THE FUTURE OF CALIFORNIA'S SACRAMENTO-SAN JOAQUIN DELTA: WATER POLICY ALTERNATIVES AND ROBUST POLITICAL VIABILITY

RACHAEL GOODHUE[†] SUSAN STRATTON SAYRE[††] LEO SIMON[†††]

JULY 29, 2011

ABSTRACT. Researchers applying predictive political economic models to specific real world policy problems typically face at least two difficult challenges. First, they have very limited information on which to assign parameters to the mappings from policies to outcomes to utilities. Second, they have very little basis for deciding which political economic model will adequately represent the political process in question. This paper presents a methodology, which we call Probabilistic Political Viability (PPV), that applies tools of political economy to make probabilistic predictions about the viability of various political alternatives in real world settings. We apply this methodology to the ongoing debate over the future of California's Sacramento-San Joaquin Delta. We construct a computer simulation model that is a highly stylized representation of the main features of this debate and study the properties of this model under a wide range of plausible parameter specifications. We define a set of alternative policy responses which includes several widely debated options and assess the likelihood that any particular option will be politically viable. We base our assessments on the fraction of our universe of parameter values for which each option meets a relatively weak viability criterion. We first establish that certain alternatives would be "robustly politically viable," if all stakeholder groups trusted that these alternatives would be implemented in accordance with negotiated guidelines. We then incorporate institutional mistrust into our model and examine how the political viability of these alternatives changes as institutional mistrust increases. Broadly speaking, mistrust in our model represents the concern of certain stakeholder groups that urban and agricultural users will fully exploit the carrying capacity of a newly constructed conveyance structure, regardless of any negotiated agreement to limit water exports and maintain through-Delta flows. When institutional mistrust is high, none of the policy alternatives we consider are robustly politically viable.

KEYWORDS: Pareto Optimality, Delta, California, Political Economy, Deep uncertainty, Robust decision making, Modeling uncertainty

JEL CLASSIFICATION NUMBERS: P48, Q25, Q34

ACKNOWLEDGEMENTS: 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 and UC Davis Ag-IO workshop for helpful comments. Kathy Edgington and Laurie Warren provided invaluable computer resources support. Goodhue and Simon are members of the Giannini Foundation of Agricultural Economics,

[†]Department of Agricultural and Resource Economics, University of California, Davis, email: goodhue@primal.ucdavis.edu
[††]Department of Economics, Smith College, email: ssayre@smith.edu
[†††]Department of Agricultural and Resource Economics, University of California, Berkeley, email: leosimon@berkeley.edu
.

1. Introduction

The task of applying political economic models to predict political outcomes is challenging for many reasons. First, it is never easy to decide which model best represents the real world political process under consideration. Second, there is rarely any solid basis for assigning values to the key parameters of whichever model is selected. As a result, it is virtually impossible for political economists to provide precise and credible predictions of the outcomes of any but the simplest political processes. To address these challenges, we present in this paper a methodology that applies the tools of political economy to make probabilistic predictions about the viability of various alternatives in real world settings. There is a real need for tools such as ours that can deliver at least modest and qualified predictions about the likelihood that an ongoing political negotiation will be resolved successfully, and if so, what kinds of properties the negotiated agreement will exhibit.

Our methodology is designed to analyze specific, one-time policy negotiations. As resource economists, the problems that interest us typically involve tradeoffs between economic vs environmental objectives, market vs non-market valuations and private vs public goods. In such contexts, the problems associated with constructing a model are particularly challenging. The complexities really matter: it is important to model the interconnected economic, social, and ecosystem impacts of the various policy options under consideration. Our approach must be more fine-grained than models that use econometric techniques to identify broad regularities linking processes to outcomes; the more we abstract from the idiosyncratic details of our problem, the less credible will be our probabilistic predictions. This leads to models that are too complex to be solved analytically. As a result, our methodology is designed to use simulation methods to predict outcomes and conduct comparative statics analysis. Moreover, trying to predict exactly what policy will emerge is too ambitious a goal; instead, we seek to identify policies that meet a coarser "political viability" criterion. Finally, when modeling complex, "one of a kind" policy debates, it is virtually impossible to assemble a database rich enough to use econometric techniques to estimate model parameters. Since we cannot have confidence in any particular parameterization of the model *and* must utilize numerical rather than analytical comparative statics methods, it is critical that we study the properties of the model under the widest possible range of plausible parameter specifications. Having done so, we can then assess the likelihood that any particular policy option will be politically viable. Specifically, we use the fraction of our universe of parameter values for which the option meets our viability criteria as a measure of political viability. In this paper, we will reserve the term *modeling uncertainty* to refer to our lack of information about how best to model, and then parameterize, the political-economic environment that we wish to study.[1]

