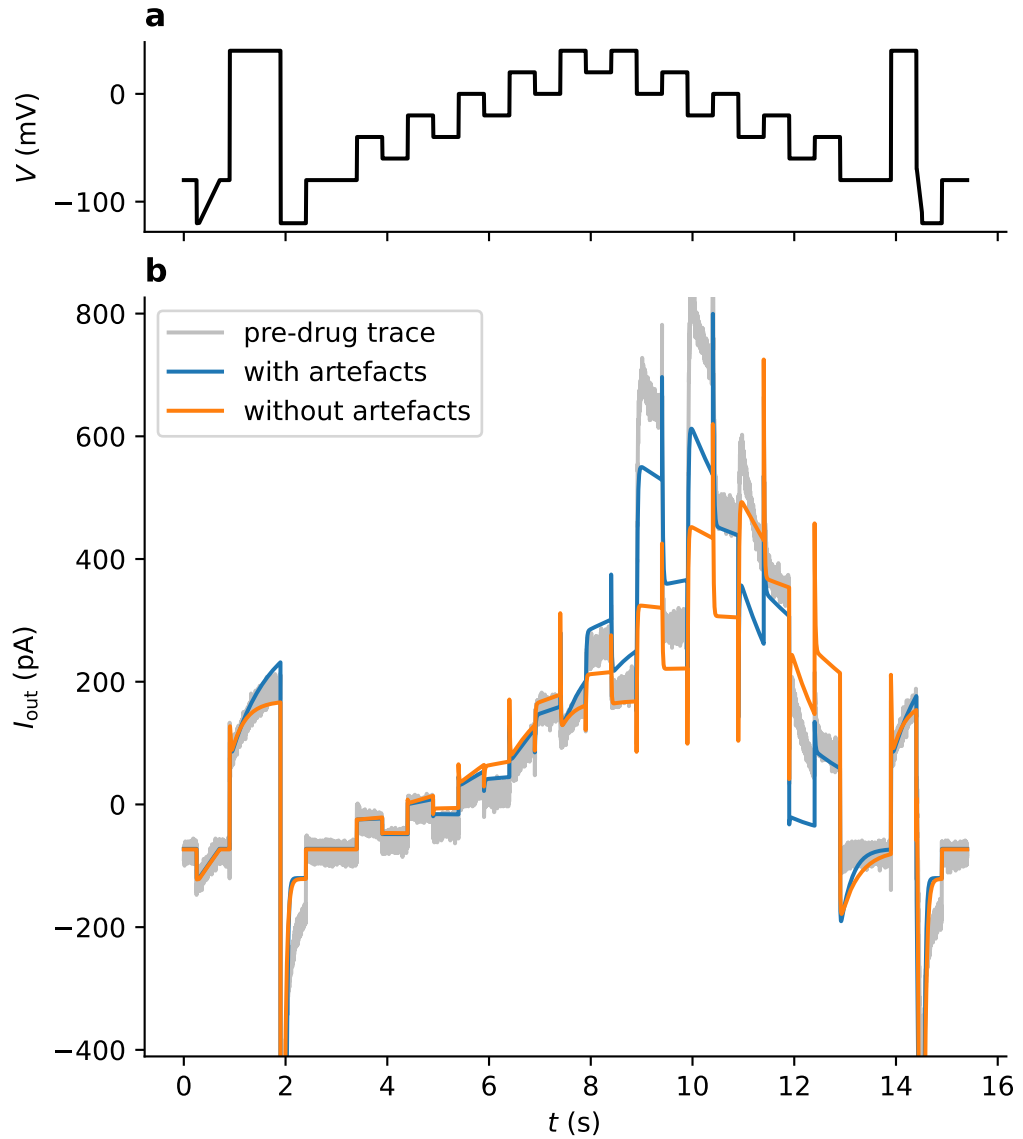


Figure 5.8: Comparison of real, experimental data, an ideal-patch model and a model including artefact effects. Both models include the Beattie et al. (2018) Model with its original parameter set (that is, the parameter set obtained from fitting to Cell #5 in its original publication). The data were collected from Well B20 under the application of the *staircase* protocol (first repeat). In both cases, the maximal conductance was fitted as to minimise the RMSE between the model and the data. The artefact parameters were inferred using the methods described in this section. Here, we can see that the inclusion of experimental artefacts drastically affects the dynamics of I_{Kr} and V_m . Model output is shown both with artefact effects (orange), and without artefacts (but, with a non-zero voltage offset) under our ideal-patch assumptions (blue).

5.4 Fitting predictive models of I_{KR} using the artefact model

As with the application of ideal-patch models (in Chapter 4), we again remove $R = 5$ ms of data after each discontinuous jump in voltage. Whilst it may be suitable to reduce R or set $R = 0$ after the inclusion of artefact effects, there are some short-timescale effects that are omitted from the artefact model. However, as we have seen previously, such short-timescale information is not absolutely necessary to parameterise our I_{KR} models.

To allow an evaluation of the artefact model, we implement each of our literature models under various assumptions. Firstly, in Case IV, we omit the artefact effects and instead use the ideal-patch assumptions, and assume that $V_{\text{off}} = E_{\text{Nernst}} - E_{\text{post}}$. This is almost identical to Case III introduced in Chapter 4, but, to allow fair comparison, we fit our models directly to pre-drug traces without leak correction (requiring the inclusion of the leak current in our models). Then, in Case V, we relax this assumption, fitting our models to the same data using the artefact model, after estimating V_{off} according to the approach outlined in the previous section. Cases IV and V are outlined in Table 5.1.

To integrate our ODE system, fit our models and make predictions, we use LSODA (Gupta et al., 1985) with absolute and relative tolerances set to 10^{-8} . The initial conditions of the artefact model are set by computing the model's stable equilibrium at the holding potential (that is, when we fix $V_{\text{cmd}} = -80$ mV). This is done by setting the Markov model component to its steady state with $V_m = -80\text{mV} + V_{\text{off}}$, and allowing the artefact model to reach equilibrium by integrating with LSODA for 5 seconds. As in the previous chapter, we use CMA-ES (Hansen, 2016) as implemented by the Python PINTS package (Clerx et al., 2019b). The number of optimisation repeats, solver tolerances, parameter transforms and initial guess sampling are performed in exactly the same way as described in Section 4.4.1.

5.5 Results

5.5.1 Model fitting

Parameter sets were obtained by fitting each of our model structures to each sweep of our data, with and without artefact effects (that is, under the assumption of both Case IV and V). Figure 5.9 shows one particular example where the Beattie et al. (2018) Model is fitted to data taken from the *staircase* protocol (d_1) under Case IV. Figure 5.10 shows the same model structure fitted to the same data under the assumption of Case V (that is, with the inclusion of experimental artefacts).

Case	parameterisation	physical interpretation
IV	$g_{\text{leak}} = \hat{g}_{\text{leak}}$, $E_{\text{leak}} = \hat{E}_{\text{leak}}$ are estimated using simple linear regression; $V_{\text{off}} =$ $E_{\text{Nernst}} - E_{\text{post}}$; and $E_{\text{Kr}} = E_{\text{Nernst}}$.	There are no artefact effects; we use the ideal-patch formulation from the previous section and fit our models directly to leak-corrected pre-drug traces (that is, without drug subtrac- tion).
V	$g_{\text{leak}} = \hat{g}_{\text{leak}}$ and $E_{\text{leak}} = \hat{E}_{\text{leak}}$ are estimated using the arte- fact model; $R_{\text{series}} = \hat{R}_{\text{series}}$, $C_{\text{m}} = \hat{C}_{\text{m}}$ from ma- chine estimates; $E_{\text{Kr}} = E_{\text{Nernst}}$ and V_{off} is estimated using the nested optimisation pro- cedure outlined in the text.	There are artefact effects, including a systematic voltage offset ($V_{\text{off}} \neq 0$). We assume $E_{\text{Kr}} = E_{\text{Nernst}}$, and in- fer E_{before} from the reversal ramp (accounting for series resistance ef- fects) and use this to calibrate V_{off} which, together with g_{L} , E_{L} , C_{m} , and R_{series} is treated as a known constant whilst fitting the maximal- conductance and transition-rate pa- rameters of the I_{Kr} model.

Table 5.1: The two ways in which we fit our Markov models of I_{Kr} in this chapter. Case IV differs from Case III (as used in the previous chapter), because the model is fit directly to the pre-drug trace without the subtraction of a post-drug trace. This is done to permit a fair comparison with the artefact model (Case V).

In the previous chapter, we saw that an apparent overcorrection of leak current resulted in differences in the obtained value of model parameter p_1 between the first and second sweeps (as shown in Figures 4.18 and 4.33). Here, in Figures 5.9 and 5.10, where we fit the Beattie et al. (2018) Model (under both Case IV and V) to the pre-drug traces, we see no such difference between sweeps. We also see that we obtain similar estimates (compared to those shown in the previous chapter) of p_1 and p_2 under both models for both Cases IV and V.

In Figures 5.9–5.10, we can see that the Beattie et al. (2018) Model fails to recapitulate the large tail current during the initial -120mV step following the leak ramp—whether or not artefact effects are included (Case V) or not (Case IV). This may be a particular shortcoming of the Beattie et al. (2018) Model, which contains only one pair of closed states—especially when fitted to room-temperature recordings (Vandenberg et al., 2012; Kemp et al., 2021). It appears that the addition of experimental artefacts is not sufficient to remedy this specific aspect of model discrepancy. More “deeper” closed states from which more transitions are required for a channel to open (such as those included in the Kemp et al. (2021) and Wang et al. (1997) models) may permit a closer fit to the data in this particular part of protocol. The properties of this tail current, and to what extent it can be recapitulated by the artefact model (together with the Wang et al. (1997) model) is discussed further in Section 5.5.4.

5.5.2 Predictive accuracy

Following the methodology of Chapter 4, we compute the normalised root-mean-square error (NRMSE) between our model predictions and the corresponding data. This was done for using every possible pair of protocols for fitting (except d_6) and validation. Figures 5.11 and 5.12 show these NRMSE values for the best and worst performing wells according to the Beattie et al. (2018) Model under Cases IV and V, respectively.

The resulting cross-validation heatmaps, averaged across all wells, are shown in Figure 5.15. Here, we can see that our Case IV models (with no artefacts) performs better than our Case V models (when averaged across wells). Though, our most accurate predictive model (in terms of $\mathcal{E}_{\text{predict}}$) for both Well B20 and Well L23 (independently) is the Kemp et al. (2021) model under our Case V assumptions—this is despite the equivalent ideal-patch model producing better fits on average (as quantified by \mathcal{E}_{fit}). This difference in predictive accuracy may be due to the larger series-resistances found in these wells (as shown in Figure 5.3), which may lead to greater discrepancy between our ideal-patch assumptions and the true DGP.

Even so, the most accurate fits for this well are found under Case IV. Overall, it appears that our models which include artefacts (or more precisely, our particular treatment of the artefact model under Case V) produce less quantitatively accurate fits

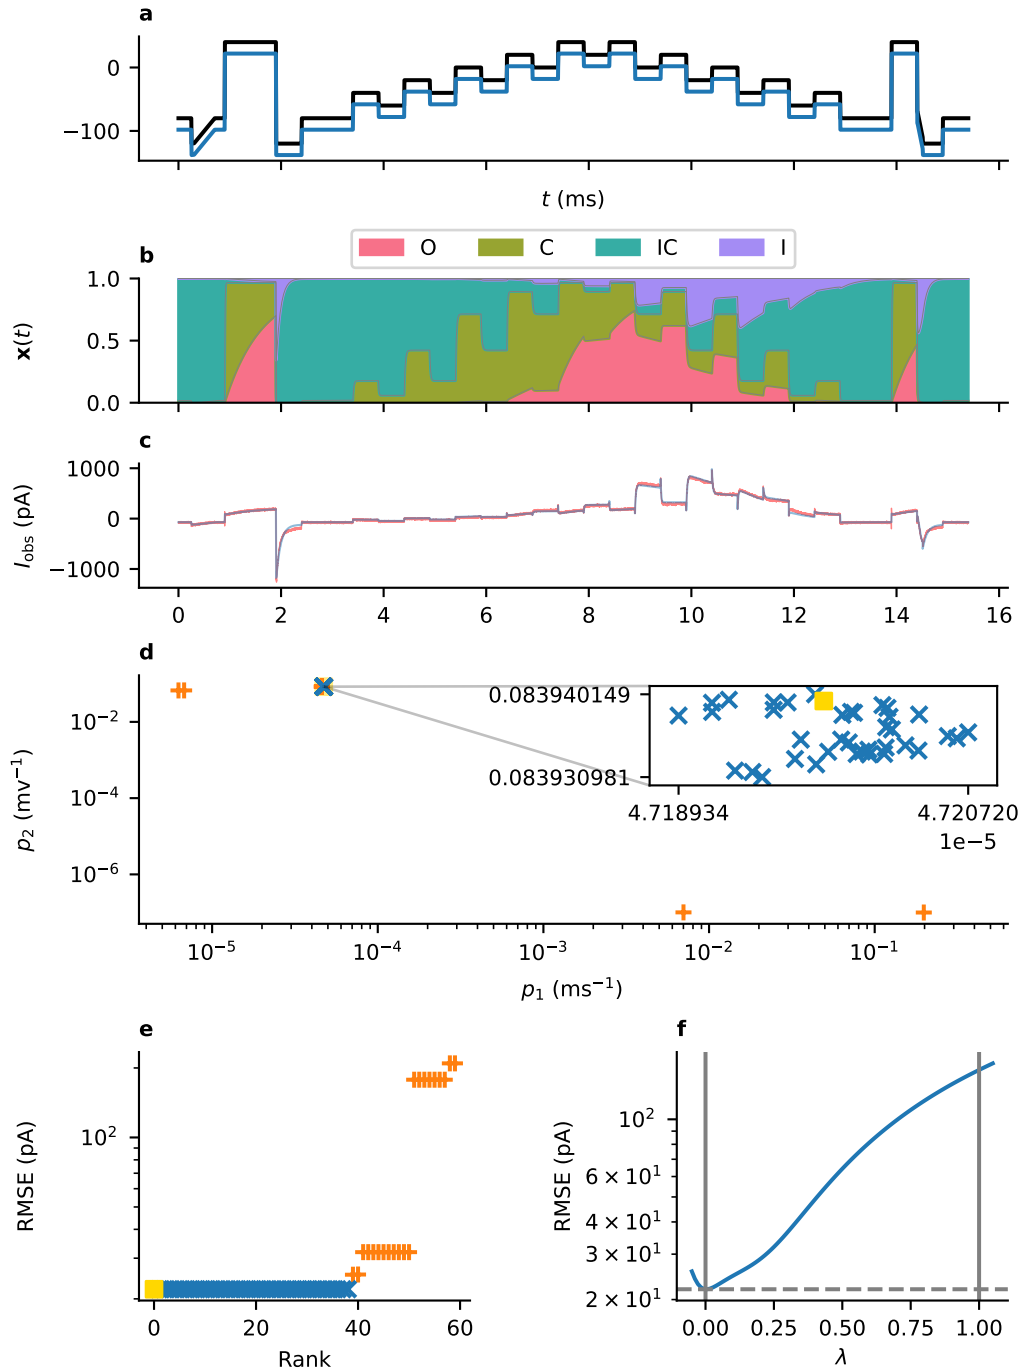


Figure 5.9: Results from fitting the Beattie et al. (2018) Model (Case IV) data from Well B20 under the *staircase* protocol (second repeat). Panel **a**: V_{cmd} (black) and V_{m} under the fitted model (blue). Panel **b**: the occupancies of each conformational state under the fitted model. Panel **c**: the data and the model fit. Panel **d**: estimates of p_1 and p_2 obtained from thirty repeats of the optimisation routine—the blue markers shown in the inset correspond to those results where the RMSE is at most 101% of the minimum value. Panel **e**: the amount of RMSE error in each fit. Panel **f**: a cross-section through the likelihood surface, starting at our best estimate of the parameters ($\lambda = 0$, blue square) and finishing (when $\lambda = 1$) at a parameter set with identical maximal conductance (g), but where the transition-rate parameters are taken from the model's original publication (Cell # 5) (Beattie et al., 2018).

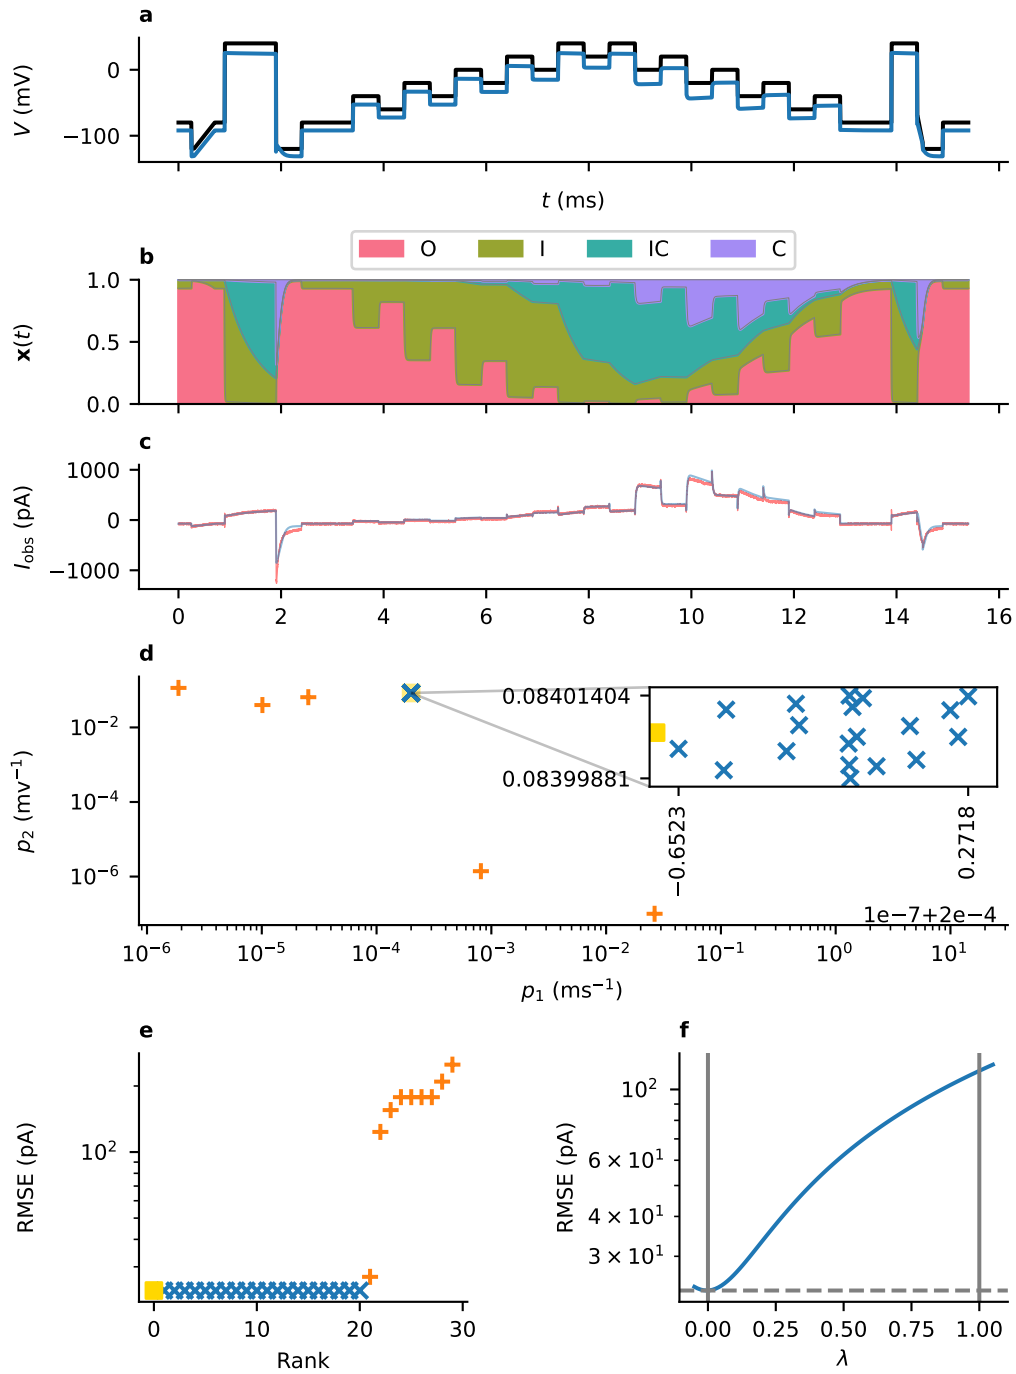


Figure 5.10: Results from fitting Beattie et al. (2018) Model (Case V) data from Well B20 under the application of the *staircase* protocol (second repeat). Panel **a**: V_{cmd} (black) and V_m under the fitted model (blue). Panel **b**: the occupancy of each conformational state under the fitted model. Panel **c**: the real-data traces together with the model fit. Panel **d**: the results of repeated optimisations—blue markers correspond to those results where the RMSE is at most 101% of the minimum value. Panel **e**: the amount of RMSE error in each fit. Panel **f**: a cross-section through the likelihood surface, starting at our best estimate of the parameters ($\lambda = 0$, blue square) and finishing (when $\lambda = 1$) at the model's original transition-rate parameters publication (Cell #5 Beattie et al., 2018).

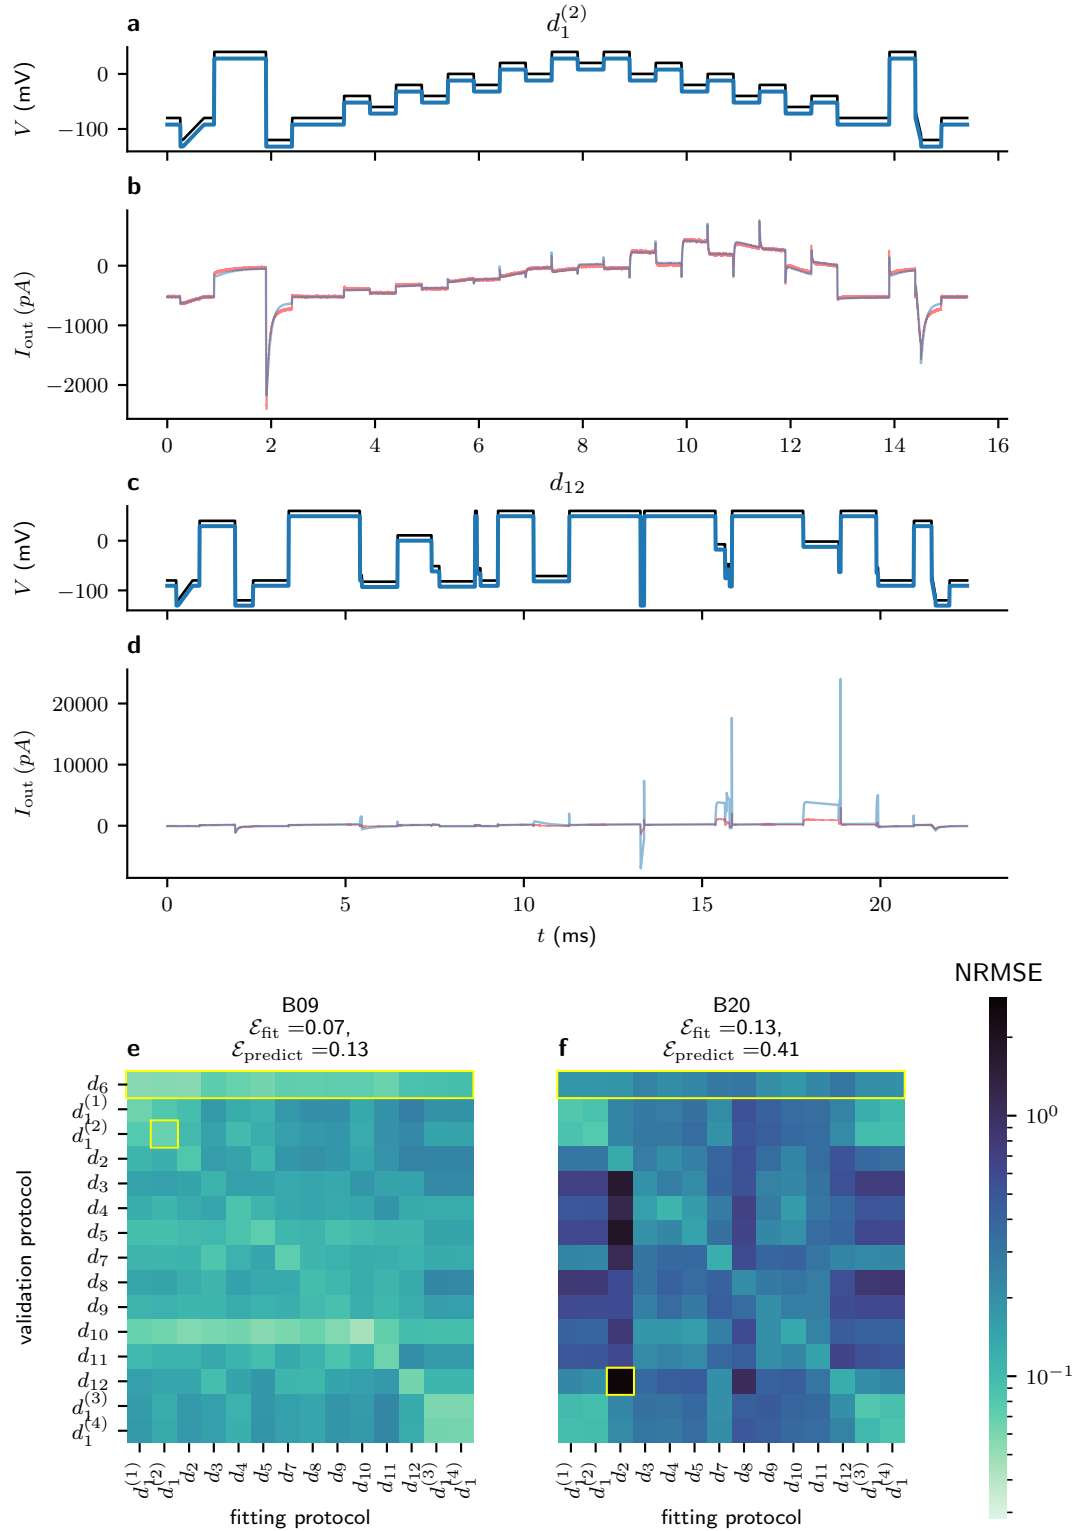


Figure 5.11: Cross-validation heatmaps for the Beattie model under Case IV (without artefacts), showing the best and worst performing wells (as quantified by $\mathcal{E}_{\text{pred}}$). The best prediction from the best performing well is shown in panels **a** and **b**, and the worst prediction from the worst performing well is shown in panels **c** and **d**. Note that in both panels **b** and **d**, the model fails to recapitulate the large tail current occurring at $t = 2$ s. Panels **a** and **c** show the corresponding voltage traces: V_{cmd} is shown in black and V_m is shown in blue.

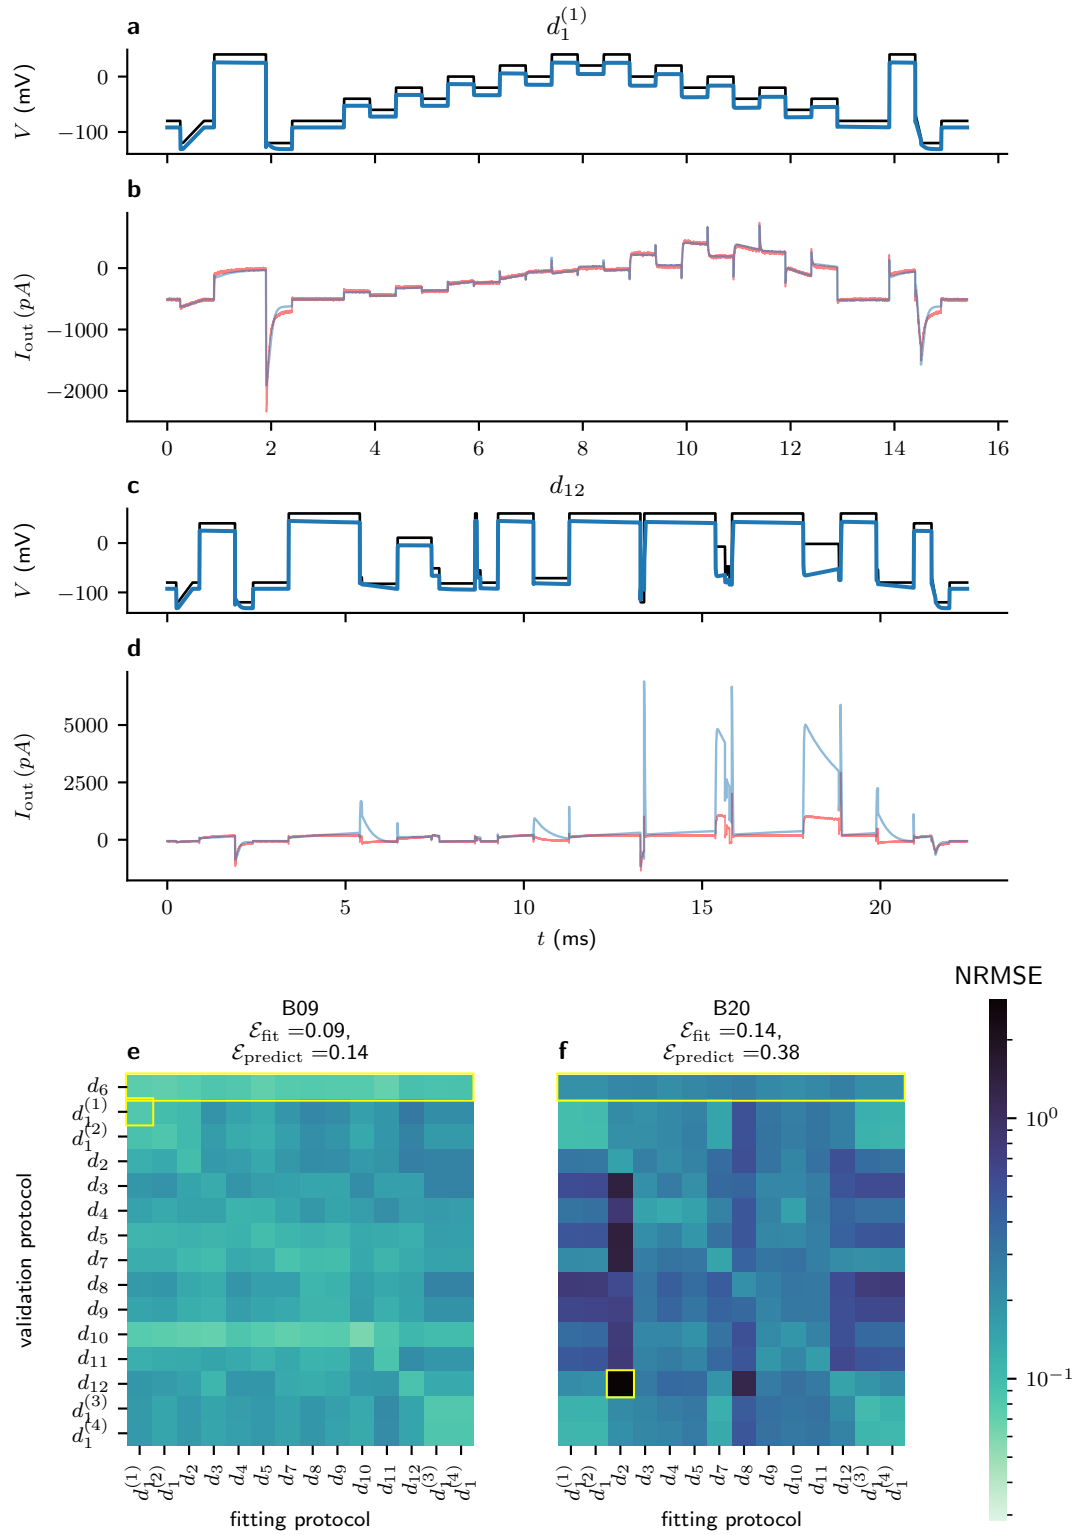


Figure 5.12: Cross-validation heatmaps for the Beattie model under Case V (with artefacts), showing the best and worst performing wells (as quantified by $\mathcal{E}_{\text{pred}}$). The best fit (and the heatmap cell exhibiting the lowest error) from the best performing well (Well B09, panel e) is shown in panels a and b, and the worst prediction from the worst performing well (Well B20, panel f) is shown in panels c and d.

