

12-15-2016

Modeling Negotiations Over Water and Ecosystem Management: Uncertainty and Political Viability

Rachael E. Goodhue

University of California, Davis

Susan Stratton Sayre

Smith College, ssayre@smith.edu

Leo K. Simon

University of California - Berkeley

Follow this and additional works at: https://scholarworks.smith.edu/eco_facpubs

Part of the [Economics Commons](#)

Recommended Citation

Goodhue, Rachael E.; Sayre, Susan Stratton; and Simon, Leo K., "Modeling Negotiations Over Water and Ecosystem Management: Uncertainty and Political Viability" (2016). Economics: Faculty Publications, Smith College, Northampton, MA.
https://scholarworks.smith.edu/eco_facpubs/1

This Article has been accepted for inclusion in Economics: Faculty Publications by an authorized administrator of Smith ScholarWorks. For more information, please contact scholarworks@smith.edu

MODELING NEGOTIATIONS OVER WATER AND ECOSYSTEM MANAGEMENT: UNCERTAINTY AND POLITICAL VIABILITY

RACHAEL E. GOODHUE, SUSAN STRATTON SAYRE, AND LEO K. SIMON

JULY 30, 2016

ABSTRACT. We present a modeling approach for generating robust predictions about how changes in institutional, economic, and political considerations will influence the outcome of political negotiations over complex water-ecosystem policy debates. Evaluating the political viability of proposed policies is challenging for researchers in these complex natural and political environments; there is limited information with which to map policies to outcomes to utilities or to represent the political process adequately. Our analysis evaluates the viability of policy options using a probabilistic political viability criterion that explicitly recognizes the existence of modeling uncertainty. The approach is used to conduct a detailed case study of the future of California's Sacramento-San Joaquin Delta. Several other possible applications of the approach are briefly discussed.

KEYWORDS: Pareto optimality; Delta; California; Political economy; Deep uncertainty; Robust decision making; Modeling uncertainty.

JEL CLASSIFICATION: P48, Q25, Q34

Goodhue is a Professor at the Department of Agricultural and Resource Economics, University of California, Davis. Sayre is an Assistant Professor at the Department of Economics, Smith College. Simon is Professor at the Department of Agricultural and Resource Economics, University of California, Berkeley and a member of the Faculty of Business and Economics, Monash University. Goodhue and Simon are members of the Giannini Foundation of Agricultural Economics.

This research was supported in part by the Giannini Foundation of Agricultural Economics. The authors thank Jim Chalfant, Richard Howitt, Jeffrey LaFrance, Anthony Millner, Christian Traeger, and Jeffrey Williams for helpful conversations and the participants of the Five College Junior Economics Faculty Seminar, the participants of the UC Davis Ag-IO workshop, the editors of SBE and two anonymous reviewers for helpful comments. Kathy Edgington and Laurie Warren provided invaluable computer resources support. Goodhue (goodhue@primal.ucdavis.edu) is the corresponding author.

1. INTRODUCTION

Conflicts between water users and ecosystem needs are common around the world. Whether and how these conflicts are resolved depends on the specific political, institutional and economic circumstances of a particular problem. In many cases, the precise nature of these circumstances is difficult to define, so isolating a predicted outcome of a negotiation is challenging. An alternative approach is to seek to predict outcomes that are very unlikely to emerge as a negotiation outcome regardless of the precise definitions of all elements of the bargaining environment. In this paper, we present and illustrate a modeling approach to generating such predictions.

Our approach is designed to address common feature of conflicts regarding water and ecosystem management. First, we focus on specific, one-time policy negotiations between several distinct stakeholder groups involving market and non-market valuations, private and public goods, and tradeoffs between economic and environmental objectives. In this setting, the complexities really matter: it is important to model the interconnected economic, social, and ecosystem impacts of the various policy options under consideration. As a result, any model which attempts to capture many of these interactions will necessarily be too complex to be solved analytically. Second, when modeling a complex, idiosyncratic policy debate of this nature, it is virtually impossible to assemble a database rich enough to estimate statistically the large set of utility function parameters that determine stakeholder preference between competing objectives. We thus need to assess the sensitivity of predictions to a wide range of uncertain parameters. Finally, we focus on situations where there is considerable uncertainty about the complex, relatively unstructured political terrain within which conflicts will have to be resolved. We propose a general approach to addressing all three challenges. This approach involves constructing a numerical political economic model that can identify policies satisfying a relatively coarse “political viability” criterion and using Monte Carlo simulation to assess the robustness of the model’s predictions with respect to our large set of imprecisely known parameters.

Our approach is closely related to the “robust decisionmaking” approach developed to evaluate problems characterized by “deep uncertainty.” Deep uncertainty refers to situations where the researcher or affected parties cannot agree on how to characterize the problem in question in one or more of the following ways: the appropriate set of conceptual relationships defining the

problem and potential solutions, the probability distributions that represent uncertainty about key relationships and parameters, and/or the desirability of alternative outcomes (Lempert; 2002).¹ In robust decisionmaking, computer simulations are used to generate a large ensemble of outcomes, each based on a specific model. Rather than interpreting the results using summary statistics of realized outcomes, as one would in a Monte Carlo setting, the results are interpreted as representing modeling uncertainty. If a potential solution performs well for a substantial share of the simulations then it is deemed robust. Lempert (2002) argues that robust decisionmaking does not need to be based on a model known to make reliable forecasts. Rather, the model must be capable of identifying key players, relationships, and potential states of the world well enough to identify which potential strategies are likely to fare well under a wide range of specifications. At the same time, the potential values of the individual elements of each specification are limited to realistic ranges (Lempert; 2002). These ranges can be defined using expert opinion or other information.²

We adopt a similar approach. In this political economic context, just as in a decision-theoretic context, the value of a single optimal solution based on a single model specification is less useful, the more sensitive is the solution to uncertainty regarding the model specification (Lempert et al. (2006)). An appropriate model may be sufficiently complex that a single specification cannot be useful because the effects of the many assumptions it incorporates cannot be disentangled from each other. Furthermore, probabilities play two distinct roles: first, the conventional one of representing the likelihood of realizations of states of the world, or known uncertainty; and second, the provision of a framework for summarizing information about the effect of modeling uncertainty on the performance of specific policies according to specific criteria (Lempert et al.; 2004). An important distinction is that in the decision-theoretic context, the emphasis is on uncertainty about how to model the impact of policies on outcomes of concern to stakeholders, while our emphasis is on uncertainty about how to model the stakeholders' preferences and the political process itself.

Our analysis is also related to the "robust control" and the "info-gap" literatures. Robust control is a means of modeling ambiguity-averse preferences (Hansen and Sargent; 2001) that has been applied to a number of environmental and natural resource problems.³ Info-gap theory, which

¹ Deep uncertainty is closely related to the distinction between situations of "risk" and of "uncertainty" introduced to economists by Frank Knight.

² Methodologically, robust decisionmaking is very closely related to multi-model analysis and perturbed physics analysis, which have been used extensively to model climate change, among other applications (For examples of this literature, see Murphy et al. (2004); Piani et al. (2005); Stainforth et al. (2005); Rougier (2007); Dessai et al. (2009)).

is designed to identify policies that decisionmakers can be confident will meet an acceptability criterion (Ben-Haim; 2006), has been applied to environmental issues by Stranlund and Ben-Haim (2008).

In the next section, we present our approach in a general setting. We then apply the approach to a detailed case study of the debate over the future of California's Sacramento-San Joaquin Delta and illustrate the types of questions and answers the approach can give. Finally, we sketch the application of the method to several other water and ecosystem policy debates to illustrate its flexibility.

2. PREDICTIVE POLITICAL ECONOMY MODEL

Our model is designed to assess the impact of particular exogenous factors on the outcome of an imperfectly understood political process. The predictive approach we use has three key steps. First, we construct a model of the political process under consideration. We incorporate our uncertainty about this process by including a number of uncertain parameters in the model and by focusing on models that yield predictions about sets of policies that *might* emerge rather than point predictions of the policy that *will* emerge. Second, we use Monte Carlo simulations to generate probabilistic predictions from our model for a large number of draws from the set of plausible parameter values. Finally, we compare the probabilistic predictions under different values of the exogenous variables of concern. Since our model yields set predictions rather than point predictions, we suggest two approaches to summarizing the probabilistic information about these sets.

We begin by developing a highly stylized description of the specific components of the model and describe the role they play in our analysis. First, we let λ describe the *exogenous factors* of interest. The political process includes several *stakeholders*, and results in the selection of a specific policy vector \mathbf{x} from a set of *policy options* \mathbb{X} . The impact of policy vectors on stakeholders' well-being is influenced by a vector \mathbf{z} of *model parameters* about which researchers have imprecise information. The vector \mathbf{z} is drawn from the set of possible parameter values, \mathbb{Z} . A *political prediction concept* describes which policies might emerge from the process as a function of how they affect individual

³ Examples include environmental and energy planning (Babonneau et al.; 2010), extractive fisheries and water allocation (Shaw and Woodward; 2008); (Woodward and Tomberlin; 2014), water management (Roseta-Palma and Xepapadeas; 2004), and biodiversity (Vardas and Xepapadeas; 2009).

stakeholders. For a particular policy process, this concept identifies the set of policy vectors $\mathbf{W}(\mathbf{z}, \lambda)$ deemed politically viable for the particular parameter vector \mathbf{z} given the exogenous factors λ . Our central goal is to assess how this set changes as λ changes.

This question is challenging to answer for two reasons. First, the answer will depend on the imprecisely known parameters in \mathbf{z} . Second, our political prediction approach yields sets of viable policies rather than point predictions, requiring more complex characterization of how our predictions change. In the settings we consider, there will not generally be a closed-form representation of $\mathbf{W}(\mathbf{z}, \lambda)$. We therefore develop a numerical model and use Monte Carlo simulation to draw \mathbf{z} vectors from \mathbb{Z} . Looking across these simulations, we present two types of results. First, we summarize by policy and compute a *probabilistic political viability* function $V(\mathbf{x}, \lambda)$ which tells us the probability computed over possible realizations of our modeling uncertainty that a particular policy is politically viable. Second, we use summary statistics $\mathbf{M}(\mathbf{W}(\mathbf{z}, \lambda))$ to describe the feasible set for different \mathbf{z} vectors. These summary statistics include the size of the feasible set and the mean values of policy vectors within the set. We can then assess the influence of the exogenous factors λ on the political process by examining how the estimates of political viability for a particular policy depend on the exogenous factors (first type of results) or on how our characterizations of the viable sets depend on these factors (second type of results).

2.1. Utility Functions. In most applications, stakeholders derive utility not from a particular policy *per se*, but from a vector \mathbf{y} of outcomes that might result if the policy was implemented. The relationship between policies and outcomes will generally be probabilistic and will depend on the future state of the world s . The outcome of policy \mathbf{x} conditional on the state of the world s and the exogenous factors λ is denoted by $\mathbf{y}(\mathbf{x}, s, \lambda)$. The distribution of possible states may be influenced by the imprecisely known model parameters; we thus write the probability density function over possible states as $h(s, \mathbf{z}, \lambda)$.

Since the outcome of any policy choice is uncertain, the representation of stakeholder utility must reflect preferences over uncertain outcomes. The model can accommodate a variety of preferences. In our applications, we assume that stakeholders are expected utility maximizers and thus define

expected utility of each individual stakeholder given \mathbf{z} and λ as

$$\mathbf{E}u(\mathbf{x}, \mathbf{z}, \lambda) = \int \mathbf{u}(\mathbf{y}(\mathbf{x}, s, \lambda,)) h(s, \mathbf{z}, \lambda,) ds.$$

Note that the uncertainty we are particularly concerned about in our modeling approach is *not* the uncertain future state of the world. Our central concern is our uncertainty as modelers regarding exactly what decisions stakeholders will make. While the decisions they make will be influenced by the uncertainty about the future state of the world, our predictions about their decisions are also influenced by our uncertainty regarding the precise structure of the model and the values of the elements of \mathbf{z} .⁴

2.2. Political Prediction Concepts. One element of our uncertainty about the stakeholders' decision is that it is often difficult to construct a model of the political process that is sufficiently detailed to yield credible predictions of the specific policy that will emerge from the process. This difficulty arises even if it were the case that all of the specific values of the elements of \mathbf{z} were known with certainty. To address this concern, we broaden our predictions from the identification of single points to the identification of sets of policies that are deemed viable. The identification of what makes a policy viable is unique to the particular setting. In our detailed application to the Delta debate, we focus on policies that Pareto dominate a default outcome.

To define the default, we specify what we expect to happen if the political process ends in a stalemate. The default could be the outcome of any non-cooperative game, such as a Prisoners' Dilemma. We then look for all Pareto improvements on the default.⁵ If each individual stakeholder has the power to derail the political process, the only feasible outcomes are the default outcome itself and the set of other policies that Pareto dominate this outcome. If the set of Pareto improvements is empty for a particular \mathbf{z} , then we predict no agreement. However, the existence of Pareto improvements does not guarantee that stakeholders will successfully agree to implement one of these solutions.

In other settings, it may be more appropriate to use different prediction concepts. For example, in one of our secondary applications, we suggest using the core of a cooperative game as the political

⁴ If the value of a parameter is not known by the stakeholders, it is part of s .

⁵ Pareto optimal outcomes are only a subset of outcomes which Pareto dominate the default; a solution may Pareto dominate the default and not be Pareto optimal.

prediction concept. If a subset of stakeholders can impose a solution on other parties, we could identify the set of policies preferred to the default by a large enough group of stakeholders to impose a solution. If any potential solution would have to pass a referendum, we could identify voters with different stakeholder groups and look at policies preferred to the status quo by at least 50% of the population.

The chosen political viability prediction concept is represented by the correspondence $\mathbf{W}(\mathbf{z}, \lambda)$ which returns the set of all possible policy vectors in \mathbb{X} that satisfy the chosen political viability criterion. Our ultimate interest is in understanding how political viability is affected by the exogenous factor(s) included in λ by comparing $\mathbf{W}(\mathbf{z}, \lambda_1)$ and $\mathbf{W}(\mathbf{z}, \lambda_2)$. However, our prediction of which policies are viable may depend in important and possibly unanticipated ways on the values of the large number of imprecisely known parameters in \mathbf{z} . We consider two possible ways to summarize the variability of $\mathbf{W}(\mathbf{z}, \lambda)$ across different values of \mathbf{z} and examine changes as λ changes.

2.3. Probabilistic Viability of Policies. Our first summary approach considers the probability that individual policies are politically viable. We define a probabilistic political viability function

$$V(\mathbf{x}, \lambda) = Pr_{\mathbb{Z}}(\mathbf{x} \in \mathbf{W}(\mathbf{z}, \lambda)). \tag{1}$$

that gives the probability computed over possible realizations of \mathbf{z} that policy \mathbf{x} satisfies our viability criterion. To evaluate $V(\cdot)$ numerically for any given policy, we draw random vectors \mathbf{z} from the set of possible parameters and calculate the set of policies $\mathbf{W}(\mathbf{z}, \lambda)$ that meet the political viability criterion. We then calculate the fraction of these draws for which a given policy is an element of $\mathbf{W}(\mathbf{z}, \lambda)$ to calculate an approximation for $V(\mathbf{x}, \lambda)$. One way to summarize the effect of λ on political viability is thus to look at how $V(\mathbf{x}, \lambda)$ varies with λ .

