

A Machine Learning Approach to Modeling Dynamic Decision-Making in Strategic
Interactions and Prediction Markets

By

John Jacob Nay

Dissertation

Submitted to the Faculty of the
Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Interdisciplinary Studies: Integrated Computational Decision Science

May, 2017

Nashville, Tennessee

Approved:

Jonathan M. Gilligan, Ph.D.

Yevgeniy Vorobeychik, Ph.D.

George M. Hornberger, Ph.D.

Mark D. Abkowitz, Ph.D.

ACKNOWLEDGMENTS

I'm extremely grateful for the guidance of Dr. Jonathan M. Gilligan. Dr. Gilligan's insights improved this entire dissertation and my understanding of science and its potential to improve policy. Chapters 2 and 4 are based on articles co-authored with Dr. Gilligan.

Chapter 3 is based on an article co-authored with Dr. Yevgeniy Vorobeychik. Dr. Howard Kunreuther of the University of Pennsylvania provided helpful comments on the work in Chapter 3.

Martin Van der Linden was a co-author of the paper that led to Chapter 4. I thank Manuel Naumovich Velasco for sharing the data from his Total Solar Irradiance prediction for the work comprising Chapter 4.

The entire ADAPT-SL and Water Conservation research teams at Vanderbilt University, and especially Dr. George Hornberger, have provided significant inspiration and support for conducting multi-disciplinary research.

U.S. National Science Foundation grants EAR-1416964 and EAR-1204685 partially supported this research.

Finally, my parents, my brother, and my wife, have all supported me and inspired me throughout this process.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	iii
LIST OF TABLES	vi
LIST OF FIGURES	viii
1 CHAPTER 1 INTRODUCTION	1
1.1 Overview	1
1.2 Organization	2
2 CHAPTER 2 A MACHINE LEARNING APPROACH TO UNDERSTANDING DYNAMIC DECISIONS	4
2.1 Abstract	4
2.2 Model Representation	4
2.3 Estimation	6
2.4 Experimental Game Data	10
2.5 Simulated Data	17
2.6 Obervational Data	20
2.7 Discussion	23
3 CHAPTER 3 A MACHINE LEARNING APPROACH TO PREDICTING DY- NAMIC DECISIONS	25
3.1 Abstract	25
3.2 Introduction	25
3.3 Data	27
3.4 Model	34
3.5 Results	37
3.5.1 Individual-level performance	37

3.5.2	Aggregate-level performance	39
3.6	Analysis	49
3.7	Conclusion	52
4	CHAPTER 4 COMPUTATIONAL POLICY DESIGN: A PREDICTION MAR-	
	KET CASE STUDY	54
4.1	Abstract	54
4.2	Introduction	54
4.3	Related Work	56
4.4	Model Design	57
4.4.1	Temperature Models	57
4.4.2	Climate Beliefs	59
4.4.3	Traders and Markets	60
4.4.4	Social Network	61
4.4.5	Model Dynamics	63
4.4.6	Model Parameters	65
4.5	Results	66
4.5.1	Historical Climate	66
4.5.2	Future Scenarios	67
4.6	Discussion	69
4.7	Conclusion	70
5	CHAPTER 5 CONCLUSION	72
	BIBLIOGRAPHY	74

LIST OF TABLES

Table	Page
3.1 Payoff table. Payoff table where one player plays from the perspective of the columns and the other from the rows. For this to be a repeated Prisoner's Dilemma, it must hold that $T > R > P > S$, and $R > (S + T)/2$ [49].	28
3.2 Summary of thirty game structures that compose the full combined data set [52, 47, 40, 41, 42, 43, 48, 13]. BR 2006 [40] and DB 2005 [47] both also conducted one-shot games; I only describe and use their repeated game data. KSBS 2009 [42] also conducted games with partial information; I only describe and use their full information data. AM 1993 [52] also conducted games that matched humans with computers; I only describe and use the games they conducted where humans played other humans. FO 2012 [48] included one-shot games and games with very different protocols for how and when to make a choice in order to study continuous choices; I only use the "Grid treatment with $n = 8$ subperiods," which they say is, "comparable to the 10-stage repeated games featured in previous laboratory studies." DO 2009 [41] also conducted random matching of opponents; I only use their fixed matching treatments.	31

3.3 Model performance comparison. Best performance for each test is *italicized*.

First four rows are performance on 32,614 predictions of period one actions and 135,772 predictions of period greater than one actions. Each evaluation is an average for how that model performed with out-of-sample predictions for each game structure. I conduct paired sample t-tests (not assuming equal variances) to determine if the thirty accuracy (Acc.) and log likelihood values (LL) for the full model are statistically greater than the values of the next best model. Accuracies for $t > 1$ of the full model ($p = 0.03$) and the likelihoods for $t > 1$ of the full model ($p < 0.001$) are significantly higher than the next best model (dynamic). Accuracies for $t = 1$ of the full model are greater than the next best model, the static model ($p = 0.07$), while the likelihoods for $t = 1$ of the full model are not significantly greater than the likelihoods of the static model ($p = 0.31$). **Last four rows** are performance on average cooperation level in each structure ($n=30$) and time series of average cooperation in each structure ($n=212$). Infinitely repeated interactions with delta set to 0.5 are on average only two periods long and there is not sufficient empirical data to extend out to eight periods so I extend to seven. Two structures are finitely repeated for two periods and two others are finitely repeated for four periods. I conducted paired sample t-tests between the full model and competitors, with a null hypothesis that the true difference in means of the 212 squared errors between predicted and real cooperation levels at all times in all game structures is equal to zero, i.e. that the full model and a competitor are statistically indistinguishable in terms of squared errors on time series predictions. I did the same for the thirty predictions of overall cooperation levels. I reject the null of no difference for all comparisons except with the static model for both tests and the dynamic model for the time series. . . . 41

LIST OF FIGURES

Figure	Page
2.1 Game theoretic strategies represented as finite state machines.	6
2.2 Genetic algorithm overview.	10
2.3 Payoff table for a Prisoner’s Dilemma game. C is cooperate and D is defect. .	12
2.4 Empirically estimated finite state machine.	14
2.5 GA predictions.	16
2.6 Role of noise in correctly identifying the player’s strategy. Points have been jittered for visualization. Lines are LOESS-curves. Noise represents the probability, in any period, that an agent will take the opposite action from its decision model. Both the GT and TFT strategies can be represented as a 10 element matrix, of which 6 elements have identifiable values. The model error represents the number of mismatched elements between the estimated and true matrices, so it ranges from 0 (best) to 6 (worst).	19
2.7 Models of cooperation over water issues for sets of countries. The italicized transitions would not be feasible outcomes if this strategy was played deter- ministically, but they are realized in the data. The inputs (c,nc) represent co- operation and non-cooperation, respectively; (f,nf) conflict and non-conflict; and (t,nt) signing and not signing a treaty.	22
3.1 Game structures ($n = 30$) with location based on the payoff variable values (A.), and delta and error values (B.). Colors represent proportion of cooper- ation observed in the game structure. Locations have been slightly randomly shifted to improve visualization.	32

3.2 A hierarchical view of the data. The top level divides the data into 30 game structures. The next level down are the interactions between two players. Within each interaction, there are T time periods. In repeated games in which termination is stochastic, T ranges up to thirty-eight. Across all interactions and structures, T is five, on average. Within each time period, player 1 takes action C_1 and player 2 takes action C_2 . 168,386 actions were taken across all the experimental data. 33

3.3 Model validation process for individual-level actions. I assign each of the thirty game structures into either training or test data. With the training data, I learn the parameters of the individual-level model, and then predict the decisions in game structures assigned to the test data. I repeat this process thirty times, including a different game structure as the held-out test each time (leave-out-one-cross-validation), until I have predictions for all the decisions for each of the game structures. 38

3.4 Model validation process for aggregate-level patterns. I tested the dynamic-only model by sampling lagged outcomes for ‘period zero’ actions from a Bernoulli distribution with equal probability of cooperation and defection, which is approximately the mean cooperation rate in the data. A subtle, but crucial, distinction between this process and the model validation process for individual-level action predictions (Fig. 3.3) is that, here, *I only pass game structures* for the test games, rather than the full behavioral data and the game structure. 40

3.5 How predictive performance varies with data splitting. RMSE (**A.**) and correlation (**B.**) for time series forecasts of play in 30 game structures, varying folds in cross-validation from 30 to 2. The full model consistently has lower prediction error and higher correlation than the baseline model and the fEWA model until there are only two folds. It is, in general, difficult to make accurate predictions when the ratio of observational units to folds is small. In the case of predicting aggregate and dynamic play, the game structure itself is the observational unit, and I only have thirty, so it's not surprising that performance can degrade at two folds depending on the particular random realization of fold assignments. 42

3.6 Model forecasts. Out-of-sample forecasts of cooperation level over time, for all game structures, conditional only on the game structure ($n = 212$). . . . 44

3.7 Variable importance scores. Variable importance scores for individual-level dynamic component of full model, i.e. for predictions of an agents' probability of cooperation in periods > 1 . Variables separated by '*' represent an interaction between those two variables. These relative importance scores are derived from the absolute values of the t-statistics for each model parameter, which correspond to the effects of the predictor variables (accounting for variability in the estimates) on the probability of cooperation, *ceteris paribus*. 46

3.8 Predicted proportions of cooperation ($n = 28$). 48

3.9 Model simulation analysis. **A.** is a partial rank correlation coefficient analysis [78] of the effects of the game parameters on average cooperation; lines are bootstrapped 95% confidence intervals ($n = 1,000$). Continuous is set to its empirical mode, 0, because there is no within experiment variation on this. **B.** shows that first period play strongly affects cooperation levels in periods greater than one. Setting the game structure variables to the mean of the empirically observed values: if I exogenously set the probability of cooperation during the first period to 0 the simulated proportion of cooperation in subsequent periods is only 0.18 ('simulated experiment 1'), and if I set the probability of cooperation during the first period to 1 the simulated proportion of cooperation in subsequent periods is 0.68 ('simulated experiment 4'). When the probability of first period cooperation is set to 0, and I use a game structure that **A.** suggests should maximize cooperation, the proportion of cooperation is 0.43 ('simulated experiment 2'); and when the probability of first period cooperation is set to 1, with the game structure that should minimize cooperation, the cooperation level is 0.35 ('simulated experiment 3'). 51

4.1 Historical measurements of temperature and a realization of possible future temperatures under two alternate models of climate physics 59

4.2 An example of the social network among traders. Yellow and blue circles represent traders who believe in the two different models of climate change. Lines indicate social network connections. In a highly segmented network (seg close to 1), most links connect like-minded traders and few connect traders with opposing beliefs. 62

4.3 Convergence over trading sequences for different degrees of social-network connection and segmentation 67

4.4 Estimated effects of model parameters on convergence of beliefs in future scenario	69
5.1 Pasteur's Quadrant	73

CHAPTER 1 INTRODUCTION

1.1 Overview

My dissertation lies at the intersection of computer science and the decision sciences. With psychology and sociology, I'm interested in models where social and cognitive factors influence human decisions, especially in social dilemmas and out-of-equilibrium dynamics such as learning and adaptation. With microeconomics and game theory, I realize that modeling human behavior as an attempt to maximize individual well-being is useful. I combine theoretical insights from the decision sciences with computational methods to understand and predict human behavior.

My work is distinct from most decision science research in two respects: (1) I emphasize prediction, and (2) I am not attempting to make simple economic decision models better describe behavior. First, the majority of economic and behavioral science research focuses on either describing in-sample phenomena or testing theories that posit causal relationships among theoretical constructs [1, 2].¹ The primary difference between prediction and causality research approaches derives from the unit of interest – causal explanation is directly concerned with theoretical population-level constructs, while prediction is directly concerned with sampled data. My research demonstrates that data-driven models with a prediction focus can be strategically designed and implemented to inform theory. As for the second point of divergence, I am not part of what constitutes most of the field of behavioral economics: the “subjective expected utility repair program” [6]. This is the active line of research adding psychological parameters to the subjective expected utility model [7] to allow it to better fit behavioral data [8].

My overarching modeling goal for my dissertation is to maximize generalization – some

¹This is done by transforming the constructs into measurable variables and their causal relationships into functional forms of (usually linear) statistical models with parameters that, once estimated on a sample, are interpreted in the terms of the theory with the use of assumptions about counterfactual states of the world [3, 4, 5].

function of data and knowledge – from one sample, with its observations drawn independently from the distribution \mathcal{D} , to another sample drawn independently from \mathcal{D} ,² while also obtaining interpretable insights from the models. The processes of collecting relevant data and generating features from the raw data impart substantive knowledge into predictive models (and the model representation and optimization algorithms applied to those features contain methodological knowledge). I combine this knowledge with the data to train predictive models to deliver generalizability, and then investigate the implications of those models with simulations systematically exploring parameter spaces. The exploration of parameter space provides insights about the relationships between key variables.

1.2 Organization

Chapter 2 describes a method I developed to generate descriptive models of strategic decision-making. I use an efficient representation of repeated game strategies with state matrices and a genetic algorithm-based estimation process to learn these models from data. This combination of representation and optimization is effective for modeling decision-making with experimental game data and observational international relations data. This chapter introduces the Prisoner’s Dilemma game, which is a focus of Chapter 3. I also demonstrate the method’s ability to recover known data-generating processes by simulating data with agent-based models and correctly deriving the underlying decision models.

Chapter 3 demonstrates that models can deliver high levels of *generalizability* with accurate out-of-sample predictions and *interpretable* scores of variable importance that can guide future behavioral research. I combine behavioral-game-theory-inspired feature design with data to train predictive models to deliver generalizability, and then investigate interactive implications of those models with optimization and sensitivity analyses.

Chapter 4 presents a computational model as a test-bed for designing climate prediction

²The distribution, \mathcal{D} , is a theoretical concept and whether samples should be considered as part of the same distribution ultimately needs to be argued for outside the data.

markets. Traders adapt their beliefs about future temperatures based on the profits of other traders in their social network. I simulate two alternative climate futures, in which global temperatures are primarily driven either by carbon dioxide or by solar irradiance. These represent, respectively, the scientific consensus and the most plausible hypothesis advanced by prominent skeptics. Then I conduct sensitivity analyses to determine how a variety of factors describing both the market and the physical climate may affect traders' beliefs about the cause of global climate change. Market participation causes most traders to converge quickly toward believing the "true" climate model, suggesting that a climate market could be useful for building public consensus.

Chapter 5 concludes with the goals I have for my work and situates my dissertation relative to basic and applied science.

CHAPTER 2 A MACHINE LEARNING APPROACH TO UNDERSTANDING DYNAMIC DECISIONS

2.1 Abstract

This chapter¹ outlines a method for automatically generating models of dynamic decision-making that have strong predictive power and are interpretable in human terms. This method provides direct insight into observed dynamic processes. I use an efficient model representation and an evolutionary algorithm-based estimation process to generate simple approximations that explain most of the structure of complex stochastic processes. This method, implemented in C++ and R (and released as free open source software), scales well to large data sets. I apply my methods to empirical data from human subjects game experiments and international relations. I also demonstrate the method's ability to recover known data-generating processes by simulating data with agent-based models and correctly deriving the underlying decision-making models for multiple agent models and degrees of stochasticity.

2.2 Model Representation

This chapter describes a modeling method designed to facilitate the process of understanding data on dynamic decision-making. We created a practical, easy-to-use software package implementing the method. Although the method is more broadly applicable, the motivation for the model representation was prediction of individual behavior in strategic interactions, i.e. games. Most *behavioral* game-theoretic treatments of repeated games use action-learning models that specify the way in which attractions to actions are updated by an agent as play progresses [9]. Action learning models can perform poorly at predicting behavior in games where cooperation (e.g. Prisoner's Dilemma) or coordination (e.g. Bach

¹Parts of this chapter appeared in Nay, J. J., Gilligan, J. M. (2015). "Data-Driven Dynamic Decision Models". *Proceedings of the 2015 Winter Simulation Conference*, 2752–2763. IEEE Press.

or Stravinsky) are key [10]. Also, action learning models often fail to account for the effects of changes in information and player matching conditions [11]. In this chapter, I model repeated game strategies as decision-making procedures that can explicitly consider the dynamic nature of the environment, e.g. if my opponent cooperated last period then I will cooperate this period. I represent decision-making with finite-state machines and use a genetic algorithm to estimate the values of the state transition tables. This combination of representation and optimization allows one to efficiently and effectively model dynamic decision-making.

Traditional game theories define strategies as complete contingent plans that specify how a player will act in every possible state; however, when the environment becomes even moderately complex the number of possible states of the world can grow beyond the limits of human cognition [12, 13]. One modeling response to cognitive limitations has been to exogenously restrict the complexity of repeated game strategies by representing them as Moore machines with a small number of states [14, 12, 15]. Moore machines are finite state machines whose outputs depend only on their current state [16], as opposed to Mealy machines, for which the output also depends on current input [17].

Moore machines can model bounded rationality, explicitly treating procedures of decision-making [18]. A machine modeling agent i responding to the actions of agent j is a four-tuple $[Q_i, q_i^0, f_i, \tau_i]$, where Q_i is the set of states, $q_i^0 \in Q_i$ is the initial state, $f_i : Q_i \rightarrow A_i$ is the output function mapping a state to an action, and $\tau_i : Q_i \times A_j \rightarrow Q_i$ (where $j \neq i$) is the transition function mapping a state and another agent's action to a state [18]. I generalize this model beyond games by allowing for more inputs in τ_i than A_j , and by providing empirical rankings of these inputs that can be used to induce sparsity in more context-rich environments. The Moore machine can have many underlying states for a single observable action, allowing it to represent arbitrarily complex decision processes. In practice the number of states needs to be limited for computational reasons. The complexity is directly controlled by the number of states, which is a tuning parameter of the method that can be

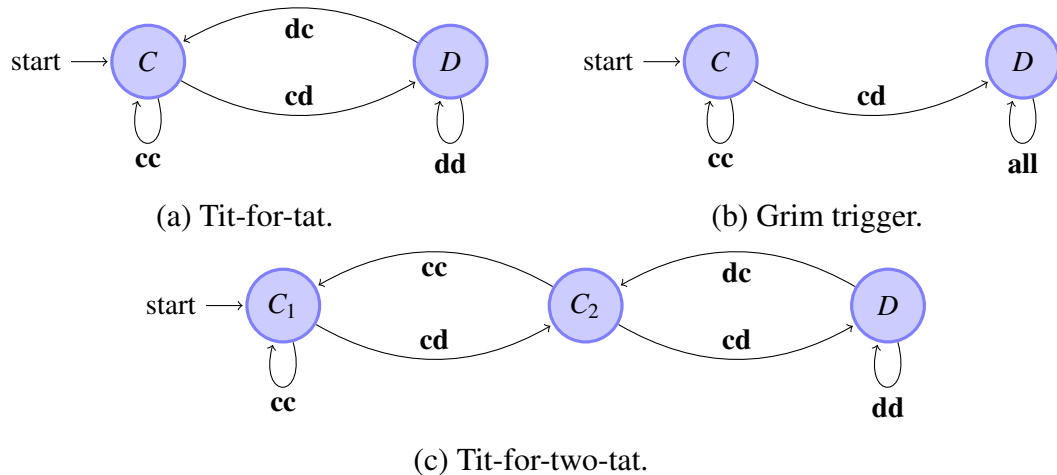


Figure 2.1: Game theoretic strategies represented as finite state machines [14, 12, 13]. Lower-case letters denote the possible outcomes of the previous period of play, e.g. **cd** means the player cooperated and her opponent defected.

optimized by Algorithm 2 for predictive performance.