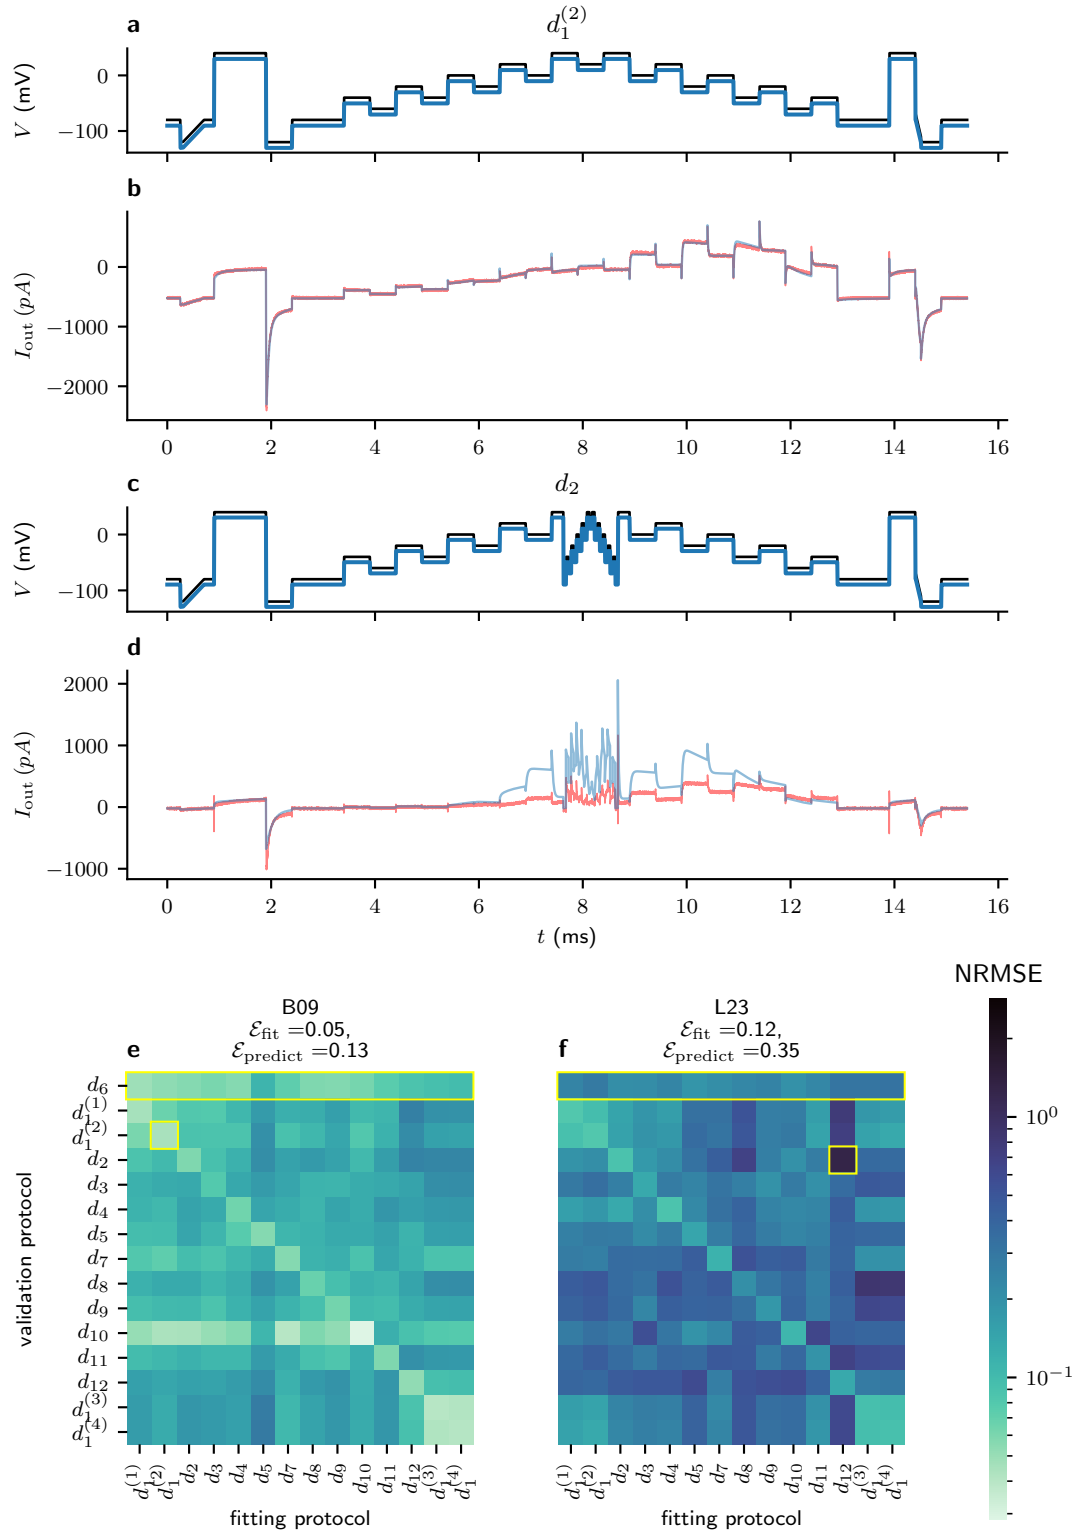


Figure 5.13: Cross-validation heatmaps for the Wang et al. (1997) model under Case IV (without artefacts), showing the best and worst performing wells (as quantified by \mathcal{E}_{pred}). The best prediction from the best performing well (Well B09, panel e) is shown in panels a and b, and the worst prediction from the worst performing well (Well L23, panel f) is shown in panels c and d. Panels a and c show the corresponding voltage traces: V_{cmd} is shown in black and V_m is shown in blue

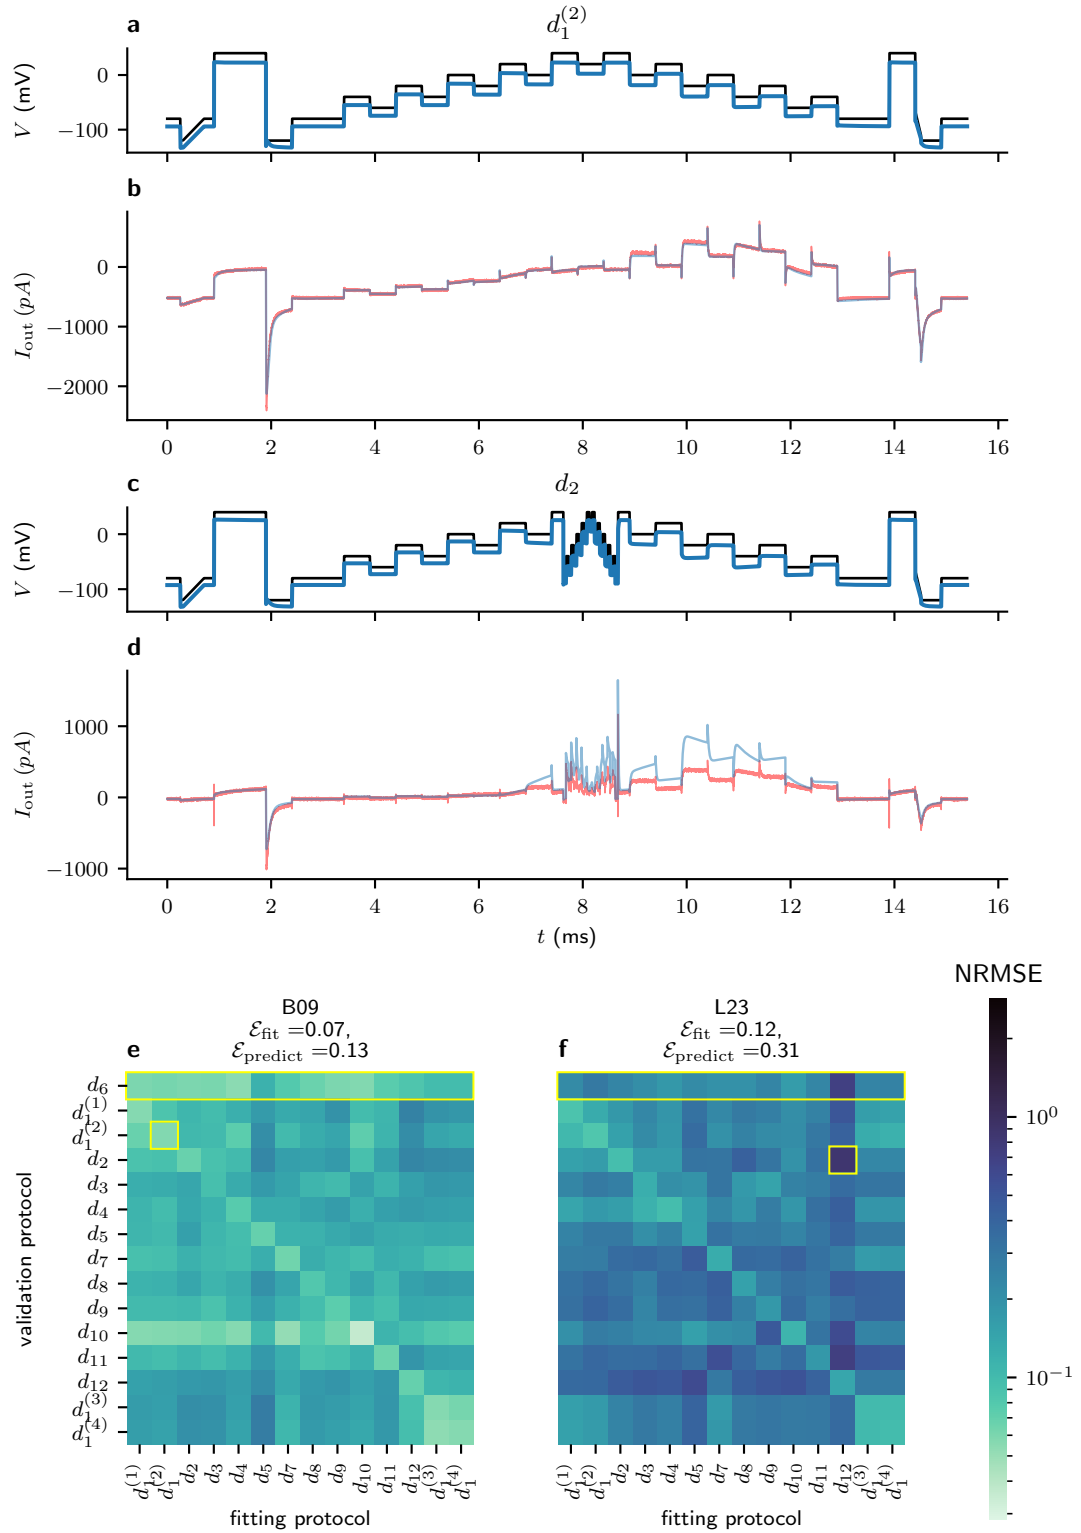


Figure 5.14: Cross-validation heatmaps for the Wang et al. (1997) model under Case V (with artefacts), showing the best and worst performing wells (as quantified by \mathcal{E}_{pred}). The best prediction from the best performing well (Well B09, panel e) is shown in panels a and b, and the worst prediction from the worst performing well (Well L23, panel f) is shown in panels c and d. Note that in both panels b and d, the model closely matches the large tail current occurring at $t = 2$ s.

and predictions than our ideal-patch models. Also, the beneficial effect of including additional closed states is also apparent because models fitted to d_2 with only a single level of closed states (the C-O-I and Beattie et al. (2018) models) are noticeably less accurate (in terms of prediction) than those include an extra layer of closed states (the Kemp et al. (2021) and Wang et al. (1997) models) as shown in Figure 5.15 and, especially, in Figures 5.16 and 5.17. These results support the inclusion of deeper closed states in our I_{Kr} model structures.

5.5.3 Variability in parameter estimates

As in the previous chapter, we consider the collection of parameter estimates obtained by fitting to each sweep of data (that is, each repeat of each protocol in each well). Accordingly, we proceed to use similar methodology to the previous section, using a simple linear statistical model (assuming independent well and protocol effects) to quantify the variability in our models' parameter estimates. As in Chapter 4, this simple linear statistical model recapitulates the variability of our parameter estimates reasonably well. These protocol-dependent effects are shown in Figures 5.18 and 5.19, for the Beattie et al. (2018) Model under Cases IV and V, respectively. Likewise, the apparent well-dependence of the same parameter estimates are shown in Figures 5.20 and 5.21 (again, for Cases IV and V, respectively). Here we see that the output of our statistical model is largely aligned with the parameter estimates computed during model fitting. As in the previous chapter, the suitability of these linear statistical models is discussed further in Appendix E.

Again, as in Chapter 4, we use log-likelihood differences to quantify the magnitude of well and protocol dependence in our parameter estimates. As in Chapter 4, we compute two log-likelihood difference statistics, **LLD-W** quantifying the degree of well dependence, and **LLD-D** quantifying the degree of protocol dependence. These values are presented in Table 5.2 for each of our candidate model structures. In any case, there remains a noticeable well and protocol dependence in our parameter estimates, indicating that our implementation of the artefact model is imperfect and does not full capture the aforementioned well-to-well variability. The size of **LLD-D** indicates model discrepancy in the dynamics of our models, because different protocols reliably result in different parameter estimates—this is similar to the synthetic data case discussed in Chapter 3.

These results suggest that, at least for the C-O-I and Beattie et al. (2018) models, artefact effects may explain some of the variability in parameter estimates. However, this difference is quite subtle, and the Case IV models seem to better match our recordings (see Figure 5.15). This well-to-well variability is also reflected in our model outputs. To show this, we produce normalised model output under d_1 and d_6 without the presence




Figure 5.15: Averaged cross-validation heatmaps for each candidate model under Cases IV and V. Generally speaking, the Case V models tend to produce worse fits than the Case IV models. Though, the best predictive model (as quantified by $\mathcal{E}_{\text{predict}}$) is the Wang model under Case V (with artefacts).

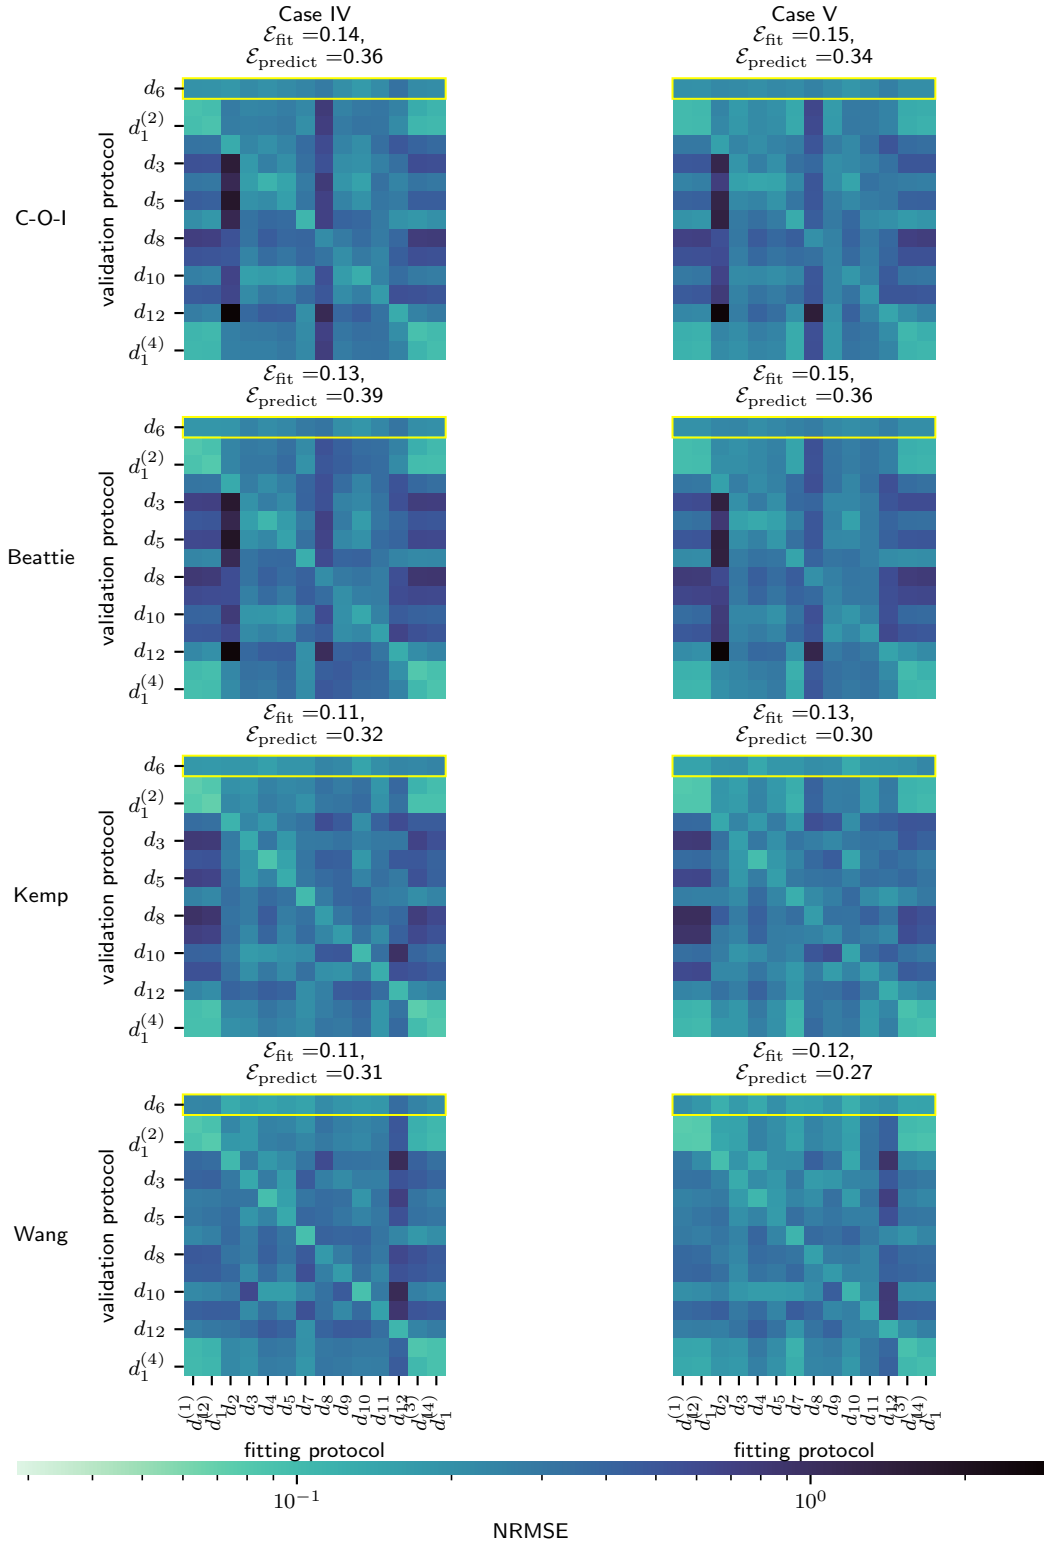


Figure 5.16: Cross-validation heatmaps showing the prediction error, in terms of the normalised root-mean-square error (NRMSE), when fitting and validating each model using data from Well B20. This well produced the worst average validation score for the Wang et al. (1997) Model under both Cases IV and V. We can see that the Case V models perform better than the Case IV models in terms of fitting, but worse in terms of prediction of unseen protocols. Under either case, the Wang model seems to produce both the best fits and predictions.

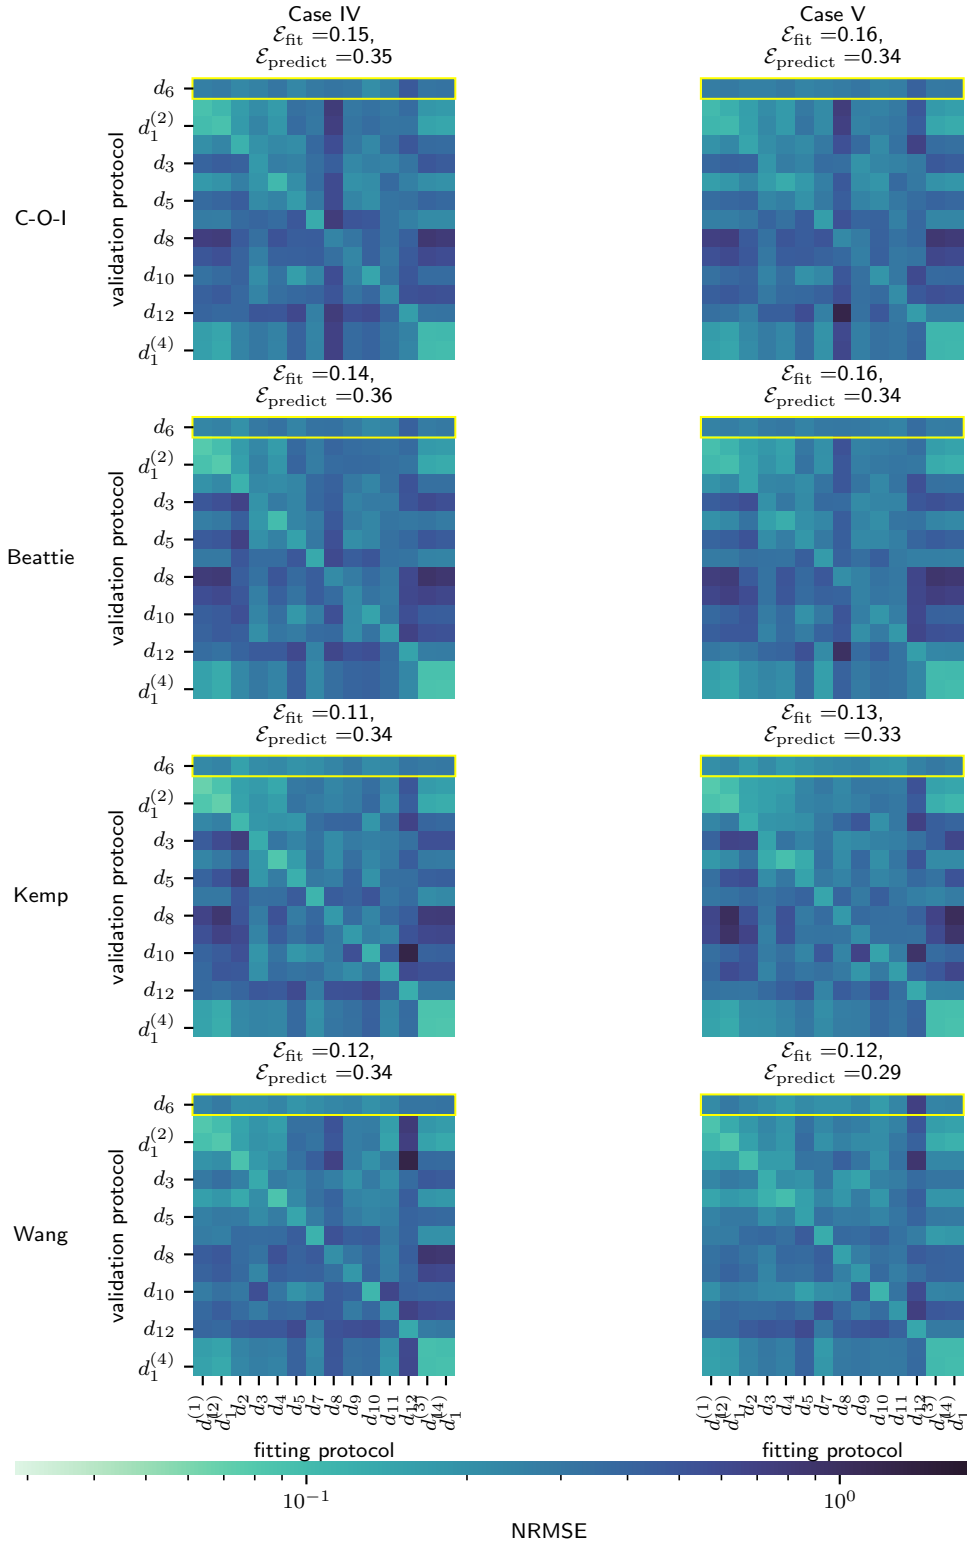


Figure 5.17: Cross-validation heatmaps showing the prediction error, in terms of the normalised root-mean-square error (NRMSE), when fitting and validating each model using data from Well L23. This well produced the worst average validation score for the Beattie et al. (2018) Model under both Cases IV and V. We can see that the Case V models perform better than the Case IV models in terms of fitting, but worse in terms of prediction of unseen protocols. Under either case, the Wang model seems to produce both the best fits and predictions.

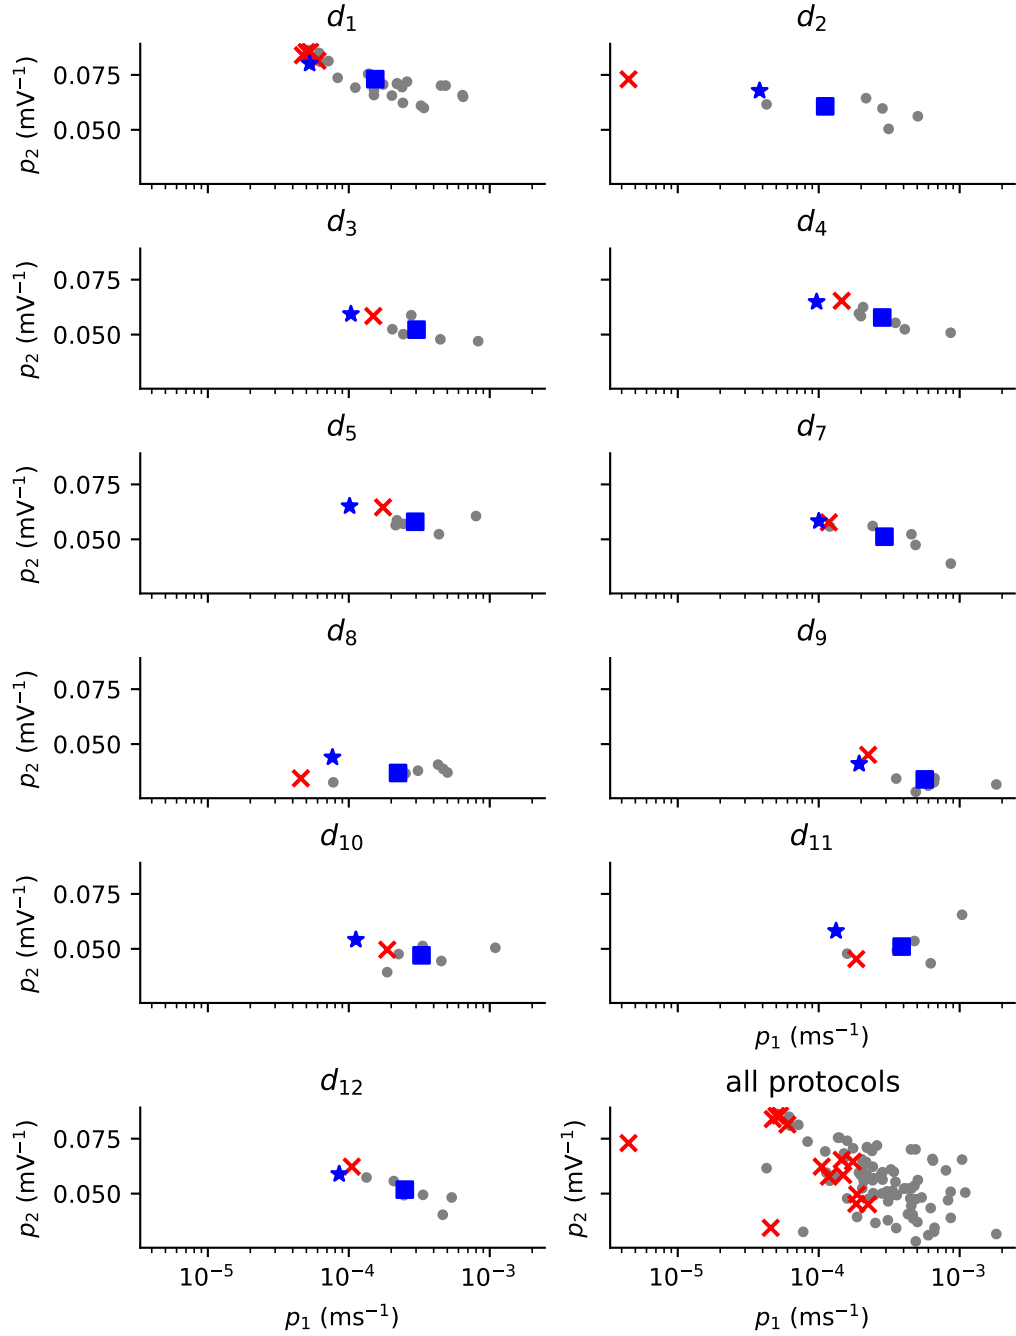


Figure 5.18: Parameter estimates obtained from each trace using the Beattie et al. (2018) Model (Case IV). Each panel shows the parameter estimates obtained from a given protocol, and highlights the parameter estimates obtained from well B20 (red crosses). The blue square shows the well-dependent effect according to \mathcal{M}_w , our linear model with no protocol-dependent effects, and the blue star shows the sum of the well and protocol effects according to $\mathcal{M}_{w,d}$. Note that p_1 , on the x-axis, is plotted on a log scale, consistent with how p_1 appears in the linear model.