2.4. Robust Characterization of Viable Sets. Our second approach looks at how characterizations of the individual viable sets $\mathbf{W}(\mathbf{z}, \lambda)$ vary across \mathbb{Z} . Let $M_j(\mathbf{W}(\mathbf{z}, \lambda))$ be a summary statistic about the set. Such statistics could include the size of the set, its expected value along some dimension of the policy space \mathbb{X} or other measures of its dispersion or shape. The numerical solution procedure described in the previous set allows us to look at the distribution $h_M(M_j; \lambda)$ of these summary statistics induced by the distribution of possible parameter vectors \mathbf{z} . Comparing these

distributions for different values of λ provides another way to look at how λ influences political viability.

2.5. Caveats and Limitations. Implementation of the approach sketched above is subject to a number of caveats and limitations. First, the modeler must have sufficient knowledge of the policy problem, stakeholders' objectives, and the underlying economic, biological, engineering and physical systems which constrain the set of potential solutions to the negotiation problem. Omitting an important stakeholder or key dimension of the policy space could result in model outcomes regarding policy viability that are fundamentally distorted. Similarly, including too many dimensions of secondary importance can add unnecessary complexity and muddy insights. Second, the political prediction concept should correspond to the nature of the negotiation problem. For example, if a simple majority is required, then the model should not require unanimity. Similarly, the major characteristics of the default outcome must be identified. Third, implementing any model of a complex policy problem invariably requires making assumptions about functional forms and parameter values. Allowing unknown parameters to vary and evaluating the robustness of the political viability of solutions to varying parameter values, as done in this approach, requires making assumptions about the distribution of possible values, although the outcome is less restrictive than selecting single values. Finally, the model is not a structural model of the negotiation process. It neither imposes a solution concept nor explicitly considers strategic behavior by stakeholders. All of these considerations suggest that designing a predictive model of a complex negotiation process is subject to many limitations. While the approach is still subject to these limitations, our more modest objective recognizes them explicitly. Rather than focusing on the structure of the negotiation process or its specific outcome, however, the objective here is to assess which potential solutions are more likely to prove viable over a wide range of specifications, as with robust decisionmaking.

Recognizing these issues, we turn now to applying the model to a specific case study: California's Sacramento-San Joaquin Delta.

3. CASE STUDY: POLITICALLY VIABLE SOLUTIONS TO CALIFORNIA'S DELTA CRISIS

In this section, we illustrate our approach by using a political economic model of the Delta debate to investigate how institutional mistrust affects the political viability of different policies. The Sacramento-San Joaquin Delta serves two critical needs for California: ecosystem services in the form of habitat for many species and water infrastructure as the Delta serves as the hub of the state's water export system, conveying fresh water from the Sacramento River watershed to diversion pumps in the southern Delta. While there has always been some competition between these needs, the conflict between them has intensified in recent years. According to scientists, the water supply system has affected the Delta ecosystem. Fish populations have crashed, and multiple species are listed as threatened or endangered. Lawsuits filed under the Endangered Species Act (ESA) have led to dramatic cuts in water exports (United States District Court; 2007), which in turn have contributed to rising unemployment rates in many agricultural regions reliant on the Delta for water. A multi-year drought has accentuated these difficulties. Moreover, the aging levees protecting the man-made Delta islands are at risk of failure, including catastrophic simultaneous failures due to earthquakes on the region's faults.

California stakeholders are actively debating the Delta's future. This debate centers around two critical questions: first, how much water can be exported from the Delta watershed without violating its economic and ecological integrity and, second, should the state build tunnels (or another isolated conveyance structure) that would deliver water from the Sacramento River directly to diversion pumps, avoiding the Delta entirely?

Several independent studies have considered how the state should respond to the Delta crisis (Bay Delta Conservation Plan; 2007, 2009; Blue Ribbon Task Force; 2007, 2008; Cooley et al.; 2008; Delta Vision Committee; 2008; Lund et al.; 2007, 2008, 2010; Michaels; 2012). There has also been a literature focusing on the regulatory institutions involved in policy issues regarding the Delta as a case study of collective governance (Holley; 2015; Kallis et al.; 2009; Fullerton; 2009; Hanemann and Dyckman; 2009).

Madani and various coauthors have studied the Delta problem from a variety of modeling perspectives, including bargaining theory, multi-criteria decision-making and social choice theory Madani

and Lund (2011); Madani et al. (2011); Madani and Lund (2012); Shalikaran et al. (2011); Mokhtari et al. (2012). With some exceptions, the conclusions of these papers are broadly consistent with ours, although the details and methods of analysis are quite different. In particular, the papers listed above are all “positive” in orientation, in the sense that they focus on identifying solutions to the Delta problem, while our orientation is “negative,” in the sense that we focus on *eliminating* potential solutions, on the grounds that they are Pareto dominated by the default. In this sense, our approaches are complementary. In Madani and Lund (2011, p. 615), the authors conclude that the status quo is unsustainable, and “is likely to be replaced by a more stable solution.” The most likely alternative to the status quo, they believe, is tunnel construction; they are, however, more optimistic than we are that a negotiated tunnel solution will be reached. We both agree that a solution involving no water exports is unlikely. The main theme of Madani and Lund (2012) is that the Delta problem, viewed through a game-theoretic lens, has evolved from a Prisoner’s Dilemma-type game to a game of Chicken:⁶ both parties would prefer to concede, and adopt an Alternative Plan (Delta restoration or a conveyance facility), rather than maintain the status quo, but each has an incentive to delay until the other party concedes, and bears the full cost of the Alternative Plan. They conclude, as do we, that government intervention will likely be necessary, in order to avert a Delta catastrophe. In Madani et al. (2011), the authors apply to the problem a novel variant of bargaining theory, known as “Fallback Bargaining:” “bargainers start by indicating their preference orders over the alternatives. Then they fall back, in lockstep, to less and less preferred alternatives until they find an alternative on which all bargainers agree” (p. 192). Once again, the tunnel solution emerges as the most likely outcome. The paper also concludes that the outcome “no agreement” is unlikely. In contrast, we find that if mistrust is high, there are many parameterizations of the model for which no agreement is possible.⁷ As in Madani and Lund (2011), Mokhtari et al. (2012) adopts a Multi-criteria decision making approach to the problem, but in this paper the focus is prescriptive rather than descriptive. They consider the problem facing a benevolent dictator tasked with finding the fairest solution, using five different criteria for fairness. Under all five criteria, the tunnel emerges as the best outcome, followed by dual conveyance. An alternative prescriptive approach is adopted in Shalikaran et al. (2011), which applies a Monte-Carlo social choice methodology to identify a socially optimal solution to the Delta problem, under

⁶ Conventionally, games of chicken (technically known as “wars of attrition”) are modeled as dynamic games (see Hendricks et al. (1988) for a classical treatment). To simplify matters, Madani and Lund (2012) models the chicken game as a one-shot game, while explaining heuristically how the game would evolve in a formal, dynamic framework.

⁷ See Figure 5 below.

a variety of social choice rules. Once again the tunnel is the best outcome, for all but one of the rules considered.

Existing analyses exclude a key component of stakeholder interactions: a lack of trust that negotiated solutions will be implemented. A lack of trust among stakeholder groups often characterizes the class of complex policy problems our modeling approach is designed to address. In this application, one of the obstacles to reaching a solution regarding water exports and the Delta is that key stakeholder groups have expressed serious mistrust in the institutions that would implement any solution. A significant source of concern among environmental groups is that the tunnels will not be operated in accordance with environmental protection laws and the ESA. These groups have expressed a fundamental lack of trust in existing water management institutions, noting that these are the same institutions that failed to prevent the current crisis. Other opponents, including in-Delta water users and Sacramento Valley water users, also express mistrust.

We now describe briefly the policy options commonly discussed for the Delta, and then describe how we implement our modeling approach for this case study. Among the studies of the Delta, Lund et al. (2008) was particularly influential, and was the technical base for Lund et al. (2010). Our analysis draws heavily on it for parameterizations. Its authors argue that there are four basic strategies available to the government: stop exporting water from the Delta altogether, invest in reinforcing the Delta's levees and continue exporting water through it, build a canal, tunnel or other isolated conveyance system to carry exports around the Delta, or combine the last two alternatives in a dual conveyance system where some water is exported through the Delta and some around it using an isolated conveyance system.

Both the tunnel only and the dual conveyance options contemplate the construction of a large isolated conveyance structure that could in principle increase the total export capacity of the system. Many groups, including in-Delta interests and some environmentalists, have repeatedly expressed concerns that if such a structure were built, its entire capacity would be maximally utilized regardless of stated intentions to lower the level of exports.⁸ This concern has been exacerbated

⁸ Blogger Dan Bacher voices a widely held view: "Although the Delta Vision Task Force's report recommended that less water be exported out of the Delta to help the estuary's collapsing ecosystem, canal opponents note that the construction of a canal with increased water export capacity would inevitably be used to export more water out of the system.... I have repeatedly asked canal advocates to give me one example, in U.S. or world history, where the construction of a big diversion canal has resulted in less water being taken out of a river system. I have also asked them to give me one example, in U.S. or world history, where the construction of a big diversion canal has resulted in a restored or improved ecosystem. None of the canal backers have been able to answer either one of these two questions" (Bacher; 2009).

by calls from engineers to choose the tunnels' capacity to match engineering constraints rather than to implement any particular export level. The reason for this approach is compelling: large tunnels would provide maximum flexibility to time export flows during the least environmentally damaging time periods. However, the approach would build in substantial excess capacity, fueling fears of exports greater than those agreed upon. More broadly, Holley (2015) documents that many actors fear that stakeholders will invalidate all or part of any agreement using the court system. These fears are a source of mistrust regarding other stakeholders' commitment to comply with agreed-upon export levels after the tunnels are constructed. Limiting the maximum capacity of the tunnels limits the maximum exports in the event stakeholders do renege on their commitments.

3.1. A Computable Political Economic Model of the Delta Debate. In this subsection, we apply the approach described in Section 2 to predict the influence of institutional mistrust on the political viability of the different approaches to resolving the Delta crisis. We capture institutional mistrust by letting the exogenous factor λ measure the probability that water managers will renege on the agreement about export levels, and an isolated conveyance will ultimately be filled to capacity. Our model includes five broadly specified stakeholder groups: urban users of exported water, the agricultural regions of the San Joaquin Valley that rely on exported water, environmentalists, state taxpayers, and in-Delta interests. These groups have conflicting concerns about the financial, ecological, and employment impacts of the possible options available to the government. Since we are interested in generating predictions that are robust to our uncertainty, the parameters in each group's utility functions are included in our uncertain parameter vector \mathbf{z} .

3.1.1. Possible Policies. Each of the four solutions described above can be represented as a vector $(x_{\text{ex}}, x_{\text{shr}})$ where x_{ex} is the total amount of water exported and x_{shr} is the share of exports routed around the Delta through the twin tunnels.

Our simulations consider values of x_{ex} between 0 and 7.5 million acre feet (maf).⁹ x_{shr} varies from zero to one. Prior to the court-mandated cutbacks, exports averaged approximately 6 maf; we refer to this as the pre-2007 export level. Tunnel size is not a policy choice in our model; we assume that if constructed tunnels will be sized based on engineering considerations as recommended by Lund

⁹ Lund et al. (2008) identify 7.6 maf as the maximum level of exports consistent with minimum flow constraints on the Sacramento River.

TABLE 1. Costs of various export regimes and how they are allocated

Cost	Allocation	Dependence on Export Regime
Reduced exports	Agricultural and urban water users	Increases as total exports decline
Water treatment	Agricultural and urban water users	Increases as total amount of water exported through Delta increases
Levee maintenance	Primarily taxpayers; also agricultural and urban water users	Increases as total amount of water exported through Delta increases
Repair	Taxpayers	Constant, but only paid if no canal is built, either initially or after disaster
Canal Construction	Mostly water users; also taxpayers	Constant for any level of canal exports; also paid if canal is built following collapse
Collapse	Water users and taxpayers	Constant, but only paid if no canal is built

et al. (2008). The parameterization is self-explanatory for each of the strategies except the dual conveyance alternative. Because the report does not include a precise description of how a dual conveyance plan would allocate exports between the isolated conveyance system and through-Delta pumping, we match their outcomes estimates with our policies by assuming their estimates apply to exports divided evenly between the two export paths.

3.1.2. *Stakeholders and utility functions.* The outcomes of interest to our stakeholders include several different costs that are allocated among stakeholders, the possibility of fish extinction, agricultural employment in the San Joaquin Valley, and inflows to the Delta. Different export regimes impose different types of costs that are shared among three stakeholder groups: agricultural users, urban users and taxpayers.¹⁰ The model includes five specific costs: costs due to reduced water exports, water treatment costs, levee maintenance costs, repair costs following a major collapse, and costs associated with a major collapse of the levee system.¹¹

Table 1 summarizes the allocation of these costs and the key pathways through which the policy vector \mathbf{x} influences them. Several of the costs are borne only in the event that tunnels either are

¹⁰ Lund et al. (2008) report upper and lower bounds on the total cost and environmental impact for each alternative. These values were derived from upper and lower bounds on a number of specific input variables. We use the same upper and lower bounds for these inputs. Because our interest is in the negotiation between stakeholders, we disaggregated total costs into costs paid by each stakeholder. Because we model continuous policy decisions regarding total exports and the share of exports routed through the twin tunnels, we made some additional assumptions to allow us to predict costs for variations in the share of exports that flow through the tunnels and for lower export levels than considered by Lund et al. (2008). These adjustments are described in detail in the appendix.

¹¹ The costs included in the cost of a levee collapse are the costs of a sudden disruption of water supplies during the transition period until either the levees are repaired or tunnels are constructed. They do not cover all potential consequences of a collapse, most importantly the costs to in-Delta interests of a catastrophic collapse. In-Delta concerns about costs associated with levee collapse are captured by including levee maintenance in their utility function as described on p. 13.

or are not constructed. As a result, the mapping from policies to outcomes is discontinuous when moving from $x_{\text{shr}} = 0$ to $x_{\text{shr}} > 0$.

The following list introduces the arguments of each stakeholder group’s utility function. We use the two letter code after each group’s title as shorthand to identify groups. Each group’s preferences over outcomes induce preferences over policy variables, in sometimes complex ways. Each stakeholder group has a CES utility function defined over the components of the outcome vector that we assume affect its utility, with the exception of environmentalists as discussed below.

State taxpayers (Tp): Taxpayers are concerned with reducing the government’s total expenditure liability and are risk neutral. The two major (variable) determinants of the government’s liability are the cost of levee maintenance, which increases with the amount of water exported through the Delta, and the costs of a major collapse, borne only if tunnels do not exist. Thus Tp’s utility is increasing in x_{shr} , decreasing in x_{ex} , and increases discontinuously when moving from $x_{\text{shr}} = 0$ to $x_{\text{shr}} > 0$.

Urban users (Ur): This group is an aggregate of urban interests in Southern California and the San Francisco Bay Area. It is concerned with minimizing the cost of meeting its water supply needs. Delta exports are cheaper than alternatives, so urban user utility increases with x_{ex} . Moreover, both water treatment costs and the probability of cutbacks for ecosystem protection decrease as water exports are shifted from the Delta to the twin tunnel, so Ur’s utility also increases with x_{shr} .

Agricultural users (Ag): This group represents agricultural interests in the San Joaquin Valley that rely on water exported through the Delta. The two arguments in Ag’s utility function are farming profits and agricultural employment. Ag’s preferences are very similar to those of Ur, although Ag’s utility decreases faster than Ur’s as x_{ex} falls because Ag’s profits and agricultural employment both decline.