In this paper, we apply our methodology to the debate over the future of California's Sacramento-San Joaquin Delta. At present, this debate centers around two critical questions: (1) how much water can be exported from the Delta without violating the economic and ecological integrity of this important region and (2) should the state build a "Peripheral Canal" that would deliver water from the Sacramento River directly to diversion pumps in the Southern Delta? In 1982, a canal construction proposal was soundly defeated via ballot initiative, but since 2007, support for building a canal or tunnel to deliver water while bypassing the Delta has grown. The plan still faces an uphill battle; a water bond that included initial funding to pave the way for a canal was withdrawn from the November 2010 ballot due to concerns that inclusion of the canal would lead to the defeat of the entire bond proposal.

Using our probabilistic political viability methodology, we investigate the viability of various possible solutions to the Delta crisis and the impact of institutional mistrust on that viability. There has been extensive analysis of the environmental and economic consequences of various Delta alternatives (Lund *et al.*, 2007, 2008; Cooley *et al.*, 2008) and some rankings of these alternatives based on a variety of financial and non-financial criteria (Lund *et al.*, 2008). Some analyses have examined the history and current situation regarding water policy and the Delta from institutional and game-theoretic perspectives. Hanemann & Dyckman (2009) and Madani & Lund (2011b) conclude that stakeholders are unlikely to agree on an alternative in the absence of credible government intervention or a substantial worsening of the current situation. Madani & Lund (2011a) analyze a simple model with two players (environmentalists and water exporters), two decision criteria (annual cost and likelihood of a viable Delta smelt population), and the four discrete policy options identified by Lund *et al.* (2008). They examine the frequency, ranging over model parameterizations, with which each of these discrete policy options is a solution under four game-theoretic solution concepts. They conclude that the construction of a conveyance facility to convey water exports around the Delta may emerge as an equilibrium. Their analysis strongly suggests, however, that their two parties may be unable to agree upon a solution, in which case the status quo will prevail. In short, existing institutional and game theoretic

[1]The very similar term "model uncertainty" is widely associated with the work of Hansen & Sargent (2001), which builds on work by Gilboa & Schmeidler (1989) and others. Gilboa *et al.* (2008) provides an accessible overview of the mathematical/statistical issues addressed in this literature. This literature is motivated by a problem very similar to the one that we confront: how to deal with situations where probabilities are unknowable. However, they focus on the question of how to optimize an objective function in this context; we eschew optimization altogether, instead attempting to isolate potential solutions that satisfy a weak necessary condition for optimality.

work suggests that there is little likelihood of stakeholder agreement unless water governance institutions are substantially redesigned or the government takes some other decisive action.

In this paper, we synthesize a selection of the existing literature into a formal mathematical model designed to evaluate political support for various alternatives. We first establish that for the range of parameterizations that we consider, certain alternatives would be "robustly politically viable," if all stakeholder groups trusted that these alternatives would be implemented in accordance with negotiated guidelines. We then incorporate institutional mistrust into our model and examine how the political viability of these alternatives changes as institutional mistrust increases. In Section 2 we develop our probabilistic political viability methodology. Section 3 presents the details of our Delta application and constructs a formal model of that political process, which is then embedded within our viability methodology. In Section 4, we present the results of our analysis and discuss their signficance for both the Delta application and our methodology. Section 5 concludes.

2. The Probabilistic Political Viability Methodology

The political economic model we construct in this paper has four basic components: a set of *policy options*; a set of *stakeholders* or participants in the process; a 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.