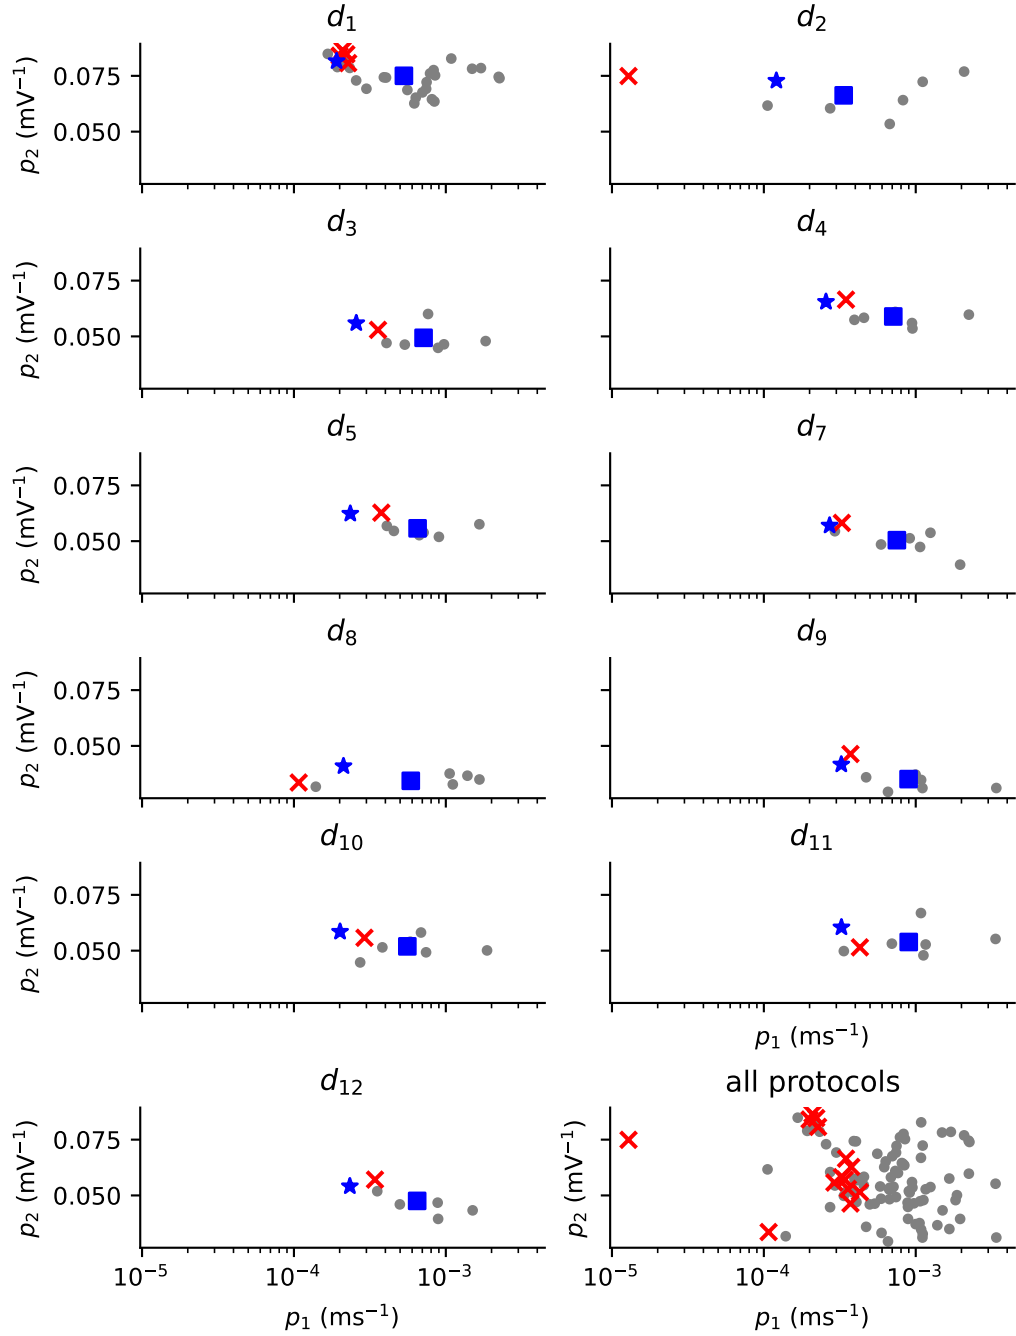


Figure 5.19: Parameter estimates obtained from each trace using the Beattie et al. (2018) Model (Case V). Each panel shows the parameter estimates obtained from a given protocol, and highlights the parameter estimates obtained from well B20 (red crosses). The blue square shows the well-dependent effect according to \mathcal{M}_w , our linear model with no protocol-dependent effects, and the blue star shows the sum of the well and protocol effects according to $\mathcal{M}_{w,d}$. Note that p_1 , on the x-axis, is plotted on a log scale, consistent with how p_1 appears in the linear model.

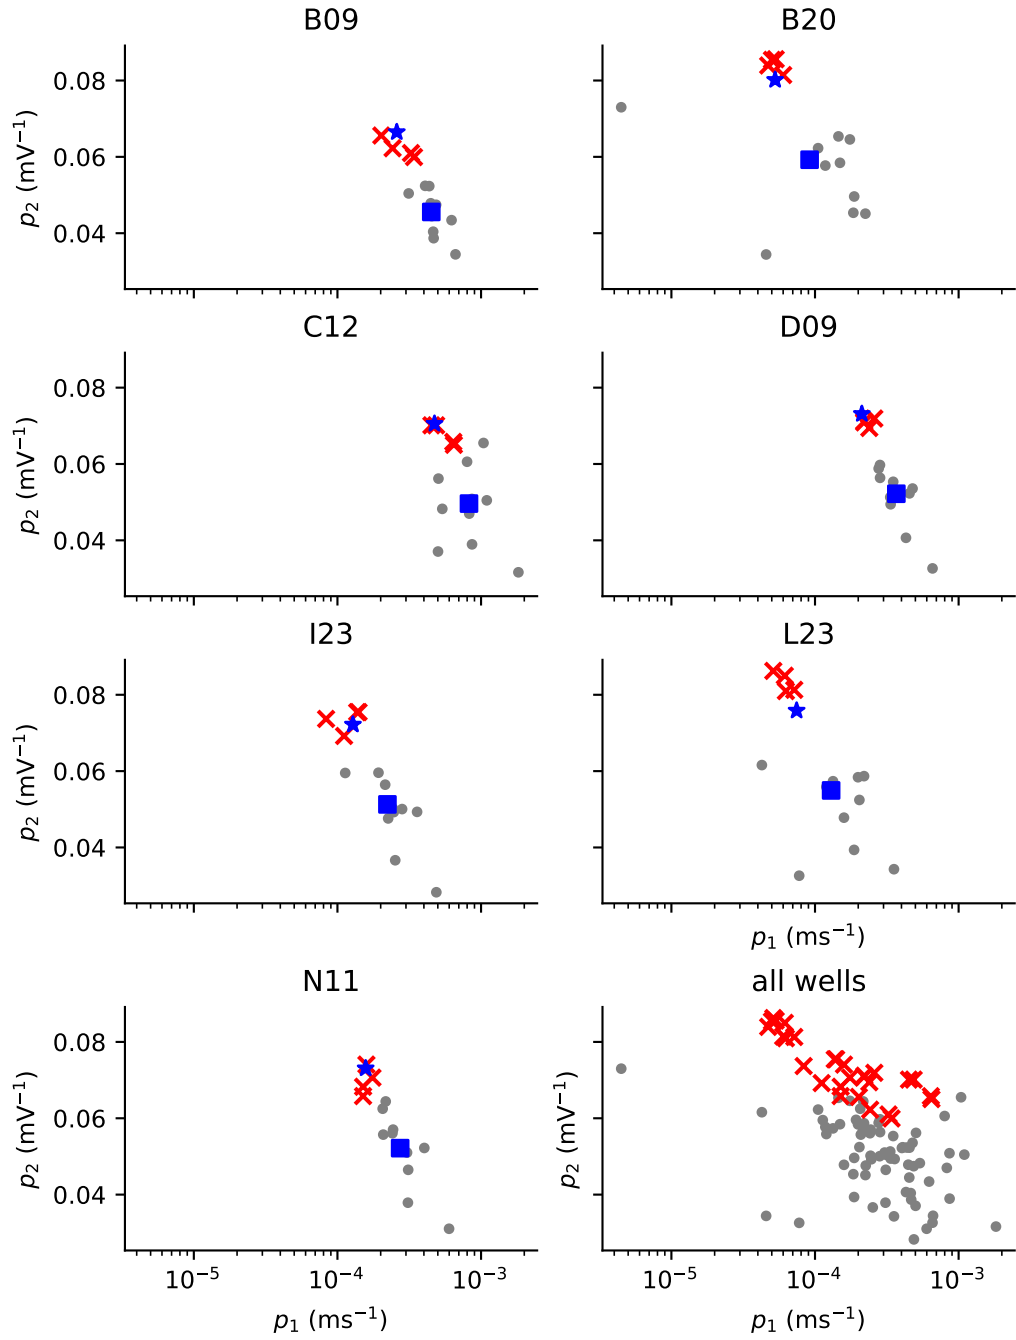


Figure 5.20: Parameter estimates obtained from each trace using the Beattie et al. (2018) Model (Case IV). Each panel shows the parameter estimates obtained from a given well, and highlights the parameter estimates obtained from the *staircase* protocol (d_1). Data points arising from repeats of this protocol are highlighted with red crosses. The blue square shows the well-dependent effect according to \mathcal{M}_w , our linear model with no protocol-dependent effects, and the blue star shows the sum of the well and protocol effects according to $\mathcal{M}_{w,d}$. Note that p_1 , on the x-axis, is plotted on a log scale, consistent with how p_1 appears in the linear model.

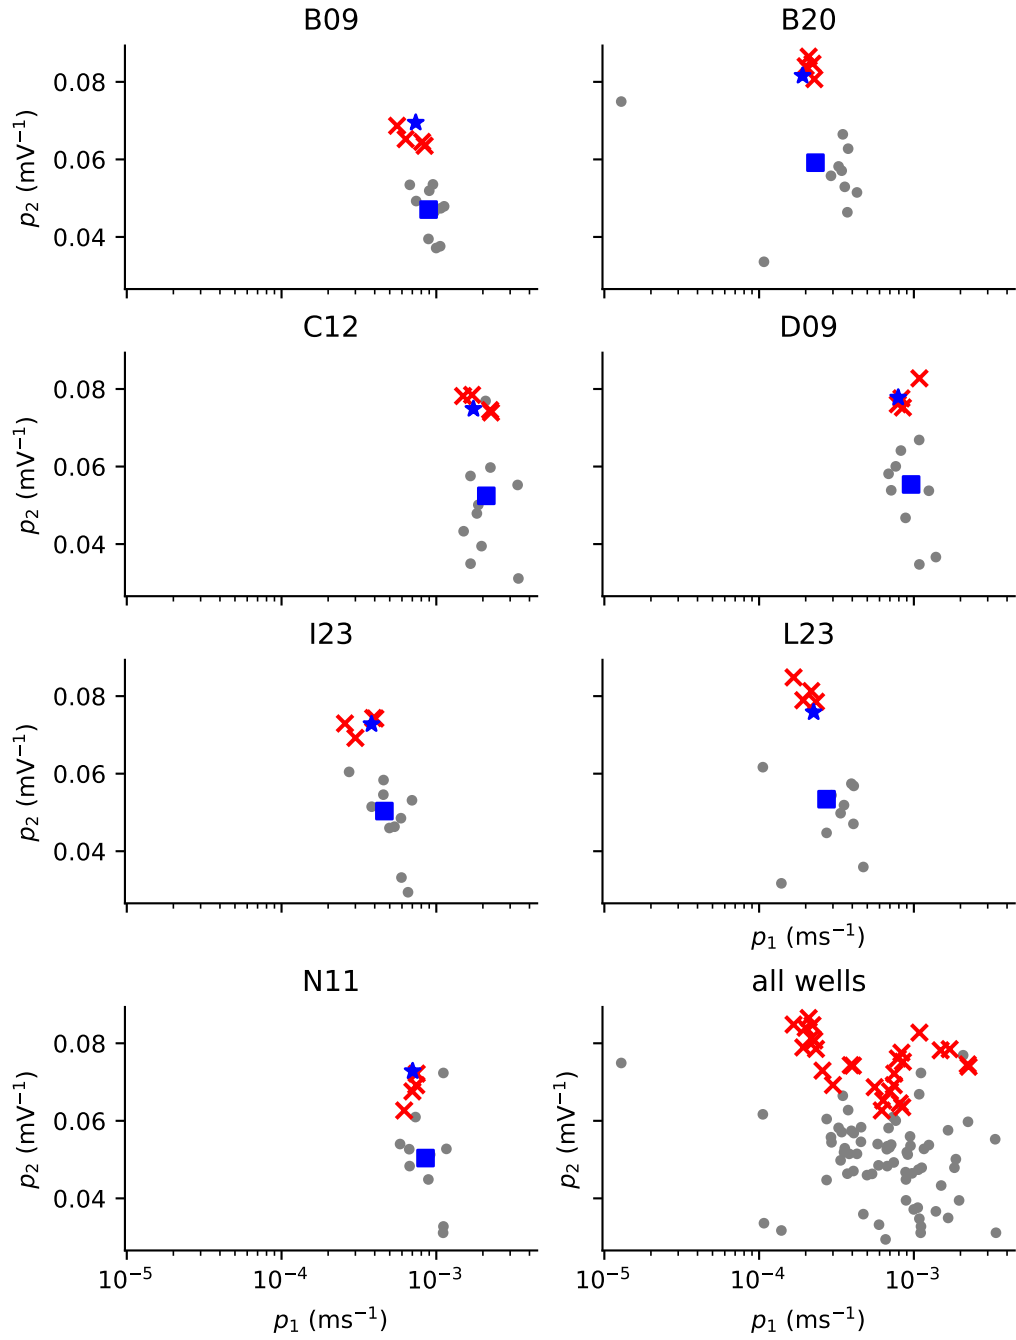


Figure 5.21: Parameter estimates obtained from each trace using the Beattie et al. (2018) Model (Case V). Each panel shows the parameter estimates obtained from a given well, and highlights the parameter estimates obtained from the *staircase* protocol (d_1). Data points arising from repeats of this protocol are highlighted with red crosses. The blue square shows the well-dependent effect according to \mathcal{M}_w , our linear model with no protocol-dependent effects, and the blue star shows the sum of the well and protocol effects according to $\mathcal{M}_{w,d}$. Note that p_1 , on the x-axis, is plotted on a log scale, consistent with how p_1 appears in the linear model

Model	Case	\mathcal{M}_0	\mathcal{M}_w	\mathcal{M}_p	$\mathcal{M}_{w,d}$	LLD-W	LLD-D
C-O-I	Case IV	1313.2	1512.1	1553.0	1996.8	443.8	484.7
	Case V	1104.9	1249.0	1320.6	1576.6	256.0	327.6
Beattie	Case IV	1372.3	1578.8	1622.2	2070.1	447.9	491.2
	Case V	1394.9	1620.8	1661.5	2099.1	437.7	478.4
Kemp	Case IV	849.5	999.8	1039.8	1266.0	226.2	266.2
	Case V	915.9	1085.5	1127.6	1408.0	280.5	322.5
Wang	Case IV	−362.3	−226.8	−88.0	127.6	215.6	354.3
	Case V	−222.0	−95.9	104.7	290.8	186.1	386.7

Table 5.2: Log-likelihoods and likelihood ratios for each of our model structures fitted using Cases IV and V. The statistic **LLD-W** quantifies the significance of our parameter estimates’ well-dependence, whereas **LLD-D** quantifies the significance of protocol dependence.

of artefacts and with $V_{\text{off}} = 0$. As shown in Figure 5.22, the well-to-well variability in model outputs under this regime are noticeably reduced when models are fitted with the artefact model. This suggests that, under Case IV, the transition-rate parameters in our models absorb some of the well-dependent artefact effects, whereas, when artefact effects are included, the resulting channel kinetics appear much more uniform across wells. Note, however, that the assumptions underpinning this linear model of parameter-estimate vectors may be somewhat less suitable for the Wang et al. (1997) model (under Case V) than the other models due to the apparent non-normality of residuals as discussed in Appendix E.

Nevertheless, the reduction in well-to-well variability of model output is corroborated by the values in Table 5.3 which show the average spread in predictions across the protocol for parameter sets obtained from different wells, that is,

$$\mathcal{B}_{\max}^{(i)} - \mathcal{B}_{\min}^{(i)}, \quad (5.14)$$

where $\mathcal{B}_{\max}^{(i)}$ and $\mathcal{B}_{\min}^{(i)}$ are the maximum and minimum model output at time t_i across our ensembles of parameter sets. However, unlike in Chapter 3, these ensembles of parameter sets are obtained from the same protocol (the first sweep of d_1 , to be precise), but from different wells.

5.5.4 Variability in time constants of decay

In Section 4.3, we saw that there is noticeable variability in the time-constant of decay during our protocols’ second -120 mV step. The current during this step was fitted

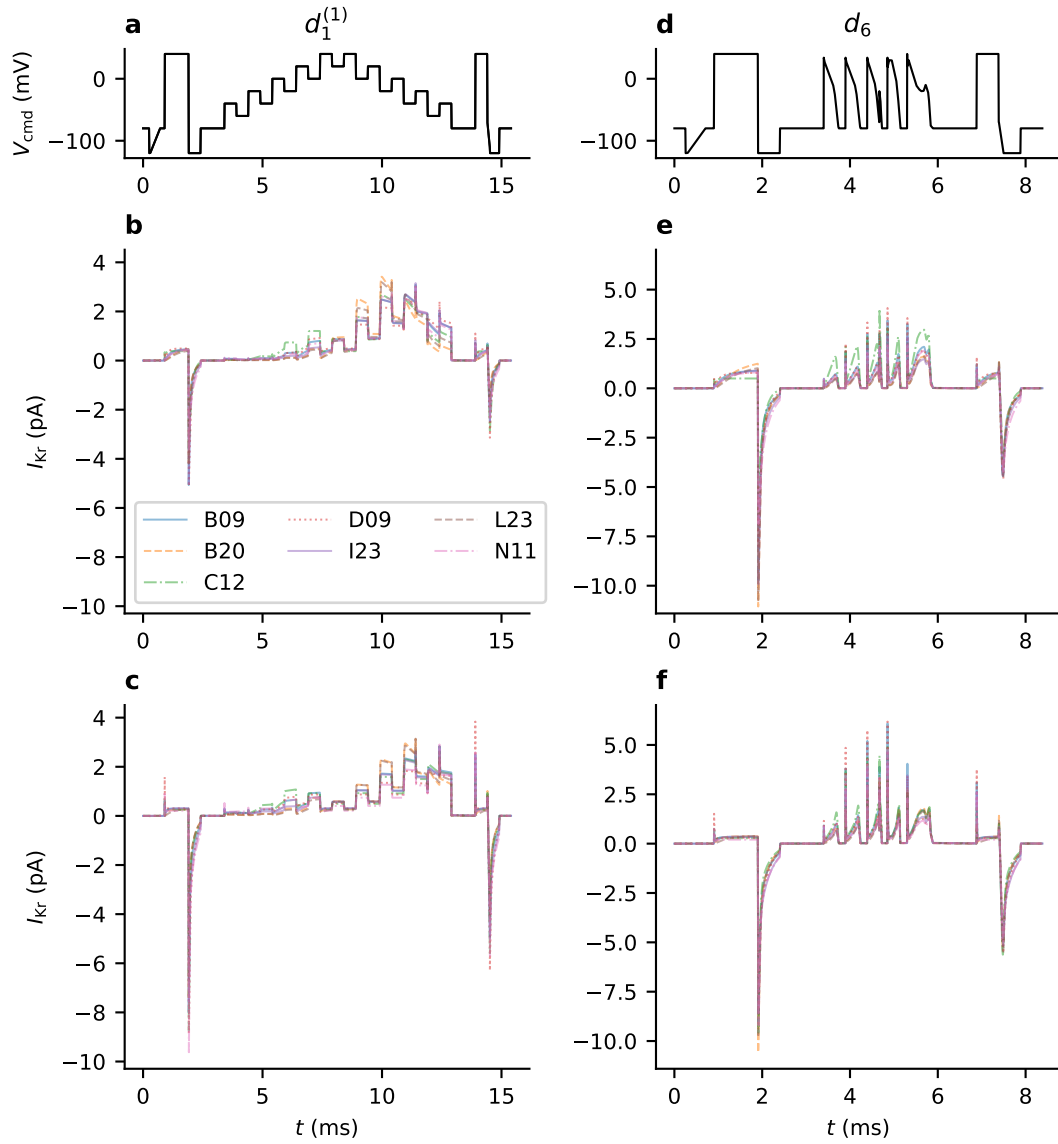


Figure 5.22: Fitting our models to data from different wells results in a range of model output (using the Wang et al. (1997) Model). Here, for d_1 (panels **a**, **c** and **e**), an ensemble of model output is generated using the parameter estimates obtained from each well (when fitted using d_1). This model output is then normalised by dividing it by its standard deviation. Panels **b**, **d** and **f** show protocol d_6 , for which we use the parameter estimates obtained from d_1 to produce an ensemble of predictions. The second row, panels **c** and **d** shows output arising from our Case IV parameter estimates, whereas panels **e** and **f** are produced using our Case V parameter estimates.

Protocol	Case IV	Case V
d_1	0.25	0.22
d_2	0.26	0.22
d_3	0.23	0.18
d_4	0.30	0.17
d_5	0.29	0.21
d_6	0.35	0.20
d_7	0.17	0.13
d_8	0.22	0.16
d_9	0.21	0.13
d_{10}	0.21	0.11
d_{11}	0.21	0.13
d_{12}	0.30	0.26

Table 5.3: The spread between maximum and minimal predictions under each of our protocols. For each protocol, an ensemble of model output is generated using the Wang et al. (1997) Model parameter estimates obtained from each well under that protocol. In the case of d_6 (which is not used for model fitting) we use the parameter estimates obtained from d_1 instead. This model output is then normalised by dividing it by its standard deviation. The values shown are the average spread of this normalised model output, averaged over the duration of the protocol.

using a simple, single-exponential decay model of the form,

$$I_{\text{post}}(t_i) = \alpha \exp \left\{ \frac{t_i - t_{\text{peak}}}{\tau} \right\}, \quad (5.15)$$

where τ and α where subject to optimisation, and t_{peak} is the time at which the (negative) peak current was observed. We showed that there was noticeable variability in τ . However, under our ideal-patch assumptions, such variability is not possible unless there is large variability in the gating kinetics of different cells. An alternative hypothesis (Lei et al., 2020a) is that this variability is largely caused by the variability in artefact effects between wells.

To investigate to what extent our artefact model explains this variability, we choose one set of parameter values for the Kemp et al. (2021) model, those obtained from the first sweep of the *staircase* protocol in Well B09 under Case III. We assume these are the correct kinetic parameters for I_{Kr} , and for each trace in our data, simulate the artefact model using our estimates of R_{series} , C_m , g_L , E_L , and V_{off} that are computed as described in Section 5.4. Here, we calibrate the maximal conductance, g , by minimising the sum-of-square errors between the model output and the data from the start of the trace to the end of each protocol's second -120 mV step (until 50 ms after this step, to be precise). We then fit a single time-constant of decay as described in Section 4.3.

The results of this exercise are shown in Figure 5.23. Here, we can see that the artefact effects included in our model do induce changes in τ , and these changes largely coincide with those found in the data. This correlation exists within the estimates of τ taken from a particular well (which is most noticeable in Well C12), but also between wells. Note, however, that our artefact model does not perfectly recapitulate these values, especially in the case of Well N11, suggesting significant discrepancy between our artefact model and the true data-generating process. It is noteworthy that Well N11 has the highest estimates of C_m and the lowest estimates of R_{series} .

5.6 Discussion

In this chapter, we developed new methods allowing the multiprotocol, cross-validation approach (as introduced in Chapters 3 and 4) to be applied to models augmented with artefact effects. This is the first use of a suite of information-rich protocols to fit and validate Markov models of I_{K_r} . An important feature of this approach is that it enables artefact-model parameters (namely, g_L , E_L and V_{off}) to be estimated in a way that is independent of the particular I_{K_r} model under investigation (albeit using a reference model for the estimation of V_{off}). This enabled us to provide a comparison of models with and without the addition of Lei et al. (2020a)’s artefacts effects model.

Though our analysis of our ensembles of parameter estimates may suggest that the well-to-well variability of the resultant parameter estimates is lessened for the C-O-I and Beattie et al. (2018) models, the same was not found to be true of the Kemp et al. (2021) model. It should be noted that the predictive performance of our ideal-patch models may be substantially reduced if capacitive spikes are included in the data for fitting and prediction, because this is where transient effects of the experimental artefacts (in particular the capacitive charging and discharging of the cell membrane) are most pronounced. Note that capacitive spike removal (removing 5 ms of data after each voltage discontinuity, as described in Appendix C), was performed to enable a fair comparison between the models as models of I_{K_r} . It is likely that our artefact models would have out-performed our ideal-patch models had this data been included for training and validation.

The reduction in fitting accuracy shown in Section 5.5.1, together with the limited ability of our models to recapitulate basic tail currents (as discussed in Section 5.5.4) indicates some deficiencies in our artefact model. It appears that model discrepancy is more noticeable in particular wells, such as Well N11 (as shown in Figure 5.23). These problems may arise due to certain limitations of our methodology. Firstly, our models rely on particular estimates of C_m and R_{series} provided by the automated patch-clamp platform. Whilst these estimates have enabled accurate predictive models to be fitted, the robustness and accuracy should be improved, if possible. In particular, the

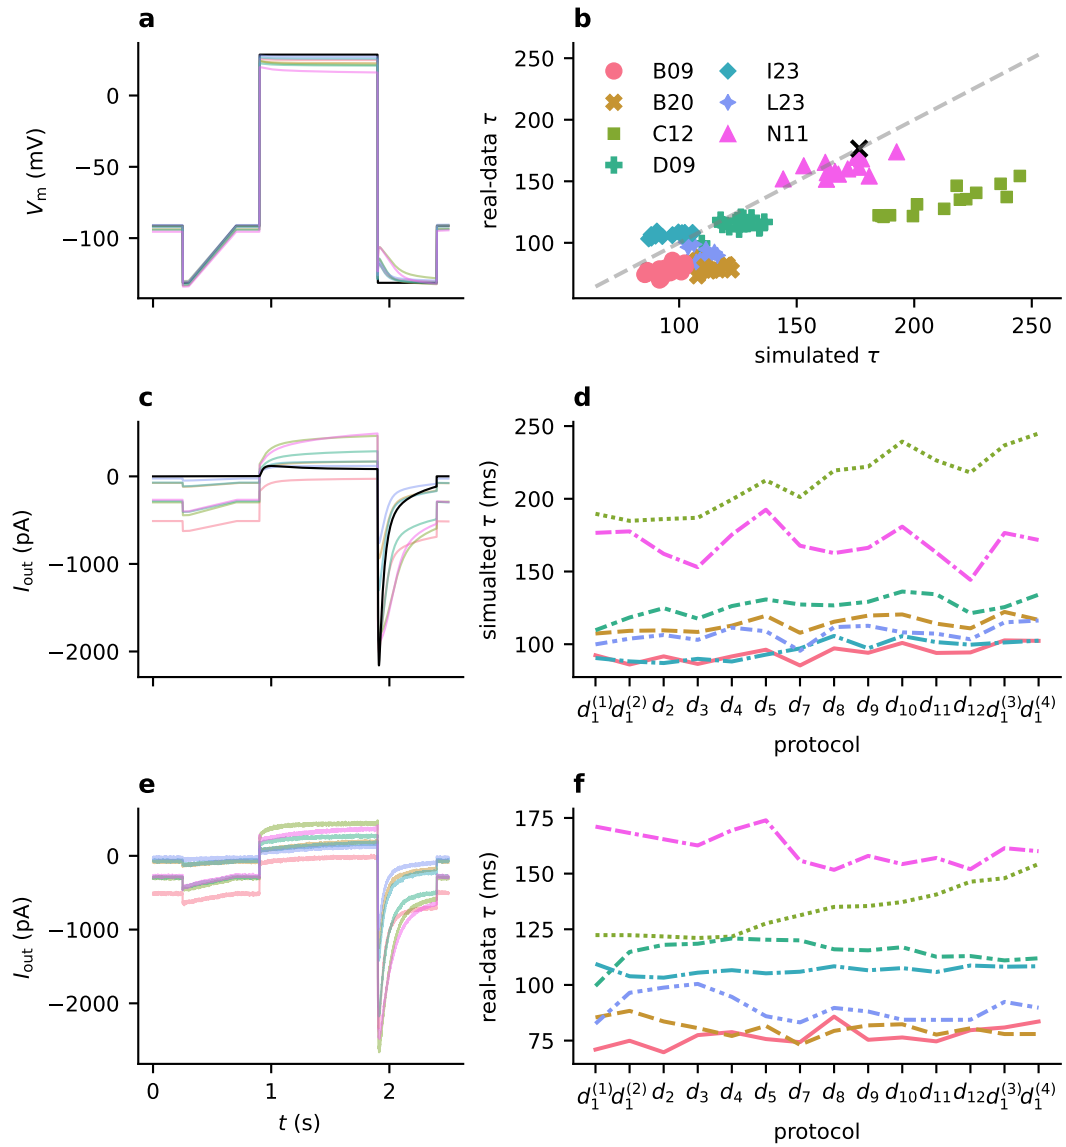


Figure 5.23: With fixed kinetic parameters and trace-specific artefact parameters, the artefact model somewhat recapitulates the variability found in the time-constant of decay for the second -120 mV step. Panel **a**: the predicted membrane potential for each of the wells under the application of the first portion of the *staircase* protocol. Here, the output from our ideal-patch model (with no leak current) is shown in black. Panel **b**: a scatter plot of the time-constant of decay for our protocols' second -120 mV step for a fixed set of kinetic parameters (but with fitted artefact parameters) compared with that of the data. Panel **c**: the real-data traces together with the model fit. Panel **b**: the time-constants of decay obtained from our real data and those obtained from our model output. These values are plotted chronologically (for each sweep) in Panels **d** and **f**.

robustness of these estimates for large values of C_m and R_{series} should be ensured, as some such wells (Wells B20, N11 and C12, for example) seem to exhibit greater model discrepancy. Nevertheless, some amount of discrepancy is unavoidable when fitting to a real, experimental dataset due to the presence of endogenous currents, and possible data pollution.