Fig. 2.1 shows examples of finite state machines (FSMs) representing strategies for the Iterated Prisoner’s Dilemma game (see Section 2.4 for game details): The possible states are cooperate (C) and defect (D), and after initialization the current state is determined by the history of the player and her opponent cooperating or defecting (**cc**, **cd**, **dc**, **dd**) in the previous period.

2.3 Estimation

There have been other machine learning approaches to estimating finite state machines but most have relied on recurrent neural networks, which can be very difficult to train [19, 20]. A more traditional statistical machine learning approach to estimating finite state machine models from data is to model the finite state machine probabilistically. The predictions could be probabilistic, e.g. 0.75 probability of cooperation, rather than deterministic, e.g. cooperation. This would allow the optimization procedure to use a differentiable loss function and thus gradient descent, the most common optimization process used for estimating machine learning models, could be employed. However, I was interested in

a process to estimate deterministic models because they are more similar to the existing theoretical repeated game models and are more interpretable.

For the deterministic finite state machine estimation, we have a vector that is length I , where each x_i is 0 or 1. Our objective is to find the values of x_i that minimize 0-1 loss on the training data, i.e. the misclassification rate. We can add a second term to the objective function to promote sparsity. In addition to finding values of x_i that minimize 0-1 loss on the training data, the objective could include minimizing the number of transitions that move a finite state machine from one state to another, s . When there are fewer s , the model is simpler and easier to understand. s can be simply computed by counting the number of states that cause a move to another state. The objective would then be to simultaneously minimize both the 0-1 loss and s . The current work only optimizes 0-1 loss but future research could include another term of this nature in the objective function that regularizes the model and explicitly pushes the search process involved in estimating models toward parameter spaces that contains models that are both effective predictors of data and simple.

The problem of estimating a model is NP-hard to directly optimize because the 0-1 loss function is non-convex and non-differentiable [21]. The objective function can be optimized by a search algorithm such as branch-and-bound, simulated annealing, tabu search, particle swarm optimization, or evolutionary algorithms.

A popular evolutionary algorithm, a genetic algorithm (GA), has been used to model agents updating beliefs based on endogenously determined variables in a general equilibrium environment [22], and agents learning to make economic decisions [23, 24, 25, 26]. In contrast to investigations of GAs as models of agent learning and behavior, I use GAs to search for interpretable agent decision models that predict empirical data. This is similar to work by [27, 12, 28] in which GAs evolved FSMs based on their interactions with one another in simulated games, but whereas these were theoretical exercises, I am estimating models to explain and predict observed interactions among real agents. I use GAs as optimization routines for estimation because they perform well in rugged search spaces to

quickly solve discrete optimization problems, are a natural complement to the binary string representation of FSMs [29], and are easily parallelized.

Duffy and Engle-Warnick [30] combined empirical experimental data with genetic programming (GP) to model behavior. GP, with the same genetic operations as most GAs [31], is a process that can evolve arbitrary computer programs [32]. I apply genetic operations to FSM representations rather than to all predictor variables and functional primitives because I am interested in deriving decision models with a particular structure: FSMs with latent states, rather than models conditioning on observable variables with any arbitrary functional form. With data-driven modeling, it is desirable to impose as many constraints as can be theoretically justified on the functional form of the model (see [28] for interesting theoretical results related to FSM agents interacting in games). This avoids overfitting the model to a particular training data set by constraining the model to a functional form that is likely generalizable across contexts, allows genetic selection to converge better, and reduces the computational effort required to explore parameter space. An additional challenge in implementing GP is specifying the genetic operations on the function primitives while ensuring that they will always produce syntactically valid programs that represent meaningful decision models. This requires fine-tuning to specific problems, which I avoid because I am designing a general method applicable across domains.

Using Moore machines as the building blocks of the decision modeling method ensures that estimation will produce easily interpretable models with latent states that can be represented graphically (see Fig. 2.1 for examples). My process represents Moore machines as Gray-encoded binary strings consisting of an action vector followed by elements that form the state matrix [33]. For details, see Fig. 2.4a and the `build_bitstring`, `decode_action_vec`, and `decode_stat_mat` functions. This way, genetic operators can have free reign to search the global parameter space guided by the ability to predict provided data with the decoded binary strings.

The vast majority of computation time for Algorithm 1 is the evaluation of the predic-

Data: Time series of actions taken by agents and the relevant predictors of each action.

Result: Finite state machine (FSM) with highest predictive performance.

Set convergence criteria (maximum number of generations or number of generations without improvement in performance of best FSM) based on the number of parameters to estimate;

Create initial population at step $k = 0$ of p individuals $\{\theta_1^0, \theta_2^0, \dots, \theta_p^0\}$ (FSMs, encoded as binary strings);

while *convergence not satisfied* **do**

- Decode each individual's string into an FSM, evaluate its predictive performance on training data, and set this as the individual's fitness for step k , $f(\theta_i^k)$;
- Assign each individual a probability for selection proportional to its fitness, $p_i^k \propto f(\theta_i^k)$;
- Select individuals by sampling from the population with replacement;
- Create next generation, $\{\theta_1^{k+1}, \theta_2^{k+1}, \dots, \theta_p^{k+1}\}$, by applying random *crossover* and *mutation* to the selected sub-population;

end

Return $\operatorname{argmax}_{\theta_i^k} f(\theta_i^k)$: the individual with the greatest predictive power;

Check each element of the solution's state transition matrix for contribution to predictive performance, and evaluate solution on test data, if supplied.

Algorithm 1: Evolving finite state machines with a genetic algorithm.

tive accuracy of the FSMs (not the stochastic generation of candidate FSMs). To improve performance I implement this evaluation in C++ using the Rcpp package [34], and, because it is embarrassingly parallel, distribute it across processor cores. I have incorporated the code into an R package with an API of documented function calls and using the GA package [35] to perform the GA evolution. A user can generate an optimized FSM by calling `evolve_model(data)`, where `data` is an R `data.frame` object with columns representing the time period of the decision, the decision taken at that period, and any predictor variables. There are many additional optional arguments to this `evolve_model` function, but they have sensible default values and are not likely to need to be tuned. My package then generates C++ code for a fitness function and uses it to evaluate automatically generated candidate models. Once the convergence criteria of this iterative search process is satisfied, the best FSM is identified, and each predictor variable is assessed by checking its identifiability and computing its importance in that decision model. The return value

contains a descriptive summary of all results, including those shown in Fig. 2.4.

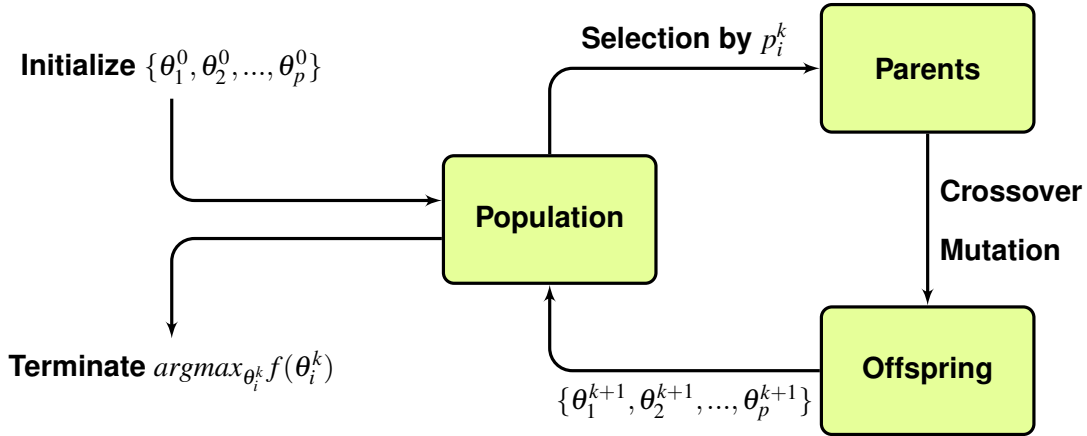


Figure 2.2: Genetic algorithm overview.

The number of states in the FSM and the number of predictor variables to include are hyper-parameters that control the complexity of the model. Beginning with the simplest possible model and increasing complexity by adding states and variables, we often observe that at first, out-of-sample predictive accuracy grows because bias falls more quickly than variance rises; but eventually, adding further complexity reduces bias less than it increases variance so accuracy decreases [36]. Cross-validation on the training data can facilitate discovering values for the hyper-parameters that maximize predictive accuracy (Algorithm 2). I assess the out-of-sample predictive accuracy of the final model with a hold-out test set of data, distinct from the cross-validation test-sets in Algorithm 2. Increasing complexity to optimize predictive accuracy introduces a new trade-off because more complex decision models are harder to interpret in human terms, so the “best” solution will depend on the goals of the analysis.

2.4 Experimental Game Data

The Iterated Prisoner’s Dilemma (IPD) is often used as a model of cooperation [37]. A one-shot PD game has a unique equilibrium in which each player chooses to defect even though both players would be better off if they cooperated. Suppose two players

<p>Data: A dataset, a performance metric (e.g. accuracy or area under the ROC curve), and possible hyper-parameter values (e.g. number of states or predictor variables to include).</p> <p>Result: An estimate of the hyper-parameters that lead to best predictive performance.</p> <p>for <i>each row in design matrix of hyper-parameter sets</i> do</p> <p style="padding-left: 2em;">Sample data into k (e.g., 10) groups;</p> <p style="padding-left: 2em;">for <i>each k</i> do</p> <p style="padding-left: 4em;">Set group k as a testing set and everything else as a training set;</p> <p style="padding-left: 4em;">Evolve model on training set with Algorithm 1;</p> <p style="padding-left: 4em;">Predict testing set and compare to actual values based on performance metric;</p> <p style="padding-left: 2em;">end</p> <p style="padding-left: 2em;">Calculate average performance across all predictions on all k testing sets;</p> <p>end</p> <p>Return hyper-parameter set with best average performance;</p>

Algorithm 2: Use cross-validation to optimize FSM hyper-parameters for predictive performance.

play the simultaneous-move PD game in Fig. 2.3, observe the choice of the other person, and then play the same simultaneous-move game again. Even in the (finitely) repeated version, no cooperation can be achieved by rational income maximizers. This tension between maximizing collective and individual gain is representative of a broad class of social situations [18], e.g. the “tragedy of the commons” [38].² One example of a repeated Prisoner’s Dilemma game is global emissions reductions choices made by the United States and China, two of the largest emitters of the pollutants that cause the climate to warm. In the game, there are four outcomes corresponding to the two players each choosing whether to cooperate (reduce emissions) or defect (keep emissions high). Reducing emissions costs a country economic growth and *not* reducing emissions incurs the costs associated with the

²Ininitely repeated games may be a more appropriate model for some real-world cooperation situations. “A model should attempt to capture the features of reality that the players perceive; it should not necessarily aim to describe the reality that an outside observer perceives, though obviously there are links between the two perceptions. Thus the fact that a situation has a horizon that is in some physical sense finite (or infinite) does not necessarily imply that the best model of the situation has a finite (or infinite) horizon. A model with an infinite horizon is appropriate if after each period the players believe that the game will continue for an additional period, while a model with a finite horizon is appropriate if the players clearly perceive a well-defined final period” [18], p. 135. However, in infinitely repeated games, analytically deriving equilibria lacks explanatory and predictive power because there are so many potential equilibria [18].

	C	D
C	(3,3)	(0,4)
D	(4,0)	(1,1)

Figure 2.3: Payoff table for a Prisoner’s Dilemma game. C is cooperate and D is defect.

negative impacts of future climate change due to the emissions. Global warming is a logarithmic function of the concentration of greenhouse gases in the atmosphere, so if we have two equal emitters and one cuts emissions 100%, while the other does not cut emissions, that only cuts future warming by roughly 30%, which will not solve the climate change problem. The costs of cutting emissions are high and could contribute to a trade imbalance because manufactured goods in the country reducing emissions are more expensive than from the country maintaining the status quo, due to the higher cost of energy inputs in an economy that is reducing emissions. This could cause increased emissions from the country that is not reducing emissions, which could further offset the cuts by the country taking emissions reductions. The combination of the cost to the economy of reducing emissions could be greater than the economic benefits from reducing future environmental consequences of global warming by 30%. Therefore, either country is worse off if they reduce their emissions while the other country does not.

Another important strategic situation is a coordination game. Coordination games are similar to Prisoner’s Dilemma games but the important difference is that there are two equilibria in a coordination game and the social dilemma is how to coordinate between the players to choose the equilibrium. The Prisoner’s dilemma has only one equilibrium – mutual defection – which has everyone worse off than they would be if they cooperated. My method for estimating decision models based on behavior in games applies to both games but I focus on the Prisoner’s Dilemma in this dissertation because there is much more data publicly available on play in the Prisoner’s Dilemma.

I applied the procedure to data from laboratory experiments on human subjects playing IPD games for real financial incentives. I gathered and integrated data from many experiments [39], conducted and analyzed by Bereby-Meyer, Roth, Duffy, Ochs, Kunreuther, Silvasi, Bradlow, Small, Dal Bo, Frechette, Fudenberg, Rand, and Dreber [40, 41, 42, 43, 13]. All of the experiments share the same underlying repeated Prisoner’s Dilemma structure, although the details of the games differed. My data set comprises 135,388 cooperation decisions, which is much larger than previous studies of repeated game strategies.

Fudenberg, Rand, and Dreber [13] and Dal Bo and Frechette [43] modeled their IPD experimental data with repeated game strategies; however, they applied a maximum likelihood estimation process to estimate the prevalence of a relatively small predefined set of strategies. In contrast, my estimation process automatically searches through a very large parameter space that includes all possible strategies up to a given number of states and does not require the analyst to predefine any strategies, or even understand the game.

I used 80% of the data for training and reserved the other 20% as a hold-out test set. Fig. 2.4 shows different representations of the the fittest two-state machine of a GA population evolved on the training data: The raw Gray-encoded and binary string (Fig. 2.4a), the bitstring decoded into state matrix and action vector form (Fig. 2.4b), and the corresponding graph representation (Fig. 2.4c). I measure variable importance (Fig. 2.4d) by switching each value of an estimated model’s state matrix to another value in its feasible range, measuring the decrease in goodness of fit to the training data, normalizing the values, then summing across each column to estimate the relative importance of each predictor variable (in this case, the moves each player made in the previous turn).


Fig. 2.5 illustrates the GA run that evolved the FSM of Fig. 2.4 by predicting cooperation decisions in IPD training data games. This GA run, which only took a few seconds on a modest laptop, used common algorithm settings: a population of 175 FSMs initialized with random bitstrings. If the analyst has an informed prior belief about the subjects’ decision models, she can initialize the population with samples drawn from that prior distribution,

raw string: 0, 1, 0, 0, 1, 0, 0, 0, 0, 0 Action vector: 1, 2
 encoded av, sm: 0, 1 0, 0, 1, 0, 0, 0, 0, 0
 decoded av, sm: 0, 1 0, 0, 1, 1, 1, 1, 1, 1 State matrix:


State	cc	dc	cd	dd
C (1)	1	2	2	2
D (2)	1	2	2	2

(a) Raw string; string split into Gray-encoded action vector (av) and state matrix (sm); and decoded binary representations. Decoded sm consists of column-wise elements that index into the action vector to determine the action for each state.

(b) Bitstring decoded into action vector and state matrix. For this action vector, 1 corresponds to cooperation (C) and 2 to defection (D). Columns of state matrix correspond to the observed behaviors at $t - 1$ and rows correspond to the state.



(c) Corresponding graph representation.



(d) Relative variable importance. Most important is 100.

Figure 2.4: A finite state machine estimated on 108,305 decisions with 82% accuracy on a hold-out test data set of 27,083 decisions. Transitions that would be accessible in strictly deterministic play are represented in boldface and inaccessible transitions in italic. Because the human players did not follow exact deterministic strategies, and the italicized transitions were taken in simulating this model with actual game play, the values of these transitions were identifiable.

but this chapter focuses on deriving useful results from random initializations, corresponding to uniform priors, where the analyst only provides data. A linear-rank selection process used the predictive ability of individuals to select a subset of the population from which to create the next generation. A single-point crossover process was applied to the binary values of selected individuals with 0.8 probability, uniform random mutation was conducted with probability 0.1, and the top 5% fittest individuals survived each generation without crossover or mutation, ensuring that potentially very good solutions would not be lost [35]. These are standard GA parameter settings and can be adjusted if convergence is taking particularly long for a given dataset, which would likely only be an issue when the GA run is repeated within a cross-validation procedure searching through a large number of

hyperparameter sets.

This process could be embedded in a more general behavioral estimation process with two hyperparameters, α and β : Instead of selecting the single best FSM from the final iteration of the GA, I could select a set of α solutions from the entire population of the last generation that had the highest α values of the fitness function, which adds no additional computational cost. α controls *breadth of cognitive capacity*: the number of unique FSMs an agent utilizes to make a decision. Thus, I could model an agent as using multiple FSMs to make recommendations about what action to choose. For implementation in a long-run agent-based simulation, choices between FSMs could be made as in [15], where agents use reinforcement learning to update attractions to repeated game strategies. The parameter β would control *depth of cognitive capacity*, the maximum complexity of each FSM, by setting the number of states (rows in the matrices).

Using theoretical agent-based simulations and a fitness measure that is a function of simulated payoffs, Axelrod [44] demonstrated the fitness of the tit-for-tat (TFT) strategy. Using a fitness measure that is a function of the ability to explain human behavior, I discovered a hybrid of TFT and grim trigger (GT), which I call noisy grim (NG). TFT's state is determined solely by the opponent's last play. GT will never exit a defecting state, no matter what the opponent does.

With traditional repeated game strategies such as TFT and GT, the player always takes the action corresponding to her current state (boldface transitions in Fig. 2.4c), but if I add noise to decisions so the player will sometimes choose the opposite action from her current state (italic transitions in Fig. 2.4c), then the possibility arises for both the player and opponent to cooperate when the player is in the defecting state (i.e. to reach the second row first column position of the state matrix in Fig. 2.4b). This would return the player to the cooperating state (see, e.g., Chong and Yao [45]). Noisy grim's predictions on the hold-out test data are 82% accurate, GT's are 72% accurate, and TFT's are 77% accurate. I also tested 16 other repeated game strategies for the IPD from Fudenberg, Rand, and Dreber




Figure 2.5: Proportion of correct predictions of human decisions in IPD games by an evolving population of decision models (strategies). On average, a random choice strategy will have 50% accuracy. Shaded area is the difference between the median of the population of strategies and the best individual strategy.