In-Delta interests (Dt): This group is a composite of local residents, farmers, and recreational users within the Delta. The two arguments of Dt’s utility function are Delta inflows and levee maintenance. The first argument proxies the quality of water in the Delta, which is highly correlated with Delta inflows. In the absence of tunnels, Delta inflows are determined by factors exogenous to the model—hydrological variables and upstream diversions.¹² If

tunnels are built, then any water exported through them would reduce inflows into the Delta. The second argument, levee maintenance, is a function of the amount of water exported through the Delta: any agreed-upon policy package will allocate funds for levee maintenance according to a formula that increases with through-Delta exports. Both impacts imply that Dt's utility decreases with x_{shr} . The impact of increasing x_{ex} depends on the value of x_{shr} . At high values of x_{shr} increasing exports reduces Dt's utility due to reduced inflows; at low values of x_{shr} , Dt's utility increases with exports due to increased levee maintenance.

Environmentalists (Ev): Ev is concerned exclusively with the survival of two fish species: Delta smelt and salmon. We define four state-dependent utility levels for Ev, similar to Woodward and Shaw (2008). If both species survive, Ev's utility is 1; if neither survive it is zero. If only one of the two species survive, Ev's utility is some number between 0 and 1, depending on which species survives. Ev's expected utility thus increases as the probabilities of survival increase and hence decreases with x_{ex} . Following Lund et al. (2008), we assume that for a given level of x_{ex} , survival probabilities are higher when a relatively large share of those exports flows through an isolated conveyance. The precise dependence of survival on x_{shr} is described in more detail in the results section.

All of the outcomes defined above (various costs, the possibility of fish species extinction, agricultural employment and Delta inflows) are contingent on the state of the world s , which incorporates all of the uncertainty that stakeholders confront. This uncertainty includes uncertainty about how fish populations will respond to changes in water export regimes (Hanak et al.; 2013), uncertainty about the level and/or allocation of various costs, the unknown occurrence and timing of a major levee collapse, and the possibility that the export regime will eventually diverge from the negotiated regime for one of three reasons: possible changes in exports following a levee collapse, court mandated cutbacks due to the ESA, or political pressure from water user groups.

Each participant in the political process has a utility function defined over outcomes that depends on the parameter vector \mathbf{z} and makes decisions based on the expected utility over possible states of the world. In the appendix, we provide a detailed description of all the uncertainty facing

¹² California's multi-year drought suggests that additional elements of stakeholder uncertainty could include expectations regarding the effects of climate change on average inflows and inflow variability. We omit those considerations here; the likely primary impact on the analysis would be to increase users' willingness to accept lower water exports, but the qualitative results would be unaffected because groups' fundamental interests would not change.

stakeholders, how we map a particular policy to possible outcomes, and the precise specification of each group’s utility function.

3.1.3. *Political Prediction Concept.* The political negotiations over the future of the Delta are complex. We do not attempt to explicitly model the strategic interactions between parties. Instead, we focus on identifying policies that are or are not “politically viable.” *Sacramento Bee* political reporter Dan Walters has commented on the “unwritten rule” in California that “any major policy decree must have virtually unanimous support from every stakeholder group or it will ultimately fail because opponents have so many political ways to kill it” (Walters; 2010). This implies that any negotiated solution to the Delta crisis must be acceptable to all stakeholder groups.¹³ Given a “default outcome” that will be implemented if the participants in the political process cannot negotiate an agreement, the set of politically viable policies $\mathbf{W}(\mathbf{z}, \lambda)$ is thus the set of alternatives that Pareto dominate this outcome when the model is parameterized by \mathbf{z} given mistrust level λ . If $\mathbf{W}(\mathbf{z}, \lambda)$ is non-empty, it is possible, but no means certain, that the stakeholders will reach an agreement.¹⁴

The default outcome in our model has deterministic and random components. Like Madani and Lund (2011), we assume that if stakeholders cannot agree on a policy alternative, then no major policy change will be implemented. As a result, no tunnels will be built and no money will be spent on maintaining the levees, so the probability of a massive levee failure will increase. Apart from a massive levee failure, the issue of primary concern to stakeholders under the default outcome is the level of water exports. This level is uncertain. As in the case of agreement, the actual level of

¹³ An alternative, non-cooperative game-theoretic approach to the Delta negotiations is provided by Madani and Lund (2011), who consider six “stability” concepts: one is Nash equilibrium, the other five extend Nash’s notion in ways that are more realistic in various ways. Their payoffs assume that movement away from the status quo requires agreement from all parties. Their analysis predicts that in many cases, stakeholders will fail to reach an agreement despite the existence of Pareto improving policies.

¹⁴ As a criterion for political viability, Pareto dominance has an obvious shortcoming: each stakeholder in the model is assumed to have veto power over the decision-making process. In this respect, it is a flawed representation of virtually every actual political process: either it endows some modeled stakeholders with more power than they actually have, or it excludes from the model stakeholders who, though lacking veto power, may have considerable political influence. Our model may overstate the power of in-Delta interests to veto an agreement and omits another interest group with some degree of power in the negotiations – agricultural users upstream of the Delta – due to lack of data. In the former instance, the set of politically viable options will be underestimated; an option can fail to meet our criterion because it is unacceptable to some stakeholder that in the real world would lack the political clout to block it. In the latter instance, the set will be overestimated; it will include policy options that are acceptable to all of the stakeholders with veto power, but in the real world would not survive the combined opposition of multiple stakeholders, none of whom had the political power to veto the outcome unilaterally. It is nonetheless a helpful exercise to identify the Pareto dominant set. In particular, as section 4 will demonstrate, it can be especially instructive to learn that certain highly publicized possibilities fail to satisfy even this relatively modest selection criterion.

default exports will depend on whether or not fish populations show signs of recovery and whether export reductions are imposed if they do not.

The vector of default outcomes, contingent on the state of the world s and the parameters vector \mathbf{z} , is given by $\mathbf{y}^d(s, \mathbf{z})$ and the default distribution function over possible outcomes is given by $h^d(s, \mathbf{z})$. By definition, both functions are independent of any policy vector and of the level of λ since no tunnels will be built in the default. The vector of expected default utilities is given by

$$\mathbf{Eu}^d(\mathbf{z}) = \int \mathbf{u}(\mathbf{y}^d(s, \mathbf{z}), \mathbf{z}) h(s, \mathbf{z}) ds$$

and the Pareto dominant (PD) set $\mathbf{W}(\mathbf{z}, \lambda)$ for a given \mathbf{z} and λ is thus the set of all policy options that satisfy

$$Eu_i(\mathbf{x}, \mathbf{z}, \lambda) \geq Eu_i^d(\mathbf{z})$$

for every stakeholder i .

There is no closed form solution for $\mathbf{W}(\mathbf{z}, \lambda)$, so we evaluate it numerically. For a given parameter vector \mathbf{z} , we compute players' payoffs for each policy and for the default outcome for a large number of draws of the state of nature s . We then take expectations over these draws to identify the PD set.

3.2. Modeling Uncertainty. If \mathbf{z} were known with certainty, we could simply repeat the procedure described in the last paragraph for different levels of λ and compare the resulting sets to determine the influence of mistrust on possible negotiated solutions. Our challenge is that these sets vary in important ways on the precise parameterization of stakeholders' utility functions and the default distribution over states of the world. In Table 2, we list the elements of \mathbf{z} for our application along with the upper and lower bounds we considered plausible. Because we have no basis for specifying an informative prior over these intervals, the principle of insufficient reason dictates that each element should be independently and uniformly distributed over the interval we specify as its support.

Specifying an interval for the potential value of an unknown parameter is always arbitrary, to some degree. Similarly, better information regarding values and intervals can be incorporated into a

simulation analysis if it becomes available. Using outside information when available, we set the bounds reported in Table 2 to allow utility functions to exhibit a variety of behaviors while seeking to mitigate the influence of extreme values. Ag and Dt have utilities dependent on two factors. In the absence of specific information on their relative importance to the stakeholders, a natural base point is to assume equal weights, or 0.5 for each factor. We specify a lower bound of 0.2 and an upper bound of 0.8 for the weight placed on each factor, with pairs of weights summing to 1. These bounds encompass a broad set of possible weights and it seems unlikely that the weights would be outside these limits for Ag, a player concerned with the value of agricultural production and employment. Dt values levee maintenance and inflows to the Delta because they both determine his (expected) water supply, and there is no reason to expect one to be dramatically more important than the other. The ranges of elasticities of substitution between the two factors in Ag's and Dt's utilities include both elastic and inelastic values. We set the upper bound on the risk aversion coefficient for all stakeholders to 1 to include the possibility of risk neutral (but not risk-loving) stakeholders. We set the lower bound of 0.2 to include the possibility of very risk averse stakeholders because we know very little about the risk preferences of the different stakeholder groups. In the results, we briefly discuss how tightening these bounds would affect our results. Ev's utility if only one species survives varies from 0.25 to 0.75, compared to his utility of 1 if both survive. This range allows for his utility function to be either submodular or supermodular in the survival of the two species, limiting the conceptual restrictiveness of the selected range. Regarding variation around the mean ("spread" in the table) for the default export distribution, we chose a lower bound of zero in order to allow for the case of complete certainty. The maximum bound allows for exports in the default to be as much as 2 maf above or below the mean. This bound is specified based on the maximum level of exports consistent with flow constraints by Lund et al. (2008), 7.6 maf. Allowing 2 maf as the maximum difference means that realized exports in the default could be above or below the mean by as much as slightly more than 25 percent of the maximum export level.

4. DELTA RESULTS

The model described in the previous section yields predictions about the impact of institutional mistrust on the viability of policies. We first discuss our predictions under the counterfactual assumption that we know for certain that each element of the parameter vector \mathbf{z} is equal to its

TABLE 2. Elements of modeling uncertainty and their distributions

Variable	Lower Bound	Upper Bound
Weight on jobs vs money in Ag utility	0.2	0.8
Weight on maintenance vs inflows in Dt utility	0.2	0.8
Constant in Dt utility	0.05	0.15
Ag elasticity of substitution	0.5	1.5
Dt elasticity of substitution	0.5	1.5
Ag risk aversion coefficient	0.2	1
Dt risk aversion coefficient	0.2	1
Ur risk aversion coefficient	0.2	1
Ev utility if only smelt survive*	0.25	0.75
Ev utility if only salmon survive*	0.25	0.75
Spread of default export distribution above and below mean (maf)	0	2

* Ev utility is scaled so that 0 represents the utility if neither species survives and 1 represents the utility if both survive.

expected value. The results are counterintuitive and potentially misleading. We then incorporate our uncertainty about the true parameter vector and use the probabilistic assessments described in our model section to demonstrate that the impact of mistrust is highly dependent on the precise parameterization and that mistrust itself decreases the robustness of our viability predictions.


4.1. Known Parameters. We begin by setting \mathbf{z} equal to its expected value $\bar{\mathbf{z}}$ and specifying perfect trust ($\lambda = 0$) and discuss the resulting PD set in considerable detail. The set is illustrated in Figure 1. The boundaries of Figure 1 coincide with the boundaries of the policy space. Moving from left to right in the diagram, the total amount of water exports, x_{ex} , increases; moving from the bottom to the top, the percentage of the exports flowing through the tunnels, x_{shr} , increases.¹⁵ The filled stars in the figure depict the four policy alternatives discussed in detail in Lund et al. (2008).

Each contour line depicted in Figure 1 identifies the set of policy options which yield that group expected utility equal to its expected utility from the default outcome. The arrows attached to each constraint line are gradient vectors, pointing into the region of the policy space which the stakeholder group prefers to the default. $\mathbf{W}(\bar{\mathbf{z}}; \lambda = 0)$ is the shaded area.

The shapes of the participation constraints in Figure 1 reflect the preferences of the various stakeholder groups. Ag and Ur will veto policies that involve low levels of exports. Both groups will

¹⁵ Note that the volumetric distinction between vertically differentiated points shrinks to zero as their horizontal location moves to the left of the diagram: in the limit, obviously, there is no distinction between different fractions of zero.

FIGURE 1. The PD set with perfect trust conditional on \bar{z}



trade slightly smaller total exports for larger shares through the tunnels, but are concerned primarily with achieving a high base level of exports; consequently, their participation constraints slope steeply downward. They will also veto policies with very low shares conveyed through the tunnels because they will not find it worth the cost of construction. Tp will veto policy vectors in the lower right corner of the space because the levels of exports through the Delta, and hence expenditures on levee maintenance, will be the highest in this region.

Ev will veto policies that involve high levels of exports, although this group is more willing to accept exports if they are routed at least partially through the tunnels. This reflects the conclusion in Lund et al. (2008) that fish populations are more likely to recover under either a dual conveyance or isolated conveyance system option than if exports are pumped exclusively through the Delta.¹⁶ For Dt , the two outputs which matter—freshwater inflows to the Delta and expenditures on levee maintenance—both decrease with exports through tunnels; hence Dt will veto policies in the uppermost region of the policy space. Dt will accept even very high levels of x_{ex} provided x_{shr} is sufficiently low because levee maintenance expenditures increase with total exports. As x_{ex} increases, there is a decline in the maximum level of x_{shr} that is acceptable to Dt . It is surprising that Dt is willing to accept such a large fraction of the available alternatives, since Delta interests

¹⁶ The precise shape of Ev 's participation constraint is due to the specification of the fish survival probabilities and the assumption that the isolated conveyance option splits exports equally between the conveyance structure and the Delta. It does not drive the results of the analysis.

have always vociferously opposed an isolated conveyance structure. The reason is that these alternatives are being compared to a default outcome that is extremely unsatisfactory in expectation: unless some kind of agreement is negotiated, expenditures on levee maintenance will be minimal, increasing the probability of a major levee collapse, which would be devastating to Dt.

This analysis has three implications for the policy question we consider. First, $\mathbf{W}(\bar{\mathbf{z}}; \lambda = 0)$ is nonempty for the parameter vector $\bar{\mathbf{z}}$ of modeling uncertainty; there do exist policies that Pareto dominate the default. This suggests that if the model, when parameterized by $\bar{\mathbf{z}}$, is a reasonable stylization of the actual political process, *some* negotiated solution might emerge from the political process, although our model does not rule out the possibility of a stalemate leaving us at the default. Second, for a policy to be acceptable to all stakeholders under this parameterization, at least half of all exports will flow through the tunnels. Finally, under this parameterization at least three of the four options discussed by Lund et al. (2008) would be vetoed by some group. Whether or not the fourth option—dual conveyance—would be vetoed depends on how the option would be implemented. It would be vetoed for a large majority of x_{shr} values.