For the majority of protocols, our models produced similarly accurate model predictions. However, parameter estimates obtained from the d_1 protocol seem to produce noticeably worse model predictions, despite producing accurate fits to the data. The same can be said of parameter estimates obtained from protocol d_2 under the C-O-I and Beattie et al. (2018) models (Figure 5.15). Perhaps this suggests that the additional activation/deactivation transition rates included in the Kemp et al. (2021) model (k_5 and k_6) are necessary to build an accurate I_{K_r} model under the dynamics induced during the fast, central portion of the d_2 protocol (at room temperature, at least). Similarly, the staircase protocol seems to yield noticeably less accurate predictions in all but the Wang et al. (1997) model, perhaps indicating that the independence of activation and inactivation, together with deeper closed states, are necessary properties of an accurate I_{K_r} model.

The work in this chapter is limited by the consideration of only four of many possible model structures of I_{K_r} . Our methodology may be repeated for any other possible model structures, possibly leading to more accurate predictive models. An obvious extension would be to the Mazhari et al. (2001) model, which has a similar topology to the Wang et al. (1997) model but includes a direct inactive-closed transition which bypasses the open state. The inclusion of such a pathway is supported by single-channel results (Kiehn et al., 1999), and may lead to improved predictive accuracy, and perhaps a reduction in protocol-to-protocol and well-to-well variability in parameter estimates and model output.

In fact, this work could, in theory, be extended to any number of potential model structures, as in Mangold et al. (2021). However, the computational demands of these analyses are high—requiring models to be fit individually to numerous sweeps of data with each individual model structure being fitted to 12 protocols of data (with some repeats) for each of our 8 wells, with and without the inclusion of artefact effects. To this end, our work in this chapter, and the previous chapter, may be seen as proof-of-concept, introducing the methods necessary to repeat this analysis with other model structures, datasets and perhaps other ion-channel currents.

Regarding further data collection, we found that, with and without the presence of experimental artefacts, a large voltage offset is necessary to explain the apparent discrepancy between E_{post} and E_{Nernst} . Further datasets (perhaps using different experimental setups) should be analysed to ensure that these results hold when such a large voltage-offset is not present. With the aim of identifying the most accurate models

of I_{Kr} , such future work should focus on the inclusion of more model structures in the analysis, improvements to the artefact model (including new methodology for the estimation of relevant parameters), and the analysis of other, similar datasets (perhaps using newly designed protocols).

Chapter 6

Discussion

The selection of a suitable mathematical model for I_{K_r} remains a pertinent and open question. In this thesis, we have used information-rich, time-series data to fit and validate a small cohort of candidate models using similar methodology to that laid out in Fink and Noble (2009); Beattie et al. (2018); Clerx et al. (2019a). That is, we fit model parameters to cell-specific, time-series recordings of currents obtained under a short, information-rich protocol (referred to as whole-current fitting), rather than summary statistics gathered from a range of experiments performed on multiple cells (Clerx et al., 2019a). However, this previous work has used information-rich protocols only in synthetic-data studies (Fink and Noble, 2009), or used only a single information-rich protocol for model fitting (Beattie et al., 2018; Clerx et al., 2019a; Lei et al., 2019a). In this thesis, we present the first analysis of a multiprotocol experiment where many information-rich protocols have been used for both fitting and validation. The resulting observations of I_{K_r} under unexplored dynamics provides insight into the nature of model discrepancy. Using the whole-current approach, these data allow for more thorough model validation, as shown in Chapters 4 and 5. However, as these chapters demonstrate, the whole-current approach is sensitive to the presence of experimental artefacts and small amounts of data pollution—such as nonlinear or time-dependent leak (Lei et al., 2020b).

Nevertheless, we were able to mitigate the presence of artefacts and nonlinear leak through the use of a fluoride-free set-up and the application of stringent quality-control criteria. A key improvement in our work is the application of multiple voltage-clamp protocols during a single experiment, allowing a more thorough validation of dynamic models of I_{K_r} , demonstrating that our cell-specific models can be used to make largely accurate quantitative predictions of I_{K_r} for a broad ranges of voltage signals. Additionally, most of these protocols have not been used previously, meaning that our data promises to provide further information for selection of suitable I_{K_r} model structures.

We first presented a synthetic data study, demonstrating this multiprotocol approach and the anticipated impact of model discrepancy on the resulting parameter estimates and model predictions. Here, we showed that model discrepancy can be exhibited as protocol-to-protocol variability in parameters estimates, despite low apparent uncertainty in each individual estimate (under the assumption that the discrepant model is correctly specified). The tendency of discrepant models to produce accurate fits to training data, but inaccurate predictions under unseen protocols, is also demonstrated here.

In Chapters 4 and 5 we demonstrated that these symptoms of model discrepancy were also found when fitting our models to real, experimental data. Considerable attention was given to our postprocessing methodology presented in Chapter 4. By augmenting methods used in previous work (Lei et al., 2019b), we were able to use a fully-automated procedure to perform quality control, resulting in a small selection of wells exhibiting remarkably clean recordings—where we see little endogenous current or changes in our recordings of I_{K_r} over the course of the experiment. Then, later in Chapter 4, we fitted and validated our candidate models to these data under the assumption of an ideal patch-clamp setup. Whilst these modelling assumptions result in the omission of known experimental artefacts (Lei et al., 2020a), they resulted in largely accurate, well-specific, predictive models of I_{K_r} . It was our opinion, however, that our Markov models were overfitted to the unaccounted-for, well-specific, artefact effects, resulting in significant well-to-well dependence in the resulting parameter estimates—even though we obtained largely accurate fits and predictions for each cell, individually. The fact that these parameter estimates are, to a large degree, confounded with uncompensated artefact effects undermines their suitability for general applications (in action-potential models and studies on drug block, for instance).

It did seem possible that inclusion of these artefact effects, by way of Lei et al. (2020a)’s simplified artefact model, could remedy this situation. However, the integration of such effects within our multiprotocol approach entailed some challenges—namely the estimation of newly introduced artefact parameters. Here, we introduced a method to infer the leak-model parameters (g_L and E_L), and a voltage offset, V_{off} , using the reversal and leak ramps included in each protocol. This allowed the fitting and validation of our I_{K_r} models with artefact effects, without the need to include additional parameters in our model-fitting routines (during which only the maximal conductance and parameters relating to our models’ transition rates are inferred).

The imperfect estimation of our artefact-model parameters, including the machine estimates of C_m and R_{series} may have contributed to model discrepancy somewhat. Nevertheless, the inclusion of these artefact effects seemed to yield a more accurate predictive model—our most accurate predictive model being the Wang et al. (1997) Model augmented with artefacts. However, the improvement in terms of predictive

accuracy was rather slight, and the quality of our model fits actually reduced with the addition of the artefact model. Notably, the well-to-well variability of Wang et al. (1997) Model parameter estimates appears to have reduced with the addition of artefact effects. We showed that this model, the Wang et al. (1997) Model with additional artefact effects, was, to some extent, able to explain the noticeable well-to-well variability in our -120 mV tail currents across wells under the assumption of uniform kinetic parameters. This may suggest that artefact effects are responsible for much of the well-to-well variability we have seen, both in our data, and in our parameter estimates throughout Chapters 4 and 5. Perhaps improved estimation of key artefact-model parameters (namely, C_m and R_{series}) may yield even more convincing results.

Here, our focus on general predictive accuracy (across a range of fitting and validation protocols) contrasts with more typical model selection criteria. Perhaps the most commonly used is the Akaike Information Criterion (AIC) (Akaike, 1998; Vandenberg and Bezanilla, 1991), which concerns only the accuracy of model fits and the number of model parameters. However, as we demonstrate in Chapters 4 and 5, the model which produces the most accurate fits does not necessarily produce the most accurate predictions (even when the same number of model parameters are used), and may not be the most suitable.

Our results highlight the important role that artefact effects play in patch-clamp recordings of I_{Kr} . Here, we used two key aspects of our models for validation: the variability of our parameter estimates between different wells and protocols (which we would expect to be largely uniform across provided a correctly specified model), and the accuracy of our model predictions under unseen (during model fitting) validation protocols. In this thesis, both aspects of model validation have led us to the same conclusion: that it is necessary to include artefact effects in our models. However, in other situations where the model with the most accurate predictions is not the most interpretable, it may be necessary to compromise on one or the other: either prioritising quantitative accuracy or the interpretability of models (Lema-Perez et al., 2019).

Future directions

The consideration of more I_{Kr} model structures is a natural extension of our work. In particular, the fact that the Wang et al. (1997) model produced the most accurate predictions (on average) is surprising given that a direct closed-open pathway has been suggested somewhat more recently (Kiehn et al., 1999; Mazhari et al., 2001). Perhaps the model error induced by the omission of this transition in the Wang et al. (1997) model is much less significant than the assumption of independent activation and inactivation assumed by both the Beattie et al. (2018) and Kemp et al. (2021) models.

Moreover, the extension of this work to further model structures could be achieved through an iterative computational approach (Menon et al., 2009; Mangold et al., 2021).

Many general approaches to the identification of ODE systems have been presented in the literature. Sparse identification of Nonlinear Dynamics (SiNDy) (Brunton et al., 2016; Massonis et al., 2023) and Graph Reconstruction by Additive Differential Equations (GRADE) (Chen et al., 2017) are two such examples which are used to consider a wide range of possible model structures. The direct application of these methods to ion-channel recordings may produce cell-specific models that produce accurate predictions across a range of voltage-clamp protocols. However, such a model may be less interpretable than an equivalent Markov model (even one that produces worse predictions) because any model we fit to our patch-clamp datasets necessarily describes a combination of artefact effects and underlying I_{Kr} dynamics. It is, therefore, unclear how such a “grey-box” approach could disentangle artefact effects from the underlying dynamics of I_{Kr} and allow the latter to be integrated into a whole-cell action potential model or used to model drug block, for example. For these purposes, the integration of domain-specific knowledge (regarding our experimental methodology and underlying mechanisms of channel gating, for instance) is essential.

Whilst our results suggest that the Wang et al. (1997) model is our most accurate predictive model, and that inclusion of artefact effects led to an improvement in predictive accuracy, these conclusions would be bolstered by the analysis of repeated runs of our multiprotocol experiments. Future work could also include new voltage-clamp protocols—ensuring that our results still hold under other choices of fitting and validation protocols. Here, care should be taken to ensure that whatever new protocols are used permit practical identifiability for every model under consideration. Because of the observed well-to-well variability in recorded currents, a probabilistic approach, such as Bayesian or pseudo-Bayesian optimal experimental design (Ryan et al., 2016; Overstall et al., 2018; Dette et al., 2023), may be considered to ensure that models are practically identifiable under any plausible instantiation of artefact-effect and kinetic parameters.

Under synthetic data generation, the generation of experimental designs which allow precise parameter estimation (practical identifiability) is possible by using the standard approaches discussed in Section 2.3.6. However, there is no guarantee that such designs lead to practical identifiability once data is collected—model discrepancy and data pollution, for example, may render some experimental designs unsuitable even if they are optimal under a given utility function. For this reason, future work should strive to ensure the robustness of voltage-clamp protocols by testing for practical identifiability under a range of model structures, artefact effects and, perhaps, the addition of small amounts of synthetically added, time-dependent leak current.

Experimental designs that permit discrimination between our candidate model structures are desirable, but, in practice, may be challenging to produce such designs in

the presence of model discrepancy. Standard approaches Atkinson and Federov (1975) may magnify model differences but fail to explore interesting, novel dynamics, as shown by the *squarewave* protocol (Lei et al., 2024). The combination of a “T-optimal” utility function (Atkinson and Federov, 1975; Dette et al., 2012; Tommasi, 2008) with the space-filling curve approach (Mirams et al., 2024) may succeed at highlighting model differences across a larger selection of our models’ state spaces. However, the computational requirements for the computational design of T-optimal are intense, requiring the solution of a nested optimisation problem.

Whilst we have aimed to validate our models through cross-validation with a range of voltage-clamp protocols, further validation of the resulting I_{Kr} models is necessary. This may be done, for example, by ensuring that integration of our I_{Kr} models with existing *in silico* action-potential models yields expected output, as has been done previously (Beattie et al., 2018). Further model validation may also make use of drug studies, where the interaction of a drug with I_{Kr} is recorded and then modelled. Here, we may seek to validate whether our chosen model structure and parameter estimates yield a plausible model of drug-channel interactions.

Finally, the experimental and mathematical methodology presented in this thesis may be extended to the study of other voltage-dependent ion-channel currents. The combination of rapid inactivation/recovery-from-inactivation and markedly slower activation/deactivation is a distinguishing feature of I_{Kr} . The dynamics exhibited by other voltage-dependent ion-channel currents differ from this, and, as such, our particular choice of voltage protocols may not be suitable. If not, new voltage-clamp protocols should be designed to account for the particular dynamics of the current under investigation. In this way, it may be possible to apply these methods to other important cardiac currents, such as the *fast sodium current*, I_{Na} (Ebihara and Johnson, 1980; Pathmanathan et al., 2015), for example.

Similarly to drug-channel interactions, genetic mutations affect the function of ion-channel currents like I_{Kr} , and consequently have an effect on the function of the heart as a whole (Perry et al., 2019; Whittaker et al., 2020b; Kemp et al., 2021). Markov models, such as those used in this thesis, could also be used to model the impact of such mutations, and any potential increase in risk under the application of a given drug. Here, our multiprotocol approach could make apparent important changes in the dynamics of channel gating and used to inform *precision medicine* approaches, where an individual’s characteristics (such as their genetics) are considered when clinical decisions are made regarding certain interventions (Niederer et al., 2019; Nakajima et al., 2021).

In any case, the methods presented herein not only provide novel insights into I_{Kr} , but provide a blueprint for assessing and improving the modelling of many other important ion-channel currents.

Appendix A

Fundamental properties of Markov models

In this appendix, we review some basic properties of Markov models, as introduced in Section 2.2.3. We begin by directly demonstrating how a Hodgkin-Huxley style gating variable may be reformulated as a simple chain-like Markov model. As discussed in Section 2.2.3, this shows that Markov models may be viewed as a generalisation of Hodgkin-Huxley style gating variables. Consequently, we do not need to consider Hodgkin-Huxley style gating variables separately when selecting model structures.

Then, in Section A.2 we discuss the behaviour of Markov models of macroscopic ion-channel currents when the command voltage, V_m is held constant. This often occurs under our ideal-patch assumptions (that is, for the models presented in Chapters 3 and 4), both at the start of our protocols, where V_{cmd} is held at -80mV , and during each “voltage step”. In these setting, the governing equation (see Equation (2.18)) becomes a linear system of ODEs. As discussed in Appendix F, these results also prove useful in the analysis of our artefact models (introduced in Chapter 5). Then in Section A.3, we discuss the connection between our deterministic ODE-based macroscopic current models and the stochastic models used to model the gating of individual channels.

A discussion of *microscopic reversibility* and its implications for the dynamics of Markov models follows in Section A.4. Here, we review the Kolmogorov Theorem (Kelly, 2011), which allows the reversibility of a Markov model to be easily checked via a simple graph-based approach. Finally, in Section A.5, we explain the method used to simplify our Markov models’ systems of ODEs. Whilst this method has been described (briefly) elsewhere (such as in Colquhoun and Hawkes (1995)), we present a general linear algebra approach that may be modified to improve numerical stability.

A.1 Hodgkin-Huxley style gating variables as Markov models

To show that we may express any Hodgkin-Huxley style gating variable as a Markov model, we consider a single variable, y , governed by,

$$\frac{dy}{dt} = (1 - y)\alpha(V_m) - y\beta(V_m). \quad (\text{A.1})$$

Examples of such gating variables appear in the Hodgkin-Huxley action-potential model (Equation (2.1)–(2.9)).

Let \mathbf{x} be a $(\mu + 1)$ -vector with i^{th} element,

$$x_i = \binom{\mu}{i} y^i (1 - y)^{\mu - i}, \quad (\text{A.2})$$

where $\binom{\mu}{i}$ denotes the number of i -combinations in a set of μ elements, $\binom{\mu}{i} = \frac{\mu!}{(\mu - i)!i!}$. Then x_i is the probability that a given channel has i independent, open gates. The channel only allows current to flow when all gates are open, hence the conductance of the channel is given by,

$$\bar{g}y(V_m - E) = \bar{g}y^\mu(V_m - E), \quad (\text{A.3})$$

where \bar{g} is the maximal conductance and E is the reversal potential of the channel. To show the equivalence of the Markov model formulation with the use of a Hodgkin-Huxley style gating variable, we seek to construct a Markov model, with state vector, \mathbf{x} , such that $x_O(t) = y(t)^\mu$.

Differentiation of Equation (A.2) yields,

$$\frac{dx_i}{dt} = \binom{\mu}{i} (iy^{i-1}(1 - y)^{\mu - i} - (\mu - i)y^i(1 - y)^{\mu - i - 1}) \frac{dy}{dt} \quad (\text{A.4})$$

$$= (\alpha(1 - y) - \beta y) \binom{\mu}{i} (iy^{i-1}(1 - y)^{\mu - i} - (\mu - i)y^i(1 - y)^{\mu - i - 1}), \quad (\text{A.5})$$

where we temporarily omit the dependence of the transition rates, $\alpha(V_m)$ and $\beta(V_m)$, on the transmembrane potential, V_m . Then, by expanding the right-hand side of this

expression, we see,

$$\begin{aligned} \frac{dy}{dt} = & \alpha i \binom{\mu}{i} y^{i-1} (1-y)^{\mu-(i-1)} \\ & - \alpha (\mu-i) \binom{\mu}{i} y^i (1-y)^{\mu-i} \\ & - \beta i \binom{\mu}{i} y^i (1-y)^{\mu-i} \\ & + \beta (\mu-i) \binom{\mu}{i} y^{i+1} (1-y)^{\mu-i-1}. \end{aligned} \quad (\text{A.6})$$

Note how some terms disappear when $i = \mu$ or $i = 0$ which correspond to $y_0 = (1-x)^\mu$ and $y_\mu = x^\mu$, respectively. Now, when $i > 0$, we have the identity,

$$i \binom{\mu}{i} = (\mu+1-i) \binom{\mu}{i-1}, \quad (\text{A.7})$$

and when $i < N$,

$$(\mu-i) \binom{\mu}{i} = (i+1) \binom{\mu}{i+1}. \quad (\text{A.8})$$

Hence, for $0 < i < N$, we may rewrite Equation (A.6),

$$\begin{aligned} \frac{dx_i}{dt} = & \alpha (\mu - (i-1)) \binom{\mu}{i-1} y^{i-1} (1-y)^{\mu-(i-1)} \\ & - \alpha (\mu-i) \binom{\mu}{i} y^i (1-y)^{\mu-i} \\ & - \beta i \binom{\mu}{i} y^i (1-y)^{\mu-i} \\ & + \beta (i+1) \binom{\mu}{i+1} y^{i+1} (1-y)^{\mu-i-1}, \end{aligned} \quad (\text{A.9})$$

We then have a system of coupled ODEs,

$$\frac{dx_0}{dt} = x_1 \beta - x_0 \alpha \mu, \quad (\text{A.10})$$

$$\frac{dx_i}{dt} = (\mu - i + 1) x_{i-1} \alpha - (\mu - i) x_i \alpha + (i + 1) x_{i+1} \beta - i x_i \beta \quad \text{for } 1 \leq i \leq \mu - 1, \quad (\text{A.11})$$

$$\text{and, } \frac{dx_\mu}{dt} = x_{\mu-1} \alpha - \mu x_\mu \beta. \quad (\text{A.12})$$

In fact, this system of ODEs may be expressed as a Markov model,

$$\frac{d\mathbf{x}}{dt} = \mathbf{Q}(V_m)^\top \mathbf{x}, \quad (\text{A.13})$$

where,

$$\mathbf{Q}(V_m) = \begin{pmatrix} -k_0 & k_0 & 0 & \dots & 0 & 0 & 0 \\ l_1 & -l_1 - k_1 & k_1 & \dots & 0 & 0 & 0 \\ & \vdots & & \ddots & & \vdots & \\ 0 & 0 & 0 & \dots & k_{\mu-1} & -k_{\mu-1} - l_{\mu-1} & l_{\mu-1} \\ 0 & 0 & 0 & \dots & 0 & l_{\mu-1} & -l_{\mu} \end{pmatrix}, \quad (\text{A.14})$$

is a $(\mu + 1) \times (\mu + 1)$ transition-rate matrix with,

$$k_i = (\mu - i)\alpha(V_m),, \quad (\text{A.15})$$

$$\text{and } l_i = i\beta(V_m). \quad (\text{A.16})$$

This demonstrates that any Hodgkin-Huxley style gating variable, raised to a positive integer power, may be expressed as a Markov model. As explained in the main text, a model may consist of the product of two or more such Hodgkin-Huxley style gating variables. The Beattie et al. (2018) is such an example (Clerx et al., 2019a).

The Markov model system of ODEs shown above has more state variables compared to Equation (A.1) ($\mu + 1$ state variables as opposed to 1), However, the relation, $y^\mu = x_O$, holds provided suitable initial conditions are chosen. In this thesis, our models are always initialised at equilibrium—a unique equilibrium point shared by both formulations, and so, when mapped to our observations through the observation function (Equation (2.19) in the main text) Hodgkin-Huxley style models (which may consist of one or more gating variables) and the equivalent Markov models provide identical model output. This steady-state behaviour is discussed in the following section.

A.2 Steady states of Markov models

In this section, we review some fundamental results regarding the dynamics of Markov models when V_{cmd} is held constant. Again, for concision, we omit the voltage dependence and consider linear systems of the form,

$$\frac{d\mathbf{x}}{dt} = \mathbf{Q}^\top \mathbf{x}, \quad (\text{A.17})$$

where, as in Chapter 2, we have:

$$\mathbf{Q} \text{ is a real-valued } (N \times N) \text{ matrix for all } V; \quad (\text{A.18})$$

$$\mathbf{Q}\mathbf{1} = \mathbf{0}; \quad (\text{A.19})$$

$$Q_{i,j} > 0 \text{ if and only if } Q_{j,i} > 0; \quad (\text{A.20})$$

$$\text{and for all } 1 \leq i \leq N, 1 \leq j \leq N, \text{ with } i \neq j, \text{ we have } Q_{i,j} \geq 0. \quad (\text{A.21})$$

Provided these conditions hold, it is simple to show that the total occupancy of the system is conserved. Differentiating Equation (A.19) constraint yields,

$$\frac{d}{dt} \mathbf{1}^\top \mathbf{x} \quad (\text{A.22})$$

$$= \mathbf{1}^\top \frac{d}{dt} \mathbf{x} \quad (\text{A.23})$$

$$= \left(\mathbf{1}^\top \mathbf{Q}^\top \right) \mathbf{x} \quad (\text{A.24})$$

$$= (\mathbf{Q}\mathbf{1})^\top \mathbf{x} \quad (\text{A.25})$$

$$= \mathbf{0}. \quad (\text{A.26})$$

Thus, the total occupancy remains constant, and we may, without loss of generality, assume that, the initial conditions of our governing equations are such that $\mathbf{1}^\top \mathbf{x} = 1$. This condition is chosen such that the i^{th} state-variable, $x_i(t)$, represents the proportion of channels in conformation i at time t .

Connectedness and Irreducibility

It is natural to associate \mathbf{Q} with a directed graph (digraph) D (Bondy and Murty, 2008) where an edge between states i and j exists, with weight $Q_{i,j}$, if and only if $Q_{i,j} > 0$. Then, some concepts from graph theory prove useful in characterising the behaviour of Equation (A.17). Moreover, given Property (A.20) holds, we need only consider the underlying (undirected) graph instead, which we denote G . In particular, G has a node for every state in the model, and an edge between states i and j if $Q_{i,j} > 0$. Note that a weighted digraph may also be called a *network* (Kelly, 2011).

Let G be an (undirected) graph. A subgraph $H \leq G$ is said to be connected if, for any two nodes, $u, v \in H$ there is a walk from u to v . In this way, a graph can be partitioned into *connected components*, each of which is a distinct, connected subgraph. A graph which has only one connected component is said to be connected. A matrix, say \mathbf{A} , is said to be reducible if there exists some *permutation* matrix \mathbf{M} such that \mathbf{MAM}^{-1} is upper triangular (Berman and Plemmons, 1979).

These concepts may be related to properties of our matrix \mathbf{Q} ; It is known that \mathbf{Q} is irreducible if and only if G is connected (Berman and Plemmons, 1979). As we will see

in the following section, the connectedness of G (and, equivalently, the irreducibility of \mathbf{Q}) proves useful in the following sections where we characterise the behaviour of Markov models. The main focus of this appendix is the following result.

Proposition 1 *Let \mathbf{Q} satisfy conditions (A.18)–(A.21). Then, if \mathbf{Q} is irreducible, there is exactly one equilibrium point of the system*

$$\frac{d}{dt}\mathbf{x} = \mathbf{Q}^\top \mathbf{x}, \quad (\text{A.27})$$

satisfying the constraint $\mathbf{1}^\top \mathbf{x}(0) = 1$.

In other words, provided the underlying graph of our Markov model is connected, there exists a unique stable equilibrium point. That is, we have $\mathbf{x}(t) \rightarrow \mathbf{x}_\infty(V)$ for some limit, $\mathbf{x}_\infty(V)$ which does not depend on $\mathbf{x}(0)$. This is important because, when fitting our models, we assume that at when $t = 0$ the system is at equilibrium, and the uniqueness of this equilibrium allows us to set the initial condition of a given model.

We first focus on the existence of an equilibrium point for Equation (A.27). We begin by may be considering the eigenvalues of \mathbf{Q}^\top . Firstly, we can see that the condition,

$$\mathbf{Q}\mathbf{1} = \mathbf{0},$$

means precisely that the pair $(0, (\mathbf{1}))$ is an eigenpair of \mathbf{Q} . Hence, we see that,

$$\mathbf{0}^\top = (\mathbf{Q}\mathbf{1})^\top \quad (\text{A.28})$$

$$= \mathbf{1}^\top \mathbf{Q}^\top, \quad (\text{A.29})$$

and so $\mathbf{1}$ is a left-eigenvector of \mathbf{Q}^\top with eigenvalue 0. This implies that there is a (right) eigenpair of \mathbf{Q}^\top of the form $(0, (\mathbf{v}))$. That is, $\mathbf{Q}^\top \mathbf{v} = \mathbf{0}$ and \mathbf{u} is an equilibrium point of \mathbf{Q} . Also, we have,

$$\mathbf{Q}^\top \mathbf{v} = \mathbf{0}, \quad (\text{A.30})$$

and so, if $\mathbf{x} = \mathbf{v}$,

$$\frac{d\mathbf{x}}{dt} = \mathbf{0}. \quad (\text{A.31})$$

However, this equilibrium point is not unique. We can see that for any $\alpha \in \mathbb{R}$, we also have $\mathbf{Q}^\top(\alpha\mathbf{v}) = \mathbf{0}$. The only equilibrium point on this line that satisfies Equation (A.19) is $\frac{\mathbf{u}}{\mathbf{1}^\top \mathbf{u}}$.

Next, we consider the uniqueness of this steady state by supposing that there is a distinct eigenpair of the form $(0, \mathbf{u})$. That is, we assume that \mathbf{u} and \mathbf{v} linearly dependent, but $\mathbf{Q}^\top \mathbf{u} = \mathbf{0}$. This implies the existence of a plane of equilibrium points,

$$\text{span}(\mathbf{u}, \mathbf{v}) = \{\lambda\mathbf{u} + \mu\mathbf{v} : \lambda, \mu \in \mathbb{R}\}.$$

The following proposition, which we proceed to prove, concerns the uniqueness of the steady states of Equation (A.30)

Proposition 2 *Let \mathbf{Q} satisfy conditions (A.18)–(A.21). If \mathbf{Q} is irreducible, then 0 is an eigenvalue of \mathbf{Q}^\top with geometric multiplicity 1.*

Proof: Let \mathbf{u} and \mathbf{v} be such that $\mathbf{Q}^\top \mathbf{u} = \mathbf{Q}^\top \mathbf{v} = \mathbf{0}$. Then \mathbf{u} is also an eigenvector of \mathbf{Q} with corresponding eigenvalue, 0. Then for any linear combination $\lambda_1 \mathbf{v} + \lambda_2 \mathbf{u}$, with $\lambda_1, \lambda_2 \in \mathbb{R}$, we have $\mathbf{Q}^\top (\mu \mathbf{v} + \nu \mathbf{u}) = \mu \mathbf{Q}^\top \mathbf{v} + \nu \mathbf{Q}^\top \mathbf{u} = \mathbf{0}$. Then,

$$\mathbf{w} = \mathbf{v} - \frac{v_1}{u_1} \mathbf{u}, \quad (\text{A.32})$$

is a vector satisfying $\mathbf{Q}\mathbf{w} = \mathbf{0}$ and $w_1 = 0$. Now,

$$\mathbf{Q}^\top \mathbf{w} = \mathbf{Q}^\top \left(\mathbf{v} - \frac{v_1}{u_1} \mathbf{Q}^\top \mathbf{u} \right), \quad (\text{A.33})$$

$$= \mathbf{Q}^\top \mathbf{v} - \frac{v_1}{u_1} \mathbf{Q}^\top \mathbf{u}, \quad (\text{A.34})$$

$$= \mathbf{0}. \quad (\text{A.35})$$

Hence,

$$\sum_{j=1}^N Q_{j,1} w_j = 0, \quad (\text{A.36})$$

and so, since the first component of \mathbf{w} is zero, that is, $w_1 = 0$, we have,

$$\sum_{j=2}^N Q_{j,1} w_j = 0. \quad (\text{A.37})$$

By assumption, the underlying graph is strongly connected, and so, there must be a simple walk between l and state 1, that is, (s_1, s_2, \dots, s_n) where $s_1 = l$ and $s_n = 1$, where each edge corresponds to a positive transition rate (Bondy and Murty, 2008). Therefore, we have $Q_{s_m, s_{m+1}} > 0$ for each $1 \leq m \leq n-1$. Hence, $w_1 = w_{s_n} = 0$ implies that $w_{s_{n-1}} = 0$, and $w_{s_m} = 0$ implies that $w_{s_{m-1}} = 0$ for $m > 1$. It follows by induction $w_{s_1} = w_l = 0$. Hence, for any k , we have $w_k = 0$ that is, $\mathbf{w} = \mathbf{0}$. In other words, \mathbf{v} and \mathbf{u} are linearly independent. This proves Proposition 2.