[13]. Their accuracy on the test set ranged from 46% to 77%. My method uncovered a deterministic dynamic decision model that predicts IPD play better than all of the existing theoretical automata models of IPD play that I am aware of and has interesting relationships to the two most well-known models: TFT and GT.

This process has allowed us to estimate a highly interpretable decision model (fully represented by the small image of Fig. 2.4c) that predicts most of the behavior of hundreds of human participants, merely by plugging in the dataset as input. I address the potential concern that the process is too tuned to this specific case study by inputting a very different dataset from the field of international relations and obtaining useful results. However, before moving to more empirical data—where the data-generating process can never be fully known—to test how robustly the method can estimate a known model, I repeatedly

simulate a variety of known data-generating mechanisms and then apply the method to the resulting choice data.

2.5 Simulated Data

In the real world, people rarely strategically interact by strictly following a deterministic strategy [45]. Strategic randomization, inattention, or error may induce a player to choose a different move from the one dictated by her strategy. To study whether the method could determine an underlying strategy that an agent would override from time to time, I followed the approach of Fudenberg, Rand, and Dreber [13] and created an agent-based model of the IPD in which agents followed deterministic strategies, such as TFT and GT, but made noisy decisions: At each time period, the deterministic strategy dictates each agent's preferred action, but the agent will choose the opposite action with probability p , where p ranges from 0 (perfectly deterministic play) to 0.5 (completely random play). The noise parameter, p , is constant across all states of a strategy of a particular agent for any given simulation experiment conducted.


When a player follows an unknown strategy, characterized by latent states, discovering the strategy (the actions corresponding to each state and transitions between the states) requires observed data that explores as much as possible of the state transition matrix defined by all possible combinations of state and predictor values (for these strategies the predictors are the history of play). Many deterministic strategy pairings can quickly reach equilibria in which players repeat the same moves for the rest of the interaction. If the player and opponent both use TFT and both make the same first move, every subsequent move will repeat the first move. If the opponent plays GT, then after the first time the player defects the opponent will defect for the rest of the session and the data will provide little information on the player's response to cooperation by the opponent. However, if the opponent plays with noise, the play will include many instances of cooperation and defection by the opponent, and will thus sample the accessible state space for the player's strategy more

thoroughly than if the opponent plays deterministically. Indeed, this is why Fudenberg, Rand, and Dreber [13] added noise to action choices in their human subjects experimental games.


In general, the ability of the analyst to discover the underlying decision model of an agent is directly affected by the play of the other agents because it is the interactive play that determines the emergent realized data. The implication is that even if a researcher has a lot of perfectly measured data on individual decisions an agent has taken and perfectly measured covariates that deterministically predict those decisions, the researcher's ability to estimate the underlying decision model is limited by the situations that other agents put the agent in.

I simulated approximately 17 million interactions, varying paired decision models of each agent [(*TFT*, *TFT*), (*TFT*, *GT*), (*GT*, *TFT*), (*GT*, *GT*)] and also varying the noise parameter p (0, 0.025, ... , 0.5) for each of two noise conditions: where both players made equally noisy decisions, and where only the opponent made noisy decisions while the player under study strictly followed a deterministic strategy. I ran 25 replicates of each of the 168 experimental conditions, with 4,000 iterations of game play for each replicate, and then applied the FSM estimation method to each replicate of the simulated choice data to estimate the strategy that the agent player under study was using.

Being in state/row k (e.g. 2) corresponds to the player taking action k (e.g. D) in the current turn. All entries in row k corresponding to the player taking action k in the current period (e.g. columns 2 and 4 for D) are identifiable. Entries in row k that correspond to not taking action k in the current period (e.g. columns 1 and 3 for row 2) represent transitions that cannot occur in strictly deterministic play, so their values cannot affect play and thus cannot be determined empirically. I take this into account when testing the method's ability to estimate underlying deterministic models: this is why only 6 elements of a 10-element TFT or GT matrix can be identified (Fig. 2.6). I also take this into account when estimating models from empirical data, where the data-generating process is assumed to be stochastic:



(a) Effect of decision models and noise in choice on ability to estimate true model when a noisy player faces an equally noisy opponent. Color indicates the pair of decision models used by the player and opponent. For noise probabilities between 0.025 and 0.325 the estimated model always exactly matches the true model.



(b) Comparison of a purely deterministic player facing a noisy opponent versus a noisy player facing a noisy opponent. As noise approaches 0.5, the decision model's impact on actual choices vanishes, so the model error rises sharply when the player's decisions are noisy.

Figure 2.6: Role of noise in correctly identifying the player's strategy. Points have been jittered for visualization. Lines are LOESS-curves. Noise represents the probability, in any period, that an agent will take the opposite action from its decision model. Both the GT and TFT strategies can be represented as a 10 element matrix, of which 6 elements have identifiable values. The model error represents the number of mismatched elements between the estimated and true matrices, so it ranges from 0 (best) to 6 (worst).

each element of the matrix that would be inaccessible under deterministic play is identified, and the fitness is calculated with a strategy matrix in which that element is changed to its complement (“flipped”). If flipping the element does not change the fitness, then the two complementary strategies are indistinguishable and the element in question cannot be determined empirically. If each element decreases the fitness when it is flipped, then the strategy corresponds to a deterministic approximation of a stochastic process and all of the elements of the state matrix can be identified.

When the noise parameter was zero, most of the models estimated by the GA had at least two incorrect elements. However, for moderate amounts of noise ($p = 0.025$ – 0.325), all of the models estimated by the GA were correct (see Fig. 2.6a). For noise levels above $p = 0.325$ in the player, the amount of error rose rapidly with p , as expected because at $p = 0.5$ the action the player chooses moves completely at random so there is no strategy to discover. When a strictly deterministic player faced a noisy opponent, the GA correctly identified the player’s strategy for all noise levels above $p = 0.025$ (see Fig. 2.6b).

2.6 Observational Data

In order to extend this method to more complex situations, the predictor variables (columns of the state matrices) can include any time-varying variable relevant to an agent’s decision. In context-free games such as the IPD, the only predictor variables are the moves the players made in the previous turn, but models of strategic interactions in context-rich environments may include other relevant variables.

It is difficult to interpret graphical models with more than four predictors, but an analyst who had many potentially relevant predictor variables and was unable to use theory alone to reduce the number of predictors sufficiently to generate easily interpretable models with the method could take four courses of action (listed in order of increasing reliability and computation time):

1. Before FSM estimation, apply a (multivariate or univariate) predictor variable selec-

tion method.

2. Before FSM estimation, estimate an arbitrary predictive model that can produce variable importance rankings and then use the top $p < 4$ predictors for FSM estimation.
3. After FSM estimation with $p \geq 4$ predictors, inspect the returned predictor variable importance ranking, and remove all but the top $p < 4$ from her dataset and re-run estimation.
4. Conduct FSM estimation with all combinations of $p < 4$ predictors out of all relevant predictors and choose the estimated model with the best performance (usually highest out-of-sample accuracy).

I illustrate the use of extra predictor variables by applying the method to an example from international relations involving repeated water management-related interactions between countries that share rivers. I use data compiled by Brochmann [46] on treaty signing and cooperation over water quality, water quantity, and flood control from 1948–1999 to generate a model for predicting whether two countries will cooperate. I used three lagged variables: whether there was water-related conflict between them in the previous year, whether they cooperated around water in the previous year, and whether they had signed a water-related treaty during any previous year. This data set was too small to divide into training and hold-out subsets for assessing predictive accuracy, so I report models' accuracy in reproducing the training data (a random choice model is 50% accurate). A two-state decision model (Fig. 2.7a) is 73% accurate, a three-state model (Fig. 2.7c) is 78% accurate, and a four-state model is 82% accurate, but its complexity makes it difficult to interpret visually so it is not shown.

Accuracy can be a problematic measure when the classes are imbalanced, i.e. if a class the model is trying to predict is rare. Many alternatives to accuracy are available that illuminate different aspects of predictive power. For instance, precision is the proportion of (cooperation) event signals predicted by the models that are correct and recall is the pro-

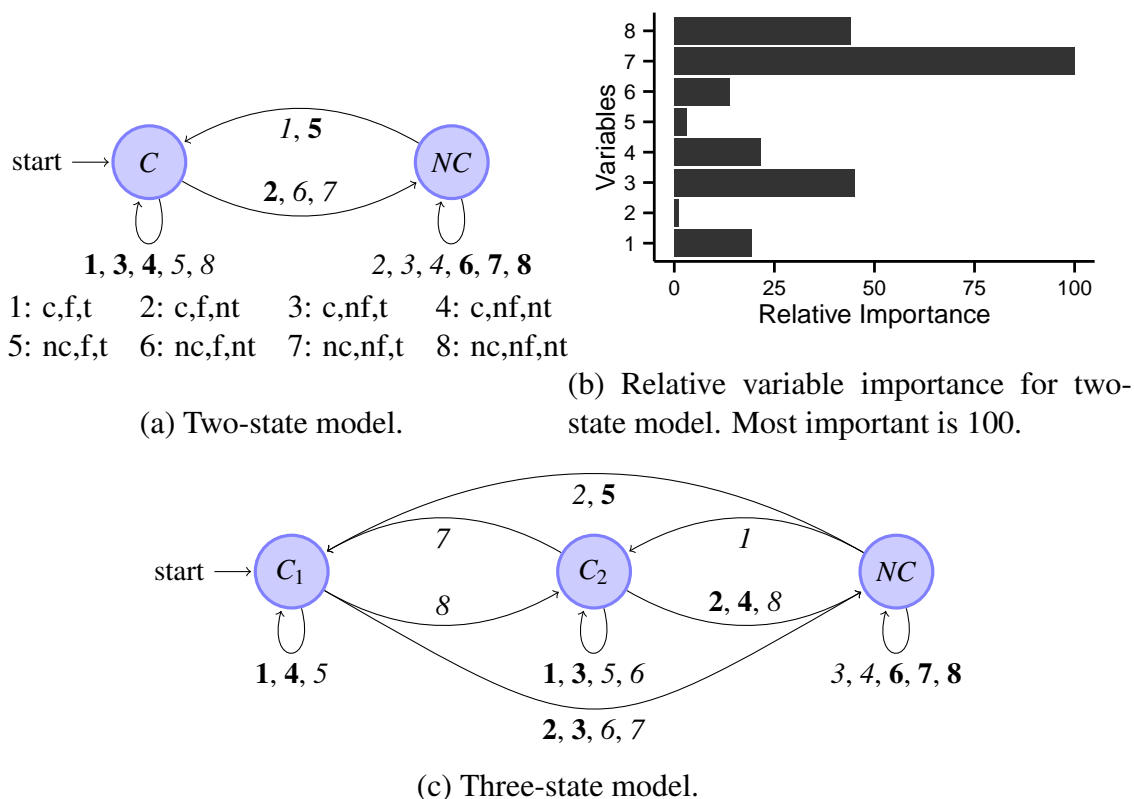


Figure 2.7: Models of cooperation over water issues for sets of countries. The italicized transitions would not be feasible outcomes if this strategy was played deterministically, but they are realized in the data. The inputs (c,nc) represent cooperation and non-cooperation, respectively; (f,nf) conflict and non-conflict; and (t,nt) signing and not signing a treaty.

portion of events that are predicted by the models. For this subset of the dataset, cooperate and defect were almost evenly distributed and to maintain a comparison to the experimental and simulated data I used accuracy as the fitness measure.

In the two-state model, whether or not the countries cooperated in the previous year, the combination of conflict and treaty-signing in the previous year always produces cooperation, whereas conflict without treaty-signing in the previous year always produces non-cooperation. In the three-state model, three of the four outcomes that include conflict lead to a transition from non-cooperation to cooperation, and four of the six outcomes that cause transitions from cooperation (states C_1 and C_2) to non-cooperation are non-conflict outcomes. While this does not tell us something decisive about the role of conflict, it sug-

gests that there may be a counter-intuitive role of conflict in promoting cooperation. [46], using a bivariate probit simultaneous equation model, has a similar finding: “In the aftermath of conflict, states may be particularly eager to solve important issues that could cause future problems” (p. 159).

2.7 Discussion

This chapter outlined a method for estimating interpretable models of dynamic decision-making. By estimating a global, deterministic, simple function for a given dataset, imposing constraints on the number of predictor variables, and providing options for reducing the number of predictor variables, the process facilitates capturing a significant amount of information in a compact and useful form. The method can be used for designing empirically grounded agent models in agent-based simulations and for gaining direct insight into observed behaviors of real agents in social and physical systems. Combining state matrices and a genetic algorithm has proven effective for simulated data, experimental game data, and observational international relations data. With the simulated data, I successfully recovered the exact underlying models that generated the data. With the real data, I estimated simple deterministic approximations that explain most of the structure of the unknown underlying process. I discovered a theoretically interesting dynamic decision model that predicted IPD play better than all of the existing theoretical models of IPD play that I’m aware of.

Going forward, I will explore in greater depth the conditions under which we can recover known data-generating processes. For instance, I would like to map out the multi-dimensional limits of the ability of the analyst to estimate strategic models from observed data by exploring interactions between dozens of different decision models beyond GT and TFT, and continuing to vary randomness. It will be interesting to see if there are types of strategies that, when playing other types of strategies, are more easily estimated. There may be particular classes of strategies that are empirically indistinguishable under the majority

of conditions, but have important differences when placed in particular strategic environments.

I have released an open-source software package (<https://github.com/JohnNay/datafsm>) that implements the methods described here to estimate any time series classification model that uses a small number of binary predictor variables and moves back and forth between the values of the outcome variable over time. Larger sets of predictor variables can be reduced to smaller sets by applying one of the four methods outlined in Section 2.6. Although the predictor variables must be binary, a quantitative variable can be converted into binary by division of the observed values into high/low classes. Future releases of the package may include additional estimation methods to complement GA optimization.

CHAPTER 3 A MACHINE LEARNING APPROACH TO PREDICTING DYNAMIC DECISIONS

3.1 Abstract

The Prisoner's Dilemma has been a subject of extensive research due to its importance in understanding the ever-present tension between individual self-interest and social benefit.¹ A strictly dominant strategy in a Prisoner's Dilemma (defection), when played by both players, is mutually harmful. Repetition of the Prisoner's Dilemma can give rise to cooperation as an equilibrium, but defection is as well, and this ambiguity is difficult to resolve. The numerous behavioral experiments investigating the Prisoner's Dilemma highlight that players often cooperate, but the level of cooperation varies significantly with the specifics of the experimental predicament. I present the first computational model of human behavior in repeated Prisoner's Dilemma games that unifies the diversity of experimental observations in a systematic and quantitatively reliable manner. The model relies on data integrated from many experiments, comprising 168,386 individual decisions. The model is composed of two pieces: the first predicts the first-period action using solely the structural game parameters, while the second predicts dynamic actions using both game parameters and history of play. The model is successful not merely at fitting the data, but in predicting behavior at multiple scales in experimental designs not used for calibration, using only information about the game structure. I demonstrate the power of my approach through a simulation analysis revealing how to best promote human cooperation.

3.2 Introduction

The Prisoner's Dilemma game has been a subject of extensive research due to its importance in understanding the ever-present tension between individual self-interest and so-

¹Parts of this chapter appeared in Nay, J. J., Vorobeychik, Y. (2016) "Predicting Human Cooperation." *PLoS ONE* 11(5): e0155656.

cial benefit. From a theoretical perspective, a strictly dominant strategy (defection), when played by both players, is mutually harmful: cooperation by both yields significant mutual benefits relative to defection. For example, local maintenance of shared drinking water systems in rural communities represents a Prisoner’s Dilemma that can result in a “tragedy of the commons” [38]. From each community member’s perspective, they are better off if someone else invests in maintaining the infrastructure. If the majority of the community adopts this strategy, everyone is worse off because the system breaks down and no longer provides clean water.

In most social dilemma settings, however, interactions are repeated. Thus, for example, community members must repeatedly make water infrastructure investment decisions. Repetition of the Prisoner’s Dilemma, a more realistic model of human interaction than a one-shot game, can theoretically give rise to cooperation as an equilibrium if players are sufficiently “patient”; still, defection remains an equilibrium as well, and this ambiguity is difficult to resolve. In particular, theoretical treatment of repeated Prisoner’s Dilemma games is not instructive in identifying when cooperation or defection emerges as the predominant outcome. Given the limitations of theory in explaining repeated cooperation, researchers have turned to experiments to better understand behavior and the effects of institutional structure on social outcome by considering different game structures and investigating associated cooperation proclivities of human subjects. The experiments highlight that humans often cooperate, but the overall level and temporal evolution of cooperation vary significantly with the specific design.

I develop a predictive model of dynamic cooperation that reliably forecasts behavior across heterogeneous game designs, and then analyze this model to tease apart the magnitude and direction of the effects of game design variables on cooperation. For this purpose I compiled data from previously analyzed repeated Prisoner’s Dilemma experiments. This dataset was expanded beyond the data used in Chapter 2. I created standardized measures of the game and individual behavior across these games, and used machine learning tech-

niques to calibrate and evaluate computational models. My model is successful in predicting individual decisions, average cooperation levels, and cooperation dynamics *in games not used for model calibration*. Moreover, I demonstrate that this model can predict the high-level quantitative and qualitative findings of the human subject experiments.

The long-term goal of this research program is to map the experimental variables onto real-world policy design factors and use model analyses to inform policies that facilitate cooperation where the underlying social structure would otherwise lead to a breakdown. For instance, how can we best design development programs that lead to sufficient voluntary maintenance of shared water systems? Is it more important to increase the potential benefits of mutual cooperation over mutual defection, or to increase the benefits of mutual cooperation over losing out by being the sole cooperator?

3.3 Data

The data are from human subjects experiments that used real financial incentives and transparently conveyed the rules of the game to the subjects, which is standard procedure in experimental economics. Subjects anonymously interact and their decisions to cooperate or defect at each time period of each interaction are recorded. They receive payoffs proportional to the outcomes in a specified payoff table similar to Table 3.1. From the description of the experiments in the published papers and the publicly available data sets, I was able to build a comprehensive collection of game structures and individual decisions.