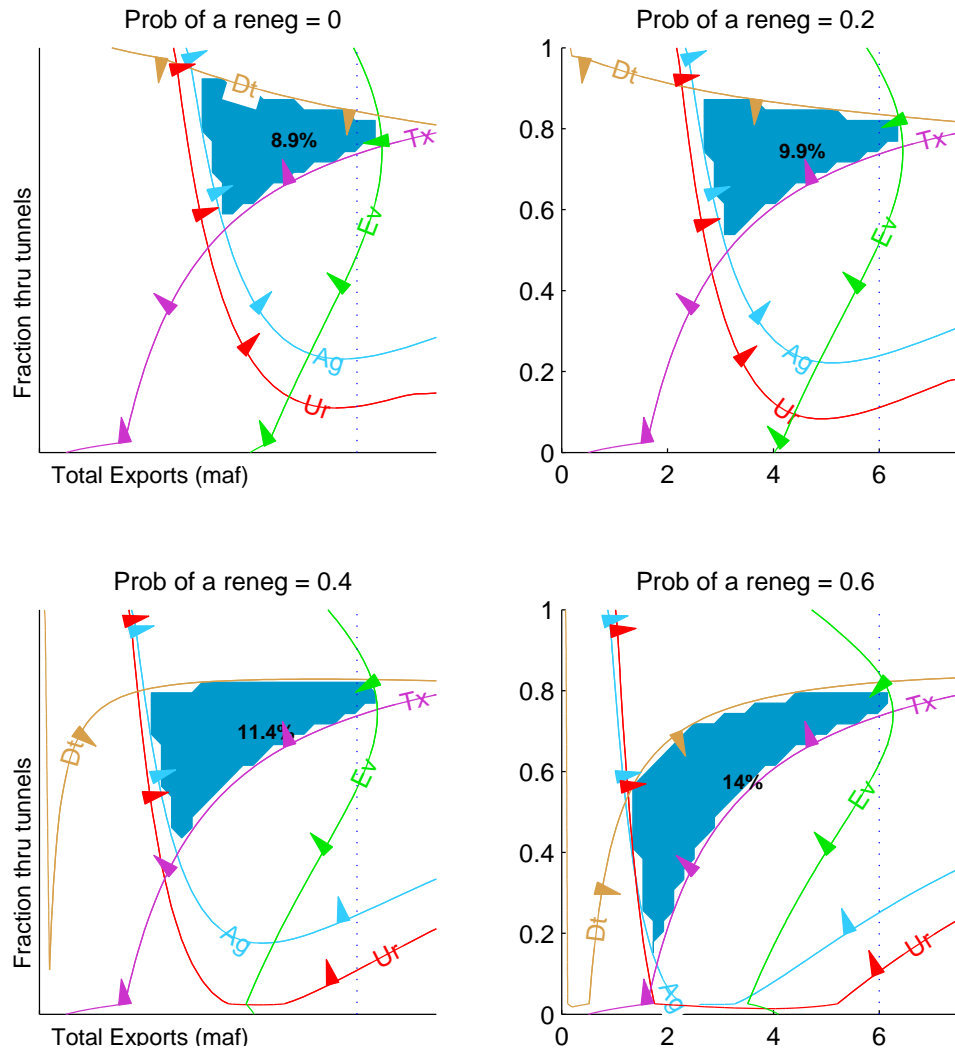
We then ask whether the existence of policies that Pareto dominate the default is contingent upon perfectly trustworthy institutions by progressively increasing the value of λ , again restricting attention to one parameterization of the model, $\bar{\mathbf{z}}$. Figure 2 is the analog of Figure 1 for all four levels of mistrust. The first panel replicates Figure 1. The percentage number printed inside of each set $\mathbf{W}(\bar{\mathbf{z}}; \lambda)$ indicates the size of this set relative to the entire policy space.

All of the costs for which Tp is responsible are independent of the level of mistrust, so the group’s participation constraint is independent of the level of mistrust. The successive increases in mistrust induce shifts in all other stakeholders’ participation constraints and, hence, the location and size of the PD set. The PD set shifts down and left, and increases in size. The acceptable policies are characterized by fewer total exports and a smaller share of those exports through the tunnels.¹⁷

The leftward shift in the PD set is due to participation constraints for Ag, Ur and Ev all moving to the left. Consequently, the PD set excludes outcomes with relatively high total exports that were

¹⁷ The shift toward lower agreed upon export levels depicted in the figure is not surprising, but this figure tells only part of the story. As mistrust increases, so does the probability that water exports will fill the tunnels to capacity. Hence the reduction in *expected actual* exports and the actual share of those exports flowing through the tunnels associated with alternatives in the PD set is less than the figure would suggest; expected values of both may in fact increase.

FIGURE 2. Impact of mistrust on the location of PD set with parameterization \bar{z}



part of the PD set when there was less mistrust, and includes outcomes with smaller total exports. At higher levels of mistrust, Ag and Ur are willing to accept an agreement with lower total exports because they believe they may be able to use political or legal pressure to increase exports through the tunnels in the future despite the agreement. Ev's participation constraint moves to the left for the same reason.¹⁸

The downward shift (reduction in share through the tunnels) is driven primarily by Dt. Greater mistrust reduces the declared share of water transported through the tunnels that Dt is willing to accept given declared total exports. Consequently, Dt's constraint increasingly limits the set of

¹⁸ It also exhibits an increase in curvature that is dependent on the specific functional form for fish survival and has little impact on the results.


possible agreements. When trust is not an issue, Dt is a relatively obliging negotiating partner: the fraction of possible alternatives that this group is willing to accept is higher than that of any other group. As mistrust increases, this fraction declines more than proportionately. Once the probability of a trust violation reaches 0.6, the fraction of alternatives that Dt will accept is strikingly small. The model thus suggests that a critical factor driving Dt’s highly publicized opposition to the proposed tunnels is its strong belief that agreed upon restrictions on exports are unlikely to be honored in practice.

In addition to shifting down and left, the PD set’s size increases monotonically with mistrust, from 8% to 14% of the entire space. While this result seems counter-intuitive at first glance, it can be explained by two factors. First, Ag’s and Ur’s constraints are shifting left at a significantly faster rate than Ev’s, increasing the width of the PD set. Second, the interaction between Tp’s and Dt’s participation constraints causes the PD set to be much narrower at its right-hand edge than at its left-hand edge; at high levels of total exports, the interval of export shares that are acceptable to all parties is much smaller than at low levels of total exports. As a consequence, even if Ev’s and Ag and/or Ur’s constraints were to shift left with mistrust at the same rate, the PD set would increase in size: a “short” column would be eliminated from the set, while a “tall” column would be added. Both factors are clearly dependent on the interplay among a large number of utility function parameters.

4.2. The Effect of Parameter Uncertainty. To determine whether we can be confident that increasing mistrust will increase the size of the viable set and shift it to the left and downward, we repeat the analysis above for 1,000 draws from the set of possible parameter vectors \mathbb{Z} and summarize the results in two ways. We first evaluate the probabilistic political viability function at each element of the policy space, \mathbb{X} . We then measure the size of each PD set and compute the average levels of total exports and tunnel shares within each PD set. We first present the results assuming perfect trust and then discuss how mistrust affects these measures.

4.2.1. Perfect trust. Figure 3 partitions \mathbb{X} into regions depending on the probability that each policy belongs to the PD set. A policy is termed RPV if it Pareto dominates the default for at least 80% of the realizations of modeling uncertainty; such policies are marked in the figure with the largest solid circles. Policies that are politically viable for at least one parameterization are

FIGURE 3. Probability of Pareto dominance with perfect trust



referred to as *politically viable* (PV). The PV set includes all of the points identified with some marker in the figure. Finally, policies in the white (unmarked) region are said to be NPV. (For ease of comparison the solid line is the boundary of the PD set for the parameterization \bar{z} plotted in Figure 1.) Figure 3 is in some sense similar to Figure 1; both suggest that most of the policies that have some chance of emerging from the political process involve export levels less than pre-2007 levels which are routed primarily, but not exclusively, through the twin tunnels. Yet, Figure 3 contains far more information than Figure 1 because it considers the implications of modeling uncertainty for the political feasibility of the examined policies.

One critical difference between the two figures is the interpretation of the unmarked regions of the policy space. Policies in the unmarked region of Figure 1 are Pareto dominated by the default for the single parameterization \bar{z} . By contrast, unmarked policies in Figure 3 are NPV; they are Pareto dominated by the default for *all* realizations of modeling uncertainty in our sample. As discussed in our model specification, we intentionally set very wide bounds on the modeling uncertainty in our analysis to be conservative in identifying policies as NPV. Narrowing the bounds on our uncertainty could expand the white space defining the NPV policies, while expanding the ranges could decrease the white space by expanding the set of PV policies.


A striking property of the figure is that all four policy options described in Lund et al. (2008) are NPV as shown by the locations of the four stars in the unshaded area. All points on the graph's left edge (corresponding to ceasing all exports) and its bottom edge (corresponding to routing all

exports through the Delta) are NPV, not just the two proposed solutions. Lund et al. (2008) were similarly skeptical of all alternatives lacking an isolated conveyance structure, noting that a policy of stopping all exports is simply too expensive for the state despite its environmental benefits, while continuing to rely exclusively on through-Delta exports carries unacceptable risks to both water supply reliability and the ecosystem. Moreover, all points on the top edge of the graph (corresponding to pure isolated conveyance alternatives) are also NPV. These alternatives are always vetoed by at least one stakeholder. If export levels are too low to justify the cost of construction, water users are unwilling to pay for the tunnels. On the other hand, if export levels are too high, Dt will veto any pure tunnel alternative because reduced Delta inflows cause two negative consequences: water quality in the Delta will decline relative to the default and expenditures on levee maintenance will remain at zero. Finally, the analysis suggests that a dual conveyance alternative with pre-2007 exports evenly split between the tunnels and the Delta is NPV, although other dual conveyance configurations are PV.

The set of RPV policies in Figure 3 is considerably smaller than the shaded set $\mathbf{W}(\bar{\mathbf{z}}; 0)$ in Figure 1. While just under 9% of the policy space belongs to $\mathbf{W}(\bar{\mathbf{z}}; 0)$, less than 1% of the policy space is RPV. Put another way, less than 10% of the policies inside the solid line are RPV, although almost all of them satisfy the Pareto criterion with probability at least 40%. There is no policy which Pareto dominates the default for more than 85% of the realizations of \mathbf{z} . This illustrates the obvious point that inclusion in the PD set for the mean realization of modeling uncertainty (or any other realization) is no guarantee of robust political viability.

4.2.2. *Impact of mistrust.* The first panel of Figure 4 replicates Figure 3; the remaining panels show the impact of increasing mistrust, i.e., increasing probability that water managers will renege on the negotiated agreement about water export levels. Legends for these figures are the same as for Figure 3. In each panel, we overlay for reference the boundaries of the shaded set $\mathbf{W}(\bar{\mathbf{z}}; \lambda)$ in the corresponding panel of Figure 2. As mistrust increases, the set of PV policies increases; this effect is driven primarily by the participation constraints for Ag and Ur slackening. As discussed on pp. 21-22, the more likely it is that the tunnels will be utilized to capacity, the more willing Ag and Ur will be to accept somewhat lower values of x_{ex} and x_{shr} . The size of this effect depends on how their utility functions are parameterized. Mistrust thus decreases the number of policies


FIGURE 4. Impact of mistrust on the probability of Pareto dominance



we can be confident will *not* emerge from the political process. This conclusion is consistent with the increasing size of the PD sets in Figure 2. On the other hand, the set of RPV policies shrinks dramatically, virtually disappearing even for $\lambda = 0.2$. In other words, as mistrust increases, it becomes increasingly difficult (and eventually impossible) to identify policies that we can be confident are Pareto improvements on the default. Unsurprisingly, none of the four policy options from Lund et al. (2008) becomes RPV as mistrust increases.

We next explore the impact of mistrust by examining the distribution $h_M(\mathbf{M}, \lambda)$ of three summary statistics about the individual PD sets: their size, the average total export level and the average share exported through the twin tunnels. Box plots of these distributions are shown in Figure 5. In each panel, for each value of λ , the solid horizontal line indicates the median value across our sample from \mathbb{Z} for the measure being plotted. The thick, squat rectangles denote 95% confidence intervals for the population medians. The thin, elongated rectangles denote the interquartile ranges

FIGURE 5. Distribution of set measures across modeling uncertainty



of the sample data, and the whiskers (thin, dashed lines) indicate the support of the sample data.¹⁹ At the bottom of each panel, the filled ovals corresponding to each λ indicate the probability that the PD set is empty; the area of the oval is proportional to the percentage of parameterizations for which the PD sets are empty, given that level of mistrust.

Several conclusions emerge from this figure. First, as mistrust increases, there is a striking increase in the percentage of parameterizations for which the PD set is empty. At low levels of mistrust, we are fairly confident that there are policies that PD the default, suggesting the potential for a negotiated solution. At high levels of mistrust, there is a striking decrease in this confidence as 30% of our parameterizations have an empty PD set when λ reaches 0.6. Yet, even at this very high level of mistrust, our model predicts a non-empty PD set nearly 70% of the time.

The left panel of Figure 5 also tells us that the median size of the Pareto set $\mathbf{W}(\cdot; \lambda)$ varies little with λ , while both the inter-quartile range and the support of the entire sample increase dramatically. This tells us two things: first, the conclusion from Figure 2 that the PD set increases in size is *not* robust to our uncertainty about the parameterization of our model because we see no consistent trend in the size of the PD set. Moreover, mistrust exacerbates our uncertainty about the size of the

¹⁹ As with the probabilistic viability results, the support of these measures is determined by the support of the modeling uncertainty. Contracting the modeling uncertainty ranges could decrease the ranges in Figure 5, while expanding the modeling uncertainty ranges could expand them.

PD set. At low levels of mistrust, the size of the PD sets varies within a relatively small range. As mistrust increases, the size of the PD set becomes more dependent on the specific parameterization. Some of our parameterizations suggest that over 30% of the policy options in our grid are PD, while a large fraction suggest an empty PD set.

The center and right panels tell us that we can be reasonably confident that the PD sets do indeed move down and to the left as suggested by Figure 2. Moreover, the middle panel shows that mistrust has little impact on the horizontal dispersion of the PD sets; the lengths of the interquartile ranges and the sample data support remain more or less constant as mistrust increases. Mistrust thus has little impact on our confidence in our ability to predict the total export levels included in the PD set. However, in the right panel, we see that the vertical dispersion of the PD sets increases with mistrust, especially at high levels. As mistrust increases, we become much less confident in how much of the negotiated exports will flow through the canal. It thus becomes much harder for us to identify specific policies we are confident PD the default, even as we continue to predict the existence of a PD set in most of our parameterizations.

5. OTHER APPLICATIONS

The Delta case study presented in the previous two sections offers a detailed look at how our approach works in practice. The core elements of the approach can be extended to a variety of debates. In this section, we sketch how our method could be applied to several other policy negotiations involving tradeoffs between water supplies, ecosystem benefits, and other objectives that have been modeled by previous authors. In each case, the authors focused on outcomes that satisfy a particular solution concept and asked how some change(s) in political, economic, or institutional factors would change the solution.

Applying our approach to these same negotiations would complement the existing analyses in two important ways. First, by including a large number of model parameters in the uncertain vector \mathbf{z} , we would generate more robust predictions of the impact of the changes considered by the authors. Second, by focusing on a coarser prediction concepts that identify sets of policies that might emerge from negotiations, we generate predictions that are consistent with a broader set of game theoretic models.

5.1. Mekong Basin. Conflicts between countries, water users, hydropower and ecosystems are intense in the Mekong river basin in southeast Asia. Houba et al. (2013) construct a regional hydrological model whose policy space includes the operation and construction of a number of dams, as well as water use by different sectors within China, Thailand, Laos, and Vietnam. Their model traces the impact of these choices on four different groups within each country (hydropower, irrigation, environmental users, and households). Pham Do and Dinar (2014) build on this analysis by considering the impact of expanding the policy space to also include trade liberalization between the countries. Both studies investigate the impact of strengthening the governance of the Mekong River Commission (MRC), an organization of four lower basin states (Thailand, Laos, Cambodia and Vietnam).

Our approach could be applied to generate robust predictions of the impact of strengthening the governance of the MRC. In the terminology of our paper, the governance of the Mekong River Commission is an example of an exogenous factor λ whose impact on the political process we want to assess. It is a somewhat unique exogenous factor in that changes in governance would map to changing the political prediction concept $\mathbf{W}(\mathbf{z}, \lambda)$. Under weak governance, individual policies and extractions would be viable as long as they provided positive net benefits for the country or countries constructing each individual dam and extracting given water amounts. Under strong governance, policies would only be viable if all dams and extractions within the MRC collectively provided positive net benefits for the MRC as a whole.