Whilst Proposition 2 proves the uniqueness of a steady state, it provides little information about the properties of this steady state. It does not, for example, guarantee the stability of this steady state. Moreover, the interpretation of our models' state variables as "state occupancies" is only valid if we have $v_i > 0$ for any i , and Proposition 2 provides no such guarantee. The following result concerns the properties of the non-zero eigenvalues. These results allow us to assess the stability of the equilibria described above. For example, if \mathbf{Q} has an eigenvalue, $\lambda > 0$, then the system is unstable.

Proposition 3 *Let \mathbf{Q} be a matrix satisfying conditions (A.18–A.21). Then, for any eigenvalue of \mathbf{Q} we have $\text{Re}(\lambda) \leq 0$.*

Our proof of this proposition relies on a result regarding the loci of eigenvalues of square matrices due to Gershgorin: the *Gershgorin circle theorem*, the proof of which lies beyond the scope of this thesis.

Theorem 1 (Gershgorin Circle theorem, Horn and Johnson (2013)) *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ then each eigenvalue of \mathbf{A} , λ , there is an i such that,*

$$\lambda \in \{x : |x - A_{i,i}| \leq r_i\},$$

where $r_i = \sum_{j \neq i} |A_{i,j}|$.

We defer the proof of this theorem to Horn and Johnson (2013).

Proof of Proposition 3: by Conditions (A.18)–(A.21), we have

$$|Q_{i,i}| = \left| \sum_{j \neq i}^N Q_{i,j} \right| = \sum_{j \neq i} |Q_{i,j}|, \quad (\text{A.38})$$

because $Q_{i,j} \geq 0$ for all i, j with $i \neq j$. In other words, \mathbf{Q} is (weakly) *diagonally dominant*. We may then directly apply the Gershgorin Circle Theorem to see that the eigenvalues must satisfy $\text{Re}(\lambda) \leq 0$ because our conservation condition tells us that $Q_{i,i} = -r_i$. Then, because the eigenvalues of \mathbf{Q} and \mathbf{Q}^\top are identical, all non-zero eigenvalues of \mathbf{Q}^\top satisfy $\text{Re}(\lambda) \leq 0$, also. This proves Proposition 3.

We demonstrate these eigenvalue properties for two cases, the first, a Markov model where \mathbf{Q} is irreducible, and the second, an example of a Markov model with a reducible \mathbf{Q} , for which there is no unique equilibrium.

A.2.1 Example I: A connected Markov model

Consider the three-state Markov model with transition rate matrix,

$$\mathbf{Q} = \begin{pmatrix} -k_1 - k_6 & k_1 & k_6 \\ k_2 & -k_2 - k_3 & k_3 \\ k_5 & k_4 & -k_4 - k_5 \end{pmatrix}. \quad (\text{A.39})$$

The corresponding weighted digraph is shown in Figure A.1. Firstly, \mathbf{Q} is irreducible, there is no permutation matrix \mathbf{P} such that \mathbf{PQP}^{-1} is in *triangular block form* (Berman and Plemmons, 1979). This can be clearly seen from the fact that the corresponding digraph is strongly connected. For this reason, we know that there is exactly one eigenvector of \mathbf{Q}^\top corresponding to a vanishing eigenvalue. We proceed to show this by direct computation.




Figure A.1: A three-state Markov model. This graph is connected, so the corresponding transition rate matrix must be irreducible, and the corresponding system has a unique equilibrium point (under the constraint $\mathbf{1}^\top \mathbf{x}(0) = 1$).

The characteristic polynomial of \mathbf{Q}^\top is

$$\begin{aligned} \det(\mathbf{Q}^\top - \lambda \mathbf{I}) &= \lambda \left(-k_1 k_3 - k_1 k_4 - k_2 k_4 - k_1 k_5 - k_2 k_5 - k_3 k_5 - k_2 k_6 - k_3 k_6 - k_4 k_6 \right. \\ &\quad \left. - k_1 \lambda - k_2 \lambda - k_3 \lambda - k_4 \lambda - k_5 \lambda - k_6 \lambda - \lambda^2 \right). \end{aligned} \quad (\text{A.40})$$

Hence, 0 is an eigenvalue of multiplicity 1. Meaning that there is an eigenvector \mathbf{v} satisfying $\mathbf{Q}^\top \mathbf{v} = \mathbf{0}$ which is the unique in the sense that any other vector, \mathbf{u} , such that $\mathbf{Q}\mathbf{u} = \mathbf{0}$, may be expressed as a scalar multiple of \mathbf{v} , that is, $\mathbf{u} = \alpha \mathbf{v}$. Therefore, if we include the constraint that $\mathbf{1}^\top \mathbf{x} = 1$, the governing equation (Equation (A.13)) has exactly one equilibrium point.

A.2.2 Example II: A disconnected Markov model

Consider the four-state Markov model,

$$\mathbf{Q} = \begin{pmatrix} -k_1 & k_1 & 0 & 0 \\ k_2 & -k_2 & 0 & 0 \\ 0 & 0 & -k_3 & k_3 \\ 0 & 0 & k_4 & -k_4 \end{pmatrix}. \quad (\text{A.41})$$

We can see from the corresponding graph (Figure A.2) that this Markov model is disconnected. Alternatively, we can see that \mathbf{Q} is reducible because it is already in block upper triangular form.

Again, we may compute the eigenvectors and eigenvalues of \mathbf{Q}^\top directly. For this example, the characteristic polynomial is

$$\det(\mathbf{Q}^\top - \lambda \mathbf{I}) = k_1 \lambda^2 (\lambda + k_3 + k_4) + \lambda^3 (\lambda + k_2 + k_3 + k_4) + \lambda^2 (k_2 k_3 + k_2 k_4).$$




Figure A.2: A four-state Markov model with a disconnected graph. Since the graph is disconnected, the corresponding transition-rate matrix is reducible. The two connected components (A and B, and C and D) may be seen wholly independent processes, each with a unique stable equilibrium. Whilst such models do not meet satisfy our Markov model definition, the connected components may be combined to form a Markov model as discussed in Chapter 2.

Hence, 0 is an eigenvalue of \mathbf{Q}^\top with algebraic multiplicity 2. By considering the linear system

$$\mathbf{Q}^\top \mathbf{w} = 0, \quad (\text{A.42})$$

we find two linearly independent eigenvectors (corresponding to the vanishing eigenvalue):

$$\mathbf{v} = \left(\frac{k_1}{k_2}, -1, 0, 0 \right), \quad \mathbf{u} = \left(0, 0, \frac{k_3}{k_4}, -1 \right). \quad (\text{A.43})$$

We may decompose this Markov model into its two connected components, and write

$$\frac{d\mathbf{y}}{dt} = \mathbf{R}^\top \mathbf{y}, \quad (\text{A.44})$$

$$\frac{d\mathbf{z}}{dt} = \mathbf{S}^\top \mathbf{z}, \quad (\text{A.45})$$

where,

$$\mathbf{x} = \begin{pmatrix} \mathbf{y} \\ \mathbf{z} \end{pmatrix}$$

and,

$$\mathbf{R} = \begin{pmatrix} -k_1 & k_1 \\ k_2 & -k_2 \end{pmatrix}, \quad (\text{A.46})$$

$$(\text{A.47})$$

$$\mathbf{S} = \begin{pmatrix} -k_3 & k_4 \\ k_3 & -k_4 \end{pmatrix}. \quad (\text{A.48})$$

We may then identify the eigenvectors of \mathbf{Q}^\top , \mathbf{u} and \mathbf{v} with equilibrium points of the Markov models described by transition-rate matrices \mathbf{R} and \mathbf{S} :

$$\left(\frac{k_1}{k_2}, 1\right)^\top \quad (\text{A.49})$$

for the former; and, for the latter we have,

$$\left(\frac{k_3}{k_4}, 1\right)^\top. \quad (\text{A.50})$$

In this case, the equilibrium that the system reaches will depend on the initial conditions, for example if we choose,

$$\mathbf{x}(0) = \left(\frac{1}{2}, \frac{1}{2}, 0, 0\right)^\top,$$

we will find that $x_3(t) = x_4(t) = 0$ for all $t > 0$, and if we choose,

$$\mathbf{x} = \left(0, 0, \frac{1}{2}, \frac{1}{2}\right)^\top,$$

then $x_1(t) = x_2(t) = 0$ for all $t > 0$.

Provided the Markov model has only one connected component, we know that it has an irreducible transition-rate matrix, which (using Proposition 3) means that it has no positive eigenvalues and a zero-eigenvalue with multiplicity 1. However, can also apply the same result to the connected components of a disconnected Markov model, and such models may be treated as the combination of two separate Markov models. It is easy to check for connectedness either visually, or by using an algorithm such as Depth-First Search (Gibbons, 1985).

A.2.3 Application of the Perron-Frobenius Theorem

The following result provides some more properties of our models' steady states.

Theorem 2 (Perron-Frobenius Theorem (Berman and Plemmons, 1979)) *Let $\mathbf{A} = (A_{i,j})$ be an $N \times N$ matrix with strictly positive entries. Then there exists an eigenpair (r, \mathbf{v}) such that:*

- $v_i > 0$ for all i ;
- $r \in \mathbb{R}$;
- and $r > \text{Re}(\lambda)$ for all eigenvalues $\lambda \neq r$.

Such an eigenvalue, r , may be referred to as the *leading* eigenvalue. This is an incomplete statement of the full Perron-Frobenius Theorem, which is provided in Berman and Plemmons (1979).

Now, for some irreducible \mathbf{Q} taken from a connected Markov model, \mathbf{Q}^\top is not a positive matrix (its diagonal elements are negative) and so the Perron-Frobenius Theorem is not directly applicable. Instead, we let $\mathbf{A} = \mathbf{Q}^\top + k\mathbf{I}$ where $k = \max_i \{-\mathbf{Q}_{i,i}\} + 1$. Then, $A_{i,j} > 0$ for all i, j and so, we may apply the Perron-Frobenius Theorem to \mathbf{A} .

Now, let (λ, \mathbf{v}) be an eigenpair of \mathbf{A} . Then,

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{v}, \quad (\text{A.51})$$

and so,

$$(\mathbf{Q}^\top + (k+1)\mathbf{I})\mathbf{v} = \lambda\mathbf{v}, \quad (\text{A.52})$$

$$\mathbf{Q}^\top\mathbf{v} + (k+1)\mathbf{v} = \lambda\mathbf{v}, \quad (\text{A.53})$$

$$\mathbf{Q}^\top\mathbf{v} = (\lambda - k - 1)\mathbf{v}. \quad (\text{A.54})$$

That is, if λ is an eigenvalue of \mathbf{A} , then $\lambda - k - 1$ is an eigenvalue of \mathbf{Q}^\top , and the corresponding eigenvectors are equal. In other words, there is a one-to-one mapping from eigenpairs of \mathbf{A} onto eigenpairs of \mathbf{Q}^\top given by,

$$(\lambda, \mathbf{v}) \mapsto (\lambda - k - 1, \mathbf{v}). \quad (\text{A.55})$$

We know from Proposition 3 that all eigenvalues of \mathbf{Q}^\top satisfy $\text{Re}(\lambda) \leq 0$ and from Proposition 2, that 0 is one of these eigenvalues and so, $1 + k$ must be the leading eigenvalue of \mathbf{A} . Hence, applying the Perron-Frobenius theorem to \mathbf{A} tells us that for each eigenvalue of \mathbf{A} , λ , satisfies

$$|\lambda| \leq k + 1, \quad (\text{A.56})$$

and so, each eigenvalue of \mathbf{Q}^\top satisfies,

$$|\lambda - (k + 1)| \leq k + 1. \quad (\text{A.57})$$

These loci are shown in Figure A.3.

The Perron-Frobenius theorem also tells us that the leading-eigenvalue, r corresponds to some eigenvector \mathbf{v}^+ with $v_i^+ > 0$ for all i . Since 0 is the value of \mathbf{Q}^\top with maximal real part, the eigenpair of \mathbf{A} , (r, \mathbf{v}^+) must correspond to an eigenpair of \mathbf{Q}^\top , $(0, \mathbf{v}^+)$ (using Equation (A.55)). This proves Proposition 1.




Figure A.3: The Perron-Frobenius theorem provides information about the eigendecomposition of \mathbf{A} . In particular, there exists a leading, positive eigenvalue, greater in absolute value than all other eigenvalues, which corresponds to a wholly positive eigenvector. The map $\mathbf{A} \mapsto \mathbf{Q}^\top$ induces a simple translation of eigenvalues and tells us that each eigenvalue of \mathbf{Q}^\top has non-positive real part.

We conclude by demonstrating the results introduced in this section by computing the eigenvalues of \mathbf{Q}^\top for a three-state model. We set transition rates as,

$$(k_1, k_2, k_3, k_4, k_5, k_6) = (1, 2, 3, 4, 5, 10). \quad (\text{A.58})$$

The resulting transition rate matrix is,

$$\mathbf{Q} = \begin{pmatrix} -11 & 1 & 10 \\ 2 & -5 & 3 \\ 5 & 4 & -9 \end{pmatrix}. \quad (\text{A.59})$$

To one decimal place, the eigenpairs of \mathbf{Q}^\top are,

$$\lambda_1 = -17.1, \quad \mathbf{v}_1 = (-0.7, 0.3, 1)^\top, \quad (\text{A.60})$$

$$\lambda_2 = -7.9, \quad \mathbf{v}_2 = (0.6, -1.6, 1)^\top, \quad (\text{A.61})$$

$$\text{and } \lambda_3 = 0, \quad \mathbf{v}_3 = (0.59, 0.88, 1). \quad (\text{A.62})$$

Also, we have $k = \max_{i,j} \{-Q_{i,j}\} = 11$ so,

$$\mathbf{A} = \mathbf{Q}^\top + (k+1)\mathbf{I}, \quad (\text{A.63})$$

$$= \mathbf{Q}^\top + 12\mathbf{I}, \quad (\text{A.64})$$

$$= \begin{pmatrix} 1 & 2 & 5 \\ 1 & 7 & 4 \\ 10 & 3 & 3 \end{pmatrix}. \quad (\text{A.65})$$

Then, to one decimal place, the eigenpairs of \mathbf{A} are,

$$\lambda_1 = -5.1, \quad \mathbf{v}_2 = (-0.7, -0.3, 1)^\top, \quad (\text{A.66})$$

$$\lambda_2 = 4.1, \quad \mathbf{v}_3 = (0.6, -1.6, 1), \quad (\text{A.67})$$

$$\lambda_3 = 12, \quad \mathbf{v}_1 = (0.59, 0.88, 11)^\top. \quad (\text{A.68})$$

As we would expect, the eigenvectors of \mathbf{Q}^\top and \mathbf{A} are equal, but the eigenvalues are different. Additionally, no eigenvalues of \mathbf{Q}^\top have positive real part, and the vanishing eigenvalue corresponds to the eigenvector $(33, 49, 53)^\top$ which has only positive components. We have hence demonstrated that the linear system of ODEs defined by Equation (A.30) has a unique, global steady equilibrium.

A.3 Single-channel models and macroscopic currents

Markov models have also been used to model the behaviour of single channels (Colquhoun and Hawkes, 1981). Here, the conformation of the channel is modelled stochastically by a *continuous-time Markov chain* (CTMC) on a finite state space. Such a CTMC may be seen as the combination of a collection of distributions of *waiting-times* for each state, and an embedded random jump process, a discrete-time Markov chain with probability matrix, \mathbf{P} . The holding-time distributions for each state i follow an exponential distribution, $\mathbf{w}_i \sim \exp(\mu_i)$ for some $\mu_i > 0$. Then, the *infinitesimal generator matrix* or “Q-matrix” is defined as,

$$\mathbf{Q} = (\mathbf{I} - \mathbf{P})\boldsymbol{\mu}^{-1}. \quad (\text{A.69})$$

where $\boldsymbol{\mu} = (\mu_i)_{i=1}^N$ (Colquhoun and Hawkes, 1995). This is a matrix which satisfies conditions (A.18–A.21), which we defined for our Markov model. Hence, for any fixed voltage, V_m , every Markov model corresponds to a CTMC with infinitesimal generator matrix $\mathbf{Q}(V_m)$ (and vice versa). Let $\mathbf{p}(t) = (p_i(t))_{i=1}^N$ be a row vector, where $p_i(t)$ is the probability that the Markov chain is in state i at time t . Then, the *forward*

Chapman-Kolmogorov equation (Kelly, 2011) tells us that,

$$\frac{d\mathbf{p}^\top(t)}{dt} = \mathbf{p}(t)\mathbf{Q}, \quad (\text{A.70})$$

which is simply the system described by Equation (A.17) written with row vectors instead of column vectors. Row vectors are commonly used to describe Markov chains (both discrete and continuous), and, subject to the condition that $\mathbf{p}\mathbf{1} = 1$ (that is, the elements of \mathbf{p} sum to 1), the row-vector \mathbf{p} is known as a *probability vector* (Kelly, 2011). Note that the definition of \mathbf{p} as a probability vector is analogous to the conservation constraint in our Markov models Equation (A.19). Furthermore, any \mathbf{Q} -matrix satisfies conditions (A.18)–(A.21) and so, there is a correspondence between Markov models and CTMCs (Kelly, 2011).

To demonstrate this correspondence, we consider n_{channels} IID channels such that the proportion of channels in each state at each time is a random process, $\{X_i(t)\}_{t \geq 0}$. For each t , $n_c X(t)$ has a multinomial distribution with probability vector $\mathbf{p}(t)$. In particular, the expected number of channels in state i is

$$\mathbb{E}[X_i(t)] = \frac{1}{n_{\text{channels}}} \mathbb{E}[n_{\text{channels}} p_i(t)]. \quad (\text{A.71})$$

and the corresponding variance is

$$\mathbb{V}[X_i(t)] = \frac{1}{n_{\text{channels}}^2} \mathbb{V}[n_{\text{channels}} X(t)] \quad (\text{A.72})$$

$$= \frac{1}{n_{\text{channels}}} p_i(t)(1 - p_i(t)) \rightarrow 0, \quad (\text{A.73})$$

as $n_{\text{channels}} \rightarrow \infty$. Hence, the deterministic model (characterised by the governing equation in Equation (A.13)) becomes a reasonable approximation to this stochastic process when there is a large number of channels. Furthermore, since $p_i(1 - p_i)$ is maximal when $p_i = 0.5$, we know that,

$$\mathbb{V}[X_i(t)] \leq \frac{1}{4n_{\text{channels}}}, \quad (\text{A.74})$$

and so if $n_{\text{channels}} = 10,000$, for example, we have

$$\mathbb{V}[X_i(t)] \leq 2.5 \times 10^{-5}. \quad (\text{A.75})$$

A practical demonstration of the increasing accuracy of the mass-action approximation is provided in Figure 2.7 (Chapter 2), from which we see that there is little difference between the ODE solution and the stochastic process when n_{channels} is sufficiently large.

Without making strong assumptions regarding the independence and identical behaviour of individual channels, ODE-based Markov models are commonly used in cardiac electrophysiology (Beattie et al., 2018; Tomek et al., 2019), and our results in Chapter 4 show that such models can recapitulate experimental data, and produce accurate predictions. Though we use ODE-based models, results arising from the study of CTMCs provide useful intuition. For example, it is known that for a given CTMC, *equilibrium distribution* exists if and only if the CTMC is irreducible (Kelly, 2011) that is, if the matrix \mathbf{Q} is irreducible (Berman and Plemmons, 1979). This result is equivalent to Proposition 2. However, in the probability literature, this result may instead be obtained by considering the properties of the embedded random jump process, and not by examining the eigenvalues of \mathbf{Q} (Whitt, 2006).

A.4 Reversible Markov models

In addition to satisfying conditions (A.18–A.21), *reversibility* is another property that is often required of Markov models. This property is sometimes referred to as *microscopic reversibility*, and is equivalent to the existence of *detailed balance* and *thermodynamic equilibrium* (Colquhoun et al., 2004).

Definition 1 Consider a Markov model which satisfies Conditions (A.18–A.21) with transition-rate matrix, \mathbf{Q} such that \mathbf{Q}^\top is irreducible. Then by Proposition 1, a unique equilibrium point, \mathbf{x}_∞ , exists. We say that this Markov model is reversible if it satisfies,

$$Q_{i,j}x_\infty^{(i)} = Q_{j,i}x_\infty^{(j)}. \quad (\text{A.76})$$

for all i and j .

Since reversibility is a property wholly concerned with properties of \mathbf{Q} , we extend this definition to the transition-rate matrix and call \mathbf{Q} reversible if the corresponding Markov model is reversible. The following result shows that there may be qualitative differences between reversible Markov models and non-reversible Markov models.

Proposition 4 Let \mathbf{Q} be a matrix satisfying conditions (A.18)–(A.21). Additionally, suppose \mathbf{Q} is reversible. Then, the eigenvalues and eigenvectors of \mathbf{Q}^\top are real.

Proof: Assuming that our model is reversible, we may construct the diagonal matrix \mathbf{D} with $D_{i,i} = \sqrt{x_\infty^{(i)}}$ for all i . \mathbf{D} is a diagonal matrix with no zeroes on its diagonal, so it is

invertible. Then,

$$(\mathbf{DQ}^\top \mathbf{D}^{-1})_{j,i} = \sqrt{x_\infty^{(j)}} Q_{j,i} \frac{1}{\sqrt{x_\infty^{(i)}}}, \quad (\text{A.77})$$

$$= \sqrt{x_\infty^{(j)}} Q_{i,j} \frac{x_\infty^{(i)}}{x_\infty^{(j)}} \frac{1}{\sqrt{x_\infty^{(i)}}}, \quad (\text{using Equation (A.76)}) \quad (\text{A.78})$$

$$= \frac{x_\infty^{(i)}}{\sqrt{x_\infty^{(i)}}} Q_{i,j} \frac{\sqrt{x_\infty^{(j)}}}{x_\infty^{(j)}}, \quad (\text{A.79})$$

$$= \sqrt{x_\infty^{(i)}} Q_{i,j} \frac{1}{\sqrt{x_\infty^{(j)}}}, \quad (\text{A.80})$$

$$= (\mathbf{DQ}^\top \mathbf{D}^{-1})_{i,j}. \quad (\text{A.81})$$

Which means that \mathbf{DQD}^{-1} is a symmetric matrix. In other words, \mathbf{Q}^\top is *similar* to a symmetric matrix, and so the characteristic polynomials of $\mathbf{DQ}^\top \mathbf{D}^{-1}$ is,

$$\det(\mathbf{DQ}^\top \mathbf{D}^{-1} - \lambda \mathbf{I}) = \det(\mathbf{D}) \det(\mathbf{Q}^\top \mathbf{D}^{-1} - \lambda \mathbf{D}^{-1} \mathbf{I}), \quad (\text{A.82})$$

$$= \det(\mathbf{D}) (\mathbf{Q}^\top - \lambda \mathbf{D}^{-1} \mathbf{D}) \det(\mathbf{D}^{-1}), \quad (\text{A.83})$$

$$= \det(\mathbf{D}) \det(\mathbf{D})^{-1} \det(\mathbf{Q}^\top - \lambda \mathbf{I}), \quad (\text{A.84})$$

$$= \det(\mathbf{Q}^\top - \lambda \mathbf{I}), \quad (\text{A.85})$$

which is exactly the characteristic polynomial of \mathbf{Q}^\top . Therefore, the eigenvalues of \mathbf{Q}^\top are exactly the eigenvalues of $\mathbf{DQ}^\top \mathbf{D}^{-1}$. Then, since $\mathbf{DQ}^\top \mathbf{D}^{-1}$ is symmetric, it has real eigenvalues, and so, \mathbf{Q}^\top has real eigenvalues, also (Meyer, 2000).

This provides some insight into the implications of reversibility on the dynamics of Markov models. Eigenvalues with $\text{Im}(\lambda) \neq 0$ describe oscillatory behaviour (similar to a damped oscillator when $\text{Re}(\lambda) \leq 0$), but such eigenvalues are not possible when \mathbf{Q} is reversible. Complex eigenvalues can arise from non-reversible \mathbf{Q} are demonstrated in the following example.

Example III: A non-reversible Markov model

We return to the three-state model shown in Figure A.1. The equilibrium points of this system satisfy,

$$\mathbf{x}_\infty = \alpha \begin{pmatrix} \frac{k_2k_4 + k_2k_5 + k_3k_5}{k_1k_3 + k_2k_6 + k_3k_6} \\ \frac{k_1k_4 + k_1k_5k_4k_6}{k_1k_3 + k_2k_6 + k_3k_6} \\ 1 \end{pmatrix}, \quad (\text{A.86})$$

for $\alpha \in \mathbb{R}$. We can see that this Markov model is not reversible in general because the flux between states 1 and 2 at steady state is,

$$x_\infty^{(1)} Q_{1,2} = k_1 \frac{k_2k_4 + k_2k_5 + k_3k_5}{k_1k_3 + k_2k_6 + k_3k_6}, \quad (\text{A.87})$$

whilst, the flux in the opposite direction is,

$$x_\infty^{(2)} Q_{2,1} = k_2 \frac{k_1k_4 + k_1k_5 + k_4k_6}{k_1k_3 + k_2k_6 + k_3k_6}. \quad (\text{A.88})$$

Hence, this Markov model is reversible if and only if,

$$k_1k_3k_5 = k_2k_4k_6. \quad (\text{A.89})$$

In Section A.4, we will see that a similar condition may be used, in general, to check the reversibility of any Markov model.

When our model is not reversible, or equivalently, $k_1k_3k_5 \neq k_2k_4k_6$, the transition-rate matrix, \mathbf{Q} , may have complex eigenvalues. For example, this is the case when we choose,

$$(k_1, k_2, k_3, k_4, k_5, k_6) = (1, 100, 1, 100, 1, 100). \quad (\text{A.90})$$

In this case the eigenpairs of \mathbf{Q} are (to one decimal place),

$$\begin{aligned} \lambda_1 &= -16.5 + 7.8i, & v_1 &= (-0.5 - 0.9i, -0.5 + 0.9i, 1.0, 1.0), \\ \lambda_2 &= -16.5 + 7.8i, & v_2 &= (0.5 + 0.9i, -0.5 - 0.9i, 1.0, 1.0), \\ \lambda_3 &= 0.0, & v_3 &= (1.0, 1.0, 1.0). \end{aligned}$$

This demonstrates that if non-reversible models are considered, computational methods relying on the eigendecomposition of \mathbf{Q}^\top must account for non-real eigenvalues. This is not the case when reversible models are used, in which case, the eigenvalues and eigenvectors are guaranteed to be real.

In the following section, we introduce a useful method for ensuring the reversibility of Markov models. As we will see, our three-state Markov model is reversible if and only if,

$$k_6 = \frac{k_1 k_3 k_5}{k_2 k_4}. \quad (\text{A.91})$$

Moreover, we can find similar conditions which guarantee the reversibility of any other Markov model.

Kolmogorov's Criterion

It is possible to ensure reversibility by checking that the detailed-balance equation (Equation (A.76)) holds. However, this requires us to compute the steady state. Kolmogorov's criterion provides an alternative method which is far more convenient.

Theorem 3 (Kolmogorov's Criterion (Kelly, 2011)) *A Markov model with transition-rate matrix $\mathbf{Q}(V)$, with corresponding graph, G , in which transition rates are positive for all V , or zero for all V is reversible if and only for every cycle in G if the products of transition rates in each direction are equal. In other words, Equation (A.76) holds if and only if,*

$$\prod_i^{C_n} Q_{s_i, s_{i+1}} = \prod_i^{C_n} Q_{s_{i+1}, s_i} \quad (\text{A.92})$$

for every cycle of G , $C = (s_1, \dots, s_{C_{n+1}})$ where $s_{C_{n+1}} = s_1$.

This result is due to the study of continuous-time Markov chains (CTMCs) which are closely related to the system described by Equation (A.13). The relationship between the ODE formulation used in this section, and CTMCs is discussed in Appendix A.3. Nevertheless, Kolmogorov's result can be directly applied to our ODE formulation. For example, regarding Example I, we saw that reversibility was equivalent to the condition $k_1 k_3 k_5 = k_2 k_4 k_6$ which is an equality between the product of transition rates around the only cycle in each direction (as illustrated in Figure A.1).

This result is proved in Kelly (2011) for continuous time Markov chains (CTMCs), but an identical proof can be applied to ODEs of the form of Equation (A.30), because, for such models, \mathbf{Q} may be realised as the *infinitesimal generator* matrix of a CTMC. As explained by Colquhoun et al., we need only check that each of a subset of cycles satisfies Equation (A.92)—the *cycle basis* of G (Yang et al., 2006) .

A.5 Simplifying the governing equation

We conclude this appendix to explain how our Markov models' conservation constraint (Equation (A.19)) can be used to simplify our models' governing equation (Equation (A.30)). As discussed in (Fink and Noble, 2009), this operation reduces the number

of state variables by 1, which reduces the amount of computation required to integrate our governing equation. This ultimately allows us to fit our models and produce predictions more efficiently. If we solve the governing equation directly with numerical methods, this constraint may be violated somewhat due to the accumulation of numerical errors. On the other hand, when numerical methods are applied to the simplified form of the system of ODEs, this is no longer a concern because the conservation constraint is satisfied implicitly.

Such an simplification of the system of ODEs is commonly applied to particular Markov models of ion-channel currents (Fink and Noble, 2009; Beattie et al., 2018). Colquhoun and Hawkes (1995) provide a description of this procedure, which may be applied to any Markov model (that is, a system of ODEs satisfying Equations (A.18–A.21)). In this section, we provide an explicit derivation of this method. We then show that many similar simplifications are possible and suggest future work for the selection of more suitable schemes for simulation of Markov models.