The thirty game structures that I compiled varied substantially across a number of dimensions, aside from player payoffs. In some structures, payoffs were deterministic, whereas others featured stochastic payoffs (in this case, the expected payoffs constituted the payoff structure). In some structures, players imperfectly observed their counterparts' past actions. Another key distinction was whether or not a game had a fixed time horizon, or would terminate independently after each iteration with a fixed probability. Finally, while most games were played over a discrete sequence of iterations, some were in contin-

uous time. I use nine variables to quantify game structure along these salient dimensions. *Risk* is an indicator of whether there is stochasticity in the payoffs [40, 42]. *Error* is the probability that the choice a player makes will be exogenously flipped [13]. *Infinite* is an indicator of whether interactions are indefinitely repeated or have a fixed length [47]. δ is the probability that the next period of the current paired interaction will occur in an infinitely repeated game [43]. I used a formula, $E[\text{InteractionLength}] = \frac{1}{1-\delta}$, to compute δ for finitely repeated interactions; for instance, the finitely repeated interactions in [42] were all ten periods long so $\delta = 0.9$. *Continuous* is an indicator of whether interactions are played in “continuous time,” rather than the standard discrete rounds [48]. R is the reward received if both players cooperate; P is the punishment received if both defect; T is the temptation to defect on the other; and S is the payoff for being a sucker by cooperating as the other defects (Table 1 illustrates the way in which the four payoff values map onto the Prisoner’s Dilemma bi-matrix representation).

	C	D
C	(R,R)	(S,T)
D	(T,S)	(P,P)

Table 3.1: Payoff table. Payoff table where one player plays from the perspective of the columns and the other from the rows. For this to be a repeated Prisoner’s Dilemma, it must hold that $T > R > P > S$, and $R > (S + T)/2$ [49].

To create standardized payoff measures from the R, S, T, P values, I used two differences between payoffs associated with important game outcomes, both normalized by the difference between the temptation to defect and being a sucker when cooperating as the other defects [50].

$$r_1 = \frac{R - P}{T - S} \tag{3.1}$$

is the normalized difference between the reward received if both players cooperate and the

punishment received if both defect.

$$r_2 = \frac{R - S}{T - S} \quad (3.2)$$

is the normalized difference between the reward received if both players cooperate and the payoff for being a sucker when cooperating as the other defects. Because

$$\frac{\partial r_1}{\partial R} > 0 \quad (3.3)$$

$$\frac{\partial r_1}{\partial P} < 0 \quad (3.4)$$

$$\frac{\partial r_1}{\partial T} < 0 \quad (3.5)$$

$$\frac{\partial r_1}{\partial S} > 0 \quad (3.6)$$

r_1 has been used as an index of the cooperativeness of a Prisoner's Dilemma [50, 51], while r_2 is descriptive of how much better off a player will be if their opponent cooperates, rather than defects, while they themselves cooperate. Table 3.2 summarizes the game structures from the data sets I standardized and combined.

For the fEWA model (see details of this model below), which uses the raw payoffs, I converted payoffs to the equivalent amount in U.S. dollars that the participant in the experiment would have received: DF 1 payoff = \$0.006; DO 1 payoff = \$0.01; FR 1 payoff = \$0.033; KS 1 payoff = \$0.10; BR 1 payoff = \$1.0; AM 1 payoff = \$1.0; and FO 1 payoff = \$0.006 ($\frac{\$0.05}{8}$).

Error	Delta	Infinite	Continuous	Risk	r1	r2	Avg. Cooperation	Dataset
0.0000	0.900	0	0	0	0.18	0.590	0.60	BR
0.0000	0.900	0	0	1	0.18	0.590	0.35	BR
0.0000	0.900	1	0	0	0.33	0.670	0.56	DO

Error	Delta	Infinite	Continuous	Risk	r1	r2	Avg. Cooperation	Dataset
0.0000	0.900	0	0	1	0.33	0.830	0.31	KS
0.0000	0.900	0	0	0	0.33	0.830	0.57	KS
0.0000	0.500	1	0	0	0.18	0.530	0.10	DF
0.0000	0.750	1	0	0	0.18	0.530	0.20	DF
0.0000	0.500	1	0	0	0.39	0.740	0.18	DF
0.0000	0.750	1	0	0	0.39	0.740	0.59	DF
0.0000	0.750	1	0	0	0.61	0.950	0.76	DF
0.0000	0.500	1	0	0	0.61	0.950	0.35	DF
0.1250	0.875	1	0	0	0.20	0.600	0.34	FR
0.1250	0.875	1	0	0	0.33	0.660	0.49	FR
0.1250	0.875	1	0	0	0.43	0.710	0.59	FR
0.0000	0.875	1	0	0	0.60	0.800	0.74	FR
0.0625	0.875	1	0	0	0.60	0.800	0.78	FR
0.1250	0.875	1	0	0	0.60	0.800	0.57	FR
0.0000	0.900	0	0	0	0.25	0.583	0.43	AM
0.0000	0.875	0	1	0	0.11	0.560	0.27	FO
0.0000	0.875	0	1	0	0.14	0.710	0.33	FO
0.0000	0.875	0	1	0	0.33	0.560	0.54	FO
0.0000	0.875	0	1	0	0.43	0.710	0.62	FO
0.0000	0.500	0	0	0	0.33	0.610	0.12	DB
0.0000	0.750	0	0	0	0.33	0.610	0.24	DB
0.0000	0.750	0	0	0	0.33	0.720	0.25	DB
0.0000	0.500	0	0	0	0.33	0.720	0.13	DB
0.0000	0.500	1	0	0	0.33	0.610	0.23	DB
0.0000	0.750	1	0	0	0.33	0.610	0.35	DB

Error	Delta	Infinite	Continuous	Risk	r1	r2	Avg. Cooperation	Dataset
0.0000	0.750	1	0	0	0.33	0.720	0.36	DB
0.0000	0.500	1	0	0	0.33	0.720	0.31	DB

Table 3.2: Summary of thirty game structures that compose the full combined data set [52, 47, 40, 41, 42, 43, 48, 13]. BR 2006 [40] and DB 2005 [47] both also conducted one-shot games; I only describe and use their repeated game data. KSBS 2009 [42] also conducted games with partial information; I only describe and use their full information data. AM 1993 [52] also conducted games that matched humans with computers; I only describe and use the games they conducted where humans played other humans. FO 2012 [48] included one-shot games and games with very different protocols for how and when to make a choice in order to study continuous choices; I only use the “Grid treatment with $n = 8$ subperiods,” which they say is, “comparable to the 10-stage repeated games featured in previous laboratory studies.” DO 2009 [41] also conducted random matching of opponents; I only use their fixed matching treatments.

Fig. 3.1 plots game structures based on their values of the four quantitative game structure variables, r_1 , r_2 , δ , and *error*, illustrating the broad empirical support in the combined data across the values of these variables, and that there are no *simple* relationships between these variables and the proportion of cooperation that can be detected without, at least, controlling for variables not included in each plot.

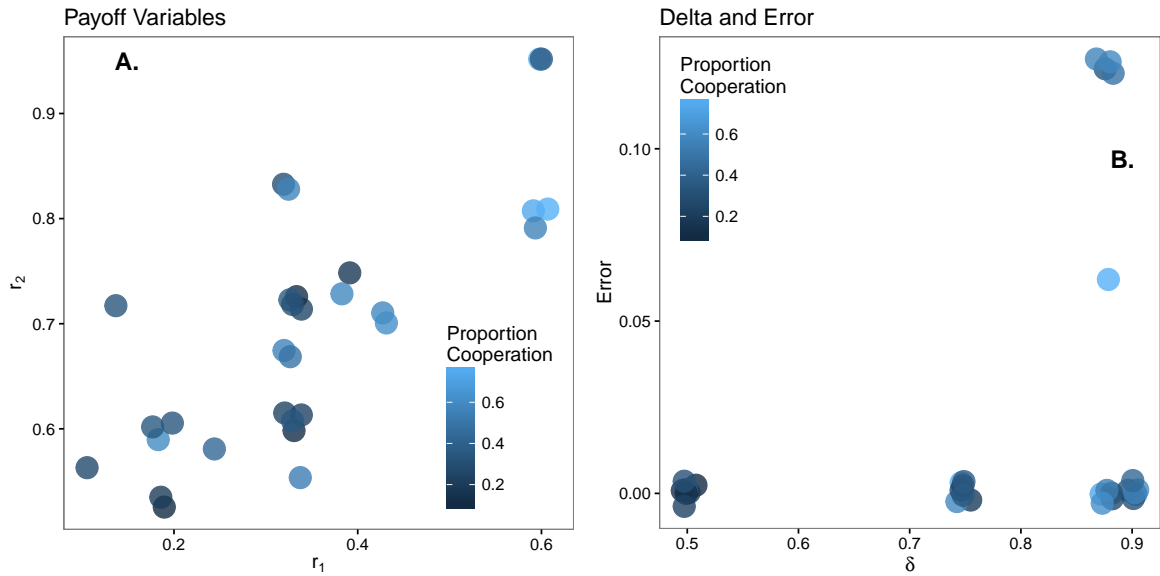


Figure 3.1: Game structures ($n = 30$) with location based on the payoff variable values (A.), and delta and error values (B.). Colors represent proportion of cooperation observed in the game structure. Locations have been slightly randomly shifted to improve visualization.

The combined data set can be organized hierarchically (Fig. 3.2). Within each game structure, there are interactions between pairs of players; these are repetitions of the same “stage-game” between the same two players. Repeating the game with past behavior as common knowledge can theoretically increase cooperation by bringing players’ reputation concerns into play. Within each interaction, there are time periods. Finally, in each time period, both players simultaneously take a single action.

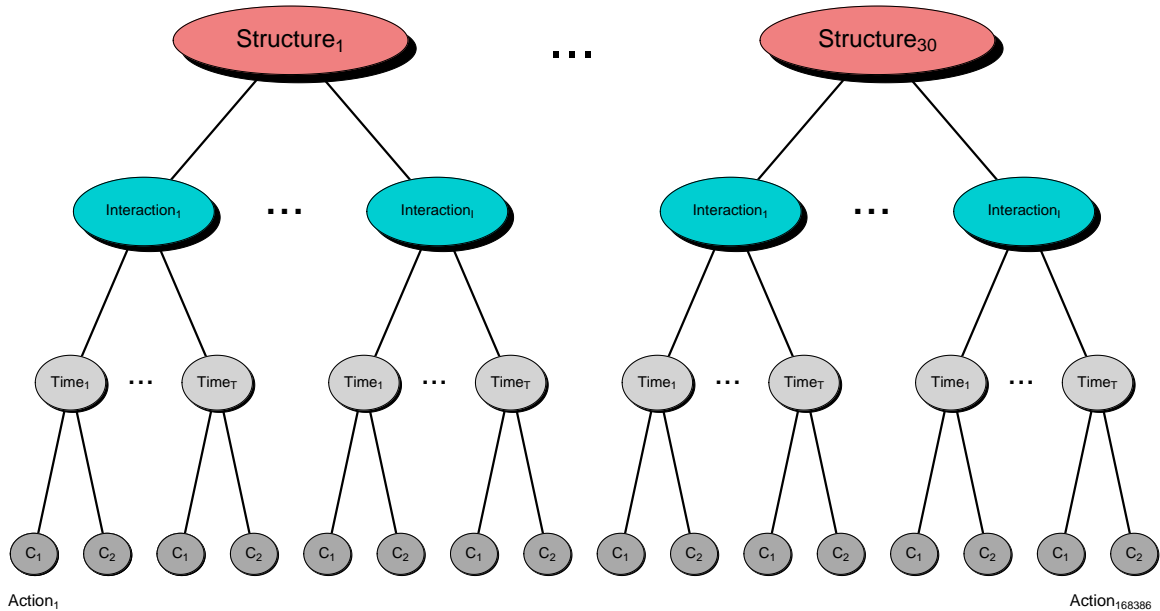


Figure 3.2: A hierarchical view of the data. The top level divides the data into 30 game structures. The next level down are the interactions between two players. Within each interaction, there are T time periods. In repeated games in which termination is stochastic, T ranges up to thirty-eight. Across all interactions and structures, T is five, on average. Within each time period, player 1 takes action C_1 and player 2 takes action C_2 . 168,386 actions were taken across all the experimental data.

My goal was to predict behavioral patterns simultaneously at several levels within this hierarchy. Specifically, my goal is to predict the effects of the game structure on average cooperation (the highest level in the hierarchy), the temporal dynamics of cooperation as a function of structure (second lowest level), and individual-level actions (lowest level). The impact of structure on cooperation has been the primary subject of experimental investigations, with the natural goal of understanding how to design institutions that promote cooperation. Understanding both short-term and long-term impacts of institutions, however, necessitates looking at behavior dynamics, rather than simply aggregate levels of cooperation. Indeed, cooperation may well be high early, but degrade with time, particularly close to the final period of the game, if it is known [40]. Finally, understanding individual behavior enables us to understand aggregate cooperation dynamics in terms of micro decision processes. If my model can successfully predict behavior at all levels in the hierarchy, we can have confidence in the ability of the resulting model to generalize

experimental findings to new institutional structures, allowing us to achieve the ultimate goal: a validated computational framework for understanding and designing institutions that promote cooperation.

3.4 Model

The behavioral model has two parts: a “static” component that predicts a player’s first period action, and a “dynamic” component that predicts a player’s actions in subsequent times of the same interaction. Both components are logistic regressions mapping a vector of predictor variables into the probability of cooperation, with parameters learned through maximum likelihood estimation on training data. The predictor variables for first period play include the game structure, \vec{Game} , $(r_1, r_2, risk, error, \delta, infinite, continuous)$, and the predictors for all other time periods (i.e., the dynamic model) include the game structure, the actions of both players from the previous period, $\vec{History}_{t-1}$, and the current time period, t . The inclusion of the history of the interaction is motivated by evidence that most participants in repeated cooperation games condition their actions on previous play [53]. In mathematical terms, the probability of cooperation, $p(C_t)$, can be expressed as follows:

$$p(C_t) = \begin{cases} t = 1 & f_{static}(\vec{Game}) \\ t > 1 & f_{dynamic}(\vec{Game}, \vec{History}_{t-1}, t) \end{cases}$$

where f is determined by the logistic regression model (distinct in both cases), calibrated on behavioral data.

Specifically, I used the following equation for first period cooperation, $C_{t=1}$:

$$r_1 + r_2 + risk + error + \delta + r_1 \times \delta + r_2 \times \delta + infinite + continuous$$

I used the following equation for cooperation in periods greater than one, $C_{t>1}$:

$$r_1 + r_2 + risk + error + \delta + r_1 \times \delta + r_2 \times \delta + infinite + continuous + \delta \times infinite +$$

$$my.decision_{t-1} + other.decision_{t-1} + error \times other.decision_{t-1} + t$$

Both equations use all the structural game features, $r_1 + r_2 + risk + error + \delta + r_1 \times \delta + r_2 \times \delta + infinite + continuous$, and because I hypothesized that δ and the payoff variables may have difference effects depending on the values of the other, I interacted them. For the dynamic model, I added an interaction term between δ and $infinite$ to capture the different effect that δ may have when it actually determines the length of the game probabilistically. In finite games, δ represents a rational expectation of the length of the game from a first period perspective. In infinite games, δ represents a rational expectation of the length of the game for all periods. Therefore, δ is used as a feature in the model of first period play, and for all periods of play beyond period one there is an interaction term that multiplies the indicator variable for whether a game is infinite by the value of δ . I interacted $error$ with $other.decision_{t-1}$ because the greater the value of $error$, the less sure the player is of the actual decision of the other player in the previous time period. The model is then a logistic sigmoid function,

$$\sigma(w^T X) = \frac{1}{1 + \exp(-w^T X)} \quad (3.7)$$

acting on a linear function of these features, \mathbf{X} , with a vector of weights, \vec{w} , the length of the feature set. The computational implementation I used was the base R “`stats::glm`” function and the caret “`train`” function [54, 55].²

²Full Software Information:

R version 3.2.3 (2015-12-10) Platform: x86_64-apple-darwin15.2.0 (64-bit) Running under: OS X 10.11.3 (El Capitan)
 attached base packages: grid stats graphics grDevices utils datasets methods base
 other attached packages: eat_0.0.1 pander_0.6.0 ggplot2_2.0.0 knitr_1.12.3
 loaded via a namespace (and not attached): Rcpp_0.12.3 gtools_3.5.0 digest_0.6.9 plyr_1.8.3
 gtable_0.1.2 formatR_1.2.1 magrittr_1.5 evaluate_0.8
 scales_0.3.0 stringi_1.0-1 rmarkdown_0.9.2 labeling_0.3
 tools_3.2.3 stringr_1.0.0 munsell_0.4.2 yaml_2.1.13

I compare my model’s performance to three alternative logistic regression models: “static-only,” which just uses the first component of the “full” model; “dynamic-only,” which just uses the second component; and “baseline,” which uses the observed average level of cooperation. Comparing a full model to its components allows one to understand the relative contributions of the components to its predictive power. I also compare my model to a state-of-the-art behavioral game theory model designed for forecasting play in out-of-sample games: functional experience-weighted attraction learning (fEWA) [56]. The actions available to agent i , which are indexed by j , are assumed to have numerical attractions for each time t , $A_i^j(t)$, and fEWA updates the attractions based on functions of i ’s experience up to time t and the payoffs of the game (i ’s chosen strategy is $s_i(t)$, i ’s opponent’s chosen strategy is $s_{-i}(t)$, i ’s payoffs are $\pi_i(s_i^j(t), s_{-i}(t))$, and I yields $I(x, y) = 0$ if $x \neq y$, and $I(x, y) = 1$ if $x = y$).

$$A_i^j(t) = \frac{\phi_i(t)N(t-1)A_i^j(t-1) + [\delta_{ij}(t) + (1 - \delta_{ij}(t))I(s_i^j, s_{-i}(t))]\pi_i(s_i^j, s_{-i}(t))}{N(t-1)\phi_i(t) + 1}$$

Then, attractions are mapped into probabilities of choosing Cooperate or Defect the next time period with a logistic stochastic response function.

$$P_i^j(t+1) = \frac{e^{\lambda A_i^j(t)}}{\sum_{k=1}^{m_i} e^{\lambda A_i^k(t)}}$$

In order to use the empirical models of individual behavior to predict interactive outcomes of new experimental designs, I simulate discrete-time dynamic systems comprised of autonomous decision algorithms (agents) that interact with each other. This allows us to simulate the play of an experiment without any behavioral data from that experiment. Player behavior is endogenous to the simulation model, which only needs to be initialized with a game structure specification. There have been a number of studies using simula-

colorspace_1.2-6.htmltools.0.3

tions to investigate cooperation games [44, 57, 58, 59, 60, 61, 62], and simulations have been used to inform institutional design of strategic interactions more broadly [63, 64, 65]. There has been research on cooperative equilibria models for predicting aggregate cooperation patterns [66], and a significant amount of work on individual-level behavioral models [67, 68, 69, 11, 9, 70, 15, 71]. My work diverges from most such research in two respects: (i) agent behavior is derived solely from individual-level empirical data, and (ii) I rigorously validate my model’s ability to predict behavior by measuring performance on many unseen game structures. Wunder, Suri, and Watts [62] also derive agent behavior solely from individual-level game data. However, I utilize data from many more experimental designs and from a different game.