As in the Delta example, the impact of changing the strength of the MRC's governance will depend on exactly how the stakeholders in the region weight different objectives. In particular, the parameter vector \mathbf{z} would include the weight each country places on the concerns of different stakeholder groups within its borders. Moreover, both Houba et al. (2013) and Pham Do and Dinar (2014) discuss the role that uncertainty about impacts might have on stakeholders. If the underlying model were extended to include uncertainty, the level of risk aversion of the different countries would be included in the parameter vector \mathbf{z} . Applying our approach would yield robust predictions about how the set of possibly viable policy outcomes would change if the governance of the Mekong River Commission were strengthened.

5.2. Upper Guadiana Basin in Spain. Rapid expansion of groundwater fed irrigation in the Upper Guadiana Basin in Spain has led to dramatic drops in water levels throughout the region. Despite legal restrictions on the amount of drilling and pumping, many farmers throughout the region rely on illegal wells to irrigate their farms. Falling water levels threaten the survival of an important wetlands containing a national park. Marchiori et al. (2012) develop a game theoretic model of local negotiations between environmentalists and several different groups of farmers, distinguished by size and water rights, over policies to address reduce water use and illegal pumping. They conclude that strong government action in the form of increased enforcement of existing policies will be required to save the wetlands and provide numerical estimates of the size of fines consistent with stabilization of water levels.

Our methodology could be used to provide a more robust answer to a somewhat broader question: how does the likelihood of meeting a water level target vary with enforcement efforts? In the language of our model, λ would be a measure of enforcement severity. Since Marchiori et al. (2012) model local negotiations using Nash bargaining, the Pareto dominance prediction criteria used in our application in this paper would remain appropriate. We would measure the likelihood of meeting the target by looking at the fraction of the policies in the Pareto dominant set for which water levels meet or exceed the target level. Finally, the parameter vector \mathbf{z} could include all the parameters in the various stakeholders utility functions and many of the parameters describing how both farmers and water levels respond to various actions.

5.3. Jucar Basin in Spain. Kahil et al. (2016) use cooperative game theory to model negotiations over water policies in the Jucar Basin in Spain. They conclude that cooperation among water users reduces the negative impact of droughts but also find that government action will likely be required to protect ecosystems in cases of severe drought.

Our approach could be used to ask how likely it is that ecosystem flows will meet a specified target or to the look at the distribution of plausible flow results, while explicitly recognizing considerable uncertainty about model parameters and the political situation. Since the authors' emphasis was on the role of government action to protect ecosystems, λ would be a measure of the stringency of government protections for ecosystems. The uncertain parameter vector \mathbf{z} would include the parameters governing the different users' payoff functions.

While Pareto dominance would be a reasonable prediction concept, we could also apply a somewhat stronger criterion. In their analysis, Kahil et al. (2016) compare the predictions of several cooperative models but also emphasize that only solutions belonging to the Core of the cooperative game are stable, since there is a coalition that would block any other solution. This observation suggests using the Core as the political prediction concept. We could then ask either how likely it is that a particular flow level is part of the Core or how the likelihood of meeting a particular target varies across plausible parameterizations. Finally, we would vary the stringency parameter λ and look at how the answers to these questions change.

5.4. Adour River Basin. Under the national French water law passed in 1992, stakeholders were required to negotiate rules for implementing water development plans at the catchment level. In the upper Adour River basin, this required negotiation over water quotas and quota prices for three groups of farmers, as well as whether or not one or more of three potential dams would be built. Simon et al. (2007) develop a non-cooperative multilateral bargaining model of this negotiation process utilizing the theoretical framework developed in Rausser and Simon (1999). Each player is a representative member of some stakeholder class. The classes are: the three groups of farmers, whose utilities depend on their quotas and prices for quotas; an environmentalist, who is concerned with both downstream water flows and the negative effects of dams on the landscape; a downstream user concerned strictly with flows; a water system manager responsible for balancing the budget of the water system, who prefers to administer a larger system; and a taxpayer who is concerned with the capital cost of dams.

One of the questions examined by the authors was whether or not a common negotiator (“spokesman”) for the three farmers can improve the utilities they obtain in the non-cooperative equilibrium. The spokesman’s objective was to maximize the average utility of the three farmers. Under this assumption, average farmer welfare generally was higher with a spokesman. Two factors drove this result: first, when he proposed a solution in the negotiation process, each individual farmer would allocate himself a large quota at a low price and charge other farmers high prices for small quotas to meet the system’s operating costs; while this “beggar-thy-neighbor” behavior was individually rational, it was collectively self-defeating; the spokesman, by “internalizing the negative externalities” that individual farmers were imposing on each other, could negotiate a better outcome for

farmers collectively; second, the spokesman rationalized prices and quotas across the three farmers in order to equate marginal returns to water quotas. However, the increase in average farmer payoff came at the expense of at least one farmer, suggesting that side payments would be necessary in order for the farmers to be willing to negotiate as one using a spokesman.

Our approach could be used to evaluate the robustness of this finding to the parameters of the environmentalist's utility function and to the weights the spokesman assigns to the utilities of the three farmer groups. To do this, our exogenous factor λ would determine whether farmers participate individually or are represented by a spokesperson. Since Pareto dominance is a necessary condition for a policy to be a solution of the Rausser-Simon model, it is the natural political prediction concept to use in our analysis. In Simon et al. (2007), the spokesman's weights are equal. In practice, it may be more appropriate for the spokesman to assign different weights, depending on, say, relative group size or political power. By including the weights in our parameter vector \mathbf{z} , we could investigate how changing the weights would change the set of side-payments that would be required, in order for for all farmers to be willing to delegate negotiating power to a spokesman.

One would also expect the outcome variables to differ, depending on the parameters determining the relative importance of downstream flows and rural landscapes in the environmentalist's utility function. The more important downstream flows are relative to rural landscapes for the environmentalist, the more dam storage capacity will be built. The resulting increase in the supply of water will be allocated across multiple uses, not just downstream flows. In turn, the allocation will depend in part on how the spokesman weights the utilities of the three farmers.

Both the parameters in the environmentalist's utility function and the weights the spokesman places on the utilities of the three farmers are unlikely to be known to the modeler, so using our approach to evaluate the sensitivity of the results to these three parameters would enable identification of the set of possible outcomes when there is a spokesman. This set could be compared to the set of possible outcomes when the three farmers act independently in the negotiations. Potentially, differences in the set of potential outcomes with and without a spokesman could induce the government to specify which groups are permitted (or required) to negotiate independently. In other words, the

government could evaluate the set of possible outcomes when deciding which stakeholders have a seat at the negotiating table.

5.5. Piave River Basin. Sgobbi and Carraro (2007) also utilize the Rausser and Simon (1999) non-cooperative bargaining model to develop a multi-player model of bargaining over summer and winter water allocations in the Piave River basin when water supplies are uncertain. Players include two regional institutions representing farmers, an upstream province interested in maintaining tourism dependent on lakes and reservoirs, a hydroelectric power company, and downstream municipalities which value river flows. The authors conduct single-variable comparative statics exercises. The two most interesting to consider using our approach are redistributing access between the two players representing farmers' interests, and varying the subjective belief of one player, the hydroelectric power company, regarding summer water resources.

As in the previous example, our approach could be used to evaluate the sensitivity of the set of possible solutions variations in players' access probabilities. In addition, Sgobbi and Carraro (2007) assume that the parameters governing the uncertainty of water supplies are exogenously specified. Our approach could investigate the effect on the set of feasible solutions of varying these parameters.

The authors calibrate their utility functions using input from the stakeholders themselves, which suggests using either a tighter interval for the parameter space than was plausible in our other applications and/or a non-uniform distribution function $f(\mathbf{z})$. At the same time, as the authors acknowledge, the conclusions of the model are still sensitive to these parameters and our approach would assess how confident we can be of the impact of changes in a player's beliefs, given that some parameters are uncertain.

6. CONCLUSION

In this paper, we have presented an approach to predictive political economy in the presence of substantial uncertainty about model parameterization. We then applied it to California's Sacramento-San Joaquin Delta in a detailed case study and discussed several other water and ecosystem management negotiations where it could be used.

Our Delta case study analysis suggests that it is extremely unlikely that several highly publicized policies will emerge as solutions, because they fail to Pareto dominate the default outcome. Because our approach considers a wide range of parameterizations for each stakeholder's utility functions, this conclusion is much stronger than it could be with only a single parameterization, even if we conducted sensitivity analysis on each element of that parameterization in turn. There is an important caveat to this conclusion: because our analysis focused on export policy alone, it does not rule out the possibility that policies we identify as non-viable could be coupled with additional policy changes outside our policy space to craft a compromise policy. Such efforts have been made in the Delta debate. For instance, earlier plans called to couple the construction of canals with substantial expenditures on habitat that are not considered in our analysis. Moreover, like Madani and Lund (2011), we abstract away from the possibility of monetary transfers between stakeholder groups.

In the case study we also find that there *are* policies that Pareto dominate the default. Strikingly, all such policies involve dual conveyance. Except under perfect trust, none of these policies meet our strong RPV standard. However, there is a large set that satisfy our Pareto dominance criterion for more than 40% of the parameterizations we consider, even with substantial mistrust of institutions. Our result is consistent with public opinion polls finding that Californians consider drought and water supplies a critical problem for the state, regardless of whether the respondents support (in a survey by Californians for Water Security Herdt (2014)) or oppose (in a survey by the Natural Resources Defense Council Herdt (2015)) investing in north-south water conveyance. We infer from these polls that while there *is* wide support for investment in some solution to California's water problems, it is unlikely that sufficient public support could be generated for a solution involving a single method for North-South conveyance. This suggests that evaluating the potential performance of a wide variety of dual conveyance options would facilitate future negotiations.

Mistrust of other stakeholders is a feature of many negotiations involving water and ecosystem management. In our case study, the impact of mistrust on the prospects for agreement is highly dependent on the precise parameters of the stakeholders' utility functions. It is thus hard to make robust predictions about the impact that institutional changes might have on the debate. When we looked at the most natural individual parameterization of the model, comparative statics on the impact of mistrust gave us surprising results. We found that there were more potentially viable

policies because a belief that they might be able to use political pressure to increase their export allotments in the future increased the utility of water users more than the risk of higher exports hurt environmentalists. When we looked across specifications, we found a more nuanced story about mistrust. Mistrust increases the probability that there are no win-win policies, a conclusion that corresponds with our intuition. At the same time, mistrust decreases our ability to predict the number of win-win policies or the average level of exports flowing through the tunnel. If being confident in the prediction of policies that PD the default is important, there are therefore two ways to improve predictions: learn more about stakeholders' utility or reduce mistrust.

One particularly useful aspects of our model is that it suggests a shift in perspective. As we note on p. 8, our paper complements other approaches in the sense that while most analyses ask a positive question, we ask a negative one. Specifically, each of the game theoretic papers we discuss in this papers specifies one or more solution concepts, and then asks "what outcome(s) satisfies our solution concept(s)?" On the other hand, one of the questions that we ask in our paper is: "what outcomes *cannot* be a solution to the problem, for a wide set of game-theoretic solution concepts, and for (almost) any possible parameterization of the problem?" Our approach thus can be viewed as a practical first step for negotiating an agreement. We find intuitive support for this in approaches to negotiations discussed in the popular business literature. Finding common ground, or "getting to yes," is sometimes presented as the basis for successful negotiations (Fisher et al.; 2011). Our approach aids in eliminating options which are highly unlikely to be acceptable to all stakeholders.

REFERENCES

- Babonneau, F., Vial, J. and Apparigliato, R. (2010). Robust Optimization for Environmental and Energy Planning, in J. Filar and A. Haurie (eds), *Uncertainty and Environmental Decision Making*, chapter 3, pp. 79–126.
- Bacher, D. (2009). California's Ghastly Peripheral Canal: Panama Canal North?, *Counterpunch*.
- Bay Delta Conservation Plan (2007). *The Bay Delta Conservation Plan: Points of Agreement for Continuing into the Planning Process*, Vol. 2009.
- Bay Delta Conservation Plan (2009). *An Overview of the Draft Conservation Strategy for the Bay Delta Conservation Plan*, Vol. 2009.
- Ben-Haim, Y. (2006). *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edn, Academic Press, London.
- Blue Ribbon Task Force (2007). *Delta Vision: Our Vision for the California Delta*.
- Blue Ribbon Task Force (2008). *Our Vision for the California Delta*, 2 edn.

- Cooley, H., Christian-Smith, J. and Gleick, P. H. (2008). *More with Less: Agricultural Water Conservation and Efficiency in California, A Special Focus on the Delta*, Pacific Institute, Oakland, CA.
- Delta Vision Committee (2008). Delta Vision Committee Implementation Report, *Technical report*.
- Dessai, S., Hulme, M., Lempert, R. and Pielke Jr, R. (2009). Climate Prediction: a Limit to Adaptation, in W. Adger, I. Lorenzoni, K. O'Brien and W. Adger (eds), *Living with Climate Change: Are There Limits to Adaptation?*, Cambridge UP, pp. 49–57.
- Fisher, R., Ury, W. L. and Patton, B. (2011). *Getting to yes: Negotiating agreement without giving in*, Penguin.
- Fullerton, D. (2009). CALFED: Tinkering at the Edges, *Environmental Science and Policy* **12**(6): 733–736.
- Hanak, E., Phillips, C., Lund, J., Durand, J., Mount, J. and Moyle, P. (2013). Scientist and Stakeholder Views on the Delta Ecosystem, *Public Policy Institute of California* .
- Hanemann, M. and Dyckman, C. (2009). The San Francisco Bay-Delta: A Failure of Decision-making Capacity, *Environmental Science and Policy* **12**(6): 710–725.
- Hansen, L. and Sargent, T. (2001). Robust Control and Model Uncertainty, *The American Economic Review* **91**(2): 60–66.
- Hendricks, K., Weiss, A. and Wilson, C. (1988). The war of attrition in continuous time with complete information, *International Economic Review* pp. 663–680.
- Herdt, T. (2014). Project Aimed at Bolstering Water Imports to Southern California.
URL: <http://www.sfexaminer.com/poll-finds-few-in-favor-of-delta-tunnel-project-aimed-at-bolstering-water-imports-to-southern-california/>
- Herdt, T. (2015). Private Poll Shows Strong Statewide Support for Delta Tunnels, 95 Percent Accurate (Online Blogs).
URL: <http://95percent.blogs.vcstar.com/2015/10/02/private-poll-shows-strong-statewide-support-for-delta-tunnels/>
- Holley, C. (2015). Crafting Collaborative Governance: Water Resources, California's Delta Plan, and Audited Self-Management in New Zealand, *Environmental Law Reporter* **45**(4).
- Houba, H., Pham Do, K. H. and Zhu, X. (2013). Saving a river: a joint management approach to the mekong river basin, *Environment and Development Economics* **18**(01): 93–109.
- Kahil, M. T., Dinar, A. and Albiac, J. (2016). Cooperative water management and ecosystem protection under scarcity and drought in arid and semiarid regions, *Water Resources and Economics* **13**: 60–74.
- Kallis, G., Kiparsky, M. and Norgaard, R. (2009). Collaborative Governance and Adaptive Management: Lessons from California's CALFED Water Program, *Environmental Science and Policy* **12**(6): 631–643.
- Lempert, R. (2002). A New Decision Sciences for Complex Systems, *Proceedings of the National Academy of Sciences of the United States of America* **99**, **Suppl 3**: 7309–7313.
- Lempert, R., Groves, D., Popper, S. and Bankes, S. (2006). A General, Analytic Method for Generating Robust Strategies and Narrative Scenarios, *Management Science* **52**(4): 514.
- Lempert, R., Nakicenovic, N., Sarewitz, D. and Schlesinger, M. (2004). Characterizing Climate-Change Uncertainties for Decision-Makers. An Editorial Essay, *Climatic Change* **65**(1): 1–9.
- Lund, J., Hanak, E., Fleenor, W., Bennett, W. and Howitt, R. (2010). *Comparing futures for the Sacramento-San Joaquin delta*, Vol. 3, Univ of California Press.
- Lund, J., Hanak, E., Fleenor, W., Bennett, W., Howitt, R., Mount, J. and Moyle, P. (2008). *Comparing Futures for the Sacramento-San Joaquin Delta*, Public Policy Institute of California, San Francisco.
- Lund, J. R., Hanak, E., Fleenor, W., Howitt, R., Mount, J. and Moyle, P. (2007). *Envisioning Futures for the Sacramento-San Joaquin Delta*, Public Policy Institute of CA.