First, to derive the scheme suggested by Colquhoun and Hawkes (1995), we seek a system of ODEs,

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{A}\tilde{\mathbf{x}} + \mathbf{b}, \quad (\text{A.93})$$

where $\tilde{\mathbf{x}}$ is the first $N - 1$ components of \mathbf{x} , \mathbf{A} is a $(N - 1) \times (N - 1)$ matrix, and \mathbf{b} is a vector of size $(N - 1)$. Then, by setting $x_n = 1 - \sum_{j=1}^{N-1} x_j$, our constraint, $\mathbf{1}^\top \mathbf{x}$ is always satisfied.

Let \mathbf{A} be a $(N - 1) \times (N - 1)$ matrix, \mathbf{b} be a $N - 1$ vector, and c be a scalar such that,

$$\mathbf{Q}(V) = \begin{pmatrix} \tilde{\mathbf{Q}} & \mathbf{a} \\ \mathbf{b}^\top & c \end{pmatrix}, \quad (\text{A.94})$$

where $\tilde{\mathbf{Q}}$ is an $(N - 1) \times (N - 1)$ matrix containing the elements in the upper-left rows and columns of \mathbf{Q} . Then, by multiplying both sides of Equation (A.30) by the $(N - 1) \times N$ matrix,

$$\begin{pmatrix} \mathbf{I}_{N-1} & \mathbf{0} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ \vdots & & & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & 0 \end{pmatrix}, \quad (\text{A.95})$$

which is simply an identity matrix with a column of zeros appended to its right side, we have,

$$\begin{pmatrix} \mathbf{I}_{N-1} & \mathbf{0} \end{pmatrix} \frac{d}{dt} \mathbf{x} = \begin{pmatrix} \tilde{\mathbf{Q}}^\top & \mathbf{I}_{N-1} \mathbf{b} \end{pmatrix} \mathbf{x}, \quad (\text{A.96})$$

$$= \begin{pmatrix} \tilde{\mathbf{Q}}^\top \tilde{\mathbf{x}} \\ 0 \end{pmatrix} + x_N \begin{pmatrix} \mathbf{b} \\ 0 \end{pmatrix}. \quad (\text{A.97})$$

Hence, substituting $x_N = 1 - \tilde{\mathbf{x}}^\top \mathbf{1}$ gives,

$$\frac{d}{dt} \tilde{\mathbf{x}} = \left(\tilde{\mathbf{Q}} - \mathbf{b} \mathbf{1}^\top \right) \tilde{\mathbf{x}} + \mathbf{b}. \quad (\text{A.98})$$

After solving this system, it is simple to use the relation $x_N = 1 - \mathbf{1}^\top \tilde{\mathbf{x}}$, to recover \mathbf{x} from $\tilde{\mathbf{x}}$.

However, Equation (A.98) is but one of many ways in which we use the conservation constraint to simplify Equation (A.30). Suppose \mathbf{T} is a non-singular matrix such that each element of the bottom row is 1. Then, we have the block matrix,

$$\mathbf{T} = \begin{pmatrix} \mathbf{U} \\ \mathbf{1}^\top \end{pmatrix}. \quad (\text{A.99})$$

Now, we have $\mathbf{s} = \mathbf{T}\mathbf{x}$, and so, from Equation (A.13), we have,

$$\frac{d\mathbf{s}}{dt} = \frac{d\mathbf{T}\mathbf{x}}{dt} \quad (\text{A.100})$$

$$= \mathbf{T} \frac{d\mathbf{x}}{dt} \quad (\text{A.101})$$

$$= \mathbf{T} \mathbf{Q}^\top \mathbf{x} \quad (\text{A.102})$$

$$= \mathbf{T} \mathbf{Q}^\top \mathbf{T}^{-1} \mathbf{s}, \quad (\text{A.103})$$

which is possible because we chose \mathbf{U} to be non-singular. If we let $\mathbf{W} = \mathbf{T} \mathbf{Q}^\top \mathbf{T}^{-1}$, we may write,

$$\frac{d\mathbf{s}}{dt} = \mathbf{W} \mathbf{s}. \quad (\text{A.104})$$

Since each element of the bottom row of \mathbf{T} is 1, we have, for the bottom component of \mathbf{S} , $s_N = \mathbf{1}^\top \mathbf{x} = 1$, and so, s_N is constant. Hence, we may simplify Equation (A.30), by writing,

$$\frac{d\tilde{\mathbf{s}}}{dt} = \mathbf{C} \tilde{\mathbf{s}} + \mathbf{d}, \quad (\text{A.105})$$

where $\tilde{\mathbf{s}}$ is the first $N - 1$ rows of \mathbf{s} , \mathbf{C} is the first $N - 1$ rows and columns of \mathbf{W} , and $\mathbf{d} = y_n \mathbf{w}_N$, where \mathbf{w}_N denotes the rightmost column of \mathbf{W} .

There are many choices for such a matrix, \mathbf{U} . Perhaps the most obvious choice is the matrix the $N - 1$ -dimensional identity matrix, \mathbf{I}_{N-1} , with a zero vector, $\mathbf{0}$, appended

to its right-hand side. In this case we have,

$$\mathbf{T} = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & 0 \\ 1 & 1 & 1 & \dots & 1 & 1 \end{pmatrix}, \quad (\text{A.106})$$

from which obtain Equation (A.98). Instead of eliminating the N^{th} state, \mathbf{x}_N , we could eliminate any other state. For example, to eliminate the first state, x_1 instead of x_N , we may define,

$$\mathbf{T} = \begin{pmatrix} 0 & 0 & 0 & \dots & 0 & 1 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & 0 \\ 1 & 1 & 1 & \dots & 1 & 1 \end{pmatrix}. \quad (\text{A.107})$$

In this way, it may be possible to choose \mathbf{T} such that the resultant system of ODEs has some desirable properties. For example, we may wish to ensure that the condition number of \mathbf{W} is as large as possible—reducing the error in matrix exponential computations (Moler and Van Loan, 2003). One simple approach would be to iterate through each of our N states, removing each state individually, and checking the resulting condition numbers.

We suggest that such a method may allow more efficient (and more accurate) simulation of Markov models of ion channel currents, as well as other similar equations. A more greedy approach, considering a larger set of possible choices of \mathbf{T} , may also be plausible. Further work could explore how the choice of \mathbf{T} affects the numerical error in numerical solutions of Equation (A.30) across various solution methods. However, it should be noted that \mathbf{Q}^\top depends on the transmembrane potential, and our model parameters, and so, the suitability of a given \mathbf{T} may vary over the course of the experiment. Hence, an easily computable heuristic for a suitable choice of \mathbf{T} may prove more useful than a harder-to-compute, more optimal method.

Appendix B

Supplementary material for Chapter 3

B.1 I_{Kr} model equations

Beattie model

In full, the system of ODEs is,

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} -k_1 - k_3 & 0 & k_4 & k_2 \\ 0 & -k_2 - k_4 & k_1 & k_3 \\ k_3 & k_2 & -k_1 - k_4 & 0 \\ k_1 & k_4 & 0 & -k_2 - k_3 \end{bmatrix} \mathbf{x}, \quad (\text{B.1})$$

where

$$k_1 = p_1 e^{p_2 V}, \quad (\text{B.2})$$

$$k_2 = p_3 e^{-p_4 V}, \quad (\text{B.3})$$

$$k_3 = p_5 e^{p_6 V}, \quad (\text{B.4})$$

$$k_4 = p_7 e^{-p_8 V}. \quad (\text{B.5})$$

Hence, the corresponding parameter set is,

$$\boldsymbol{\theta} = \left[p_1, \dots, p_8, g \right]^T, \quad (\text{B.6})$$

and,

$$\mathbf{x} = \left[C, I, IC, O \right]^T. \quad (\text{B.7})$$

Note that this model may also be written in terms of two Hodgkin-Huxley style gating variables (Clerx et al., 2019a).

Wang model

We may write this model's governing system of ODEs as

$$\frac{d\mathbf{x}}{dt} = \begin{bmatrix} -\alpha_{a0} & \beta_{a0} & 0 & 0 & 0 \\ \alpha_{a0} & -\beta_{a0} - k_f & k_b & 0 & 0 \\ 0 & k_f & -k_b - \alpha_{a1} & \beta_{a1} & 0 \\ 0 & 0 & \alpha_{a1} & -\beta_{a1} - \alpha_1 & \beta_1 \\ 0 & 0 & 0 & \alpha_1 & -\beta_1 \end{bmatrix} \mathbf{x}, \quad (\text{B.8})$$

where

$$a_1 = q_1 e^{q_2 V}, \quad (\text{B.9})$$

$$a_{a0} = q_3 e^{q_4 V}, \quad (\text{B.10})$$

$$a_{a1} = q_5 e^{q_6 V}, \quad (\text{B.11})$$

$$b_{a1} = q_7 e^{-q_8 V}, \quad (\text{B.12})$$

$$b_1 = q_9 e^{-q_{10} V}, \quad (\text{B.13})$$

$$b_{a0} = q_{11} e^{-q_{12} V}. \quad (\text{B.14})$$

$$(\text{B.15})$$

The corresponding parameter set is,

$$\boldsymbol{\theta} = \left[q_1, \dots, q_{12}, k_f, k_b \right]^T, \quad (\text{B.16})$$

and

$$\mathbf{x} = \left[C_1, C_2, C_3, O, I \right]^T. \quad (\text{B.17})$$

The default parameter values for both models are presented in Table B.1.

B.2 Further Case I results

The predictive accuracy of the model used in Section 3.4.2 (under the validation protocol, d_0), trained using each protocol, for a range of values of λ is shown in Fig. B.1. Here, we see that as there is more model discrepancy (when λ moves away from 1) our predictions become less accurate.

The parameter estimates obtained in Sections 3.4.1 and 3.4.2 are summarised in Table B.2, Table B.4 and Table B.5, respectively. Here, we can see that when the model is misspecified, the standard deviation in our estimates (across fits to different samples of our DGP) is small compared to the differences between estimates obtained from

Wang model			Beattie Model		
Parameter	Value	Units	Parameter	Value	Units
g	1.52×10^{-1}	μS	g	1.52×10^{-1}	μS
k_b	3.68×10^{-2}	ms^{-1}	p_1	2.26×10^{-4}	ms^{-1}
k_f	2.38×10^{-2}	ms^{-1}	p_2	6.99×10^{-2}	mV^{-1}
q_1	9.08×10^{-2}	ms^{-1}	p_3	3.45×10^{-5}	ms^{-1}
q_2	2.34×10^{-2}	mV^{-1}	p_4	5.46×10^{-2}	mV^{-1}
q_3	2.23×10^{-2}	ms^{-1}	p_5	8.73×10^{-2}	ms^{-1}
q_4	1.18×10^{-2}	mV^{-1}	p_6	8.91×10^{-3}	mV^{-1}
q_5	1.37×10^{-2}	ms^{-1}	p_7	5.15×10^{-3}	ms^{-1}
q_6	3.82×10^{-2}	mV^{-1}	p_8	3.16×10^{-2}	mV^{-1}
q_7	6.89×10^{-5}	ms^{-1}			
q_8	4.18×10^{-2}	mV^{-1}			
q_9	6.50×10^{-3}	ms^{-1}			
q_{10}	3.27×10^{-2}	mV^{-1}			
q_{11}	4.70×10^{-2}	mV^{-1}			
q_{12}	6.31×10^{-2}	ms^{-1}			

Table B.1: The default parameter sets we use for the Wang et al. (1997) and Beattie et al. (2018) models, with parameter sets taken from the corresponding publications. In the case of the Beattie et al. (2018) model, these parameters values are taken from Cell #5. For the Wang et al. (1997) model, we take those parameters obtained using a 2 mM concentration of potassium chloride. The same maximal conductance (g) a representative value for a CHO hERG expression system, is used for both models.




Figure B.1: Case I: predictive accuracy (under our validation protocol, d_0) decreases as $\lambda \rightarrow 1$. For 17 values of λ ($\frac{1}{4} \leq \lambda \leq 4$), the predictive error (averaged over repeats) is shown for each fitting protocol (d_1 – d_5)

different protocols—the choice of fitting protocol is less important when there is no model discrepancy.

Table B.2 details the distribution of each parameter estimate (under repeated samples of the DGP) for each protocol as λ varies (as described in Section 3.4.1). Whereas, Table B.3 shows how our spread-of-prediction intervals change under different values of λ . Here, we can see that each parameter estimate, as well as \mathcal{B} itself, show little variability under repeated samples of the DGP.

B.3 Further Case II results

Table B.4 and Table B.5 summarise the parameter estimates obtained in Section 3.4.2 using the Beattie et al. (2018) model and Wang et al. (1997) model, respectively. Here, the Wang et al. (1997) model was chosen as the DGP and so, the Wang et al. (1997) model is an example of a correctly specified model, whereas the Beattie et al. (2018) model is a discrepant model. This is reflected by the parameter estimates which show that when the Wang et al. (1997) model is fitted to the data, we obtain similar parameter estimates from each protocol, whereas our parameter estimates for the Beattie et al. (2018) model are dependent on the protocol used for training.

Table B.6 shows the behaviour of our spread-of-prediction intervals (Equation (3.19)) for both the Wang et al. (1997) Model and the Beattie et al. (2018) Model, as described in Section 3.4.2. Here, we see that the average width of this interval (averaged over the length of the protocol) is much larger for the Beattie et al. (2018) Model (a discrepant model) when compared with the Wang et al. (1997) Model (the same model used for data generation).

λ		d_1	d_2	d_3	d_4	d_5
0.25	p_1	1.8E-04±2E-07	2.4E-04±7E-07	9.0E-05±1E-06	1.9E-04±4E-07	4.7E-05±8E-08
	p_2	9.1E-02±5E-05	9.7E-02±9E-05	1.0E-01±2E-04	9.3E-02±8E-05	1.0E-01±4E-05
	p_3	3.0E-05±2E-08	1.9E-05±3E-08	1.2E-05±4E-08	1.7E-05±4E-08	9.3E-06±2E-08
	p_4	4.8E-02±6E-06	5.2E-02±2E-05	5.5E-02±3E-05	5.3E-02±2E-05	5.8E-02±2E-05
	p_5	5.0E-02±5E-05	5.5E-02±6E-05	5.0E-02±5E-05	7.9E-02±3E-04	5.1E-02±8E-05
	p_6	1.5E-02±2E-05	1.0E-02±4E-05	1.0E-02±2E-05	1.8E-02±5E-05	1.3E-02±3E-05
	p_7	1.5E-02±2E-05	1.3E-02±2E-05	9.8E-03±2E-05	2.3E-02±7E-05	1.0E-02±5E-05
	p_8	4.4E-02±2E-05	4.6E-02±5E-05	5.1E-02±4E-05	3.6E-02±4E-05	5.8E-02±1E-04
0.50	p_1	2.2E-04±2E-07	2.5E-04±6E-07	2.0E-04±6E-07	2.1E-04±4E-07	6.9E-05±2E-07
	p_2	7.6E-02±3E-05	8.5E-02±8E-05	8.0E-02±6E-05	8.8E-02±7E-05	9.7E-02±6E-05
	p_3	3.7E-05±2E-08	3.4E-05±5E-08	3.6E-05±6E-08	3.6E-05±6E-08	2.5E-05±3E-08
	p_4	5.0E-02±4E-06	5.1E-02±1E-05	5.0E-02±2E-05	5.1E-02±1E-05	5.3E-02±9E-06
	p_5	7.5E-02±7E-05	7.5E-02±6E-05	6.8E-02±6E-05	8.4E-02±2E-04	7.7E-02±7E-05
	p_6	9.4E-03±1E-05	9.0E-03±3E-05	9.7E-03±2E-05	1.1E-02±3E-05	9.9E-03±1E-05
	p_7	9.3E-03±1E-05	8.2E-03±7E-06	6.8E-03±8E-06	1.0E-02±3E-05	8.5E-03±1E-05
	p_8	3.7E-02±1E-05	3.8E-02±2E-05	4.0E-02±1E-05	3.5E-02±3E-05	3.8E-02±2E-05
1.00	p_1	2.3E-04±2E-07	2.3E-04±4E-07	2.3E-04±6E-07	2.3E-04±3E-07	2.3E-04±6E-07
	p_2	7.0E-02±3E-05	7.0E-02±7E-05	7.0E-02±5E-05	7.0E-02±5E-05	7.0E-02±6E-05
	p_3	3.4E-05±1E-08	3.4E-05±4E-08	3.4E-05±5E-08	3.4E-05±5E-08	3.4E-05±2E-08
	p_4	5.5E-02±5E-06	5.5E-02±1E-05	5.5E-02±1E-05	5.5E-02±1E-05	5.5E-02±5E-06
	p_5	8.7E-02±7E-05	8.7E-02±8E-05	8.7E-02±6E-05	8.7E-02±3E-04	8.7E-02±6E-05
	p_6	8.9E-03±9E-06	8.9E-03±2E-05	8.9E-03±1E-05	8.9E-03±3E-05	8.9E-03±1E-05
	p_7	5.2E-03±6E-06	5.2E-03±4E-06	5.2E-03±4E-06	5.1E-03±1E-05	5.2E-03±4E-06
	p_8	3.2E-02±9E-06	3.2E-02±2E-05	3.2E-02±8E-06	3.2E-02±3E-05	3.2E-02±1E-05
2.00	p_1	2.3E-04±2E-07	2.1E-04±3E-07	2.3E-04±7E-07	2.1E-04±3E-07	3.4E-04±6E-07
	p_2	6.8E-02±3E-05	6.1E-02±7E-05	7.0E-02±6E-05	5.8E-02±5E-05	6.0E-02±4E-05
	p_3	3.0E-05±1E-08	3.0E-05±4E-08	2.5E-05±4E-08	2.8E-05±5E-08	3.4E-05±2E-08
	p_4	5.9E-02±5E-06	5.8E-02±1E-05	6.0E-02±1E-05	5.8E-02±2E-05	5.8E-02±6E-06
	p_5	9.0E-02±7E-05	8.9E-02±8E-05	9.5E-02±7E-05	9.5E-02±3E-04	9.0E-02±5E-05
	p_6	1.0E-02±6E-06	9.6E-03±2E-05	9.0E-03±1E-05	8.6E-03±4E-05	9.7E-03±8E-06
	p_7	2.6E-03±3E-06	2.9E-03±2E-06	2.9E-03±2E-06	2.8E-03±1E-05	2.7E-03±2E-06
	p_8	2.7E-02±8E-06	2.7E-02±2E-05	2.6E-02±6E-06	2.8E-02±3E-05	2.7E-02±8E-06
4.00	p_1	2.3E-04±2E-07	2.0E-04±3E-07	2.3E-04±8E-07	2.1E-04±2E-07	3.8E-04±5E-07
	p_2	6.6E-02±3E-05	5.7E-02±6E-05	7.0E-02±6E-05	5.5E-02±4E-05	5.7E-02±3E-05
	p_3	2.9E-05±1E-08	2.8E-05±4E-08	2.3E-05±4E-08	2.6E-05±5E-08	3.3E-05±2E-08
	p_4	6.1E-02±5E-06	6.0E-02±1E-05	6.2E-02±1E-05	6.0E-02±2E-05	5.9E-02±6E-06
	p_5	9.4E-02±7E-05	9.1E-02±9E-05	9.7E-02±7E-05	9.9E-02±4E-04	9.6E-02±5E-05
	p_6	1.1E-02±5E-06	9.7E-03±2E-05	9.1E-03±1E-05	8.6E-03±3E-05	1.0E-02±6E-06
	p_7	1.4E-03±1E-06	1.6E-03±1E-06	1.5E-03±1E-06	1.5E-03±5E-06	1.5E-03±1E-06
	p_8	2.5E-02±8E-06	2.5E-02±1E-05	2.4E-02±5E-06	2.6E-02±3E-05	2.5E-02±7E-06

Table B.2: The mean and standard deviation of the least-squares parameter estimates (Equation (2.32)) used in Case I, where the maximal conductance, is misspecified by scaling it with λ . These values were obtained from each fitting protocol (d_1 – d_5) for multiple repeats of synthetically generated data.

λ	Mean interval width (nA)	DGP in interval (%)	Midpoint RMSE (nA)
0.25	7.4E-02 \pm 1.1E-04	3.4E01 \pm 2.3E-02	1.6E-01 \pm 8.9E-05
0.30	6.6E-02 \pm 6.5E-05	3.7E01 \pm 5.0E-02	1.4E-01 \pm 8.9E-05
0.35	6.1E-02 \pm 5.8E-05	4.2E01 \pm 4.3E-02	1.3E-01 \pm 9.1E-05
0.42	5.6E-02 \pm 6.4E-05	5.0E01 \pm 4.2E-02	1.1E-01 \pm 9.2E-05
0.50	4.9E-02 \pm 7.1E-05	5.1E01 \pm 5.1E-02	8.5E-02 \pm 9.4E-05
0.59	4.0E-02 \pm 8.4E-05	5.3E01 \pm 6.7E-02	6.5E-02 \pm 9.7E-05
0.71	2.8E-02 \pm 9.1E-05	5.5E01 \pm 8.8E-02	4.7E-02 \pm 9.0E-05
0.84	1.4E-02 \pm 9.3E-05	5.5E01 \pm 1.7E-01	3.4E-02 \pm 8.4E-05
1.00	2.1E-04 \pm 3.7E-05	9.4E01 \pm 9.2E-00	3.0E-02 \pm 6.7E-05
1.19	1.2E-02 \pm 8.9E-05	5.6E01 \pm 6.9E-01	3.3E-02 \pm 5.1E-05
1.41	2.1E-02 \pm 8.5E-05	5.7E01 \pm 6.7E-01	3.7E-02 \pm 5.3E-05
1.68	2.8E-02 \pm 8.5E-05	5.7E01 \pm 1.7E-00	4.2E-02 \pm 6.2E-05
2.00	3.3E-02 \pm 8.6E-05	5.6E01 \pm 1.7E-00	4.7E-02 \pm 6.9E-05
2.38	3.6E-02 \pm 8.6E-05	5.5E01 \pm 1.4E-00	5.1E-02 \pm 7.3E-05
2.83	3.9E-02 \pm 8.4E-05	5.4E01 \pm 1.4E-00	5.4E-02 \pm 7.5E-05
3.36	4.1E-02 \pm 8.6E-05	5.4E01 \pm 1.7E-01	5.7E-02 \pm 7.6E-05
4.00	4.3E-02 \pm 8.4E-05	5.3E01 \pm 1.5E-01	5.9E-02 \pm 7.8E-05

Table B.3: A summary of showing how the spread-of-predictions interval (Equation (3.19)) behaves under Case I. Here we show: the mean width of the interval (averaged over each observation time); the proportion of observations for which the underlying DGP (minus noise) lies within the interval; the RMSE between the data and the midpoint prediction (Equation (3.20)). By considering ten randomly sampled datasets (each containing a repeat each protocol d_1 – d_5), we show the mean and standard deviation of these values.

	d_1	d_2	d_3	d_4	d_5
g	1.5E-01 \pm 3E-05	1.5E-01 \pm 5E-05	1.6E-01 \pm 3E-05	1.5E-01 \pm 5E-05	1.6E-01 \pm 2E-05
p_1	1.6E-03 \pm 7E-07	1.6E-03 \pm 8E-07	1.7E-03 \pm 1E-06	1.7E-03 \pm 1E-06	2.0E-03 \pm 5E-07
p_2	7.3E-02 \pm 2E-05	7.9E-02 \pm 2E-05	3.7E-02 \pm 2E-05	8.4E-02 \pm 4E-05	5.4E-02 \pm 2E-05
p_3	1.9E-05 \pm 2E-08	2.2E-05 \pm 2E-08	4.3E-05 \pm 6E-08	2.2E-05 \pm 2E-08	3.0E-05 \pm 2E-08
p_4	5.2E-02 \pm 7E-06	5.1E-02 \pm 1E-05	4.6E-02 \pm 1E-05	5.1E-02 \pm 1E-05	4.9E-02 \pm 6E-06
p_5	1.1E-01 \pm 7E-05	9.6E-02 \pm 2E-05	9.3E-02 \pm 2E-05	1.2E-01 \pm 2E-04	9.6E-02 \pm 6E-05
p_6	2.3E-02 \pm 5E-06	2.4E-02 \pm 8E-06	2.2E-02 \pm 5E-06	2.7E-02 \pm 2E-05	2.3E-02 \pm 7E-06
p_7	8.9E-03 \pm 5E-06	7.1E-03 \pm 3E-06	7.4E-03 \pm 2E-06	8.9E-03 \pm 1E-05	6.8E-03 \pm 5E-06
p_8	2.9E-02 \pm 8E-06	3.1E-02 \pm 1E-05	3.0E-02 \pm 7E-06	2.9E-02 \pm 2E-05	3.1E-02 \pm 7E-06

Table B.4: The parameter estimates obtained for Case II (Section 3.4.1) when using the Beattie et al. (2018) Model to fit data generated by the Wang et al. (1997) Model.

	d_1	d_2	d_3	d_4	d_5
g	1.5E-01±3E-05	1.5E-01±4E-05	1.5E-01±2E-05	1.5E-01±4E-05	1.5E-01±2E-05
k_b	3.7E-02±4E-04	3.6E-02±1E-03	3.7E-02±2E-04	3.6E-02±2E-03	3.7E-02±3E-04
k_f	2.4E-02±9E-05	2.4E-02±4E-04	2.4E-02±9E-05	2.4E-02±6E-04	2.4E-02±7E-05
q_1	9.1E-02±5E-05	9.1E-02±7E-05	9.1E-02±2E-05	9.1E-02±1E-04	9.1E-02±6E-05
q_2	2.3E-02±6E-06	2.3E-02±1E-05	2.3E-02±5E-06	2.3E-02±1E-05	2.3E-02±8E-06
q_3	2.2E-02±5E-04	2.3E-02±7E-04	2.2E-02±3E-04	2.3E-02±8E-04	2.2E-02±3E-04
q_4	1.2E-02±4E-04	1.1E-02±7E-04	1.2E-02±2E-04	1.1E-02±8E-04	1.2E-02±2E-04
q_5	1.4E-02±2E-04	1.4E-02±4E-04	1.4E-02±1E-04	1.4E-02±7E-04	1.4E-02±1E-04
q_6	3.8E-02±2E-04	3.8E-02±3E-04	3.8E-02±1E-04	3.8E-02±6E-04	3.8E-02±2E-04
q_7	6.9E-05±7E-08	6.9E-05±3E-07	6.9E-05±1E-07	6.9E-05±4E-07	6.9E-05±1E-07
q_8	4.2E-02±8E-06	4.2E-02±4E-05	4.2E-02±1E-05	4.2E-02±5E-05	4.2E-02±1E-05
q_9	6.5E-03±4E-06	6.5E-03±5E-06	6.5E-03±2E-06	6.5E-03±8E-06	6.5E-03±7E-06
q_{10}	3.3E-02±7E-06	3.3E-02±2E-05	3.3E-02±9E-06	3.3E-02±2E-05	3.3E-02±8E-06
q_{11}	4.7E-02±1E-03	4.9E-02±3E-03	4.7E-02±4E-04	4.9E-02±3E-03	4.7E-02±6E-04
q_{12}	6.3E-02±4E-04	6.3E-02±6E-04	6.3E-02±4E-04	6.3E-02±6E-04	6.3E-02±2E-04

Table B.5: The parameter estimates obtained for Case II when using the Wang et al. (1997) Model to fit synthetic data generated under the same model.

Model	Mean interval width (nA)	DGP in interval (%)	Midpoint RMSE (nA)
Beattie	7.5E-02±9E-05	34±0.07	1.1E-01±8E-05
Wang	7.0E-04±2E-04	87±20	3.0E-02±2E-05

Table B.6: A summary showing how the spread-of-predictions interval (Equation (3.19)) behaves for both the Beattie et al. (2018) Model (a discrepant model) and the Wang et al. (1997) Model (a correctly specified model). The columns show: the mean width of the interval (averaged over each observation time); the proportion of observations for which the underlying DGP (minus noise) lies within the interval; the RMSE between the data and the midpoint prediction (Equation (3.20)). By considering ten randomly sampled datasets (each containing a repeat for each protocol d_1 – d_5), we show the mean and standard deviation of these values.

Appendix C

Effects of capacitance spike removal

In an appendix of Beattie et al. (2018)[B2.1], Beattie et al. explain that small amounts of data were removed after discontinuities in the command voltage (V_{cmd}) to minimise the presence of capacitive currents. It has been suggested that spikes in current arise from the charging and discharging of the cell membrane (Beattie et al., 2018), which acts like a capacitor (this capacitive effect is included in the Hodgkin-Huxley model as shown in Figure 2.5). As these currents are not part of I_{Kr} , we seek to minimise their presence, and continue using our simple, ideal patch-clamp Markov model. Note that in the following chapter, we model these effects explicitly.

If more data is removed, the variability of our parameter estimates will increase. We chose to remove $R = 5\text{ms}$ of data as a compromise between minimising model discrepancy, and maximising the information provided by our experimental designs. In Figure C.1, we use synthetic data to show how different values of R affect the predictive accuracy of our models. In Figures C.2 and C.3, we can see how increasing R causes an increase in parameter uncertainty under our idea patch-clamp assumptions (see Chapters 4 and 5). When choosing how much data to remove after each discontinuity, that is, choosing the removal duration, R , we aim to strike a balance between the uncertainty of our parameter estimates and the bias introduced by including data polluted by capacitive currents.

In Beattie et al. (2018), the Beattie et al. (2018) model is presented as both as the product of two Hodgkin-Huxley style gating variables, and as a four-state Markov model. Using the former representation, we can see that for constant, V_{m} , the behaviour of the model is governed by two exponentials,

$$a(t) = a_0 + (a_{\infty}(V_{\text{m}}) - a_0) \exp \left\{ \frac{t}{\tau_{\text{a}}(V_{\text{m}})} \right\} \quad (\text{C.1})$$

$$\text{and } r(t) = r_0 + (r_{\infty}(V_{\text{m}}) - r_0) \exp \left\{ \frac{t}{\tau_{\text{r}}(V_{\text{m}})} \right\} \quad (\text{C.2})$$




Figure C.1: The removal of capacitive spikes under the *staircase* protocol. Panel **a**: the highlighted regions show $R = 80$ ms of data after each discontinuity in the command voltage. Panel **b**: a synthetic I_{Kr} recording (grey) is generated by adding IID Gaussian noise to the model output. Model output with $R = 80$ ms of data removed after each spike is shown in blue.