3.5 Results

3.5.1 Individual-level performance

The first investigation evaluates models’ ability to predict individual-level actions. I divide game structures into training and test groups, estimate the parameters in training game structures, and then predict actions in held-out game structures, conditioning on the game structure and the empirically observed actions of the previous period (for periods greater than one). I repeatedly execute the process, each time slightly changing the split of the data so each game structure will be in the test data once (Fig. 3.3). The end results are out-of-sample predictions of all actions in each game structure. To make predictions with the dynamic-only model, which will have missing values for the lagged action outcomes at period one, I draw cooperate/defect actions with equal probability (corresponding, approximately, to the average cooperation/defection split over all game structures). When I instead impute “period zero” outcomes as mutual cooperation the results are qualitatively the same. For this test, I measure the log-likelihood of the observed actions in the test data, given model predictions, which is a statistically proper method for evaluating the quality

of probabilistic predictions. I also discretize model outputs into Cooperate or Defect to measure accuracy, which is the proportion of actions where the predicted probability of cooperation was above (below) 0.5 when the observed action was cooperation (defection).

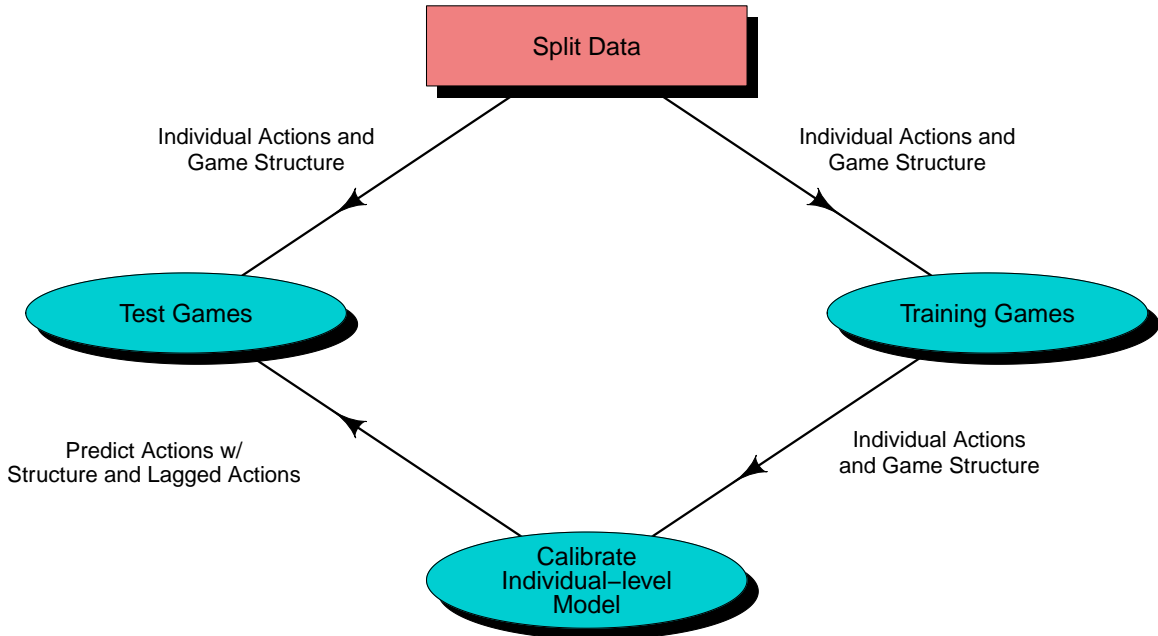


Figure 3.3: Model validation process for individual-level actions. I assign each of the thirty game structures into either training or test data. With the training data, I learn the parameters of the individual-level model, and then predict the decisions in game structures assigned to the test data. I repeat this process thirty times, including a different game structure as the held-out test each time (leave-out-one-cross-validation), until I have predictions for all the decisions for each of the game structures.

The dynamic model performs almost as well as the full model in periods greater than one, but poorly in the first period, indeed, worse than the static model (Table 3.3). Overall, my relatively simple two-piece model predicts the next action a player will take with 86% accuracy on average (a remarkably good prediction, given that human behavior is generally quite noisy). The model also significantly outperforms all alternatives in terms of the log-likelihood measure, which is more statistically appropriate in quantifying performance of stochastic forecasts, but is less intuitive.

3.5.2 Aggregate-level performance

To evaluate the model's ability to predict behavior in new game structures, I developed the following procedure (Fig. 3.4). Assign each of the thirty game structures into either training or test data. With the training data, learn the parameters of the individual-level model. Next, create a simulation in which the estimated individual-level model makes joint decisions in a repeated Prisoner's Dilemma game, and predict the probabilistic behavior in game structures assigned to the test data *using only the game structure*, i.e., using no behavioral data from the experiment. Finally, compare the predictions, $p(y_{sim}^{new} | \vec{Game}^{new})$, to actual observed cooperation dynamics, y_{obs}^{new} , using both squared error and correlation to measure success of the model in predicting behavior. Repeat this process thirty times, including a different game structure as the held-out test each time (leave-out-one-cross-validation), until I have a prediction for each of the game structures as if each prediction were made before any data had been collected for that experimental design. I also test that the results are robust to the number of folds in the cross-validation procedure (from thirty down to two), i.e. robust to the number of game structures used for training.

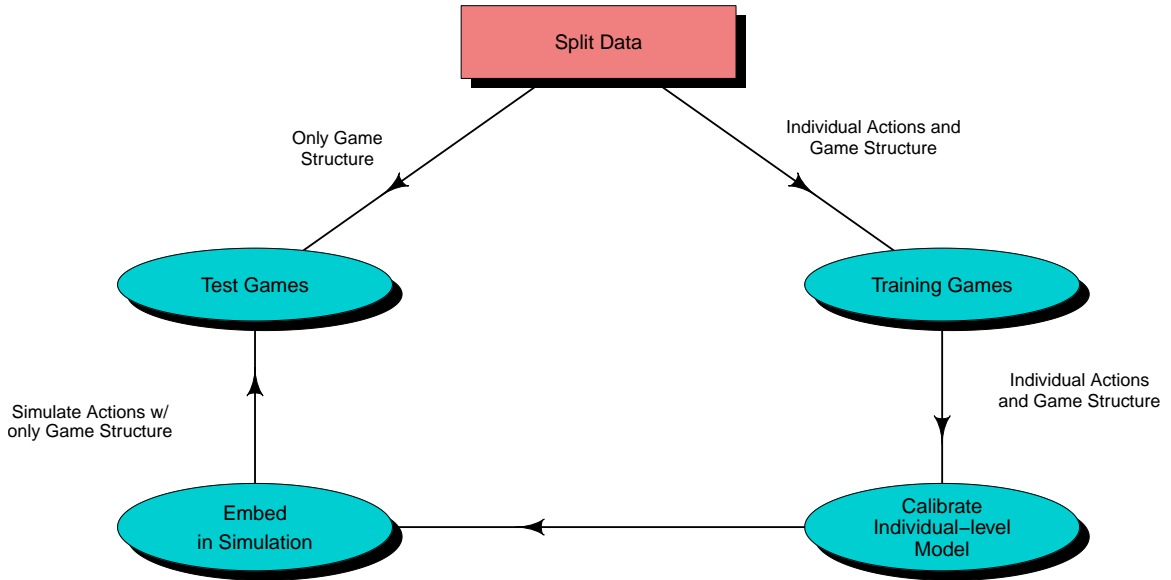


Figure 3.4: Model validation process for aggregate-level patterns. I tested the dynamic-only model by sampling lagged outcomes for ‘period zero’ actions from a Bernoulli distribution with equal probability of cooperation and defection, which is approximately the mean cooperation rate in the data. A subtle, but crucial, distinction between this process and the model validation process for individual-level action predictions (Fig. 3.3) is that, here, *I only pass game structures* for the test games, rather than the full behavioral data and the game structure.

I compare the performance of the five models’ predictions of average probability of cooperation and dynamics of cooperation (Table 3.3). My model is slightly worse at predicting overall cooperation levels than the static model, but better at predicting dynamics (neither comparison is statistically significant), and is significantly better than the other models in almost all cases. The full model outperforms the static model on six measures and the static model outperforms the full model on two measures. This may be due to noise and this demonstrates the importance of using as multiple important measures of performance.

	Full	Static	Dynamic	fEWA	Baseline
Acc. t=1	68	62	57	48	48
Acc. t>1	86	68	85	62	62
LL t=1	-656	-668	-846	-748	-761

	Full	Static	Dynamic	fEWA	Baseline
LL $t>1$	-1624	-2945	-1726	-3108	-3146
Cor-Time	0.755	0.709	0.713	0.241	-0.697
Cor-Avg.	0.774	0.819	0.721	0.106	-0.724
RMSE-Time	0.149	0.154	0.163	0.213	0.224
RMSE-Avg.	0.126	0.113	0.136	0.194	0.203

Table 3.3: Model performance comparison. Best performance for each test is *italicized*. **First four rows** are performance on 32,614 predictions of period one actions and 135,772 predictions of period greater than one actions. Each evaluation is an average for how that model performed with out-of-sample predictions for each game structure. I conduct paired sample t-tests (not assuming equal variances) to determine if the thirty accuracy (Acc.) and log likelihood values (LL) for the full model are statistically greater than the values of the next best model. Accuracies for $t>1$ of the full model ($p = 0.03$) and the likelihoods for $t>1$ of the full model ($p < 0.001$) are significantly higher than the next best model (dynamic). Accuracies for $t=1$ of the full model are greater than the next best model, the static model ($p = 0.07$), while the likelihoods for $t=1$ of the full model are not significantly greater than the likelihoods of the static model ($p = 0.31$). **Last four rows** are performance on average cooperation level in each structure ($n=30$) and time series of average cooperation in each structure ($n=212$). Infinitely repeated interactions with delta set to 0.5 are on average only two periods long and there is not sufficient empirical data to extend out to eight periods so I extend to seven. Two structures are finitely repeated for two periods and two others are finitely repeated for four periods. I conducted paired sample t-tests between the full model and competitors, with a null hypothesis that the true difference in means of the 212 squared errors between predicted and real cooperation levels at all times in all game structures is equal to zero, i.e. that the full model and a competitor are statistically indistinguishable in terms of squared errors on time series predictions. I did the same for the thirty predictions of overall cooperation levels. I reject the null of no difference for all comparisons except with the static model for both tests and the dynamic model for the time series.

Estimating the parameters of the model on a subset of the data and then evaluating the performance of the model on held-out data allows us to measure generalizability. However, randomly dividing the data increases bias of the evaluation of the predictive performance because the estimated value of the predictive power is conditional on which data were included in the training or test samples. To reduce this bias, it is common to run multiple rounds of this process and then average the resulting values of predictive performance [72].

If we do this n times, this is called leave-out-one-cross-validation (LOOCV), which has lower bias; however, LOOCV can have higher variance in the estimates compared to k -fold validation, where $k < n$ [73]. Fig. 3.5 displays the effect of the number of folds in cross-validation on model performance, demonstrating that the main results are robust to the value of k .




Figure 3.5: How predictive performance varies with data splitting. RMSE (A.) and correlation (B.) for time series forecasts of play in 30 game structures, varying folds in cross-validation from 30 to 2. The full model consistently has lower prediction error and higher correlation than the baseline model and the fEWA model until there are only two folds. It is, in general, difficult to make accurate predictions when the ratio of observational units to folds is small. In the case of predicting aggregate and dynamic play, the game structure itself is the observational unit, and I only have thirty, so it's not surprising that performance can degrade at two folds depending on the particular random realization of fold assignments.

Every panel in Fig. 3.6 is the full model's *out-of-sample forecast* for the average probability of cooperation at each time, conditional only on the game structure of that experiment. My model's time series of average cooperation is statistically significantly positively correlated ($0.76, p < 0.001$) with the observed time series. To better understand Fig. 3.6, observe, for example, Structure 14: using no data from that game structure, my model predicted the initial (high) level of cooperation almost exactly and then was perfectly cor-

related with the empirically observed mean cooperation level throughout the next seven periods of play.




Figure 3.6: Model forecasts. Out-of-sample forecasts of cooperation level over time, for all game structures, conditional only on the game structure ($n = 212$).

The dynamic-only model performed nearly as well as the full model on individual-level *period* > 1 actions, but worse on both tests of aggregate pattern predictions. By investigating the coefficients of the estimated individual-level dynamic model, I discover that the actions taken by a player and her opponent in the previous period are highly predictive of the next action (Fig. 3.7). The variable with the most predictive power is the player's own previous action: if a player cooperated (defected) in the previous period, she is very likely to cooperate (defect) in the next. There is strong inertia to Prisoner Dilemma behavior, and, therefore, accurate prediction of first period play is crucial for good performance at the aggregate level. fEWA can incorporate the payoff game structure variables but not the other variables, which prevents high first period accuracy. The full model is able to predict first period play well with a model trained only on first periods in the training data, and then use a dynamic model trained on periods > 1 in the training data, allowing for subtly different relationships between game structures and the evolution of cooperation.




Figure 3.7: Variable importance scores. Variable importance scores for individual-level dynamic component of full model, i.e. for predictions of an agents' probability of cooperation in periods > 1 . Variables separated by '*' represent an interaction between those two variables. These relative importance scores are derived from the absolute values of the t-statistics for each model parameter, which correspond to the effects of the predictor variables (accounting for variability in the estimates) on the probability of cooperation, *ceteris paribus*.

The empirical experiments varied structural game parameters to measure hypothesized differences in cooperation levels between structures. As a final validation, I compared the (out-of-sample) predicted average cooperation levels between my synthetic model of behavior to the actual observed behavior in experiments [52, 47, 40, 42, 43, 48, 13]. Overall, my model came to the same qualitative conclusions as the experiments: δ , infinite and particular payoff configurations increased cooperation, while risk reduced cooperation. I detail each paper's finding and illustrate my model's corresponding finding graphically in Fig. 3.8.

I could not include more than one game structure from two papers (game structures 3 and 18) that comprised the integrated data set [52, 41], because they were comparing either

to one-shot games or games with artificial opponents. Therefore, in my model's replication of the qualitative empirical experimental findings, I could not conduct any replication related to these two papers' findings. Dal Bo and Frechette found that: delta increases cooperation, keeping payoffs fixed (Fig. 3.8 A.); and that certain payoffs increase cooperation, while fixing delta (Fig. 3.8 B.). Bereby-Meyer and Roth found that risk reduces cooperation, where payoffs were framed as gains (Fig. 3.8 C.). Kunreuther et al. found that risk reduces cooperation, with payoffs framed as losses rather than gains (Fig. 3.8 D.). This is the only finding where I predicted marginally different cooperation levels when the empirical data indicates a larger gap. Fudenberg Rand and Dreber found that certain payoffs increase cooperation (Fig. 3.8 E.). Friedman and Oprea found that certain payoffs increase cooperation (Fig. 3.8 F.). Dal Bo found that delta increases cooperation, fixing payoffs and infinite (Fig. 3.8 G.); having an 'infinitely' repeated game increases cooperation, fixing payoffs and delta (Fig. 3.8 G.); and certain payoffs increase cooperation, fixing infinite and delta (Fig. 3.8 H.). Dal Bo also found that the cooperation levels decrease more over time within finite games (Fig. 3.6 Structures 23 - 26), compared to infinite games (Fig. 3.6 Structures 27 - 30).

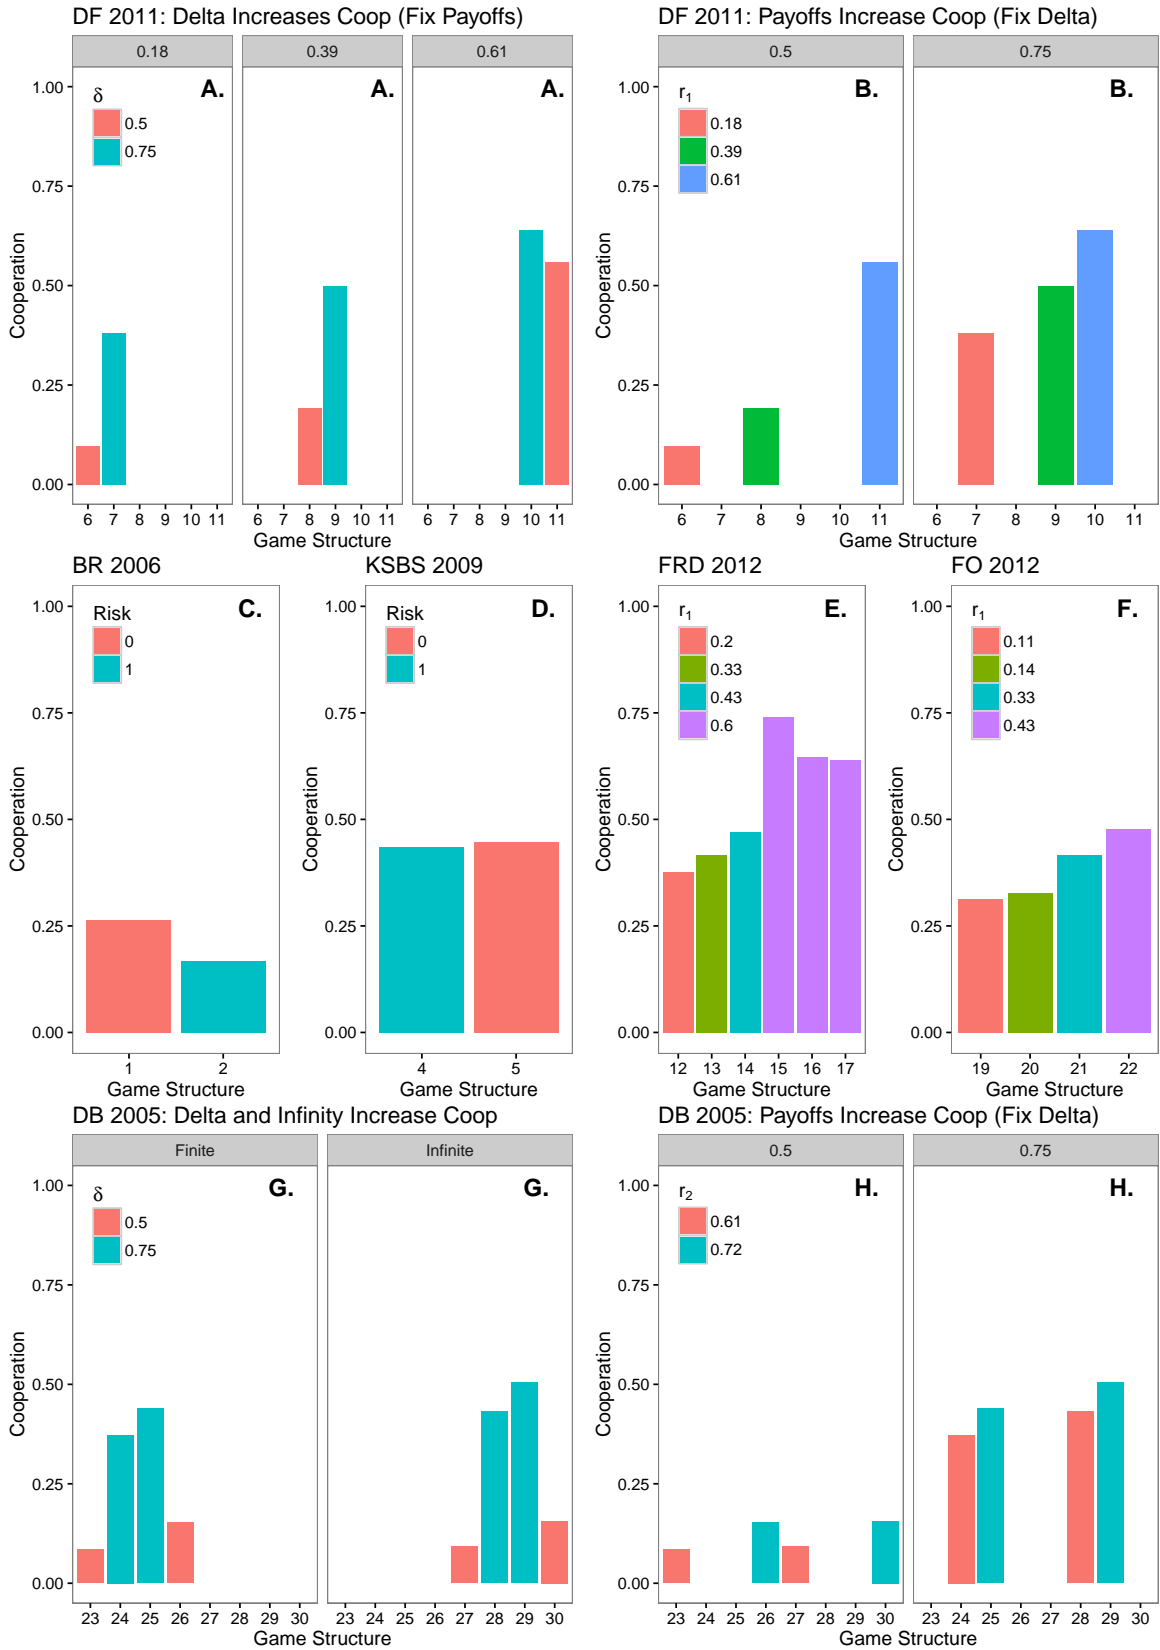


Figure 3.8: Predicted proportions of cooperation ($n = 28$).