- Madani, K. and Lund, J. (2011). A Monte-Carlo Game Theoretic Approach for Multi-criteria Decision Making under Uncertainty, *Advances in Water Resources* **34**: 607–618.
- Madani, K. and Lund, J. (2012). California's Sacramento-San Joaquin Delta Conflict: from Cooperation to Chicken, *Journal of Water Resources Planning and Management* **138**(2): 90–99.
- Madani, K., Shalikarian, L. and Naeeni, S. (2011). Resolving hydro-environmental conflicts under uncertainty using fallback bargaining procedure, *Proceeding of the 2011 International Conference on Environment Science and Engineering (ICESE 2011)*, pp. 192–196.
- Marchiori, C., Sayre, S. S. and Simon, L. K. (2012). Bargaining and devolution in the upper guadiana basin, *Environmental and resource economics* **51**(3): 453–470.
- Michaels, J. (2012). Benefit - Cost Analysis of Delta Water Conveyance Tunnels, *Technical report*, Eberhardt School of Business, Business Forecasting Center at University of the Pacific.
- Mokhtari, S., Madani, K. and Chang, N.-B. (2012). Multi-criteria decision making under uncertainty: application to the californias sacramento-san joaquin delta problem,) *2012 World Environmental and Water Resources Congress*, pp. 2339–2348.
- Murphy, J., Sexton, D., Barnett, D., Jones, G., Webb, M., Collins, M. and Stainforth, D. (2004). Quantification of Modelling Uncertainties in a Large Ensemble of Climate Change Simulations, *Nature* **430**(7001): 768–772.
- Pham Do, K. H. and Dinar, A. (2014). The role of issue linkage in managing noncooperating basins: The case of the mekong, *Natural Resource Modeling* **27**(4): 492–518.
- Piani, C., Frame, D., Stainforth, D. and Allen, M. (2005). Constraints on Climate Change From a Multi-thousand Member Ensemble of Simulations, *Geophysical Resource Letters* **32**: 23.
- Rasmusen, E. (2007). *Games and Information: An Introduction to Game Theory*, Wiley-Blackwell.
- Rausser, G. C. and Simon, L. K. (1999). A non-cooperative model of collective decision making: A multilateral bargaining approach, *Technical report*, University of California at Berkeley.
- Roseta-Palma, C. and Xepapadeas, A. (2004). Robust Control in Water Management, *The Journal of Risk and Uncertainty* **29**(1): 21–34.
- Rougier, J. (2007). Probabilistic Inference for Future Climate Using an Ensemble of Climate Model Evaluations, *Climatic Change* **81**(3): 247–264.
- Sgobbi, A. and Carraro, C. (2007). A stochastic multiple players multi-issues bargaining model for the piave river basin.
- Shalikaran, L., Madani, K. and Naeeni, S. (2011). Finding the socially optimal solution for california's sacramento-san joaquin delta problem, *World Environmental and Water Resources Congress 2011: Bearing Knowledge for Sustainability*, ASCE, pp. 3190–3197.
- Shaw, W. D. and Woodward, R. T. (2008). Why Environmental and Resource Economists should care about Non-expected Utility Models, *Resource and Energy Economics* **30**: 66–89.
- Simon, L., Goodhue, R., Rausser, G., Thoyer, S., Morardet, S. and Rio, P. (2007). Structure and Power in Multilateral Negotiations: An Application to French Water Policy, *Giannini Foundation Research Monograph* (47).
- Stainforth, D., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D., Kettleborough, J., Knight, S., Martin, A., Murphy, J. et al. (2005). Uncertainty in Predictions of the Climate Response to Rising Levels of Greenhouse Gases, *Nature* **433**(7024): 403–406.
- Stranlund, J. and Ben-Haim, Y. (2008). Price-based vs. Quantity-based Environmental Regulation under Knightian Uncertainty: An Info-gap Robust Satisficing Perspective, *Journal of Environmental Management* **87**(3): 443–449.
- United States District Court (2007). OCAP Final Interim Relief Court Order.
- Vardas, G. and Xepapadeas, A. (2009). Model Uncertainty, Ambiguity and the Precautionary Principle: Implications for Biodiversity Management, *Environmental and Resource Economics* **45**(3): 379–404.

Walters, D. (2010). Water Bond Violated Unwritten Capitol Rule, Sacramento Bee.

URL: <http://www.sacbee.com/2010/08/11/v-print/2950514/dan-walters-water-bond-violated.html>

Woodward, R. and Shaw, W. (2008). Allocating Resources in an Uncertain World: Water Management and Endangered Species, *American Journal of Agricultural Economics* **90**(3): 593–605.

Woodward, R. T. and Tomberlin, D. (2014). Practical Precautionary Resource Management Using Robust Optimization, *Environmental Management* **54**(4): 828–839.

Technical Appendix

Modeling Negotiations over Water and Ecosystem Management: Uncertainty and Political Viability

This technical appendix provides additional detail on both our general framework and the specifics of the Delta application. We begin by presenting the general model using set theory and then present the specifics of the Delta application.

A.1. PROBABILISTIC POLITICAL VIABILITY

The approach begins with a predictive political economic model that has four basic components: a set of policy options; a set of stakeholders or participants in the process; a political prediction mapping from policy options via outcomes to stakeholder expected utilities; and a prediction concept. The prediction concept selects policy options that meet a certain political viability criterion, based on the expected utilities that stakeholders assign to these options.

A.1.1. Exogenous variables. Each exogenous variable in our model is classified either as a parameter, a state-dependent variable or a policy. The term parameter refers to any exogenous variable whose value is known by stakeholders. We denote by \mathbb{Z} the space of all parameter vectors, with generic element \mathbf{z} . Our notation here diverges slightly from the main text. There we separately defined λ and \mathbf{z} although λ meets the formal definition of \mathbf{z} provided here.

The set of possible states of the world is given by $\mathbb{S} \subset \mathbb{R}$, with generic element s . For every model variable classified as state dependent, we specify a probability distribution $f(s; \mathbf{z})$ over the states of the world that represent stakeholders' uncertainty about it.²⁰ The parameters governing these distributions are included in \mathbf{z} . Conventionally, a state of the world refers to a "move by nature" (Rasmusen; 2007, p.54). Here we use the term state of the world very broadly to encompass any component of the model about which stakeholders are uncertain, including ones that are not usually thought of as being determined by nature, such as certain random aspects of the mapping from policies to outcomes and the default outcomes. Stakeholders face unpredictability in the traditional sense, i.e., Knightian risk: they know the probability distributions over which they must take expectations. In reality, however, there is no bright line distinction between Knightian risk and uncertainty. Rather, these concepts should be thought of as extreme points of a conceptual continuum, along which our stakeholders' unknowns are dispersed. $f(s; \mathbf{z})$ over the states of the world that represent stakeholders' uncertainty about it.

Finally, there is a policy space $\mathbb{X} \subset \mathbb{R}^2$, with generic element \mathbf{x} , consisting of a set of possible policy options.

A.1.2. A Political Prediction Mapping. Stakeholders derive expected utilities not from a particular policy per se, but from the range of possible outcomes that might be induced if this policy were implemented. We thus define a mapping from $\mathbb{Z} \times \mathbb{S} \times \mathbb{X}$ to the outcome space $\mathbb{Y} \subset \mathbb{R}^m$.

An element $\mathbf{y} \in \mathbb{Y}$ is called an outcome vector, and the components of \mathbf{y} will be referred to simply as outcomes. The specification of this mapping depends on the parameter vector \mathbf{z} so the outcome

²⁰ Although stakeholders are uncertain about model variables, they know the distributions of these variables. See Vercammen et al. (1990) for a political economy analysis in which stakeholders have incorrect beliefs.

of policy \mathbf{x} conditional on state of the world s and outcome parameter vector \mathbf{z}^y is denoted by $\mathbf{y}(\mathbf{x}, s, \mathbf{z})$.

Each participant in the political process has a utility function defined over outcomes that is also dependent on the parameter vector. The vector $u(\mathbf{y}(\mathbf{x}, s, \mathbf{z}), \mathbf{z})$ enumerates the utilities of all stakeholders resulting from policy in state s given \mathbf{z} . Stakeholders maximize expected utility, taking expectations over possible states of the world. The vector of stakeholders' expected utilities is

$$\mathbf{E}u(\mathbf{x}, \mathbf{z}) = \int \mathbf{u}(\mathbf{y}(\mathbf{x}, s, \mathbf{z}), \mathbf{z}) f(s, \mathbf{z}) ds. \quad (\text{A.1})$$

A predictive political economic model is represented by a political prediction mapping $\mathbf{W} : \mathbb{Z} \rightarrow \mathbb{X}$. Given a parameterization $\mathbf{z} \in \mathbb{Z}$ of the model, $\mathbf{W}(\mathbf{z})$ is the model's prediction of which element (or elements) from \mathbb{X} are politically viable, in a sense to be described below.

The specification of $\mathbf{W}(\cdot)$ is specific to the application being considered. In many cases, including the applications discussed in the paper, it will identify the set of policies meeting some threshold utility level for the individual stakeholders. As an example, we formalize the definition of the Pareto dominance political prediction criterion.

Given a “default outcome” that will be implemented if the participants in the political process cannot negotiate an agreement, $\mathbf{W}(\mathbf{z})$ is the set of alternatives that Pareto dominate this outcome when the model is parameterized by \mathbf{z} . Let $\mathbf{y}^d(s, \mathbf{z})$ denote the default outcome. The dependence of \mathbf{y}^d on s reflects the possibility that stakeholders may be uncertain about what will happen in the absence of an agreement. Moreover, \mathbf{z} will in general influence default outcomes and utilities in addition to negotiated outcomes. The vector of expected default utilities is: $\mathbf{E}u^d(\mathbf{z}) = \int \mathbf{u}(\mathbf{y}^d(s, \mathbf{z}), \mathbf{z}) f(s, \mathbf{z}) ds$. Note that by definition this vector is independent of every non-default policy \mathbf{x} in \mathbb{X} . The Pareto dominance political prediction mapping is specified as: $\mathbf{W}(\mathbf{z}) = \{\mathbf{x} \in \mathbb{X} : Eu_i(\mathbf{x}, \mathbf{z}) \geq Eu_i^d(\mathbf{z}) \text{ for all } i\}$.

A.1.3. Probabilistic Political Viability. While stakeholders know the value of \mathbf{z} , the modeler does not. To incorporate this lack of knowledge into our approach, we model the components of \mathbf{z} as stochastic; we define a random vector $\tilde{\mathbf{z}} \in \mathbb{Z}$ with density function $f(\tilde{\mathbf{z}})$, representing epistemic uncertainty about the true value of \mathbf{z} . Given this definition of parameters, constants, which are known to the modeler as well as the stakeholders, are included as elements of \mathbf{z} with a degenerate distribution.

We first study the role of modeling uncertainty through the use of a probabilistic policy viability function $V : \mathbb{X} \rightarrow [0, 1]$, where $V(\mathbf{x})$ is the probability computed over possible realizations of modeling uncertainty that policy \mathbf{x} satisfies our viability criterion, i.e., Pareto dominates the default. Formally, $V(\mathbf{x}) = Pr_{\tilde{\mathbf{z}}}(\mathbf{x} \in \mathbf{W}(\tilde{\mathbf{z}}))$. This function is a measure of our model's assessment of the political viability of the policy vector \mathbf{x} . It has no closed form solution, but can be evaluated numerically for any particular policy.

To facilitate interpretation of our results, we partition the policy space into “more likely” and “less likely” regions to summarize the information provided by our viability function. Formally, for some K , we specify a K -vector ρ of probability thresholds, where $0 = \rho_1 < \rho_k < \rho_K < 1$, and, for each k , define a “more likely” region $\mathbb{C}_k^+ = \{\mathbf{x} \in \mathbb{X} : V(\mathbf{x}) > \rho_k\}$ and a “less likely” region $\mathbb{C}_k^- = \{\mathbf{x} \in \mathbb{X} : V(\mathbf{x}) \leq \rho_k\}$. \mathbb{C}_k^+ and \mathbb{C}_k^- are, respectively, the upper- and lower-contour sets of V corresponding to ρ_k . \mathbb{C}_k^+ contains all policies that Pareto dominate the default for some fraction

exceeding ρ_k of possible realizations of modeling uncertainty. We will say that a policy in the “highest” upper-contour set \mathbb{C}_k^+ is robustly politically viable; for a policy with this designation, we can have a high degree of confidence that its political viability is not highly sensitive to specific model parameterizations. Conversely, a policy in the “lowest” lower-contour set \mathbb{C}_1^- will be called never politically viable; we can be highly confident that a policy in this category will not survive the political process, regardless of specific model parameterizations.

We also consider how descriptions of the individual viable sets vary across the distribution of \mathbf{z} . Let $M_j(\mathbf{W}(\mathbf{z}))$ be a summary statistic about the set. Such statistics could include the size of the set, its expected value along some dimension of the policy space \mathbb{X} or other measures of its dispersion or shape. The numerical solution procedure described in the previous set allows us to look at the distribution of these summary statistics induced by the distribution of possible parameter vectors. This distribution is given by

$$h_M(M_j, f(\tilde{\mathbf{z}})) = \int_{\mathbb{Z}} M_j(\mathbf{W}(\mathbf{z})) f(\tilde{\mathbf{z}}) d\mathbf{z}.$$

A.1.4. Simulation Approach. Susan: It felt out of order to me so I rewrote it fairly substantially. How does this sound? For complex policy problems it is virtually impossible to express in tractable analytical form the key elements of the predictive political model, in particular, $\mathbf{y}(\cdot)$, $\mathbf{W}(\cdot)$ and $V(\cdot)$. Instead, the model is analyzed using simulation techniques. First, specific functional forms are assigned to $\mathbf{y}(\cdot)$ and $\mathbf{u}(\cdot)$. We then specify the set of possible parameter vectors \mathbb{Z} and the set of possible states of the world \mathbb{S} . We also specify distributions over \mathbb{Z} and \mathbb{S} . The distribution over \mathbb{S} can depend on the specific realization of $\tilde{\mathbf{z}}$ since we may have modeling uncertainty over stakeholder’s beliefs about the distribution of future states.