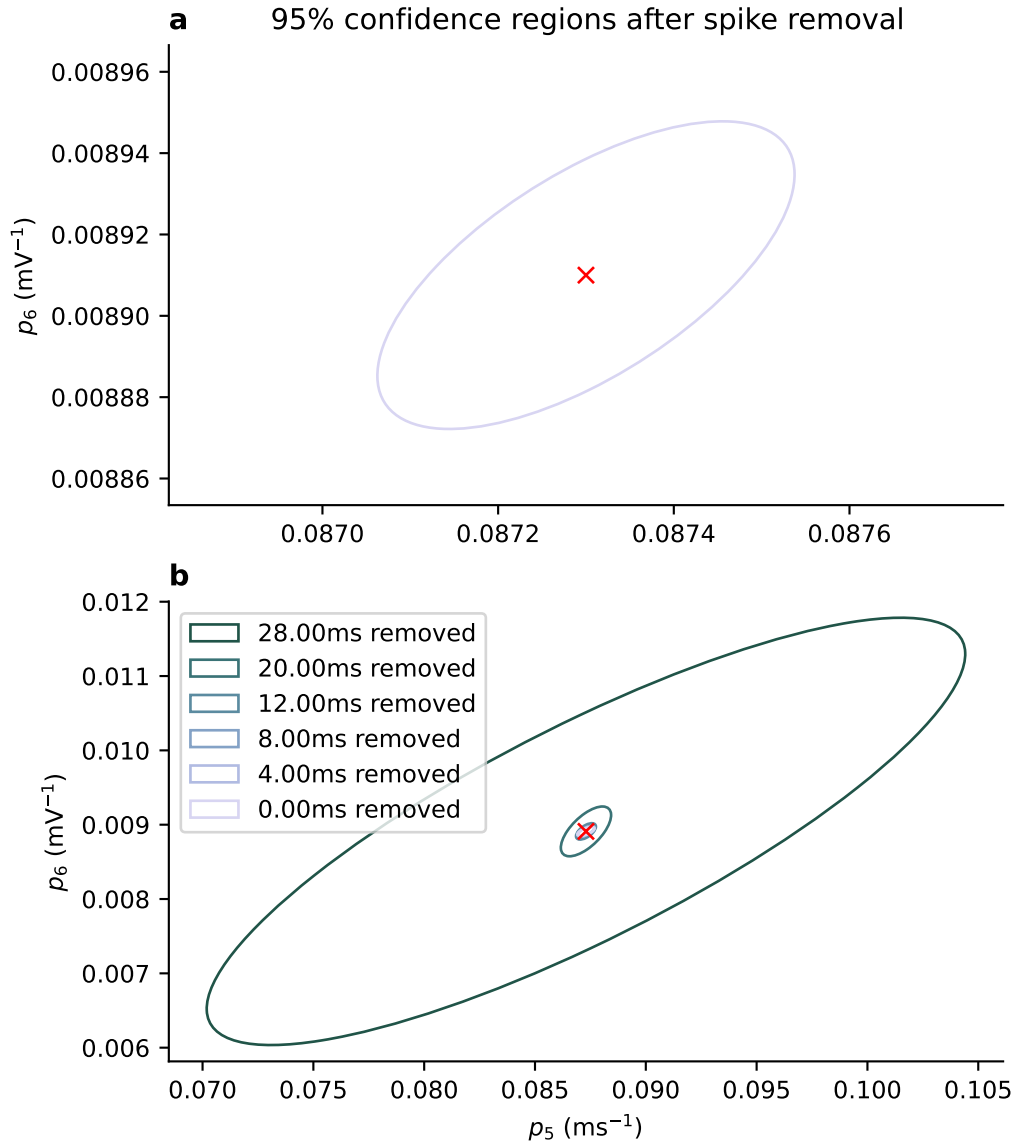


Figure C.2: Increasing the time removed after each voltage discontinuity, R , impacts the accuracy of our p_5 and p_6 parameter estimates, which determine the rate of inactivation in the Beattie et al. (2018) Model. Here, we see that parameter uncertainty increases greatly as more data is removed, that is, as R increases. Panel **a** shows only the 95% confidence region resulting from no data removal (that is, $R = 0$). Panel **b** shows similar confidence regions for various values of R .

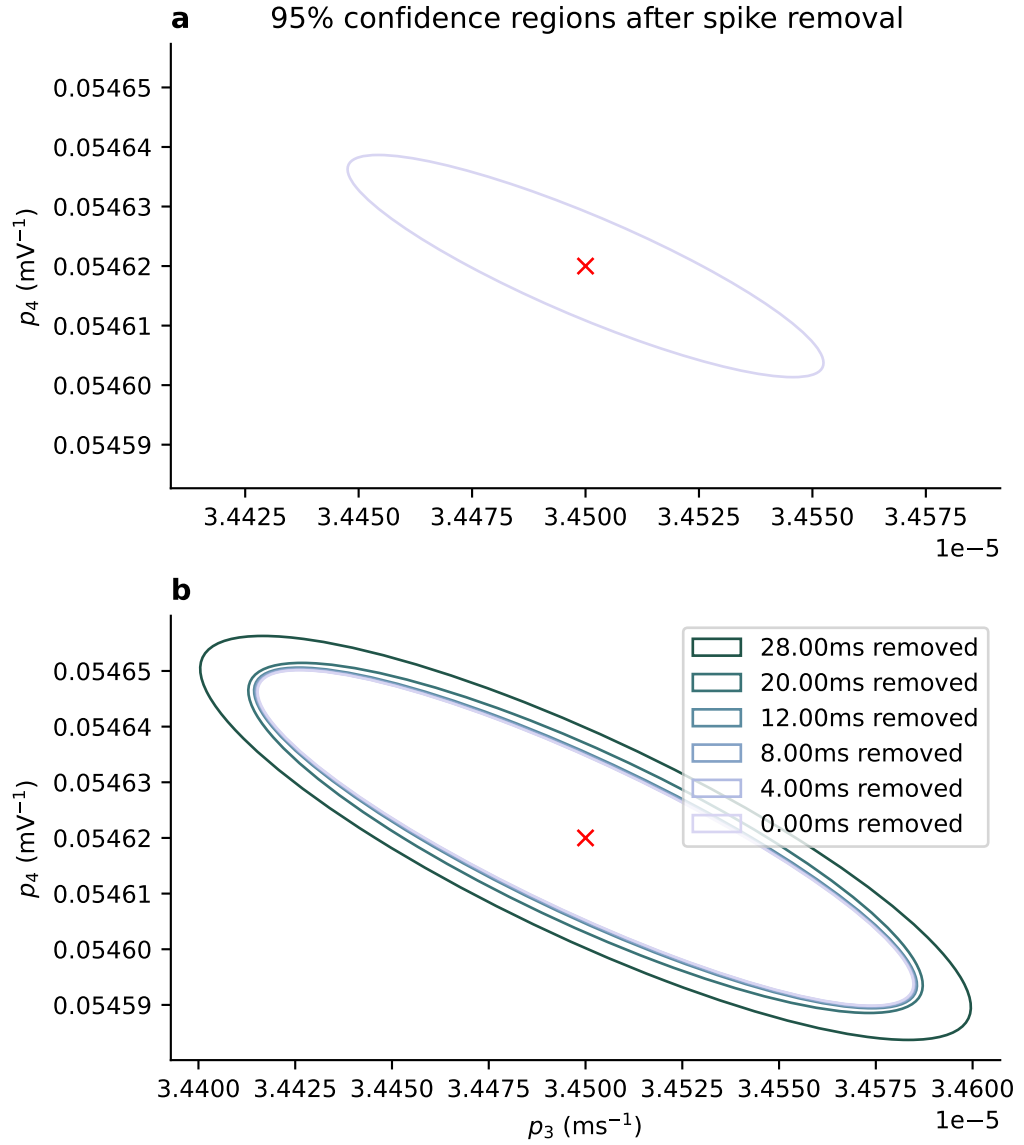


Figure C.3: The size of our parameter estimates' confidence regions grows as R increases. Approximate 95% confidence regions are shown using a linear approximation. Marginal confidence regions are shown for model parameters p_1 and p_2 , which determine the rate of activation ($k_2 = p_3 \exp\{-p_4 V_m\}$). Here, we see little increase in parameter uncertainty, even when $R \geq 20$ ms. Panel **a** shows only the 95% confidence region resulting from no data removal (that is, $R = 0$). Panel **b** shows similar confidence regions for various values of R .

where a_0 and h_0 are our initial conditions, and where $a_\infty(V_m)$ and $r_\infty(V_m)$, $\tau_a(V_m)$ and $\tau_r(V_m)$ are voltage dependent functions arising from our model.

In the Beattie et al. (2018) model, our transition rates are each of the form

$$k = A \exp\{\pm b V_m\} \quad (\text{C.3})$$

for $A, b \in \mathbb{R}$ with $A > 0$ and $b \neq 0$. In fact, we have (Clerx et al., 2019a),

$$\tau_a = \frac{1}{k_1 + k_2}, \quad (\text{C.4})$$

$$\tau_r = \frac{1}{k_3 + k_4}, \quad (\text{C.5})$$

$$a_\infty = \frac{k_1}{k_1 + k_2}, \quad (\text{C.6})$$

$$\text{and } r_\infty = \frac{k_4}{k_3 + k_4}. \quad (\text{C.7})$$

In practice, we find that τ_a (that is, the timescale relating to the $\mathbf{C} \rightarrow \mathbf{O}$ and $\mathbf{O} \rightarrow \mathbf{C}$ transitions) is much greater than the τ_r which describes the timescale of the *inactivation* and *recovery-from-inactivation* processes (the $\mathbf{O} \rightarrow \mathbf{I}$ and $\mathbf{I} \rightarrow \mathbf{O}$ transitions, respectively). This is illustrated in Figure C.4, where these variables, τ_a and τ_r , are shown for a range of voltages using the Beattie et al. (2018) model with cell #5's parameters.

If we were to only observe the model at its steady state, there would be insufficient data to identify our parameters. Therefore, we may expect that our estimates of p_5, p_6, p_7 and p_8 (which correspond to the inactivation/recovery-from-inactivation process), to be more affected by increasing the removal time, R . This is shown to be the case in Figures C.2 and C.3 where we show the uncertainty of our parameter estimates. From these figures, it seems increasing R has a much more noticeable effect on parameters pertaining r rather than a , likely because of the vast differences in the timescales that the model's two processes operate over, as shown in Figure C.4. This increase in parameter estimate uncertainty is also reflected in the A- and D-optimality measures (introduced in Section 2.3), as shown in Figure C.5. Here, the sensitivities matrix is normalised as described in Fink and Noble (2009), that is, A- and D-optimality statistics are computed using S' where, $S'_{i,j} = \theta_j S_{i,j}$ for $1 \leq i \leq n$ and $1 \leq j \leq N_{\text{parameters}}$ where n is the number of observations and $N_{\text{parameters}}$ is the number of model parameters.

From these results, we can see that, under the ideal patch-clamp assumption, decreasing R increases the accuracy of our parameter estimates. We choose retain Beattie et al.'s choice of $R = 5\text{ms}$, as this avoids excessive accuracy/information loss, yet is large enough that the impact of capacitive currents is greatly reduced. The choice of $R = 5\text{ms}$ is a somewhat subjective choice, and could be reevaluated for future protocols, models and parameter sets.




Figure C.4: Characteristic timescales of the Beattie et al. (2018) model across a physiologically relevant range of voltages. These values were computed using the Cell #5 parameter set as presented in Beattie et al. (2018). Here, τ_1 is the timescale corresponding to the slower, activation/deactivation process, and τ_2 corresponds to the much faster inactivation/recovery-from-inactivation process.




Figure C.5: A- and D-optimality utility functions computed for the Beattie et al. (2018) Model under the application of the *staircase* protocol, where increasingly more data is removed after each discontinuity in voltage. The increase in these two quantities corresponds to an increase in uncertainty in estimates of model parameters. Here, R quantifies the duration during which observations are discarded after each discontinuity in V_{cmd} .

Appendix D

Model details

Here, we provide further mathematical details of the I_{Kr} models used in Chapter 4, and later in Chapter 5—those shown in Figure 4.17. Two of these models were already used in Chapter 3, and described in Appendix B.1. The remaining two models are the three-state, Closed-Open-Inactive (C-O-I) model, and the Kemp et al. (2021) model. However, unlike in Chapter 3, there is no synthetic data generation, and so we do not require any reference parameter sets prior to model fitting. Accordingly, we only list the mathematical equations that arise from these two model structures.

All the model structures explored in chapters 3, 4 and 5 may be expressed using a transition-rate matrix $\mathbf{Q}(V_m)$, the transpose of which appears in the governing ODE (Equation (2.18)). Similarly, these models all utilise similar observation functions (as in Equation (3.28)). Here the maximal conductance, g , appears as an additional model parameter that we estimate during model fitting. Both of these models were implemented using a graph-based, symbolic manipulation approach using the Markov-builder package (see <https://github.com/CardiacModelling/markov-builder>).

D.1 Closed-Open-Inactive model

This model consists of three states, the closed state (**C**), the open state (**O**), and the inactive state (**I**). Using this ordering of states, the voltage-dependent transition-rate matrix is,

$$\mathbf{Q}(V_m) = \begin{pmatrix} -a_m & a_m & 0 \\ b_m & -a_h - b_m & a_h \\ 0 & b_h & -b_h \end{pmatrix}, \quad (\text{D.1})$$

where a_m, a_h, b_m and b_h are voltage-dependent transition rates satisfying,

$$a_m = p_1 \exp\{p_2 V_m\}, \quad (\text{D.2})$$

$$b_m = p_3 \exp\{-p_4 V_m\} \quad (\text{D.3})$$

$$a_h = p_5 \exp\{p_2 V_m\}, \quad (\text{D.4})$$

$$\text{and, } b_h = p_7 \exp\{-p_8 V_m\}, \quad (\text{D.5})$$

and $p_i > 0$ for each model parameter, $1 \leq i \leq 8$. As with all of our other model structures, the maximal conductance, g , is an additional model parameter which does not appear in \mathbf{Q} .

D.2 Kemp model

The Kemp et al. (2021) Model is a six-state model with transition-rate matrix, with twelve transition-rate parameters, p_1, \dots, p_{12} . It is similar to the (Beattie et al., 2018) model in that it consists of two independent processes—an *activation/deactivation* process and an *inactivation/recovery-from-inactivation* process. The latter, as in the Beattie et al. (2018) Model, is governed by two voltage-dependent transition rates,

$$k_3 = p_5 \exp\{p_6 V_m\}, \quad (\text{D.6})$$

$$\text{and } k_4 = p_7 \exp\{-p_8 V_m\} \quad (\text{D.7})$$

describing the rate of inactivation and recovery from inactivation, respectively. However, where activation and deactivation are described by two transition rates, k_1 and k_2 , in the Beattie et al. (2018) Model, two additional transition rates appear in the Kemp et al. (2021) Model,

$$k_1 = p_1 \exp\{p_2 V_m\}, \quad (\text{D.8})$$

$$k_2 = p_3 \exp\{-p_4 V_m\}, \quad (\text{D.9})$$

$$k_5 = p_9 \exp\{p_{10} V_m\}, \quad (\text{D.10})$$

$$\text{and,} \quad (\text{D.11})$$

$$k_6 = p_{11} \exp\{-p_{12} V_m\}. \quad (\text{D.12})$$

Each of these transition rates appears in the model's transition-rate matrix,

$$\mathbf{Q}(V_m) = \begin{pmatrix} -k_3 - k_2 & k_3 & k_2 & 0 & 0 & 0 \\ k_4 & -k_2 - k_4 & 0 & k_2 & 0 & 0 \\ k_1 & 0 & -k_1 - k_3 - k_6 & k_3 & k_6 & 0 \\ 0 & k_1 & k_4 & -k_1 - k_4 - k_6 & 0 & k_6 \\ 0 & 0 & k_5 & 0 & k_3 - k_5 & k_3 \\ 0 & 0 & 0 & k_5 & k_4 & -k_4 - k_5 \end{pmatrix}. \quad (\text{D.13})$$

where the rows and columns correspond to, the open state (**O**) the inactive state (**I**), the closed state (**C**), the inactive-closed state (**IC**), the second closed state (**C2**) and the second inactive-closed state (**IC2**), in that order.

Appendix E

Further statistical analysis of variability in parameter estimates

In Chapters 4 and 5, we use linear statistical models to characterise our collections of parameter estimates. In particular, we seek to quantify the degree of well and protocol dependence in our parameter-estimate vectors. This quantification is achieved through computing the log-likelihood difference between a model which includes well- and protocol-dependent effects and a simplified model (that is, one that omits either well- or protocol-dependent effects).

Under the assumption of Gaussian errors, our linear statistical models (including well- and/or protocol-dependent effects) are a family of nested models, as described in Chapter 4 (Figure 4.30). Likelihood ratio tests are commonly used for model selection between such families of models (Rencher and Schaalje, 2008), and concern the relative maximal likelihoods under each model. Under the hypothesis that the simpler model is true and there are sufficient observations, the distribution of the likelihood ratio is approximately a chi-squared distribution, where the degrees of freedom in the distribution is equal to the number of additional parameters included in the more complex model. Therefore, the **LLD-W** and **LLD-D** statistics discussed in Chapters 4 and 5 quantify the significance of well- and protocol-dependent effects in our parameter estimates.

However, this approach relies on the validity of the full linear model (that is, $\mathcal{M}_{w,d}$ as introduced in Chapter 4). This model is built upon the assumption of Gaussian random errors. Quantile-quantile (QQ) plots are commonly used to assess the validity of this assumption. Here, the quantiles of residuals found between the model and data are plotted against their corresponding expected values (under the assumption that the model is correct). An example of such a plot is presented in Figure E.1 for the Beattie et al. (2018) Model (Case III). Here, under a perfectly correct model, we expect the residuals to approximately lie along the line $y = x$, where the sample quantiles largely agree with the theoretical quantiles obtained under our model assumptions. The points

in Figure E.1 show some departure from this line, where the most extreme residuals are more common in the sample distribution than the theoretical distribution, indicating that the sample distribution is heavy-tailed. Whilst this figure shows that the distribution of residuals may have heavy tails, it appears that the assumption of Gaussian distributed residuals is a reasonable one to make.

Such heavy tails may be due to certain interactions between well- and protocol-dependent effects which are not accounted for in the model, leading to large errors between the model and the data for particular pairs of wells and protocols. It is also possible that such large residuals could be caused by data pollution affecting our parameter estimates. Such an example seems to be shown in the main text (Figure 4.31) for the data obtained from the second repeat of d_1 in Well B20. Additionally, the symmetry of this plot about the line $y = x$ shows that the distribution of the residuals is mostly symmetrical (that is, not skewed), as we would expect if the residuals were normally distributed.

Similar QQ plots are also shown for the C-O-I, Kemp et al. (2021) and Wang et al. (1997) models (Case III) in Figures E.2, E.3 and E.4, respectively. Furthermore, Figures E.5, E.6, E.7 and E.8 show similar plots, for the Beattie et al. (2018), C-O-I, Kemp et al. (2021) and Wang et al. (1997) models (respectively) under Case V (that is, with artefact effects). The QQ plots involving the Wang model (Figures E.4 and E.8) show heavier tails in the distribution of residuals compared to those of the other models. Also, for the Wang et al. (1997) Model under Case V (with artefacts) the distribution of residuals for some parameters appear multimodal (see Figure E.8). However, generally speaking, each figure shows that the residuals are approximately normally distributed. Overall, these figures support the validity of the linear statistical models used in Chapters 4 and 5.




Figure E.1: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the Beattie et al. (2018) Model (Case III). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate then the plotted points will approximately lie on the grey, dashed line ($y = x$).




Figure E.2: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the C-O-I Model (Case III). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate, the plotted points will approximately lie on the grey, dashed line ($y = x$).




Figure E.3: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the Kemp et al. (2021) Model (Case III). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate, the plotted points will approximately lie on the grey, dashed line ($y = x$).




Figure E.4: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the Wang et al. (1997) Model (Case III). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate, the plotted points will approximately lie on the grey, dashed line ($y = x$).




Figure E.5: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the Beattie et al. (2018) Model (Case V, with artefacts). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate, the plotted points will approximately lie on the grey, dashed line ($y = x$).




Figure E.6: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the C-O-I Model (Case V, with artefacts). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate, the plotted points will approximately lie on the grey, dashed line ($y = x$).




Figure E.7: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the Kemp et al. (2021) Model (Case V, with artefacts). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate, the plotted points will approximately lie on the grey, dashed line ($y = x$). The relative flatness of some parameters (such as q_3 , q_5 and q_6) may suggest a degree of multimodality in the distribution residuals. Whilst generally reasonable, the assumption of normally-distributed residuals may be somewhat less suitable for these parameters.




Figure E.8: A quantile-quantile plot for the full linear statistical model fitted to parameter estimates obtained under the Wang et al. (1997) Model (Case V, with artefacts). The quantiles of the residuals found in the data are compared against theoretical quantiles under the assumption that they are normally distributed (with zero mean and some fitted variance). When the model is appropriate then the plotted points will approximately lie on the grey, dashed line ($y = x$) and where we would expect all data to lie provided the model is correct and there is a large number of data.

Appendix F

Dynamics of the artefact model

As discussed in Chapter 5, when our I_{Kr} models are combined with the artefact model, we obtain a nonlinear system of ODEs where,

$$\frac{d\mathbf{x}}{dt} = \mathbf{Q}(V_m)^\top, \quad (\text{F.1})$$

and,

$$\frac{dV_m}{dt} = \frac{V_{cmd} + V_{off} - V_m}{C_m R_{series}} - \frac{I_{out}}{C_m}, \quad (\text{F.2})$$

and, as under the ideal-patch assumptions,

$$\frac{d\mathbf{x}}{dt} = \mathbf{Q}(V_m)^\top \mathbf{x}, \quad (\text{F.3})$$

where

$$I_{out} = I_L + I_{Kr}, \quad (\text{F.4})$$

and $I_L = g_L(V_m - E_L)$, and $I_{Kr} = g_{xO}(V_m - E_L)$.

In this section, we first consider the behaviour of this system when $x_O = 0$ (that is, when there is no I_{Kr} present). This is assumed to be the case when fitting our leak-model parameters, and so, this exercise provides some insight into the behaviour of the model under the leak ramp. Similarly, by assuming a fixed x_O , we may investigate the behaviour of the artefact model under the application of the reversal ramp. Note here, that under the assumption that x_O is fixed, I_{Kr} is identical to a leak current, and so, we need not include these terms explicitly.

We then briefly show how the properties of Markov models described in Appendix A can be extended to the artefact model. In particular, we show that—provided realistic Artefact Model parameters—this system has a unique steady state. This is important for our multiprotocol approach because it allows us to set the initial conditions for each protocol—we simply initialise our ODE model with the system’s steady state.

F.1 Dynamics under ramps

During the leak ramp, we assume that $g = 0$ and so,

$$I_{\text{out}} = I_{\text{L}} = g_{\text{L}}(V_{\text{m}} - E_{\text{L}}). \quad (\text{F.5})$$

Hence,

$$\frac{dV_{\text{m}}}{dt} = \frac{V_{\text{cmd}} + V_{\text{off}} - V_{\text{m}}}{R_{\text{series}}C_{\text{m}}} - \frac{I_{\text{out}}}{C_{\text{m}}} \quad (\text{F.6})$$

$$= \frac{V_0 + kt + V_{\text{off}} - V_{\text{m}} - R_{\text{series}}I_{\text{out}}}{R_{\text{series}}C_{\text{m}}} \quad (\text{F.7})$$

$$= \frac{V_0 + kt + V_{\text{off}} - V_{\text{m}} - R_{\text{series}}g_{\text{L}}(V_{\text{m}} - E_{\text{L}})}{R_{\text{series}}C_{\text{m}}} \quad (\text{F.8})$$

$$= \frac{V_0 + V_{\text{off}} - R_{\text{series}} - AV_{\text{m}} + kt + R_{\text{series}}g_{\text{L}}E_{\text{L}}}{\tau} \quad (\text{F.9})$$

where $\tau = R_{\text{series}}C_{\text{m}}$, and $A = 1 + g_{\text{L}}R_{\text{series}}$. Equation (F.9) is a linear, first-order non-homogenous ODE with solution,

$$V_{\text{m}}(t) = -\frac{V_0 + V_{\text{off}} - g_{\text{L}}R_{\text{series}}E_{\text{L}} + k\tau}{A} + \frac{kt}{A^2} + B \exp\left\{-\frac{A}{\tau}t\right\}, \quad (\text{F.10})$$

for some constant $B \in \mathbb{R}$, which we set using the initial condition, that is,

$$V_{\text{m}}(0) = V_0 + V_{\text{off}}. \quad (\text{F.11})$$

Consider the constant A . Typically, (under our QC criteria in Chapter 4, for example) we have $R_{\text{seal}} = \frac{1}{g_{\text{L}}} > 0.5\text{G}\Omega$ and $R_{\text{series}} < 25\text{M}\Omega$ in which case,

$$R_{\text{series}} \ll \frac{1}{g_{\text{L}} + \bar{g}x_{\text{O}}} \quad (\text{F.12})$$

and so,

$$A = 1 + R_{\text{series}}(g_{\text{L}} + x_{\text{O}}\bar{g}) \approx 1. \quad (\text{F.13})$$

Hence, the exponential term in Equation (F.10) decays quickly as our upper bound (according to QC.1) for this timescale of decay is approximately

$$\tau = R_{\text{seal}}C_{\text{m}} < 25\text{M}\Omega \times 100\text{pF} = 2.5\text{ms}. \quad (\text{F.14})$$

Comparatively, our leak ramp and reversal ramps much longer than this (400ms and 100ms in length, respectively). Hence, once some has elapsed after the beginning of

the ramped segment is large, ($t \gg 2.5\text{ms}$) this term becomes small, and we have,

$$V_m(t) = -\frac{k\tau}{A^2} + \frac{V_0 + V_{\text{off}} - BR_{\text{series}}}{A} + kt, \quad (\text{F.15})$$

which is simply the equation of a straight line, the gradient of which depends only on the gradient of the command voltage. We also obtain the discrepancy between the transmembrane potential and command voltage induced by artefact effects,

$$V_m(t) - V_{\text{cmd}}(t) = V_m(t) - V_0 - kt \quad (\text{F.16})$$

$$\approx -\frac{k\tau}{A^2} + \frac{V_0 + V_{\text{off}} - g_L E_L R_{\text{series}}}{A} - V_0 - kt \quad (\text{F.17})$$

$$\approx -k\tau + V_{\text{off}} - g_L E_L R_{\text{series}} \quad (\text{F.18})$$

$$\approx -k\tau + V_{\text{off}}, \quad (\text{F.19})$$

since $g_L R_{\text{series}}$ is small and $A \approx 1$. Here, for sufficiently large t , a constant-offset term becomes a reasonable approximation of $V_m(t) - V_{\text{cmd}}(t)$. The size of this offset is influenced by the gradient of the command voltage, k , the series resistance (R_{series}) and membrane capacitance (C_m). For the protocols discussed in Chapters 4 and 5, the leak ramp has a gradient of $k = 0.1 \text{ Vs}^{-1}$, and the reversal ramp has a gradient of $k = 0.4 \text{ Vs}^{-1}$. Additionally, the largest value of τ allowed by our QC criteria (**QC1**) is,

$$\tau = C_m R_{\text{series}} < 100 \text{ pF} \times 25 \text{ M}\Omega = 2.5 \text{ ms}.$$

Thus, the $-k\tau$ term is rather small.

We expect that this discrepancy between V_m and V_{cmd} has a material impact on the inference of our leak-model parameters, and on the observed reversal potentials, E_{post} and E_{before} as introduced in Sections 4.3 and 5.4, respectively.

F.2 Steady states of the artefact model

First, we reform our artefact model's system of ODEs such that each state variable is nondimensional. To do this, let $v = \frac{V_m - V_{\text{min}}}{V_{\text{max}} - V_{\text{min}}}$ where $[V_{\text{min}}, V_{\text{max}}]$ is a range of voltages that V_{cmd} is allowed to take during the experiment. For the experiments described in Chapters 4 and 5, we have $V_{\text{min}} = -120\text{mV}$ and $V_{\text{max}} = +60\text{mV}$. Note that it may be possible for the *transmembrane* potential to lie slightly outside of this range, that is, we may have $V_m \notin [V_{\text{min}}, V_{\text{max}}]$ and so $v < 0$ or $v > 1$. Then, we may formulate our governing equation in terms of v rather than V_m such that our governing equations are,

$$\frac{d\mathbf{x}}{dt} = \mathbf{R}^\top(v)\mathbf{x}, \quad (\text{F.20})$$

and,

$$\frac{dv}{dt} = \frac{1}{V_{\max} - V_{\min}} \frac{dV_m}{dt} \quad (\text{F.21})$$

$$= \frac{1}{V_{\max} - V_{\min}} \frac{V_{\text{cmd}} + V_{\text{off}} - V}{R_{\text{series}} C_m} - \frac{I_{\text{out}}}{C_m} \quad (\text{F.22})$$

$$\frac{dv}{dt} = \frac{v_{\text{cmd}} + v_{\text{off}} - v}{R_{\text{series}} C_m} - \frac{i_{\text{out}}}{C_m}, \quad (\text{F.23})$$

where $v_{\text{cmd}} = \frac{V_{\text{cmd}}}{V_{\max} - V_{\min}}$, $v_{\text{off}} = \frac{V_{\text{off}}}{V_{\max} - V_{\min}}$, $i_{\text{out}} = \frac{I_{\text{out}}}{V_{\max} - V_{\min}}$, and $\mathbf{R}(v) = \mathbf{Q}(V_m)$.

In Appendix A, we discuss the steady states of Markov models when transmembrane potential is held constant. Equivalently, using these results, we may write for any fixed $v \in \mathbb{R}$, the system,

$$\frac{d\mathbf{u}}{dt} = \mathbf{R}^\top \mathbf{u}, \quad (\text{F.24})$$

has a unique, stable, global equilibrium point. We denote the open occupancy at this equilibrium point by $u_O^\infty(v)$ —that is, when v is fixed we have, $u_O(t) \rightarrow u_O^\infty(v)$.