3.6 Analysis

After re-learning the computational model with all available data to best explore the full parameter space, I deployed it to quantify the sensitivity of cooperation to each of the structural game design parameters. To systematically explore the model, I generated thousands of collections of input values (specifications of Prisoner's Dilemma experiments) from the multi-dimensional distribution covering the feasible ranges of all input values using Latin Hypercube sampling [74, 75]. The variables are drawn from the following distributions, with the constraint that $r_1 < r_2$ because r_1 is always less than r_2 in the data: $error \sim Unif(0, 0.5)$; $\delta \sim Unif(0.45, 0.95)$; $infinite \sim Bern(0.5)$; $risk \sim Bern(0.5)$; $r_1 \sim Unif(0, 1)$; $r_2 \sim Unif(0, 1)$. Then I simulated cooperation dynamics for each experimental input set. This global sampling and simulation allows subsequent analysis to generate reliable information about the relationships between model inputs (structural game design parameters) and output (cooperation behavior) [76, 77].

Based on the results of a partial rank correlation coefficient analysis [76, 78], the six main game structure variables can be divided into three groups that contain two variables each within the 95% confidence interval of each other (Fig. 3.9A); I obtain qualitatively equivalent results with a standardized rank regression coefficient analysis. δ and r_2 have *very large positive effects* on average cooperation levels. As noted above, my δ measure is applicable to both infinite and finite games as a measure of the expected length of the game from a first period perspective, and the dynamic model has an interaction term between δ and *infinite* that allows the δ effect in periods greater than one to be different for infinite games. Surprisingly, this interaction term is the least important predictor variable in the dynamic model (Fig. 3.7), suggesting that the effect of the expected length of the game from a first period perspective is independent of whether the game is indefinitely repeated.

Infinity and r_1 have *moderately large positive effects* on cooperation. r_1 is generally used as an index of the cooperativeness of the payoff table so it is surprising that r_2 has a significantly larger impact on cooperation. The analysis suggests that we can increase

the probability of cooperation more by increasing the difference between the potential outcomes of a player and her opponent both cooperating (C,C) and only her cooperating (C,D). Increasing the difference between mutual cooperation (C,C) and mutual defection (D,D) will also increase cooperation, but less. The third group includes *error* and *risk*, which have *negative effects* on cooperation.

I empirically discovered that if a player cooperated (defected) in the previous period, she is very likely to cooperate (defect) in the next (Fig. 3.7). To explore the implications of this finding, I modified the simulation model so that I could exogenously set the probability of an agent cooperating in the first period, and found that it strongly affects cooperation levels in subsequent periods with the game structure set to the empirical mean values (Fig. 3.9B Simulated Experiments 1 and 4). However, a game structure that the sensitivity analysis indicates is very favorable to cooperation can moderate the negative effect of initial defection (Fig. 3.9B Experiment 2), and, conversely, a game structure that the analysis suggests should inhibit cooperation can moderate the positive effect of initial cooperation (Fig. 3.9B Experiment 3). The history of a particular interaction *and* the institutional structure both play important roles in determining cooperation levels.




Figure 3.9: Model simulation analysis. **A.** is a partial rank correlation coefficient analysis [78] of the effects of the game parameters on average cooperation; lines are bootstrapped 95% confidence intervals ($n = 1,000$). Continuous is set to its empirical mode, 0, because there is no within experiment variation on this. **B.** shows that first period play strongly affects cooperation levels in periods greater than one. Setting the game structure variables to the mean of the empirically observed values: if I exogenously set the probability of cooperation during the first period to 0 the simulated proportion of cooperation in subsequent periods is only 0.18 (‘simulated experiment 1’), and if I set the probability of cooperation during the first period to 1 the simulated proportion of cooperation in subsequent periods is 0.68 (‘simulated experiment 4’). When the probability of first period cooperation is set to 0, and I use a game structure that **A.** suggests should maximize cooperation, the proportion of cooperation is 0.43 (‘simulated experiment 2’); and when the probability of first period cooperation is set to 1, with the game structure that should minimize cooperation, the cooperation level is 0.35 (‘simulated experiment 3’).

I further investigated “inertia” – the probability a player will cooperate given that she cooperated last period – and its relationship to game structure. I computed an average “predicted inertia” for each of the thirty game structures by predicting the probability of cooperation after cooperating last period in a given game structure with the model, marginalizing out the effect of the time period and the opponent’s previous decision. To compute an av-

erage “actual inertia” value for each of the thirty game structures I divide the sum of the number of times all players cooperated after cooperating in the previous period by the total number of times all players cooperated in the previous period. The thirty predicted and actual inertia values have a 0.74 correlation, further evidence that the model captures the relevant patterns in the data. Game structures with longer expected length of interactions from a first-period perspective (higher δ), indefinite repetition, and higher r_2 payoff values have higher actual inertia values. δ is the strongest predictor of higher inertia and they are correlated at the 0.71 level.

3.7 Conclusion

The Prisoner’s Dilemma game is widely used to understand the tension between social and individual interests. I develop a computational model that can accurately predict human behavior in Prisoner’s Dilemma experimental games for a broad range of game structures, using only separate such structures for calibrating the model. I demonstrate that my approach can successfully predict behavior at multiple scales, yielding the most rigorously and broadly validated computational framework to date for designing institutions that promote cooperation in social dilemma scenarios. In particular, I use my model to identify variables that have the greatest impact on cooperation.

The sensitivity analysis demonstrated the importance of higher expected values of interaction length and larger differences between potential C,C and C,D outcomes (Fig. 3.9). It is more important to increase the benefits of mutual cooperation over losing out by being the sole cooperator than it is to increase the potential benefits of mutual cooperation relative to mutual defection. These insights are relevant to improving the underlying structure of new policy programs and designing new human subjects experiments. This work represents a new approach to understanding and *predicting* human interactions that will be increasingly relevant as more (experimental and observational) behavioral data is collected. With sufficient behavioral data from a variety of policy structures, my approach can be applied to

understand which factors should be prioritized to improve policy outcomes. The specifics need to be tailored to the circumstance, but models like those described herein can serve as a starting point for understanding which structural factors of a policy are most influential.

CHAPTER 4 COMPUTATIONAL POLICY DESIGN: A PREDICTION MARKET CASE STUDY

4.1 Abstract

Despite much scientific evidence, a large fraction of the American public doubts that greenhouse gases are causing global warming.¹ I present a simulation model as a computational test-bed for climate prediction markets. Traders adapt their beliefs about future temperatures based on the profits of other traders in their social network. I simulate two alternative climate futures, in which global temperatures are primarily driven either by carbon dioxide or by solar irradiance. These represent, respectively, the scientific consensus and a hypothesis advanced by prominent skeptics. I conduct sensitivity analyses to determine how a variety of factors describing both the market and the physical climate may affect traders' beliefs about the cause of global climate change. Market participation causes most traders to converge quickly toward believing the "true" climate model, suggesting that a climate market could be useful for building public consensus regardless of whether temperature changes are caused by carbon-dioxide levels or solar irradiance.

4.2 Introduction

The climate change debate has become strongly polarized over the past two decades. Although the scientific consensus on the anthropogenic nature of climate change strongly increased, beliefs about climate change did not evolve much within the public [79]. In addition, the divide on anthropogenic climate change between liberals and conservatives has grown steadily as the question is becoming increasingly politicized and potentially disconnected from scientific evidence [80]. The costs of misinformed climate policies are high. If climate change is not human-induced but is believed to be so, public resources will

¹Parts of this chapter appeared in Nay, J. J., Van der Linden, M., Gilligan, J. (2016). "Betting and Belief: Prediction Markets and Attribution of Climate Change." *Proceedings of the 2016 Winter Simulation Conference*, 1666–1677. IEEE Press.

be spent on unnecessary efforts. On the other hand, if climate change is human-induced but not recognized to be so, the costs of inaction could be high. Effective climate policies require acting quickly, so it would be valuable to bring the public to a prompt and accurate consensus on the issue.

Attempts to foster such consensus face many social and psychological challenges, some of which could be addressed by creating climate prediction markets where participants can “put their money where their mouths are” [81, 79]. The idea of using prediction markets to efficiently aggregate information about uncertain event outcomes has been widely discussed [82]. Prediction markets have interesting theoretical properties [83, 84], and perform well in terms of forecasting performance and information aggregation in experiments [85, 86] simulation models [87, 88] and the real-world [89, 90, 91]. However, to the best of my knowledge, the idea that prediction markets can generate consensus *on the factors* affecting uncertain events has never been explored quantitatively.

Bloch, Annan, and Bowle [92] have proposed using derivatives markets to reduce the scientific uncertainty in estimating the impact of future climate change. Existing prediction markets (e.g. <https://hypermind.com/hypermind/app.html>, <https://www.betfair.com/exchange/>, and <https://www.predictit.org/>) focus on near term events such as elections months away, so it is difficult to extrapolate empirical findings to the climate case. Furthermore, I am interested in investigating the unobservable beliefs of traders. Therefore, I turn to simulation modeling informed by climate and economic theory. I simulate a prediction market where traders exchange securities related to climate outcomes to explore whether, and under what social and climate conditions, prediction markets may be useful for increasing convergence of climate beliefs. This work can be extended as part of a computational design process for effective climate prediction markets and I release all the computer code to allow others to easily extend this model.

From a public policy perspective, changing the explanatory models of market participants is one of the most important roles that prediction markets might play. This is perhaps

the most important social benefit of prediction markets given the predictive power of statistical supervised learning models, which could possibly provide information at a much lower cost than creating and maintaining a market [93]. Effective climate policy not only requires an accurate consensus on future climate *outcomes*. It also requires an accurate consensus on the causal *mechanisms* influencing such outcomes. If people agree that temperature will rise, but some believe it will be due to greenhouse gases, while others believe that it will be caused by increased solar activity, inconsistent and ineffective policies may be implemented.

4.3 Related Work

Agent-based simulations of prediction markets have been studied [87, 94, 88], including some that feature communication between agents. In these models, however, beliefs about the uncertain outcomes are constructed in abstract ways. In particular, beliefs are not based on structural models from which agents could derive causal implications, so these models are not suitable for investigating the convergence of the underlying explanatory models agents employ for prediction.

Tseng, Lin, Lin, Wang, and Li [94] created an agent-based model (ABM) of a continuous double auction market with multiple market strategies, including two variants of a zero-intelligence agent. They compared the behavior of their simulation with data from a prediction market for the outcome of political elections. They found that despite their simplicity, zero-intelligence agents capture some salient features of real market data. Klingert and Meyer [87] compared the predictive accuracy of different kinds of simulated markets (continuous double auctions and logarithmic market scoring rules) and reached similar conclusions to experimental work by Hanson, Oprea, and Porter [85]. None of these models featured agents learning and updating their beliefs.

Jumadinova and Dasgupta [88] created a continuous double auction model in which agents update their beliefs about uncertain events based on newly-acquired information.

The better the information set, the more likely the agents were to put higher weight on last period's prices when revising their beliefs. However, little structure was imposed on the information set and the way it was used to generate the weights for updating beliefs, so that model does not permit studying the convergence of trader beliefs about predictive models. Ontanon and Plaza [95] studied the effect of deliberation on a simulated prediction market, in which agents used case-based reasoning [96] to debate uncertain outcomes with their neighbors in a social network. This model could in principle be used to assess convergence of predictive models, but this was not done and the process of forming and revising beliefs through argumentation is not well suited to studying stochastic time series, such as those relevant to climate prediction markets. Therefore, I created a new model based on economic theory.

4.4 Model Design

In the model, traders bet on global temperature anomalies six years in the future. During the six-year period, traders buy and sell futures. Every year, traders update their models and forecast of future temperatures based on newly available data. At the end of each six-year period, winners collect gains and traders revise beliefs about climate models, based on their ideology and the beliefs of top earners in their social network. In this section, I describe the models used to generate future temperature data, agents beliefs about those models, the market procedures, and the social network connecting agents. I conclude with details on model dynamics and an overview of the model parameters I experimentally vary.

4.4.1 Temperature Models

For climate time-series, we use the annual anomaly of global mean temperature. For years from 1880–2014 we use the GISTEMP global mean land-sea annual temperature anomalies [97, 98] and for years from 2015 onward, we project future climates under two alternative theories: In both theories, changes in global temperature are proportional to

changes in a deterministic forcing plus a stochastic noise term. For simplicity, we choose two alternative expressions for the deterministic climate forcing: one, which corresponds to conventional climate science, takes the natural logarithm of the atmospheric carbon dioxide concentration in parts per million [99], and the other, which corresponds to an alternative theory advocated by many who doubt or reject conventional climate science, takes the total solar irradiance (the brightness of sunlight, in Watts per square meter, at the top of the atmosphere) averaged over the 11-year sunspot cycle [100]. Most scientific models of the earth’s climate include many forcings, including carbon dioxide, other greenhouse gases, aerosols, total solar irradiance, and more. In these models, the changing CO₂ concentration is, by a large margin, the strongest single forcing [101]. Choosing only one forcing term for each of the competing models simplifies comparison because each model has the same number of adjustable parameters, therefore we consider only CO₂.

For CO₂ we used historical emissions through 2005, harmonized with RCP 8.5 representative concentration pathway from 2005 onward [102, 103] and for TSI we used the harmonized historical values, with a projection through 2100 from [104], which is, to my knowledge, the only prediction of TSI for the entire 21st century.

Warming coefficients for each model were determined by linear regression of historical temperatures from 1880–2014 against the historical values for each model’s forcing term. The noise model was determined by fitting an ARMA(p, q) model to the residuals from the regression, using the Stan (<http://mc-stan.org>) probabilistic programming language and the `rstan` (<https://cran.r-project.org/web/packages/rstan/>) software package [105]. Stan proved more numerically stable than the R `nlme` (<http://CRAN.R-project.org/package=nlme>) package [106] for fitting ARMA noise models. We identified the optimal model for the auto-correlated noise term by performing the regression analysis for all combinations of $p, q \in \{0, 1, 2\}$ and using the Widely Applicable Information Criterion to select the optimal noise model [107, 108]. In both cases (TSI and log CO₂), the optimal noise model was AR(1).




Figure 4.1: Historical measurements of temperature and a realization of possible future temperatures under two alternate models of climate physics.

Future climates were generated by applying the future climate forcings to Eq. 4.1:

$$T_{\text{model}}(t) = \beta_{\text{model}} F_{\text{model}}(t) + \varepsilon, \quad (4.1)$$

where $T_{\text{model}}(t)$ is the temperature at time t , under a given model of what causes warming, $F_{\text{model}}(t)$ is the forcing (TSI or $\ln(\text{CO}_2)$) at time t , β_{model} is the regression coefficient, and ε is a noise term. The coefficient β and the parameters of the ARMA noise model were fit to the historical data (1880–2014). An example realization of future climates for both the $\ln(\text{CO}_2)$ and TSI models is shown in Fig. 4.1.

4.4.2 Climate Beliefs

Traders use one of two models (temperature depends on CO_2 or TSI) to forecast future temperature. These models are interpreted as the trader's beliefs about the true climate process, the driving factor of long-term global temperatures. They represent pervasive positions on climate change in the public debate. In order to approximately match the current

configuration of beliefs in climate change in the United States, during model initialization, the CO₂ model is randomly assigned to half of the traders, while the TSI model is assigned to the other half of the traders. These random model assignments are made *before* the market initially opens.

Both when the true data generating process is CO₂ and TSI, at model initialization approximately half of the traders use the true data-generating model to make predictions. Traders using the true model do not necessarily make perfectly accurate predictions however. Although these traders believe in the correct *functional form* of the model, they still need to calibrate their model based on limited noisy data. Therefore, the values these traders assign to the parameters of the model will typically be different from the exact parameters in the data-generating process.

4.4.3 Traders and Markets

Traders are initially endowed with a single experimental currency unit (ECU). Traders use their model to forecast the distribution of future temperature and determine their reservation price for different securities. Each security pays 1 ECU at the end of the trading sequence if the temperature at the end of the sequence falls into a certain range.

Traders are risk-neutral expected utility maximizers. Therefore, their reservation price for a security is their assessment of the probability that the temperature will fall in the range covered by the security at the end of the sequence. At each time-step (year), the agents use the new year's temperature data to re-estimate the coefficients for the model they believe explains climate change, using Bayesian linear regression with an AR(1) noise model. Traders use the joint posterior probability distribution of regression and noise coefficients to estimate the probability distribution for the temperature at the end of the current trading sequence. We performed a full Bayesian analysis at each time step.

Traders then use this posterior probability distribution to assign reservation prices for buying and selling securities. Based on their reservation price, agents behave as “zero-

intelligence” traders [109]. They attempt to sell securities at a random price above their reservation, and to buy securities at a random price below their reservation. These trading strategies are simple but provide accurate approximations of behavior in prediction markets [87], and in financial markets more broadly.