In our Delta application, we define the parameter space \mathbb{Z} as a hypercube and assume that the elements of the random parameter vector $\tilde{\mathbf{z}}$ are independently and uniformly distributed, but our methodology can incorporate alternate specifications of \mathbb{Z} and $f(\tilde{\mathbf{z}})$. For each realization of $\tilde{\mathbf{z}}$, the distribution over states of the world has density $h(\cdot; \tilde{\mathbf{z}})$. Again, in our application, we assume that $h(\cdot; \mathbf{z})$ is a constant over the relevant support but alternate specifications are possible.

Once the model is specified, we use a nested loop to generate our results. For each $\mathbf{z} \in \mathbb{Z}$ (the “outer loop”), we compute players’ payoffs for each policy in \mathbb{X} and for the default outcome for each realization $s \in \mathbb{S}$ (the “inner loop”). We then take expectations over \mathbb{S} and identify $\mathbf{W}(\mathbf{z})$. This approach provides a comprehensive picture of political viability across the entire spectrum of plausible parameter configurations through the probabilistic viability function $V(\cdot)$, its associated upper and lower contour sets \mathbb{C}_k^+ and \mathbb{C}_k^- , and the distributions of our summary statistics $h_M(M_j, f(\tilde{\mathbf{z}}))$.

A.2. DELTA APPLICATION DETAILS

While the discussion in the main text used the policy \mathbf{x} as an argument for ease of exposition, the outcomes of policy choices depend on actual export regimes, not simply the chosen policy, for reasons detailed later in this appendix. The actual export regime, $\mathbf{g}(\mathbf{x})$, is described in subsection A.2.3 below. Because outcomes depend on \mathbf{g} , we have defined each of the functions below using \mathbf{g} as our argument rather than \mathbf{x} .

A.2.1. Flow Variables. Three flow variables are derived from the policy variables: the amount of water flowing into the Delta (χ_{in}), the amount of water exported through the Delta (χ_{δ}), and the amount of water exported through the twin tunnels (χ_{pc}). To calculate the Delta inflow, we specify the current inflow from the Sacramento River into the Delta using a constant (χ_{in_0}). Because the tunnel intakes would be upstream of the Delta, water that is exported through the tunnels would not flow into the Delta while water that is exported through the Delta would. Thus,

$$\begin{aligned}\chi_{\delta} &= g_{ex} (1 - g_{shr}) \\ \chi_{pc} &= g_{ex} \cdot g_{shr} \\ \chi_{in} &= \chi_{sac} - \chi_{pc}.\end{aligned}$$

A.2.2. Fish Survival Probabilities. For each fish species, we define a random variable s_i^f that takes on the value 1 if the species survives and 0 if it does not. Using information from Lund et al. (2008), we calibrate two fish survival probability functions: one for smelt and one for salmon. We assume these functions have the form:

$$\bar{\pi}_i(\mathbf{g}) = a_i + b_i \chi_{in}^{\alpha_{in}} + c_i \chi_{\delta}^{\alpha_{\delta}} + d_i \chi_{pc}^{\alpha_{pc}}$$

where the exponents α_i are elements of scientific uncertainty and are allowed to vary across terms. The precise values of the parameters a_i , b_i , c_i , and d_i are part of the state of the world; their values in a particular state of the world are generated in a two-step procedure. The four policies considered in Lund et al. (2008) correspond to four different values for the vector \mathbf{g} , expressed together in matrix form as \mathbf{G} . In the first step, for each state of the world we draw a vector $\hat{\mathbf{f}}_i$ of survival probabilities whose individual elements are the survival probabilities for each of these policies. In the second step, we compute the values of a_i , b_i , c_i , and d_i that set $\bar{\pi}_i(\mathbf{G}) = \hat{\mathbf{f}}_i$ for each of the four policies considered by Lund et al. (2008). The probabilities represent stakeholders' a priori beliefs about the distribution of s_i^f .

A.2.3. Actual Export Regime. As noted above, there is a divergence between the declared export policy \mathbf{x} and the actual export regime \mathbf{g} . The function $\mathbf{g}(\mathbf{x}; s, \mathbf{z}_y)$ gives the actual export regime. There are three distinct reasons that the actual regime may vary from the declared policy. We consider each in turn and define intermediate mappings incorporating their impacts.

Mistrust. As described in the main text, λ is the probability that exports are increased to a maximum feasible level (given by the constant χ_{max}). This implies that prior to any disaster or ecosystem related changes in exports, the actual level of exports is given by $(1 - \lambda) x_{ex} + \lambda \chi_{max}$ and the actual share through the tunnels is given by $(1 - \lambda) x_{shr} + \lambda \left(1 - \frac{\chi_{\delta}}{\chi_{max}}\right)$. We denote the resulting vector $\mathbf{g}^M(\mathbf{x}; \lambda)$, with the M indicating an adjustment for mistrust.

Ecosystem Driven Cutbacks. The possibility that managers may alter the export policy after observing a signal about fish survival based on legal requirements, environmental regulations, or other factors creates a divergence between the declared policy and the actual export level. We assume that this decision occurs after any change in exports due to mistrust has occurred. Therefore, there is a function $\mathbf{g}^E(\mathbf{x}; s, \mathbf{z}^y)$ that maps a declared policy vector \mathbf{x} into an export regime after the impacts of both mistrust and ecosystem driven cutbacks are realized.

Decisions about whether to require cutbacks (set $\mathbf{g}^E \neq \mathbf{g}^M$) must occur prior to nature determining whether the fish species survive (i.e. before nature selects a value of s_i^f). Before making a decision,

managers observe a signal that allows them to update their prior belief about the distribution of s_i^f . The realization of this signal ω_i for each species is an element of the state of the world. This variable takes on the values {good, bad} where **good** implies that managers have received positive news about fish populations leading them to believe the probability of survival is now higher than $\bar{\pi}_i$ and **bad** implies that managers have received negative information about fish populations leading them to believe the probability of survival is now lower than $\bar{\pi}_i$. Formally, the managers' updated probability distribution for s_i^f is defined as $\hat{\pi}_i$

$$\hat{\pi}_i(\mathbf{g}^E, \omega_i) = \begin{cases} \eta_i + (1 - \eta_i) \bar{\pi}_i(\mathbf{g}^E) & \text{if } \omega_i = \text{good} \\ (1 - \eta_i) \bar{\pi}_i(\mathbf{g}^E) & \text{if } \omega_i = \text{bad} \end{cases} \quad (\text{A.2})$$

where $\eta_i \in [0, 1]$ is a measure of the information content of the signal. Note that if $\eta_i = 1$, the signal perfectly predicts survival and if $\eta_i = 0$, the signal provides no added information relative to the prior. For simplicity, the distribution of ω_i conditional on declared policy \mathbf{x} is

$$\omega_i(\mathbf{g}^M) = \begin{cases} \text{good} & \text{with probability } \bar{\pi}_i(\mathbf{g}^M) \\ \text{bad} & \text{with probability } 1 - \bar{\pi}_i(\mathbf{g}^M). \end{cases} \quad (\text{A.3})$$

The formulation embodied in Equations (A.2) and (A.3) guarantees that if managers decide not change the policy in response to the signal, the initial prior and the expected fish survival after observing the signal remain consistent:

$$\bar{\pi}_i(\mathbf{g}^M) = Pr(\omega_i(\mathbf{g}^M) = \text{good}) \hat{\pi}_i(\mathbf{g}^M, \text{good}) + Pr(\omega_i(\mathbf{g}^M) = \text{bad}) \hat{\pi}_i(\mathbf{g}^M, \text{bad}).$$

Note that the signals are a function of \mathbf{g}^M while the ultimate survival probabilities are a function of \mathbf{g}^E . Because there is (and has been in the past) considerable uncertainty about whether managers will in fact reduce exports after observing a bad signal, we introduce another random variable R that takes on the value 1 if the cutbacks occur following the observation of at least one bad signal and 0 otherwise. The distribution of R is given by the variable ν , which gives the probability that $R = 1$. If cutbacks occur, managers reduce exports of all types by a constant proportion μ . Therefore,

$$\mathbf{g}^E(\mathbf{x}, \omega, R) = \begin{cases} \mu \mathbf{g}^M(\mathbf{x}; \lambda) & \text{if } R = 1 \text{ and } \omega_i(\mathbf{g}^M) = \text{bad for some } i \\ \mathbf{g}^M(\mathbf{x}; \lambda) & \text{otherwise (i.e. if } R = 0 \text{ or } \omega_i(\mathbf{g}^M) = \text{good for all } i). \end{cases}$$

Post-Collapse Exports. The final source of variation between the declared policy and the actual export regime has a different character. Based on scientific consensus, the key question regarding the probability of major levee collapse is *when* a collapse will occur, not *whether* one will occur. The random variable τ is the year in which a collapse occurs. We adopt the Lund et al. (2008) hypothesis that there is a constant annual probability of major levee collapse (p_{annFail}). This probability is calculated from a cumulative probability of failure over Yrs years given by p_{fail} according to the formula

$$p_{\text{annFail}} = 1 - (1 - p_{\text{fail}})^{\frac{1}{Yrs}}.$$

Stakeholders receive a stream of annual pre-collapse utilities and a stream of annual post-collapse utilities discounted to the present using the interest rate r .

We assume that all exports will be shifted to the twin tunnels following a major levee collapse, if they exist. If they have not been built, the state can either build one after the disaster or repair the Delta levees and continue pumping through the Delta exclusively.²¹ The decision of whether to

²¹ A third option would be to cease exports in the wake of a major levee collapse. According to Lund et al. (2008) such an outcome is unlikely because it would likely cost the state more than constructing an isolated conveyance structure would. We have no information that contradicts this statement.

build tunnels or repair as an element of the future state of the world is represented by the random variable ξ whose distribution is governed by p_ξ , the probability that the tunnels are built following a major levee collapse.

This structure implies that the export regime after a disaster is

$$\mathbf{g}^{\text{ad}}(\mathbf{x}; s, \mathbf{z}^y) = \begin{cases} \begin{pmatrix} g_{\text{ex}}^E \\ 1 \end{pmatrix} & \text{if } x_{\text{shr}} > 0 \text{ or } \xi = 1 \\ \mathbf{g}^E & \text{if } x_{\text{shr}} = 0 \text{ and } \xi = 0. \end{cases}$$

The export regime before a disaster is simply

$$\mathbf{g}^{\text{bd}}(\mathbf{x}; s, \mathbf{z}^y) = \mathbf{g}^E(\mathbf{x}; s, \mathbf{z}^y).$$

A.2.4. Agricultural Employment. The level of agricultural employment in the San Joaquin Valley depends on total water exports:

$$y_{\text{employ}}(\mathbf{g}) = \varepsilon_0 + \varepsilon_1 \frac{g_{\text{ex}}}{\chi_0}$$

where ε_0 is the number of agricultural jobs in the San Joaquin Valley with no exports and ε_1 is the increase in the number of jobs with pre-2007 export levels.

A.2.5. Costs.

Reduced water exports. Lund et al. (2008) provide estimates of the total costs to the state of reducing exports from pre-2007 levels to three levels: no exports, 25% of pre-2007 levels, and 50% of pre-2007 levels. Their detailed results in Appendix J separate these costs into specific regions which correspond roughly to urban and agricultural water users. We use that information to calibrate two functions of the form:

$$C_k^{\text{rx}}(\mathbf{g}) = C_k^{\text{nx}} e^{\vartheta_k^{\text{rx}} \frac{g_{\text{ex}}}{\chi_0}}$$

where C_k^{nx} is the cost to the water user group of no exports and ϑ_k^{rx} is the calibrated parameter.

Water treatment. Water exported through the Delta must be treated due to its high salinity. The total treatment cost is proportional to the amount of water flowing through the Delta:

$$C^{\text{treat}}(\mathbf{g}) = C_0^{\text{treat}} \frac{\chi_\delta(\mathbf{g})}{\chi_0}$$

where χ_0 is the pre-2007 level of exports (all of which flow through the Delta) and C_0^{treat} is the Lund et al. (2008) treatment cost estimate for that export level.

The water treatment costs are split between agricultural and urban water users in rough proportion to the amount of water they use. The share of total exports used by agricultural users increases with total exports, so the share of treatment costs paid by Ag is

$$\zeta_{\text{Ag}}^{\text{treat}} = \hat{\zeta}_{\text{Ag}}^{\text{treat}} \left(\frac{g_{\text{ex}}}{\chi_0} \right)^{\vartheta_{\text{Ag}}^{\text{treat}}}$$

and the share of treatment costs paid by Ur is

$$\zeta_{\text{Ur}}^{\text{treat}} = (1 - \zeta_{\text{Ag}}^{\text{treat}}).$$

Construction costs. Because the assumed size of the twin tunnels is independent of the planned level of exports (based on engineering considerations), the annualized cost of construction is

$$C^{\text{construct}} = \begin{cases} rC_0^{\text{construct}} & \text{if } \chi_{pc} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

$C_0^{\text{construct}}$ is the Lund et al. (2008) estimate of the total cost of constructing an isolated conveyance structure. These costs are allocated among Ag, Ur, and Tp according to the share vector ζ^{treat} .

Levee maintenance. The cost of maintaining the levees is

$$C^{\text{mntn}} = C_0^{\text{mntn}} \left(\frac{\chi\delta}{\chi_0} \right)^{\rho^{\text{mntn}}}$$

where C_0^{mntn} is the estimate in Lund et al. (2008) of maintenance costs with pre-2007 exports all flowing through the Delta.

Major collapse of the levee system. The cost of a major collapse is

$$C^{\text{collapse}} = \begin{cases} 0 & \text{if } t < \tau \\ rC_0^{\text{collapse}} & \text{if } t \geq \tau \text{ and } g_{\text{shr}} = 0. \end{cases}$$

These collapse costs consider only the costs that could be avoided if exports were instead routed through tunnels. Any costs of collapse that are unaffected by the export regime (e.g. ecosystem impacts or the inundation of Delta islands) are constant throughout the policy space and thus excluded.

Repair following a major collapse. The cost of repairing after a major collapse depends on whether twin tunnels are built or the levees are repaired and through-Delta exports continue. The repair cost is thus

$$C^{\text{repair}}(\mathbf{g}) = \begin{cases} 0 & \text{if } t < \tau \text{ or } g_{\text{shr}} > 0 \\ rC_0^{\text{repair}} & \text{if } t \geq \tau, g_{\text{shr}} = 0, \text{ and } \xi = 0 \\ rC_0^{\text{construct}} & \text{if } t \geq \tau, g_{\text{shr}} = 0, \text{ and } \xi = 1. \end{cases}$$

Stakeholder net benefits. Three stakeholders (Ag, Ur, and Tp) pay all of the costs in the model. Their net financial benefits are

$$B_k = B_k^0 - C_k^{\text{rx}} - \sum_c \zeta_k^c C^c$$

where $c = \{\text{treat, construct, mntn, collapse, repair}\}$.

A.2.6. Stakeholder Utility.

Taxpayers. Taxpayers are risk neutral and have utility of the form

$$u_{\mathbf{u}(\text{Tp})}(\mathbf{g}) = B_{\text{Tp}}(\mathbf{g}).$$

This utility function is not subject to modeling uncertainty apart from any effects on \mathbf{g} .

Urban users. Urban water users are only concerned about the cost of meeting their water supply needs and are risk averse, giving them a utility function of

$$u_{\mathbf{u}(Ur)}(\mathbf{g}) = (B_{Ur}(\mathbf{g}))^{\gamma^{Ur}}$$

where B_{Ur} is the benefit, net of all water supply costs, and the level of risk aversion (γ^{Ur}) is part of modeling uncertainty.