Now, if the artefact model to be at equilibrium, we must have that each of the system's state variables are at equilibrium (both our vector state occupancies, \mathbf{x} , and the transmembrane potential). That is,

$$\frac{d\mathbf{u}}{dt} = \mathbf{0}, \quad (\text{F.25})$$

and,

$$\frac{dv}{dt} = 0. \quad (\text{F.26})$$

The former condition implies,

$$v_{\text{cmd}} + v_{\text{off}} - v - R_{\text{series}} i_{\text{out}} = 0, \quad (\text{F.27})$$

and so,

$$0 = -v - v_{\text{cmd}} + v_{\text{off}} + R_{\text{series}} i_{\text{out}} \quad (\text{F.28})$$

$$= -v + v_{\text{cmd}} + v_{\text{off}} - R_{\text{series}} (g u_O^{(\infty)}(v)(v - e_{\text{Kr}}) + g_L(v - e_L)) \quad (\text{F.29})$$

$$= v_{\text{cmd}} + v_{\text{off}} - v(1 + R_{\text{series}} g u_O^{(\infty)}(v) + R_{\text{series}} g_L) + e_{\text{Kr}} R_{\text{series}} g u_O^{(\infty)}(v_m) + e_L R_{\text{series}} g_L. \quad (\text{F.30})$$

Differentiation of (F.30) with respect to v yields,

$$h(v) := -1 - R_{\text{series}} g u_O^{(\infty)}(v) - R_{\text{series}} g_L - v R_{\text{series}} g \frac{du_O^{(\infty)}(v)}{dv}, \quad (\text{F.31})$$

and provided this quantity is negative for all v there will be at most one equilibrium point. Note that we typically have $R_{\text{series}}g, R_{\text{series}}gL \ll 1$ and so, in order for two equilibrium points to exist, we would either $R_{\text{series}}, g_L, g$ or $\frac{du_O^{(\infty)}(v)}{dv}$ to be uncharacteristically large (Figure F.1 shows typical values for $x_O^{(\infty)}(V_m) = u_O^{(\infty)}(v)$ across a range of voltages). Otherwise, for realistic parameter values (that is, where $R_{\text{series}}g, R_{\text{series}}gL, \frac{du_O^{(\infty)}}{dv} \ll 1$), we have $h(v) \approx -1$ for all v , in which case we expect there to be exactly one equilibrium point.




Figure F.1: The steady-state open occupancy of the Beattie et al. (2018) model (using Cell #5 parameters) under a range of transmembrane potentials.

Bibliography

- Akaike, H. (1998). Information Theory and an extension of the maximum likelihood principle. In Parzen, E., Tanabe, K., and Kitagawa, G., editors, *Selected Papers of Hirotugu Akaike*, Springer Series in Statistics, pages 199–213. Springer, New York, NY.
- Alberts, B., Johnson, A., Lewis, J., Morgan, D., Raff, M., Roberts, K., and Walter, P. (2015). *Molecular Biology of the Cell*. Garland Science, sixth edition. edition.
- András, V., Tomek, J., Nagy, N., Virág, L., Passini, E., Rodriguez, B., and Baczkó, I. (2021). Cardiac transmembrane ion channels and action potentials: cellular physiology and arrhythmogenic behavior. *Physiological reviews*, pages 1083–1176.
- Atkinson, A. C. and Federov, V. V. (1975). The design of experiments for discriminating between two rival models. *Biometrika*, 62(1):57–70.
- Ball, F. G., Sansom, M. S. P., and Colquhoun, D. (1989). Ion-channel gating mechanisms: model identification and parameter estimation from single channel recordings. *Proceedings of the Royal Society of London. B. Biological Sciences*, 236(1285):385–416.
- Balos, C. J., Day, M., Esclapez, L., Felden, A. M., Gardner, D. J., Hassanaly, M., Reynolds, D. R., Rood, J., Sexton, J. M., Wimer, N. T., et al. (2024). SUNDIALS time integrators for exascale applications with many independent ODE systems. *arXiv preprint arXiv:2405.01713*.
- Banga, J. R. and Balsa-Canto, E. (2008). Parameter estimation and optimal experimental design. *Essays in biochemistry*, 45:195–210.
- Bates, D. M. and Watts, D. G. (1988). *Nonlinear regression analysis and its applications*. Wiley, New York.
- Bearup, D. J., Evans, N. D., and Chappell, M. J. (2013). The input–output relationship approach to structural identifiability analysis. *Computer methods and programs in biomedicine*, 109(2):171–181.

- Beattie, K. A., Hill, A. P., Bardenet, R., Cui, Y., Vandenberg, J. I., Gavaghan, D. J., Boer, T. P. d., and Mirams, G. R. (2018). Sinusoidal voltage protocols for rapid characterisation of ion channel kinetics. *The Journal of Physiology*, 596(10):1813–1828.
- Beeler, G. W. and Reuter, H. (1977). Reconstruction of the action potential of ventricular myocardial fibres. *The Journal of Physiology*, 268(1):177–210.
- Berman, A. and Plemmons, R. J. (1979). *Nonnegative matrices in the mathematical sciences*. Computer science and applied mathematics. Academic Press, New York ; London.
- Bondy, A. and Murty, U. (2008). *Graph Theory*. Number 244 in Graduate Texts in Mathematics. Springer, 3rd edition.
- Brown, A. M. (2020). *A companion guide to the Hodgkin-Huxley papers*. The Journal of Physiology.
- Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937.
- Chen, S., Shojaie, A., and Witten, D. M. (2017). Network reconstruction from high-dimensional ordinary differential equations. *Journal of the American Statistical Association*, 112(520):1697–1707.
- Chicone, C. C. (2006). *Ordinary differential equations with applications*. Texts in applied mathematics ; 34. Springer, New York, 2nd ed. edition.
- Clayton, R., Bernus, O., Cherry, E., Dierckx, H., Fenton, F. H., Mirabella, L., Panfilov, A. V., Sachse, F. B., Seemann, G., and Zhang, H. (2011). Models of cardiac tissue electrophysiology: progress, challenges and open questions. *Progress in biophysics and molecular biology*, 104(1-3):22–48.
- Clerx, M., Beattie, K. A., Gavaghan, D. J., and Mirams, G. R. (2019a). Four ways to fit an ion channel model. *Biophysical Journal*, 117(12):2420–2437.
- Clerx, M., Robinson, M., Lambert, B., Lei, C. L., Ghosh, S., Mirams, G. R., and Gavaghan, D. J. (2019b). Probabilistic Inference on Noisy Time Series (PINTS). *Journal of Open Research Software*, 7(1):23.
- Colquhoun, D., Dowsland, K. A., Beato, M., and Plested, A. J. R. (2004). How to impose microscopic reversibility in complex reaction mechanisms. *Biophysical Journal*, 86(6):3510–3518.

- Colquhoun, D. and Hawkes, A. (1981). On the stochastic properties of single ion channels. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 211(1183):205–235.
- Colquhoun, D. and Hawkes, A. G. (1995). A Q-matrix cookbook: how to write only one program to calculate the single-channel and macroscopic predictions for any kinetic mechanism. *Single-channel recording*, pages 589–633.
- Creswell, R., Lambert, B., Lei, C. L., Robinson, M., and Gavaghan, D. (2020). Using flexible noise models to avoid noise model misspecification in inference of differential equation time series models. *arXiv preprint arXiv:2011.04854*.
- Creswell, R., Shepherd, K. M., Lambert, B., Mirams, G. R., Lei, C. L., Tavener, S., Robinson, M., and Gavaghan, D. J. (2024). Understanding the impact of numerical solvers on inference for differential equation models. *Journal of the Royal Society Interface*, 21(212):20230369.
- Dette, H., Haines, L. M., and Irnhof, L. A. (2023). Maximin and Bayesian optimal designs for regression models. *Statistica Sinica*.
- Dette, H., Melas, V. B., and Shpilev, P. (2012). T-optimal designs for discrimination between two polynomial models. *The Annals of Statistics*, 40(1). arXiv: 1205.6283.
- Di Veroli, G. Y., Davies, M. R., Zhang, H., Abi-Gerges, N., and Boyett, M. R. (2013). High-throughput screening of drug-binding dynamics to HERG improves early drug safety assessment. *American Journal of Physiology-Heart and Circulatory Physiology*, 304(1):H104–H117.
- Ebihara, L. and Johnson, E. (1980). Fast sodium current in cardiac muscle. a quantitative description. *Biophysical journal*, 32(2):779–790.
- Fajrial, A. K., He, Q. Q., Wirusanti, N. I., Slansky, J. E., and Ding, X. (2020). A review of emerging physical transfection methods for CRISPR/Cas9-mediated gene editing. *Theranostics*, 10(12):5532.
- Fink, M. and Noble, D. (2009). Markov models for ion channels: versatility versus identifiability and speed. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1896):2161–2179.
- Frazier, D. T., Robert, C. P., and Rousseau, J. (2020). Model misspecification in approximate Bayesian computation: consequences and diagnostics. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(2):421–444.

- Gábor, A., Villaverde, A. F., and Banga, J. R. (2017). Parameter identifiability analysis and visualization in large-scale kinetic models of biosystems. *BMC systems biology*, 11:1–16.
- Gibbons, A. (1985). *Algorithmic graph theory*. Cambridge University Press, Cambridge.
- Giles, W. R. and Imaizumi, Y. (1988). Comparison of potassium currents in rabbit atrial and ventricular cells. *The Journal of Physiology*, 405(1):123–145.
- Gillespie, D. T. (2007). Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58(Volume 58, 2007):35–55.
- Goldwyn, J. H., Imennov, N. S., Famulare, M., and Shea-Brown, E. (2011). Stochastic differential equation models for ion channel noise in Hodgkin-Huxley neurons. *Physical Review E*, 83:041908.
- Grandi, E. (2018). Keeping it short and (not so) simple: characterizing hERG kinetics with sinusoidal waves. *The Journal of Physiology*, 596(10):1783–1784.
- Gupta, G. K., Sacks-Davis, R., and Tescher, P. E. (1985). A review of recent developments in solving ODEs. *ACM Computing Surveys (CSUR)*, 17(1):5–47.
- Hansen, N. (2016). The CMA evolution strategy: A tutorial. Number: arXiv:1604.00772 arXiv:1604.00772 [cs, stat].
- Harris, C. R., Millman, K. J., Van Der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., et al. (2020). Array programming with numpy. *Nature*, 585(7825):357–362.
- Hindmarsh, A. (1982). Large ordinary differential equation systems and software. *IEEE Control Systems Magazine*, 2(4):24–30.
- Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544.
- Horn, R. A. and Johnson, C. R. (2013). *Matrix Analysis*. Cambridge University Press, Cambridge; New York, 2nd edition.
- Hu, L., Pennington, M., Jiang, Q., Whartenby, K. A., and Calabresi, P. A. (2007). Characterization of the functional properties of the voltage-gated potassium channel Kv1.3 in human CD4+ T lymphocytes. *The Journal of Immunology*, 179(7):4563–4570.

- Hund, T. J., Kucera, J. P., Otani, N. F., and Rudy, Y. (2001). Ionic charge conservation and long-term steady state in the Luo-Rudy dynamic cell model. *Biophysical Journal*, 81(6):3324–3331.
- Jammalamadaka, S. R. (2003). *Linear Models : An Integrated Approach*. Series On Multivariate Analysis. World Scientific Publishing Company, 1st ed. edition.
- Keener, J. P. and Sneyd, J. (1998). *Mathematical Physiology*. Interdisciplinary applied mathematics. Springer, New York.
- Kelly, F. (2011). *Reversibility and stochastic networks*. Cambridge University Press.
- Kemp, J. M., Whittaker, D. G., Venkateshappa, R., Pang, Z., Johal, R., Sergeev, V., Tibbits, G. F., Mirams, G. R., and Claydon, T. W. (2021). Electrophysiological characterization of the hERG R56Q LQTS variant and targeted rescue by the activator RPR260243. *Journal of General Physiology*, 153(10):e202112923.
- Kennedy, M. C. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464.
- Kiehn, J., Lacerda, A. E., and Brown, A. M. (1999). Pathways of HERG inactivation. *American Journal of Physiology—Heart and Circulatory Physiology*, 277(1):H199–H210.
- Klimovskaia, A., Ganscha, S., and Claassen, M. (2016). Sparse regression based structure Learning of stochastic reaction networks from single cell snapshot time series. *PLOS Computational Biology*, 12(12):e1005234.
- Kreutz, C., Raue, A., and Timmer, J. (2012). Likelihood based observability analysis and confidence intervals for predictions of dynamic models. *BMC Systems Biology*, 6:1–9.
- Lambert, B., Lei, C. L., Robinson, M., Clerx, M., Creswell, R., Ghosh, S., Tavener, S., and Gavaghan, D. J. (2023). Autocorrelated measurement processes and inference for ordinary differential equation models of biological systems. *Journal of the Royal Society Interface*, 20(199):20220725.
- Lee, W., Mann, S. A., Windley, M. J., Imtiaz, M. S., Vandenberg, J. I., and Hill, A. P. (2016). In silico assessment of kinetics and state dependent binding properties of drugs causing acquired LQTS. *Progress in biophysics and molecular biology*, 120(1-3):89–99.

- Lei, C. L., Clerx, M., Beattie, K. A., Melgari, D., Hancox, J. C., Gavaghan, D. J., Polonchuk, L., Wang, K., and Mirams, G. R. (2019a). Rapid Characterization of hERG Channel Kinetics II: Temperature Dependence. *Biophysical Journal*, 117(12):2455–2470.
- Lei, C. L., Clerx, M., Gavaghan, D. J., Polonchuk, L., Mirams, G. R., and Wang, K. (2019b). Rapid characterization of hERG channel kinetics I: Using an automated high-throughput system. *Biophysical Journal*, 117(12):2438–2454.
- Lei, C. L., Clerx, M., Whittaker, D. G., Gavaghan, D. J., de Boer, T. P., and Mirams, G. R. (2020a). Accounting for variability in ion current recordings using a mathematical model of artefacts in voltage-clamp experiments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 378(2173):20190348.
- Lei, C. L., Fabbri, A., Whittaker, D. G., Clerx, M., Windley, M. J., Hill, A. P., Mirams, G. R., and de Boer, T. P. (2020b). A nonlinear and time-dependent leak current in the presence of calcium fluoride patch-clamp seal enhancer. *Wellcome Open Research*, 5:152.
- Lei, C. L., Ghosh, S., Whittaker, D. G., Aboelkassem, Y., Beattie, K. A., Cantwell, C. D., Delhaas, T., Houston, C., Novaes, G. M., Panfilov, A. V., Pathmanathan, P., Riabiz, M., dos Santos, R. W., Walmsley, J., Worden, K., Mirams, G. R., and Wilkinson, R. D. (2020c). Considering discrepancy when calibrating a mechanistic electrophysiology model. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 378(2173):20190349.
- Lei, C. L. and Mirams, G. R. (2021). Neural network differential equations for ion channel modelling. *Frontiers in Physiology*, 12:708944.
- Lei, C. L., Whittaker, D. G., Windley, M. J., Perry, M. D., Hill, A. P., and Mirams, G. R. (2024). A range of voltage-clamp protocol designs for rapid capture of herg kinetics. *bioRxiv*, page 2024.08.14.607756.
- Lema-Perez, L., Muñoz-Tamayo, R., Garcia-Tirado, J., and Alvarez, H. (2019). On parameter interpretability of phenomenological-based semiphsical models in biology. *Informatics in Medicine Unlocked*, 15:100158.
- Li, T., Lu, G., Chiang, E. Y., Chernov-Rogan, T., Grogan, J. L., and Chen, J. (2017a). High-throughput electrophysiological assays for voltage gated ion channels using syncropatch 768pe. *PLOS ONE*, 12(7):1–18.

- Li, Z., Dutta, S., Sheng, J., Tran, P. N., Wu, W., Chang, K., Mdluli, T., Strauss, D. G., and Colatsky, T. (2017b). Improving the in silico assessment of proarrhythmia risk by combining hERG (Human Ether-à-go-go-Related Gene) channel-drug binding kinetics and multichannel pharmacology. *Circulation: Arrhythmia and Electrophysiology*, 10(2):e004628.
- Li, Z., Mirams, G. R., Yoshinaga, T., Ridder, B. J., Han, X., Chen, J. E., Stockbridge, N. L., Wisialowski, T. A., Damiano, B., Severi, S., Morissette, P., Kowey, P. R., Holbrook, M., Smith, G., Rasmusson, R. L., Liu, M., Song, Z., Qu, Z., Leishman, D. J., Steidl-Nichols, J., Rodriguez, B., Bueno-Orovio, A., Zhou, X., Passini, E., Edwards, A. G., Morotti, S., Ni, H., Grandi, E., Clancy, C. E., Vandenberg, J., Hill, A., Nakamura, M., Singer, T., Polonchuk, L., Greiter-Wilke, A., Wang, K., Nave, S., Fullerton, A., Sobie, E. A., Paci, M., Musuamba Tshinanu, F., and Strauss, D. G. (2020). General principles for the validation of proarrhythmia risk prediction models: An extension of the CiPA in silico strategy. *Clinical Pharmacology and Therapeutics*, 107(1):102–111.
- Maly, J., Emigh, A. M., DeMarco, K. R., Furutani, K., Sack, J. T., Clancy, C. E., Vorobyov, I., and Yarov-Yarovoy, V. (2022). Structural modeling of the hERG potassium channel and associated drug interactions. *Frontiers in Pharmacology*, 13:966463.
- Mangold, K. E., Brumback, B. D., Angsutrararux, P., Voelker, T. L., Zhu, W., Kang, P. W., Moreno, J. D., and Silva, J. R. (2017). Mechanisms and models of cardiac sodium channel inactivation. *Channels*, 11(6):517–533.
- Mangold, K. E., Wang, W., Johnson, E. K., Bhagavan, D., Moreno, J. D., Nerbonne, J. M., and Silva, J. R. (2021). Identification of structures for ion channel kinetic models. *PLoS computational biology*, 17(8):e1008932.
- Massonis, G., Villaverde, A. F., and Banga, J. R. (2023). Distilling identifiable and interpretable dynamic models from biological data. *bioRxiv*.
- Mazhari, R., Greenstein, J. L., Winslow, R. L., Marb  n, E., and Nuss, H. B. (2001). Molecular interactions between two Long-QT syndrome gene products, HERG and KCNE2, rationalized by in vitro and in silico analysis. *Circulation Research*, 89(1):33–38.
- McAllister, R. E., Noble, D., and Tsien, R. (1975). Reconstruction of the electrical activity of cardiac Purkinje fibres. *The Journal of Physiology*, 251(1):1–59.

- Menon, V., Spruston, N., and Kath, W. L. (2009). A state-mutating genetic algorithm to design ion-channel models. *Proceedings of the National Academy of Sciences*, 106(39):16829–16834.
- Meyer, C. D. C. D. (2000). *Matrix analysis and applied linear algebra*. Society for Industrial and Applied Mathematics, Philadelphia.
- Mirams, G. R., Clerx, M., Whittaker, D. G., and Lei, C. L. (2024). Optimal experimental designs for characterising ion channel gating by filling the phase-voltage space of model dynamics. *Mathematics in Medical and Life Sciences*, 1(1):2375494.
- Mirams, G. R., Cui, Y., Sher, A., Fink, M., Cooper, J., Heath, B. M., McMahon, N. C., Gavaghan, D. J., and Noble, D. (2011). Simulation of multiple ion channel block provides improved early prediction of compounds' clinical torsadogenic risk. *Cardiovascular Research*, 91(1):53–61.
- Mirams, G. R., Davies, M. R., Cui, Y., Kohl, P., and Noble, D. (2012). Application of cardiac electrophysiology simulations to pro-arrhythmic safety testing. *British Journal of Pharmacology*, 167(5):932–945.
- Mirams, G. R., Pathmanathan, P., Gray, R. A., Challenor, P., and Clayton, R. H. (2016). Uncertainty and variability in computational and mathematical models of cardiac physiology. *The Journal of Physiology*, 594(23):6833–6847.
- Moler, C. and Van Loan, C. (2003). Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM review*, 45(1):3–49.
- Nakajima, T., Tamura, S., Kurabayashi, M., and Kaneko, Y. (2021). Towards mutation-specific precision medicine in atypical clinical phenotypes of inherited arrhythmia syndromes. *International Journal of Molecular Sciences*, 22(8):3930.
- Neher, E. and Sakmann, B. (1976). Single-channel currents recorded from membrane of denervated frog muscle fibres. *Nature*, 260(5554):799–802. Number: 5554 Publisher: Nature Publishing Group.
- Niederer, S. A., Lumens, J., and Trayanova, N. A. (2019). Computational models in cardiology. *Nature Reviews Cardiology*, 16(2):100–111.
- Noble, D. (1962). A modification of the Hodgkin—Huxley equations applicable to Purkinje fibre action and pacemaker potentials. *The Journal of Physiology*, 160(2):317.
- Noble, D. (2002). Modelling the heart: insights, failures and progress. *BioEssays*, 24(12):1155–1163.

- O'Hara, T., Virág, L., Varró, A., and Rudy, Y. (2011). Simulation of the undiseased human cardiac ventricular action potential: Model formulation and experimental validation. *PLOS Computational Biology*, 7(5):e1002061.
- Overstall, A. M., McGree, J. M., and Drovandi, C. C. (2018). An approach for finding fully Bayesian optimal designs using normal-based approximations to loss functions. *Statistics and Computing*, 28(2):343–358.
- Pathmanathan, P., Shotwell, M. S., Gavaghan, D. J., Cordeiro, J. M., and Gray, R. A. (2015). Uncertainty quantification of fast sodium current steady-state inactivation for multi-scale models of cardiac electrophysiology. *Progress in Biophysics and Molecular Biology*, 117(1):4–18.
- Perry, M. D., Ng, C.-A., Mangala, M. M., Ng, T. Y. M., Hines, A. D., Liang, W., Xu, M. J. O., Hill, A. P., and Vandenberg, J. I. (2019). Pharmacological activation of IKr in models of long QT Type 2 risks overcorrection of repolarization. *Cardiovascular Research*, 116(8):1434–1445.
- Postawa, K., Szczygieł, J., and Kułczyński, M. (2020). A comprehensive comparison of ODE solvers for biochemical problems. *Renewable Energy*, 156:624–633.
- Rampe, D. and Brown, A. M. (2013). A history of the role of the hERG channel in cardiac risk assessment. *Journal of pharmacological and toxicological methods*, 68(1):13–22.
- Rencher, A. C. and Schaalje, G. B. (2008). *Linear models in statistics*. Wiley-Interscience, Hoboken, N.J., 2nd ed. edition.
- Robertson, G. A. and Morais-Cabral, J. H. (2020). hERG function in light of structure. *Biophysical Journal*, 118(4):790–797.
- Roy, M.-L., Dumaine, R., and Brown, A. M. (1996). HERG, a primary human ventricular target of the nonsedating antihistamine terfenadine. *Circulation*, 94(4):817–823.
- Rudy, Y. and Silva, J. R. (2006). Computational biology in the study of cardiac ion channels and cell electrophysiology. *Quarterly Reviews of Biophysics*, 39(1):57–116.
- Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. (2016). A review of modern computational algorithms for Bayesian Optimal Design. *International Statistical Review / Revue Internationale de Statistique*, 84(1):128–154.
- Sakmann, B. and Neher, E. (1984). Patch clamp techniques for studying ionic channels in excitable membranes. *Annual review of Physiology*, 46(1):455–472.

- Schoening, M. E. and Silva, J. R. (2024). *Creating Computational Models of Ion Channel Dynamics*, pages 139–156. Springer US, New York, NY.
- Schoppa, N. and Sigworth, F. (1998). Activation of shaker potassium channels: I. characterization of voltage-dependent transitions. *The Journal of general physiology*, 111(2):271–294.
- Seber, G. and Wild, C. (2005). *Nonlinear Regression*. Wiley Series in Probability and Statistics. Wiley.
- Sheridan, K., Berleant, J., Bathe, M., Condon, A., and Williams, V. V. (2023). Factorization and pseudofactorization of weighted graphs. *Discrete Applied Mathematics*, 337:81–105.
- Shuttleworth, J. G., Lei, C. L., Whittaker, D. G., Windley, M. J., Hill, A. P., Preston, S. P., and Mirams, G. R. (2023). Empirical quantification of predictive uncertainty due to model discrepancy by training with an ensemble of experimental designs: An application to ion channel kinetics. *Bulletin of Mathematical Biology*, 86(1):2.
- Shuttleworth, J. G., Lei, C. L., Windley, M. J., Hill, A. P., Preston, S. P., and Mirams, G. R. (2024). Evaluating the predictive accuracy of ion channel models using data from multiple experimental designs. *bioRxiv*.
- Siekmann, I., Sneyd, J., and Crampin, E. J. (2012). MCMC can detect nonidentifiable models. *Biophysical Journal*, 103(11):2275–2286.
- Silva, J. R., Pan, H., Wu, D., Nekouzadeh, A., Decker, K. F., Cui, J., Baker, N. A., Sept, D., and Rudy, Y. (2009). A multiscale model linking ion-channel molecular dynamics and electrostatics to the cardiac action potential. *Proceedings of the National Academy of Sciences*, 106(27):11102–11106.
- Smith, R. C. (2013). *Uncertainty Quantification: theory, implementation, and applications*, volume 12. Society for Industrial and Applied Mathematics.
- Sobol, I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation*, 55(1):271–280. The Second IMACS Seminar on Monte Carlo Methods.
- Stadt, M. M., Leete, J., Devinyak, S., and Layton, A. T. (2022). A mathematical model of potassium homeostasis: effect of feedforward and feedback controls. *PLOS Computational Biology*, 18(12):e1010607.
- Städter, P., Schälte, Y., Schmiester, L., Hasenauer, J., and Stapor, P. L. (2021). Benchmarking of numerical integration methods for ODE models of biological systems. *Scientific reports*, 11(1):2696.

- Suessbrich, H., Waldegger, S., Lang, F., and Busch, A. (1996). Blockade of HERG channels expressed in xenopus oocytes by the histamine receptor antagonists terfenadine and astemizole. *FEBS Letters*, 385(1-2):77–80.
- Sung, C.-L., Barber, B. D., and Walker, B. J. (2020). Calibration of inexact computer models with heteroscedastic errors. Number: arXiv:1910.11518 arXiv:1910.11518 [stat].
- Teed, Z. R. and Silva, J. R. (2016). A computationally efficient algorithm for fitting ion channel parameters. *MethodsX*, 3:577–588.
- ten Tusscher, K. H. W. J., Noble, D., Noble, P. J., and Panfilov, A. V. (2004). A model for human ventricular tissue. *American Journal of Physiology-Heart and Circulatory Physiology*, 286(4):H1573–H1589.
- Tomek, J., Bueno-Orovio, A., Passini, E., Zhou, X., Mincholé, A., Britton, O., Bartolucci, C., Severi, S., Shrier, A., Virag, L., et al. (2019). Development, calibration, and validation of a novel human ventricular myocyte model in health, disease, and drug block. *eLife*, 8:e48890.
- Tommasi, C. (2008). Optimal designs for both model discrimination and parameter estimation. *Universita degli Studi di Milano, UNIMI - Research Papers in Economics, Business, and Statistics*, 139.
- Trayanova, N. A., Doshi, A. N., and Prakosa, A. (2020). How personalized heart modeling can help treatment of lethal arrhythmias: A focus on ventricular tachycardia ablation strategies in post-infarction patients. *WIREs Systems Biology and Medicine*, 12(3):e1477.
- Tuncer, N., Maratheva, M., LaBarre, B., and Payote, S. (2018). Structural and practical identifiability analysis of Zika epidemiological models. *Bulletin of Mathematical Biology*, 80.
- Vandenberg, C. and Bezanilla, F. (1991). A sodium channel gating model based on single channel, macroscopic ionic, and gating currents in the squid giant axon. *Biophysical Journal*, 60(6):1511–1533.
- Vandenberg, J. I., Perry, M. D., Perrin, M. J., Mann, S. A., Ke, Y., and Hill, A. P. (2012). hERG K(+) channels: structure, function, and clinical significance. *Physiological Reviews*, 92(3):1393–1478.
- Wang, S., Liu, S., Morales, M. J., Strauss, H. C., and Rasmusson, R. L. (1997). A quantitative analysis of the activation and inactivation kinetics of HERG expressed in *Xenopus* oocytes. *The Journal of Physiology*, 502 (Pt 1)(Pt 1):45–60.

- Wang, W. and MacKinnon, R. (2017). Cryo-EM Structure of the Open Human Ether-à-go-go-Related K⁺ Channel hERG. *Cell*, 169(3):422–430.e10.
- Weerakoon, P., Culurciello, E., Klemic, K. G., and Sigworth, F. J. (2009). An integrated patch-clamp potentiostat with electrode compensation. *IEEE transactions on Biomedical Circuits and Systems*, 3(2):117–125.
- Whitt, W. (2006). Continuous-time Markov chains. *Dept. of Industrial Engineering and Operations Research, Columbia University, New York*, page 164.
- Whittaker, D. G., Clerx, M., Lei, C. L., Christini, D. J., and Mirams, G. R. (2020a). Calibration of ionic and cellular cardiac electrophysiology models. *WIREs Systems Biology and Medicine*, 12(4):e1482.
- Whittaker, D. G., Kemp, J. M., Mirams, G. R., and Claydon, T. W. (2020b). Rapid characterisation of R56Q mutant hERG channel kinetics using sinusoidal voltage protocols. *Biophysical Journal*, 118(3):112a.
- Wieland, F.-G., Hauber, A. L., Rosenblatt, M., Tönsing, C., and Timmer, J. (2021). On structural and practical identifiability. *Current Opinion in Systems Biology*, 25:60–69.
- Willmott, C. J., Ackleson, S. G., Davis, R. E., Feddema, J. J., Klink, K. M., Legates, D. R., O'donnell, J., and Rowe, C. M. (1985). Statistics for the evaluation and comparison of models. *Journal of Geophysical Research: Oceans*, 90(C5):8995–9005.
- Yang, J., Bruno, W. J., Hlavacek, W. S., and Pearson, J. E. (2006). On imposing detailed balance in complex reaction mechanisms. *Biophysical Journal*, 91(3):1136–1141.
- Zhang, J., Litvinova, M., Wang, W., Wang, Y., Deng, X., Chen, X., Li, M., Zheng, W., Yi, L., Chen, X., et al. (2020). Evolving epidemiology and transmission dynamics of coronavirus disease 2019 outside Hubei province, China: a descriptive and modelling study. *The Lancet Infectious Diseases*, 20(7):793–802.
- Zotos, E. E. (2017). Determining the Newton-Raphson basins of attraction in the electromagnetic Copenhagen problem. *International Journal of Non-Linear Mechanics*, 90:111–123.