Based on traders’ sell and buy orders, traders exchange securities following a continuous-double auction (CDA) procedure (see Model Dynamics below for more details). CDAs or some close variants are common procedures to match buy and sell orders. CDA are notably used on large stock markets [94].

4.4.4 Social Network

Traders are part of a social network where each agent forms two links at random, and then forms links randomly, ensuring that each agent is connected to at least two other agents. There are other, slightly more complex but potentially more realistic, approaches to modeling social networks that I may implement in future model designs. Once these are implemented, the network topology could be varied along with the other parameters to determine the importance of the underlying topology on the convergence of beliefs caused by participation in the market. Scale-free networks and small-world networks are common network topologies used when modeling social networks [110]. These more realistic network topologies often lead to more distinct neighborhoods within the overall network, which could lead to less exposure to the beliefs and earnings of traders in other neighborhoods and thus less switching of belief models and convergence.

Every time securities are realized, each trader looks at the performance of her richest neighbor in the network. Traders start with the same initial amount of ECU and differences in ECU can only come from market interactions. Therefore, if some trader is poorer than her richest neighbor, the trader interprets it as a signal that her richest neighbor may have a better model of the climate. Then, the trader considers adopting the model of her richest neighbor. For each trader, the willingness to revise her belief is determined by how




Figure 4.2: An example of the social network among traders. Yellow and blue circles represent traders who believe in the two different models of climate change. Lines indicate social network connections. In a highly segmented network (seg close to 1), most links connect like-minded traders and few connect traders with opposing beliefs.

ideologically loaded her belief is [80], which is varied in the experiments.

An example of a snapshot of a social network is depicted in Fig. 4.2. A segmentation parameter controls the homophily of the network: the extent to which traders are preferentially linked with other traders who share their initial belief in the cause of climate change (CO_2 or TSI), as opposed to traders with the opposite initial belief. This parameter varies from 0 (no preference for like-minded traders) to 1 (traders are only connected to like-minded traders). This reflects reality where cultural and social groups are highly correlated with climate change beliefs. Although traders can change their beliefs over time, the connections between traders do not change as the market unfolds, i.e the edges are fixed. I recognize that this may be an unrealistic assumption. It is likely that the market participants would change their social network as a response to changes in the climate beliefs of those in their network in order to remain in a network of more like-minded individuals. In future work, I may explore ways to model this social network evolution. I vary the segmentation

parameter in the experiments, but it then stays fixed throughout each model run.

4.4.5 Model Dynamics

The time periods t are grouped into trading *sequences*. In a given sequence, the potential payments associated with traded securities are all based on the temperature at the end of the sequence. For instance, the third trading sequence might start in period $t = 1964$ and end in period $t = 1970$. In this case, a security traded in the third sequence pays 1 ECU if the temperature at $t = 1970$ falls into the range of temperatures covered by the security.

At each time t , traders are assumed to know the past value of the temperature $T_{0:t}$, carbon dioxide $\text{CO2}_{0:t}$, and total solar irradiance $\text{TSI}_{0:t}$. In a sequence finishing at time t^* , traders also have common knowledge of $\text{CO2}_{t:t^*}$ and $\text{TSI}_{t:t^*}$, the future values of carbon dioxide and total solar irradiance up to t^* . However, at any t , traders do not know the value of any future temperatures. Thus, the traders know the forcing terms for the rest of the trading sequence, but do not know the value of T_{t^*} . Traders can only predict T_{t^*} using their approximate model and their knowledge of $T_{0:t}$, $\text{CO2}_{0:t^*}$, and $\text{TSI}_{0:t^*}$. Notice that because $T_{0:t}$, $\text{CO2}_{0:t^*}$, and $\text{TSI}_{0:t^*}$ are common knowledge, in each period t , any two traders with the same approximate model m form the same stochastic beliefs about future temperatures $P_{t,m}(T_{t^*} \mid T_{0:t}, \text{CO2}_{0:t^*}, \text{TSI}_{0:t^*})$. The probability distribution $P_{t,m}$ incorporates both epistemic uncertainty (the trader does not know the true values of the coefficients of the climate model) and aleatory uncertainty (in addition to deterministic warming, the temperature exhibits stochastic noise that is directly modeled).

At each time t , traders:

1. re-calibrate m , their approximate model at time t , based on the new set of temperature data available at t ,
2. use the posterior probability distribution from (1) to assign beliefs about the probability distribution of future temperatures at time t^* : $P_{t,m}(T_{t^*} \mid T_{0:t}, \text{CO2}_{0:t^*}, \text{TSI}_{0:t^*})$

and use it to determine the expected value they attach to each security, and

3. trade on the CDA market as follows:

- Every trader i chooses at random a security s_i^B she will try to buy.
- Every trader i also chooses at random a security s_i^S she will try to sell among the securities she owns a positive amount of (if any).
- Traders then decide their selling price p_i^S and buying price p_i^B . To do so, traders first compute their expected values $E(s_i^B)$ and $E(s_i^S)$ for securities s_i^B and s_i^S (where expected values are with respect to i 's approximate model at time t). Then traders set p_i^S at random above $E(s_i^S)$ and p_i^B at random below $E(s_i^B)$ (see Model Parameters below for more details).
- Traders go to the market one at the time, in an order drawn randomly for each t .
- When trader i comes to the market, she places limit orders in the order book. These orders specify that i is willing to buy s_i^B at any price below p_i^B , and to sell s_i^S at any price above p_i^S .
- The market maker attempts to match i 's orders with some order which was put in the book before i came to the market.
- If there are outstanding sell offers for s_i^B at a price below p_i^B , a trade is concluded. Trader i buys one unit of s_i^B from the seller who sells at the lowest price below p_i^B , and the sell and buy offers are removed from the order book.
- If there are outstanding buy offers for s_i^S at price above p_i^S , a trade is concluded. Trader i sells one unit to the buyer who buys at the highest price above p_i^S , and the sell and buy offers are removed from the order book.
- When all traders have come to the market, any remaining offers are removed from the order book, and the trading period is concluded.

At t^* , when the sequence ends, there is only one security s^* that includes the actual temperature T_{t^*} . At t^* , traders:

1. receive 1 ECU per unit of s^* they own, and
2. consider adopting their neighbors' approximate model as described in the behavioral parameters sub-section below.

4.4.6 Model Parameters

The model depends on the following parameters, which I vary in simulation experiments to determine their effects on the convergence of beliefs. I group the parameters into climate, network and individual behavioral factors.

- Climate parameters:

- `true.model`: temperature data-generating process (CO₂ or TSI).

- Network parameters:

- `n.traders`: the number of traders.
- `n.edg`: the number of edges in the social network.
- `seg`: the segmentation parameter for the social network.

- Behavioral parameters:

- `risk.tak`: determines the distribution of risk tolerance with respect to successfully buying or selling securities. Higher risk tolerance corresponds to demanding more aggressive prices for buying and selling, and hence, a higher risk of not completing a trade that would be mutually advantageous to the buyer and seller.

For each trader i , the level of `risk.taki` is drawn uniformly at random from $[0, \text{risk.tak}]$. The higher `risk.taki`, the higher the price i will demand for selling securities, and the lower the price i will offer for buying.

Formally, in each period, trader i picks her buying or selling price for s uniformly at random in the interval $[(1 - \text{risk.tak}_i)\text{reserv}_i, \text{reserv}_{i,t}]$ for buying and $[\text{reserv}_i, (1 + \text{risk.tak}_i)\text{reserv}_{i,t}]$ for selling.

- `ideo` : determines the degree of “ideology” of traders. For each trader i , the level of `ideoi` is drawn uniformly at random from $[0, \text{ideo}]$. If `ideo` is high, traders will not revise their approximate models easily, even when faced with evidence that their richest neighbor is doing better than them. Formally, for each trader i and each sequence, `ideoi` is the probability that i adopts the approximate model of her richest neighbor if that neighbor is doing better than i at the end of the sequence (in monetary terms).

4.5 Results

4.5.1 Historical Climate

As a validation that the market is operating correctly, I ran the market using actual historical temperatures from 1880 to 2014, with market betting from 1931 to 2014. Greenhouse gas concentrations did not become high enough to begin dominating natural climate variations until the 1970s or so, and it took until the 1990s for the temperature record to show clear signs of anthropogenic interference [101]. Consistent with this, the simulated historical trading sequence (Fig. 4.3) shows convergence to belief in the TSI model until the early 1970s, after which traders converge toward believing the CO₂ model.

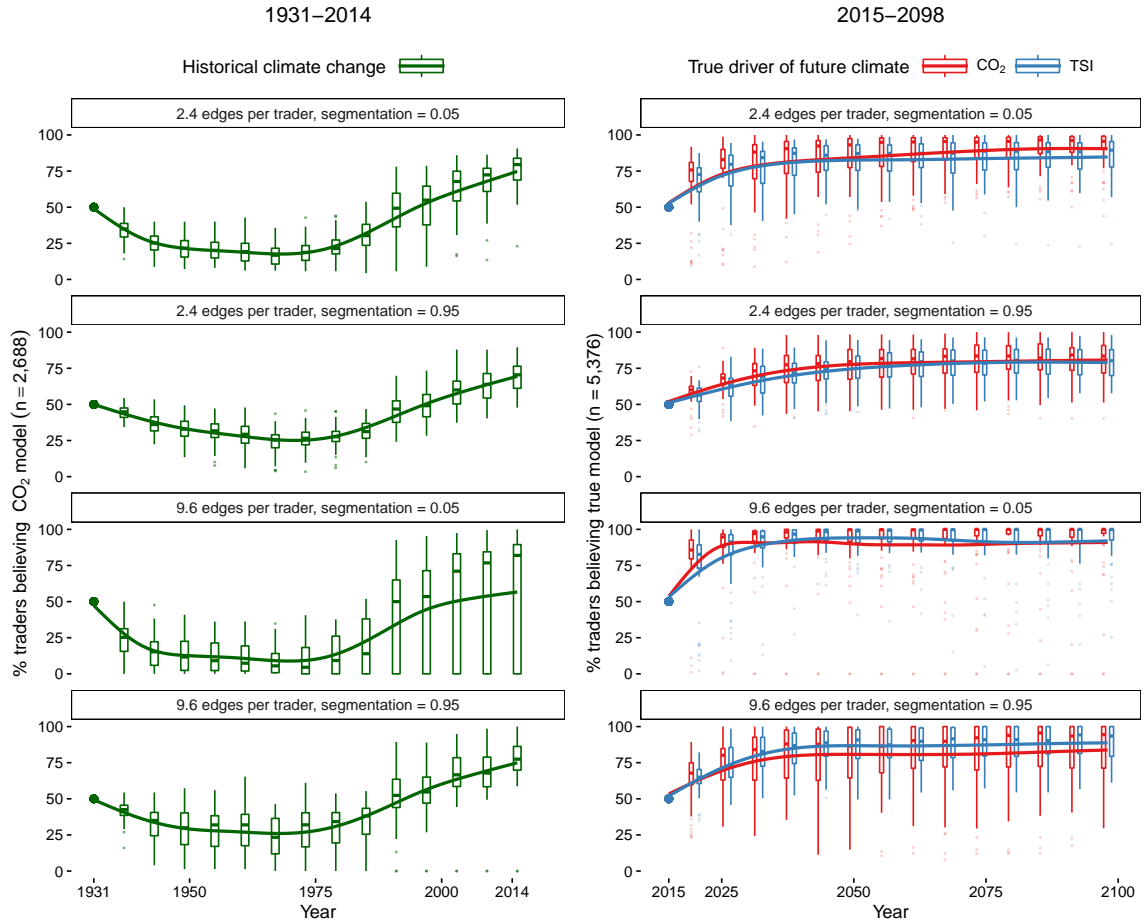


Figure 4.3: Convergence over trading sequences for different degrees of social-network connection and segmentation.

4.5.2 Future Scenarios

The primary focus is on simulating the market in future scenarios. This is more relevant for policy design and more interesting theoretically due to the increasing divergence of global temperature values under the two future scenarios. The sensitivity analysis is based on a Latin hypercube sampling of 500 parameter sets from the following distributions [74, 75].

- $ideo \sim \text{Uniform}(0, 1)$
- $n.edge \sim \text{Uniform}(100, 200)$ (mapped into integer)

- `n.traders` \sim Uniform(50, 250) (mapped into integer)
- `risk.tak` \sim Uniform(0, 1)
- `seg` \sim Uniform(0, 1)
- `true.model` \sim Bernoulli(0.5)

The future CO₂ forcing is taken from the RCP 8.5 scenario [103] and the future TSI forcing is taken from a prediction of 21st Total Solar Irradiance by [104]. Because of the stochastic noise term, the temperature time series were different in each simulation.

I used the model to perform 10 full simulations for each of the 500 input parameter sets and average the 10 convergence scores. I conducted multiple simulations for the same parameter set because there is stochasticity in the temperature time series, the social network structure, and the agent decision models. I then conducted a partial rank correlation coefficient analysis on the relationship between the input matrix, X , and the resulting simulated outcome vector of mean belief convergence scores, y [111, 78, 76]. Partial correlation computes the linear relationship between the part of the variation of X_i and y that are linearly independent of other X_j ($j \neq i$). The difference between the partial correlation and the partial rank correlation that I use here is that I first rank-transform that data in order to capture potentially non-linear relationships. I conducted 1,000 bootstrapped estimations of the partial rank correlation coefficients to obtain 95% confidence intervals.

The sensitivity analysis averages over time and thus masks time trends. I randomly drew from the above distributions for ideology, risk tolerance, and the number of traders, and conducted an experiment crossing 95th and 5th percentile values for the number of edges per trader and the segmentation of the social network and both values for the true model. I collected the time series of belief-convergence across these eight designs to visualize the distributions of convergence over time (Fig. 4.3). Under most parameterizations, the median fraction of traders believing in the true model reaches 75% in 10–20 years.




Figure 4.4: Estimated effects of model parameters on convergence of beliefs in future scenario. Positive numbers increase convergence and negative numbers decrease it. True model refers to the true climate model being CO₂, so traders converge to the true model more quickly when carbon dioxide drives climate change than when total solar irradiance does. Bars represent 95% confidence intervals.

4.6 Discussion

The sensitivity analysis, Fig. 4.4, shows that ideology, risk tolerance, the number of traders, the number of edges per trader, the segmentation of the social network, and the true model all statistically significantly affect convergence. Increasing the number of edges per trader increases the flow of information through the market, which causes traders to converge toward believing the “true” climate model.

Segmented social networks reduce convergence by creating an “echo-chamber” effect, in which lack of interaction between traders with different views reduces traders’ access to information that could persuade them to change their beliefs. This is apparent in Fig. 4.3, where the rate of convergence in highly segmented markets ($s_{seg} = 0.95$) is considerably slower, especially in the first three trading sequences (18 years), than for markets with low

segmentation ($seg = 0.05$).

When the true model is CO₂-driven there is more convergence toward the true model. We believe this is because when models are fit to the historical data, the residuals from the CO₂-driven model are considerably smaller than those from the TSI-driven model. Thus, projections of future climate change using the TSI model will exhibit significantly larger stochastic noise than the CO₂ models, and this noise makes it more difficult for traders to identify the true model when the true model is TSI.

Traders with higher risk tolerance will price buy and sell orders more aggressively, so they will earn more (or lose less) on completed trades, but will have a greater risk of failing to complete trades. Among traders with the correct model of climate, those with higher risk tolerance will earn more profit on completed trades (and their counterparties will lose more). A trader's wealth is an important source of information, and we believe this is why greater risk taking enables traders to identify the correct model more quickly. Additionally, the risk of failing to complete trades adds volatility, which also puts information in play. I have observed a similar phenomenon in a very different context, in Chapter 2, where adding stochastic noise to player decisions in iterated games improves the accuracy with which my method can identify the players' strategies [112].

What is not clear is why increasing the number of traders slows convergence, even if the number of edges per trader remains fixed. This is a topic for future research.

4.7 Conclusion

I simulate two alternative climate futures: one where CO₂ is the primary driver of global temperature and one where variations in solar intensity are the primary driver. These represent the two most plausible competing views in the public discourse and the analysis is agnostic about which is "true."

Market participation causes traders to converge toward believing the "true" climate model under a variety of model parameterizations in a relatively short time: In markets with

low segmentation, the number of traders believing the true model rises from 50% to 75% in roughly 12 years, and even in highly segmented markets, 75% convergence is achieved in 18–24 years. Ideally, I would like to compare belief convergence with and without a prediction market but because the only source of climate belief in the current model is market interactions I cannot make this comparison. However, I do use actual temperatures for the model under historical 20th century conditions and observe convergence to the CO₂ model, whereas in the real world there is no convergence of beliefs.

The model assumes that traders continue to participate in the market even if they are losing money by participating. In a future version of the model, I may explore how to model the choice of participation in the market. If a trader is losing money consistently they may leave the market with their beliefs unchanged. This would reduce the belief changing value of the market.

Both the historical and future simulation results suggest that a climate prediction market could be useful for producing broad agreement about the causes of climate change, and could have persuasive power for people who are not persuaded by the consensus among climate researchers that greenhouse gases (and especially CO₂) are responsible for the majority of observed climate change [101]. I also find that market segmentation has a large effect on the speed of convergence, so transparency and effective communication about the performance of traders with different beliefs will be important.

The fact that rapid convergence to the true model occurs regardless of which model is actually “true” may persuade those who doubt the scientific consensus that the market is ideologically neutral, and that the deck is not stacked to produce a pre-determined result.

All the code and data for the model is released as free and open source software and is available at <https://github.com/JohnNay/predMarket>. This project is a computational test-bed for public policy design: the code can be extended to test the effects of trading strategies, cognitive models, future climate scenarios, and market designs on the evolution of trader beliefs.

CHAPTER 5 CONCLUSION

Use-inspired basic research is exemplified by the work of Louis Pasteur, leading to both theoretical microbiology insights *and* the practical prevention of disease. Stokes [113] introduced the term Pasteur's Quadrant to locate use-inspired basic research in relation to basic research – guided solely by its potential to increase scientific understanding – and applied research – guided solely by its intended end-use in society (Figure 5.1). For my dissertation, I conducted basic research into computational methods for modeling dynamic decisions, but this basic research was guided by its intended application to improving public policy. I hope my work can be extended to increasingly realistic scenarios to inform policy design.

The method described in Chapter 2 can be used to analyze behavioral data to uncover decision-making models by automatically deriving simple rules that describe behaviors in large datasets. When a policy-maker needs a quick overview of how people or companies may be making decisions they can apply this method. This model representation is especially relevant for decision-making in strategic interactions and where a relatively simple model is likely to explain the underlying decision-making. The primary test case for the method was the Prisoner's Dilemma game, but I illustrated the use of observational data and extra predictor variables by applying the method to an example from international relations involving repeated water management-related interactions between countries that share rivers. I used three lagged variables: whether there was water-related conflict between them in the previous year, whether they cooperated around water in the previous year, and whether they had signed a water-related treaty during any previous year. For any decision context where history is important, my method could be very useful for teasing out potential underlying decision models of the actors involved.