Agricultural users. Agricultural interests value both agricultural employment (y_{employ}) and their benefits net of all the costs they pay (B_{Ag}). They have CES utility given by

$$u_{\mathbf{u}(Ag)}(\mathbf{g}) = [w^{\text{employ}} y_{\text{employ}}(\mathbf{g})^{e_{\text{SubExpAg}}} + (1 - w^{\text{employ}}) B_{Ag}(\mathbf{g})^{e_{\text{SubExpAg}}}]^{\frac{\gamma^{Ag}}{e_{\text{SubExpAg}}}}$$

where $e_{\text{SubExp}} = \frac{e_{\text{Sub}} - 1}{e_{\text{Sub}}}$ governs the elasticity of substitution, w^{employ} is their weighting parameter, and γ^{Ag} is their degree of risk aversion. Each of these parameters is part of modeling uncertainty.

In-Delta interests. In-Delta interests are concerned with maintenance expenditures (C^{mntn}) and Delta water quality, proxied by Delta inflow, χ_{in} . They have a CES utility function. To avoid problems when one of their utility arguments is equal to zero, we add a constant q_{mntn} to the maintenance expenditures before calculating utility. Their utility function is

$$u_{\mathbf{u}(Dt)}(\mathbf{g}) = [w^{\text{mntn}} (C^{\text{mntn}}(\mathbf{g}) + q_{\text{mntn}})^{e_{\text{SubExpDt}}} + (1 - w^{\text{mntn}}) \chi_{in}(\mathbf{g})^{e_{\text{SubExpDt}}}]^{\frac{\gamma^{Dt}}{e_{\text{SubExpDt}}}}$$

$e_{\text{SubExp}} = \frac{e_{\text{Sub}} - 1}{e_{\text{Sub}}}$ governs the elasticity of substitution, w^{mntn} is their weighting parameter, and γ^{Dt} is their degree of risk aversion. Each of these parameters is part of modeling uncertainty, as is the constant q_{mntn} .

Environmentalists. Environmentalists are concerned only with the survival of two fish species: Delta smelt and salmon. Utility is defined in terms of these species' survivals as four state-dependent utility levels. If both species survive, utility is 1 and if neither survives utility is 0. If only one species survive utility takes an intermediate value. Expected utility increases as the probability of survival increases for each species.

$$u_{\mathbf{u}(Ev)}(\mathbf{g}) = \begin{cases} v_{\text{both}}^{Ev} & \text{if } \hat{\pi}_i(\mathbf{g}) = 1 \text{ for } i = \{\text{smelt, salmon}\} \\ v_{\text{salmon}}^{Ev} & \text{if } \hat{\pi}_{\text{salmon}}(\mathbf{g}) = 1 \text{ and } \hat{\pi}_{\text{smelt}}(\mathbf{g}) = 0 \\ v_{\text{smelt}}^{Ev} & \text{if } \hat{\pi}_{\text{smelt}}(\mathbf{g}) = 1 \text{ and } \hat{\pi}_{\text{salmon}}(\mathbf{g}) = 0 \\ v_{\text{none}}^{Ev} & \text{if } \hat{\pi}_i(\mathbf{g}) = 0 \text{ for } i = \{\text{smelt, salmon}\} \end{cases}$$

Discounted utility streams. As discussed in subsection A.3, stakeholders experience one stream of utility prior to a major collapse and a second stream following a major levee collapse. Thus, their utility conditional on the state of the world is

$$= \frac{u_k(\mathbf{y}(\mathbf{x}; s, \mathbf{z}^y); \mathbf{z}^u)}{r} - \frac{u_k(\mathbf{y}^{\text{bd}}(\mathbf{x}; s, \mathbf{z}^y); \mathbf{z}^u) - u_k(\mathbf{y}^{\text{ad}}(\mathbf{x}; s, \mathbf{z}^y); \mathbf{z}^u)}{r} \left(\frac{p_{\text{annFail}}(1+r)}{p_{\text{annFail}} + r} \right).$$

This expression results from integrating over the distribution of time to failure implied by the annual hazard rate p_{annFail} .

Discontinuities. There is a discontinuity in the mapping from policies to expected utilities when $x_{\text{shr}} = 0$ because the costs and benefits of tunnels change discontinuously. We assume that if tunnels exist all exports will be routed through them in the event of catastrophic damage to Delta levees. This reflects an important discontinuity in the real-world political-economic landscape: if an isolated conveyance is built, it will have a very high option value, even if $x_{\text{shr}} \approx 0$. The discontinuity is particularly important for **Tp** and **Ev**. In the absence of an alternative conveyance option, a major levee collapse will lead to one of two outcomes: either the tunnels will be built or extensive levee repairs will be undertaken. Either would occur on an emergency basis, with a compressed schedule and the associated increased costs. The high cost of these emergency response options implies that the maximum x_{ex} (and thus the maximum regular levee maintenance expenditures) that **Tp** will accept falls when $x_{\text{shr}} = 0$. The possibility of rebuilding the Delta levees and continuing through Delta exports also creates a discontinuity for **Ev**: For any given x_{ex} , fish survival probabilities are lowest when the realized share conveyed through the tunnels is zero. In short, even if tunnels would be used only in the event of a disaster, in expectation their existence would contribute significantly to fish survival probabilities. For this reason, the maximum level of x_{ex} that **Ev** will accept falls discontinuously when $x_{\text{shr}} = 0$.

A.2.7. Default. The level of default exports is influenced by modeling uncertainty and the realized future state of the world. The variable x_{ex}^d gives the maximum level of total exports in the default. We introduce a state-contingent random variable ϕ that is distributed uniformly on the interval $(0, 1)$. The maximum level of default exports is given by

$$x_{\text{ex}}^d = \phi \left(\overline{x_{\text{ex}}^d} + v_{dx} \right) + (1 - \phi) \left(\overline{x_{\text{ex}}^d} - v_{dx} \right)$$

where ϕ is state-contingent, $\overline{x_{\text{ex}}^d}$ is a constant and v_{dx} is an element of modeling uncertainty. This formulation implies that x_{ex}^d is distributed uniformly across an interval centered on $\overline{x_{\text{ex}}^d}$ whose width varies with modeling uncertainty.

The actual default export regime is state contingent because non-zero probabilities of additional ecosystem-driven cutbacks and future levee failures exist. The variable R^d indicates whether ecosystem cutbacks occur in the default; the probability that $R^d = 1$ is ν_0 , which is set equal to 1 in our simulations as shown in Table A.3. This implies that default exports after uncertainty about ecosystem cutbacks is resolved are

$$\mathbf{g}^{E^d} \left(\mathbf{x}^d, \omega, R^d \right) = \begin{cases} \mu \mathbf{x}^d & \text{if } R^d = 1 \text{ and } \omega_i \left(\mathbf{x}^d \right) = \text{bad for some } i \\ \mathbf{x}^d & \text{if } R^d = 0 \text{ or } \omega_i \left(\mathbf{x}^d \right) = \text{good for all } i. \end{cases}$$

There are only two differences between the mapping from the actual default export regime to payoff-relevant outcomes and the mapping from the actual export regime induced by an agreement. First, in the default $C^{\text{mntn}} \left(\mathbf{g}^d \right) = 0$ regardless of the level of default exports. Second, if tunnels are constructed following a major levee collapse in the default, they will be more expensive to build due to the need to move on an accelerated schedule in the absence of any framework for doing so (and, perhaps, pre-existing enmity between the stakeholders). This consequence is captured by specifying that the construction cost in the default is $(1 + P_d) C^{\text{construct}} \left(\mathbf{g}^d \right)$.

The existence or extent of mistrust has no impact on default exports because the twin tunnels will not be built in the default, so $x_{\text{shr}}^d = 0$.

A.2.8. Model Coefficients. The model coefficients that are elements of modeling uncertainty are described in Table 2 in the main text. Table A.1 replicates this information, adding an initial

TABLE A.1. Elements of modeling uncertainty

	Variable	Lower Bound	Upper Bound
w^{employ}	Weight on jobs vs money in Ag utility	0.2	0.8
w^{mntn}	Weight on maintenance vs inflows in Dt utility	0.2	0.8
q_{mntn}	Constant in Dt utility	0.05	0.15
e_{SubAg}	Ag elasticity of substitution	0.5	1.5
e_{SubDt}	Dt elasticity of substitution	0.5	1.5
γ_{Ag}	Ag risk aversion coefficient	0.2	1
γ_{Dt}	Dt risk aversion coefficient	0.2	1
γ_{Ur}	Ur risk aversion coefficient	0.2	1
$v_{\text{smelt}}^{\text{Ev}}$	Ev utility if only smelt survive*	0.25	0.75
$v_{\text{salmon}}^{\text{Ev}}$	Ev utility if only salmon survive*	0.25	0.75
v_{dx}	Spread of default export distribution above and below mean (maf)	0	2

* Ev utility is scaled so that 0 represents the utility if neither species survives and 1 represents the utility if both survive.

column identifying the mathematical symbol used for these coefficients in this appendix. These variables are independently and uniformly distributed on their specified intervals.

Table A.2 lists the primitive variables governing uncertainty about the future state of the world. Again, the variables are assumed to be independently and uniformly distributed on the given intervals. Many of these upper and lower bounds are drawn directly from Lund et al. (2008); the source of these is listed as PPIC in Table A.2 .

The functions presented above rely on several variables that are treated as constants. These variables and their sources are listed in Table A.3.

TABLE A.2. State dependent variables and their distributions

Symbol	Variable	Lower Bound	Upper Bound	Source
π_{PC}^{Smelt}	Smelt survival prob with $\mathbf{x} = (6, 1)$.1	.4	PPIC
π_{PC}^{Salmon}	Salmon survival prob with $\mathbf{x} = (6, 1)$.2	.5	PPIC
π_{Dual}^{Smelt}	Smelt survival prob with $\mathbf{x} = (6, 0.5)$.1	.4	PPIC+assumption
π_{Dual}^{Salmon}	Salmon survival prob with $\mathbf{x} = (6, 0.5)$.2	.5	PPIC+assumption
π_{Thru}^{Smelt}	Smelt survival prob with $\mathbf{x} = (6, 0)$.05	.3	PPIC
π_{Thru}^{Salmon}	Salmon survival prob with $\mathbf{x} = (6, 0)$.1	.3	PPIC
π_{No}^{Smelt}	Smelt survival prob with $\mathbf{x} = (0, \cdot)$.3	.6	PPIC
π_{No}^{Salmon}	Salmon survival prob with $\mathbf{x} = (0, \cdot)$.4	.8	PPIC
$\zeta_{Tp}^{collapse}$	Tp share of collapse costs	.2	1	Assumption
ζ_{Tp}^{repair}	Tp share of maintenance costs	.5	1	Assumption/SU
ϑ^{mntn}	Exponent in maintenance cost function	0.5	1.5	Assumption
ϑ^{repair}	Exponent in repair cost function	.1	.2	Assumption
$C_0^{collapse}$	Total cost of collapse (\$ billion)	7.8	15.7	PPIC
C_0^{repair}	Total cost of repairs (\$ billion)	.2	2.5	PPIC
$C_0^{construct}$	Total cost of isolated conveyance construction (\$ billion)	4.75	9.75	PPIC
C_0^{treat}	Annualized treatment cost for $\chi_\delta = 6$ (\$ billion/yr)	.3	1	PPIC
ϑ_{Ag}^{fx}	Exponent in Ag scarcity cost function	-3.62	-1.58	Derived from PPIC
ϑ_{Ur}^{fx}	Exponent in Ur scarcity cost function	-6.35	-1.97	Derived from PPIC
μ	Export reduction share if fish don't recover	.25	.4	PPIC
p_{fail}	Cummulative failure probability over Yrs	.34	.95	PPIC
C_{Ag}^{nx}	Scarcity cost to ag with no exports (\$ billion/yr)	.49	.96	Derived from PPIC
C_{Ur}^{nx}	Scarcity cost to ur with no exports (\$ billion/yr)	1.1	1.54	Derived from PPIC
C_0^{mntn}	Maintenance costs for $\mathbf{x} = (6, 0)$ (\$ billion/yr)	1	2	PPIC
$\vartheta^{Ag^{treat}}$	Exponent in Ag treatment cost share	.3	.4	Derived from PPIC
ϕ	Location of x_{ex}^d in its interval	0	1	Bounds assumed
\hat{p}_d	% Increase in annual failure probability in default	0	.15	Assumption
α_δ	Exponent in the survival function on χ_δ	1.5	4	Assumption
α_{PC}	Exponent in the survival function on χ_{pc}	1.5	4	Assumption
α_{in}	Exponent in the survival function on χ_{in}	1.5	4	Assumption
P_d	Extra % post-disaster construction cost in the default	0.5	1	Assumption

Note: Additional uncertainty about state of the world incorporated in event trees (levee collapse etc.)

TABLE A.3. Constants

	Variable	Value	Source
$\zeta_{Ur}^{\text{collapse}}$	Share of the collapse cost paid by Ur	.2	Assumption
$\zeta_{Ag}^{\text{collapse}}$	Share of the collapse cost paid by Ag	.2	Assumption
$\zeta_{Tp}^{\text{repair}}$	Share of the repair cost paid by Tp	1	Assumption
$\zeta_{Ur}^{\text{construct}}$	Share of the tunnels' construction cost paid by Ur	.45	Assumption
$\zeta_{Ag}^{\text{construct}}$	Share of the construction cost paid by Ag	.45	Assumption
$\zeta_{Tp}^{\text{construct}}$	Share of the construction cost paid by Tp	.1	Assumption
ζ_{Ur}^{mntn}	Share of the costs of maintaining the Delta paid by Ur	.1	Assumption
ζ_{Ag}^{mntn}	Share of the costs of maintaining the Delta paid by Ag	.1	Assumption
B_{Tp}^0	Base benefit level for Tp (\$ billion/yr)	100	2011-12 state budget
B_{Ur}^0	Base benefit level for Ur (\$ billion/yr)	5	Assumption
B_{Ag}^0	Base benefit level for Ag (\$ billion/yr)	2	Assumption
χ_0	Pre-Wanger export level (maf)	6	PPIC
r	Interest rate used in discounting	.05	PPIC
Yrs	Years over which the prob of failure applies	42	PPIC
χ_{in_0}	Current level of inflow into the Delta (maf)	19.3	PPIC
ε_0	Ag jobs w/no exports (millions)	.5	PPIC
ε_1	Increase in Ag jobs w/ baseline exports instead of none (millions)	.1	PPIC
q_{inflow}	Scaling parameter to incorporate Delta inflows in Dt utility	$\frac{1}{15}$	Assumption
γ_{Tp}	Tp risk aversion	1	Assumption
$\hat{\zeta}_{Ag}^{\text{treat}}$	Share of treatment cost paid by Ag in base case	.4	Derived from PPIC
\underline{x}_{ex}	Minimum value of total exports in the policy space (maf)	0	Choice
\overline{x}_{ex}	Maximum value of total exports in the policy space (maf)	7.5	Choice
ν	Probability the ESA is enforced following a bad signal	.75	Assumption
η_{smelt}	Information in smelt signal	.75	Assumption
η_{salmon}	Information in salmon signal	.75	Assumption
$v_{\text{both}}^{\text{Ev}}$	Ev utility if both species survive	1	Immaterial assumption
$v_{\text{none}}^{\text{Ev}}$	Ev utility if neither species survives	0	Immaterial assumption
ν_0	Probability of ESA enforcement in the default	1	Assumption
χ_{max}	Maximum feasible export level (maf)	7.5	Downstream constraints per PPIC
p_{ξ}	Probability of isolated conveyance structure following a disaster	.5	Assumption
$\overline{x_{ex}^d}$	Mean of the x_{ex}^d distribution	4	Assumption