The approach developed for Chapter 3 can similarly benefit from the rise of massive amounts of data on decisions. Chapter 3 presents a new approach to understanding and

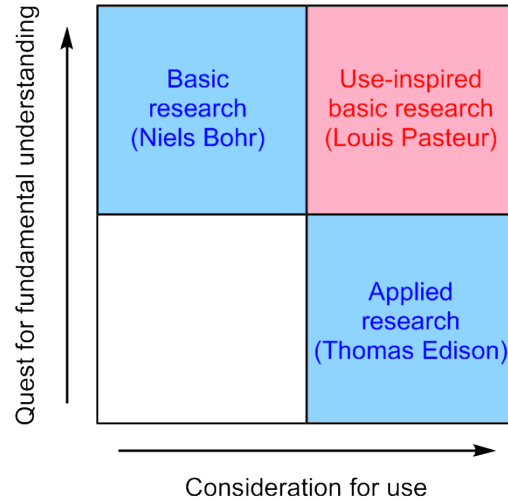


Figure 5.1: Pasteur's Quadrant [113]. Figure from [114].

predicting human interactions that will be increasingly relevant as more behavioral data is collected. With sufficient behavioral data from a variety of policy structures, the approach can be applied to understand which factors should be prioritized to improve policy outcomes. The specifics need to be tailored to the circumstance, but the results of the modeling and sensitivity analysis can serve as a starting point for understanding which structural factors of a policy are most influential.

Finally, the approach of Chapter 4, computational simulations of new markets in scenarios where the market cannot be tested easily in the real-world, can also be applied to a variety of institutional designs to allow policy-makers to stress-test new ideas before rolling out real-world policy experiments. Our simulation model of climate prediction markets, which we have released as free open source software, can easily be directly extended to test the effects of trading strategies, traders' cognitive models, alternative future climate scenarios, and market designs.

BIBLIOGRAPHY

- [1] Stephen L. Morgan and Christopher Winship. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge University Press, New York, 1 edition edition, July 2007.
- [2] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3(0):96–146, 2009.
- [3] J. J. Heckman, J. Smith, and N. Clements. Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts. *The Review of Economic Studies*, 64(4):487–535, October 1997.
- [4] Kosuke Imai and David A van Dyk. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467):854–866, September 2004.
- [5] Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press, Princeton, 1 edition edition, January 2009.
- [6] Gerd Gigerenzer and Reinhard Selten, editors. *Bounded Rationality: The Adaptive Toolbox*. The MIT Press, Cambridge, Mass., August 2002.
- [7] Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1 edition, 1954.
- [8] I. Erev and A. E. Roth. Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences*, 111(Supplement_3):10818–10825, July 2014.
- [9] Colin F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, 2003.
- [10] Nobuyuki Hanaki. Action learning versus strategy learning. *Complexity*, 9:41–50, 2004.
- [11] Richard D. McKelvey and Thomas R. Palfrey. Playing in the dark: Information, learning, and coordination in repeated games. Technical report, California Institute of Technology, Pasadena, 2001.
- [12] John H. Miller. The coevolution of automata in the repeated prisoner’s dilemma. *Journal of Economic Behavior & Organization*, 29:87–112, 1996.
- [13] Drew Fudenberg, David G Rand, and Anna Dreber. Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*, 102:720–749, 2012.
- [14] Ariel Rubinstein. Finite automata play the repeated prisoner’s dilemma. *Journal of Economic Theory*, 39(1):83–96, June 1986.

- [15] Nobuyuki Hanaki, Rajiv Sethi, Ido Erev, and Alexander Peterhansl. Learning strategies. *Journal of Economic Behavior & Organization*, 56:523–542, 2005.
- [16] Edward Moore. Gedanken-experiments on sequential machines. *Automata Studies*, 34:129–153, 1956.
- [17] George H. Mealy. A method for synthesizing sequential circuits. *Bell System Technical Journal*, 34(5):1045–1079, September 1955.
- [18] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, MA, 1994.
- [19] Axel Cleeremans, David Servan-Schreiber, and James L. McClelland. Finite state automata and simple recurrent networks. *Neural Computation*, 1(3):372–381.
- [20] Zheng Zeng, Rodney M. Goodman, and Padhraic Smyth. Learning finite state machines with self-clustering recurrent networks. *Neural Computation*, 5(6):976–990.
- [21] Shai Ben-David, Nadav Eiron, and Philip M. Long. On the difficulty of approximately maximizing agreements. *Journal of Computer and System Sciences*, 66(3):496–514.
- [22] James Bullard and John Duffy. Using genetic algorithms to model the evolution of heterogeneous beliefs. *Computational Economics*, 13:41–60, 1999.
- [23] Jasmina Arifovic. Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control*, 18:3–28, 1994.
- [24] Jasmina Arifovic and Curtis Eaton. Coordination via genetic learning. *Computational Economics*, 8:181–203, 1995.
- [25] Robert E. Marks, David F. Midgley, and Lee G. Cooper. Adaptive behaviour in an oligopoly. In Jörg Biethahn and Volker Nissen, editors, *Evolutionary Algorithms in Management Applications*, pages 225–239. Springer, New York, 1995.
- [26] David F. Midgley, Robert E. Marks, and Lee C. Cooper. Breeding competitive strategies. *Management Science*, 43:257–275, 1997.
- [27] David B. Fogel. Evolving behaviors in the iterated prisoner’s dilemma. *Evolutionary Computation*, 1:77–97, 1993.
- [28] John H. Miller and Scott E. Page. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press, Princeton, 2007.
- [29] David E. Goldberg and John H. Holland. Genetic algorithms and machine learning. *Machine Learning*, 3:95–99, 1988.
- [30] John Duffy and Jim Engle-Warnick. Using symbolic regression to infer strategies from experimental data. In *Evolutionary Computation in Economics and Finance*, pages 61–82. Springer, New York, 2002.

- [31] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Bradford, Cambridge, MA, 1992.
- [32] John Duffy. Agent-based models and human subject experiments. In *Handbook of Computational Economics*, volume 2, pages 949–1011. Elsevier, Amsterdam, 2006.
- [33] C. Savage. A Survey of Combinatorial Gray Codes. *SIAM Review*, 39(4):605–629, January 1997.
- [34] Dirk Eddelbuettel. *Seamless R and C++ Integration with Rcpp*. Springer, New York, 2013.
- [35] Luca Scrucca. GA: A package for genetic algorithms in R. *Journal of Statistical Software*, 53:1–37, 2013.
- [36] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*. Springer, New York, NY, 2nd edition, 2009.
- [37] Robert M. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
- [38] Garrett Hardin. The tragedy of the commons. *Science*, 162:1243–1248, 1968.
- [39] John Jacob Nay. Predicting cooperation and designing institutions: An integration of behavioral data, machine learning, and simulation. In *Winter Simulation Conference Proceedings*, Savannah, GA, December 2014.
- [40] Yoella Bereby-Meyer and Alvin E. Roth. The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation. *The American Economic Review*, 96:1029–1042, 2006.
- [41] John Duffy and Jack Ochs. Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior*, 66:785–812, 2009.
- [42] Howard Kunreuther, Gabriel Silvasi, Eric T. Bradlow, and Dylan Small. Bayesian analysis of deterministic and stochastic prisoner’s dilemma games. *Judgment and Decision Making*, 4:363–384, 2009.
- [43] Pedro Dal Bo and Guillaume R Frechette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101:411–429, 2011.
- [44] Robert M. Axelrod. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton University Press, Princeton, 1997.
- [45] S.Y. Chong and Xin Yao. Behavioral diversity, choices and noise in the iterated prisoner’s dilemma. *IEEE Transactions on Evolutionary Computation*, 9:540–551, 2005.

- [46] Marit Brochmann. Signing river treaties: Does it improve river cooperation? *International Interactions*, 38:141–163, 2012.
- [47] Pedro Dal Bo. Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. 95(5):1591–1604.
- [48] Daniel Friedman and Ryan Oprea. A continuous dilemma. *The American Economic Review*, 102(1):337–363, 2012.
- [49] Robert Axelrod and Douglas Dion. The further evolution of cooperation. 242(4884):1385–1390.
- [50] Anatol Rapoport and Albert M. Chammah. *Prisoner's Dilemma*. University of Michigan Press, s.l., December 1965.
- [51] Alvin E. Roth and J. Keith. Equilibrium behavior and repeated play of the prisoner's dilemma. 17(2):189–198.
- [52] James Andreoni and John H. Miller. Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. 103(418):570–585.
- [53] Anna Dreber, Drew Fudenberg, and David G. Rand. Who cooperates in repeated games: The role of altruism, inequity aversion, and demographics. *Journal of Economic Behavior & Organization*, 98:41–55, February 2014.
- [54] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014.
- [55] Max Kuhn, Steve Weston, Andre Williams, Chris Keefer, Allan Engelhardt, Tony Cooper, and Zachary Mayer and the R. Core Team. caret: Classification and regression training, August 2014.
- [56] Teck H. Ho, Colin F. Camerer, and Juin-Kuan Chong. Self-tuning experience weighted attraction learning in games. 133(1):177–198.
- [57] P. J. Deadman. Modelling individual behaviour and group performance in an intelligent agent-based simulation of the tragedy of the commons. *Journal of Environmental Management*, 56(3):159–172, July 1999.
- [58] Wander Jager and Marco Janssen. The need for and development of behaviourally realistic agents. In *Multi-Agent-Based Simulation II*, pages 36–49. Springer, 2003.
- [59] Marco A. Janssen and Toh-Kyeong Ahn. Learning, signaling, and social preferences in public-good games. *Ecology and society*, 11(2):21, 2006.
- [60] Stephen Wendel and Joe Oppenheimer. An agent-based analysis of context-dependent preferences. *Journal of Economic Psychology*, 31(3):269–284, June 2010.

- [61] Jasmina Arifovic and John Ledyard. Individual evolutionary learning with many agents. *The Knowledge Engineering Review*, 27(2):239–254.
- [62] Michael Wunder, Siddharth Suri, and Duncan J. Watts. Empirical agent based models of cooperation in public goods games. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 891–908. ACM, 2013.
- [63] James Andreoni and John H. Miller. Auctions with artificial adaptive agents. *Games and Economic Behavior*, 10(1):39–64, July 1995.
- [64] John Bower and Derek W. Bunn. Model-based comparisons of pool and bilateral markets for electricity. *The Energy Journal*, Volume21(Number 3):1–29, 2000.
- [65] Robert Marks. Chapter 27 market design using agent-based models. In L. Tesfatsion and K.L. Judd, editor, *Handbook of Computational Economics*, volume Volume 2, pages 1339–1380. Elsevier, 2006.
- [66] Valerio Capraro. A model of human cooperation in social dilemmas. 8(8):e72427.
- [67] Alvin E. Roth and Ido Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, 8(1):164–212, 1995.
- [68] Yin-Wong Cheung and Daniel Friedman. Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior*, 19(1):46–76, April 1997.
- [69] Ido Erev and Alvin E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. 88(4):848–81.
- [70] Dale O. Stahl and Ernan Haruvy. Aspiration-based and reciprocity-based rules in learning dynamics for symmetric normal-form games. *Journal of Mathematical Psychology*, 46(5):531–553, October 2002.
- [71] Ido Erev, Alvin E. Roth, Robert L. Slonim, and Greg Barron. Learning and equilibrium as useful approximations: Accuracy of prediction on randomly selected constant sum games. *Economic Theory*, 33(1):29–51, July 2007.
- [72] Ji-Hyun Kim. Estimating classification error rate: Repeated cross-validation, repeated hold-out and bootstrap. *Computational Statistics & Data Analysis*, 53(11):3735–3745, September 2009.
- [73] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning: with Applications in R*. Springer, New York, 1st ed. 2013. corr. 4th printing 2014 edition edition, August 2013.
- [74] Brian Beachkofski and Ramana Grandhi. Improved Distributed Hypercube Sampling. In *43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. American Institute of Aeronautics and Astronautics, 2002.

- [75] Rob Carnell. Latin hypercube samples. R package, Comprehensive R Archive Network, 2012.
- [76] A. Saltelli, K. Chan, and E. M. Scott. *Sensitivity Analysis*. Wiley, Chichester, 2009.
- [77] Jan C. Thiele, Winfried Kurth, and Volker Grimm. Facilitating parameter estimation and sensitivity analysis of agent-based models: A cookbook using NetLogo and r. *Journal of Artificial Societies and Social Simulation*, 17(3):11, 2014.
- [78] Gilles Pujol, Bertrand Iooss, and Alexandre Janon Lemaitre. Sensitivity: Sensitivity analysis. R package, Comprehensive R Archive Network, 2014.
- [79] Michael P. Vandenbergh, Kaitlin Toner Raimi, and Jonathan M. Gilligan. Energy and climate change: A climate prediction market. *UCLA Law Review*, 61:1962–2017, 2014.
- [80] Dan M. Kahan, Hank Jenkins-Smith, and Donald Braman. Cultural cognition of scientific consensus. *Journal of Risk Research*, 14:147–174, 2011.
- [81] Shi-Ling Hsu. A prediction market for climate outcomes. *University of Colorado Law Review*, 83:179–256, 2011.
- [82] Christian Franz Horn, Bjoern Sven Ivens, Michael Ohneberg, and Alexander Brem. Prediction markets—a literature review 2014. *The Journal of Prediction Markets*, 8:89–126, 2014.
- [83] Reinhard Set and Reinhard Selten. Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1:43–62, 1998.
- [84] Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets*, pages 3–15, 2012.
- [85] Robin Hanson, Ryan Oprea, and David Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60:449–459, 2006.
- [86] Paul J. Healy, Sera Linardi, J. Richard Lowery, and John O. Ledyard. Prediction markets: Alternative mechanisms for complex environments with few traders. *Management Science*, 56:1977–1996, 2010.
- [87] Frank MA Klingert and Matthias Meyer. Comparing prediction market mechanisms using an experiment-based multi-agent simulation. In Klaus G. Troitzsch, Möhring, and Ulf Lotzmann, editors, *Proceedings 26th European Conference on Modelling and Simulation*, pages 654–661, 2012.
- [88] Janyl Jumadinova and Prithviraj Dasgupta. A multi-agent system for analyzing the effect of information on prediction markets. *International Journal of Intelligent Systems*, 26:383–409, 2011.

- [89] Justin Wolfers and Eric Zitzewitz. Prediction markets in theory and practice. Working Paper 12083, NBER, 2006.
- [90] Deepak Pathak, David Rothschild, and Miroslav Dudik. A comparison of forecasting methods: Fundamentals, polling, prediction markets, and experts. *Journal of Prediction Markets*, 9:1–31, 2015.
- [91] Anna Dreber, Thomas Pfeiffer, Johan Almenberg, Siri Isaksson, Brad Wilson, Yiling Chen, Brian A. Nosek, and Magnus Johannesson. Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences*, 112:15343–15347, 2015.
- [92] Daniel Bloch, James Annan, and Justin Bowles. Cracking the climate change conundrum with derivatives. *Wilmott Journal*, 2:271–287, 2010.
- [93] Sharad Goel, Daniel M. Reeves, Duncan J. Watts, and David M. Pennock. Prediction without markets. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, pages 357–366, New York, NY, USA, 2010. ACM.
- [94] Jie-Jun Tseng, Chih-Hao Lin, Chih-Ting Lin, Sun-Chong Wang, and Sai-Ping Li. Statistical properties of agent-based models in markets with continuous double auction mechanism. *Physica A: Statistical Mechanics and its Applications*, 389:1699–1707, 2010.
- [95] Santi Ontañón and Enric Plaza. Argumentation-based information exchange in prediction markets. In Iyad Rahwan and Pavlos Moraitis, editors, *Argumentation in Multi-Agent Systems*, number 5384 in Lecture Notes in Computer Science, pages 181–196. Springer Berlin Heidelberg, 2009.
- [96] Agnar Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7:39–59, 1994.
- [97] GISTEMP Team. GISS surface temperature analysis (GISTEMP). Technical report, NASA Goddard Institute for Space Studies, 2016.
- [98] J. Hansen, R. Ruedy, M. Sato, and K. Lo. Global surface temperature change. *Reviews of Geophysics*, 48:RG4004, 2010.
- [99] David Archer. *Global Warming: Understanding the Forecast*. Wiley, 2nd edition, 2012.
- [100] Willie W.-H. Soon. Variable solar irradiance as a plausible agent for multidecadal variations in the arctic-wide surface air temperature record of the past 130 years. *Geophysical Research Letters*, 32:L16712, 2005.
- [101] IPCC. *Climate Change 2013: The Physical Science Basis*. Cambridge University Press, 2013.

- [102] Peter Kolp and Keywan Riahi. RCP database. Technical report, 2009.
- [103] Keywan Riahi, Shilpa Rao, Volker Krey, Cheolhung Cho, Vadim Chirkov, Guenther Fischer, Georg Kindermann, Nebojsa Nakicenovic, and Peter Rafaj. Rcp 8.5: A scenario of comparatively high greenhouse gas emissions. *Climatic Change*, 109:33–57, 2011.
- [104] V. M. Velasco Herrera, B. Mendoza, and G. Velasco Herrera. Reconstruction and prediction of the total solar irradiance: From the medieval warm period to the 21st century. *New Astronomy*, 34:221–233, 2015.
- [105] Bob Carpenter, Andrew Gelman, Matt Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Michael A. Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 2016. (in press).
- [106] Jose Pinheiro, Douglas Bates, Saikat DebRoy, Deepayan Sarkar, and R Core Team. *nlme: Linear and Nonlinear Mixed Effects Models*, 2016. R package version 3.1-126.
- [107] Sumio Watanabe. A widely applicable bayesian information criterion. *Journal of Machine Learning*, 14:867–897, 2013.
- [108] Andrew Gelman, Jessica Hwang, and Aki Vehtari. Understanding predictive information criteria for bayesian models. *Statistics and Computing*, 24:997–1016, 2014.
- [109] Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101:119–137, 1993.
- [110] Duncan J. Watts. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press, Princeton, N.J.; Woodstock, December 2003.
- [111] Simeone Marino, Ian B. Hogue, Christian J. Ray, and Denise E. Kirschner. A methodology for performing global uncertainty and sensitivity analysis in systems biology. *Journal of Theoretical Biology*, 254:178–196, 2008.
- [112] John J. Nay and Jonathan M. Gilligan. Data-driven dynamic decision models. In L. Yilmaz, W.K.V. Chan, I. Moon, T.M.K. Roeder, C. Macal, and M. Rosetti, editors, *Proceedings of the 2015 Winter Simulation Conference*, pages 2752–2763. IEEE Press, 2015.
- [113] Donald E. Stokes. *Pasteurs Quadrant: Basic Science and Technological Innovation*. Brookings Institution Press, Washington, D.C, August 1997.
- [114] Frank Leibfarth. Speaking frankly: The allure of pasteurs quadrant, June 2013.