As discussed in the introduction, researchers applying predictive political economic models to specific real world policy problems typically face at least two difficult challenges. First, they have very limited information on which to assign parameters to the mappings from policies to outcomes to utilities. Second, they have no firm basis for declaring that any given model, which by necessity is a highly stylized representation of reality, will adequately represent the political process applicable to the particular policy problem in question. To address these challenges, we introduce in this section a methodology that we call Probabilistic Political Viability (PPV). In subsection 2.1, we specify the components of a predictive political economic model; in later subsections we explain how we incorporate into our methodology the researcher's lack of information about how to parameterize the model.

2.1. **A Predictive Political Correspondence.** There is a policy space $\mathbb{X} \subset \mathbb{R}^n$, with generic element $\mathbf{x}$, consisting of a set of possible policy options. Players 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. For instance, if the policy were a tax level and the outcome after-tax profits, businesses would derive utility from their profits, not from the tax level itself. The first step in constructing a predictive political economic model is thus to define a mapping from the policy space $\mathbb{X}$ to the outcome space $\mathbb{Y} \subset \mathbb{R}^m$. An element $\mathbf{y} \in \mathbb{Y}$ is called an *outcome vector*, while the components of $\mathbf{y}$ will be referred to simply as *outcomes*.

The specification of the mapping from $\mathbb{X}$ to $\mathbb{Y}$ includes a number of *outcome parameters*, whose values are given by the vector $\mathbf{z}^{\mathbf{y}}$. In general, the outcome resulting from a particular policy choice will also depend on uncertain events. We refer to these events as *states of the world*. The set of possible states of the world is given by $\mathbb{S} \subset \mathbb{R}$, with generic element s. The outcome of policy $\mathbf{x}$, conditional on state of the world s and outcome parameter vector $\mathbf{z}^{\mathbf{y}}$ is denoted by $\mathbf{y}(\mathbf{x}; s, \mathbf{z}^y)$. Each participant in the political process has a utility function defined over outcomes; $\mathbf{z}^u$ is a vector of *utility parameters* specifying stakeholders' preferences. The vector $\mathbf{u}(\mathbf{y}(\mathbf{x}; s, \mathbf{z}^y); \mathbf{z}^u)$ enumerates the utilities of all stakeholders resulting from policy $\mathbf{x}$, in state s, given parameter vectors $\mathbf{z}^y$ and $\mathbf{z}^u$. Stakeholders are assumed to be expected utility maximizers, taking expectations over possible states of the world. The distribution over states of the world is parameterized by a vector $\mathbf{z}^s$ of *distribution parameters*. We combine our three parameter vectors, plus a fourth, $\mathbf{z}^d$, defined below, into a composite vector $\mathbf{z} = (\mathbf{z}^y, \mathbf{z}^u, \mathbf{z}^s, \mathbf{z}^d)$. The parameter vector $\mathbf{z}$ is known by all stakeholders; any stakeholder uncertainty about the mapping from policies to outcomes is incorporated in the description of the future state of the world s. The vector of stakeholders' expected utilities is given by

$$\mathbf{Eu}(\mathbf{x}; \mathbf{z}) = \int \mathbf{u}(\mathbf{y}(\mathbf{x}; s, \mathbf{z}^y); \mathbf{z}^u) f(s; \mathbf{z}^s) \, ds$$

where $f(s; \mathbf{z}^s)$ is a probability distribution over states of the world, parameterized by $\mathbf{z}^s$.

The space of all possible model parameter vectors is denoted by $\mathbb{Z}$. A predictive political economic model is represented by a mapping $\mathbf{W} : \mathbb{Z} \to \mathbb{X}$, i.e., from the parameter space into the policy space. We call such a mapping a *political prediction mapping*. 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 typical approach to political economic modeling is to isolate an alternative or set of alternatives that *solves* the model using the specified solution concept. As noted above, however, a starting point for this paper is the infeasability of isolating a single model that best represents a given complex real-world political process. Therefore, our political prediction correspondence maps not to the outcome identified by applying any one particular solution concept, but rather to a set of policies that satisfy a relatively weak condition for political viability, specifically Pareto dominance.[2] To define this correspondence, we first identify a "default outcome" to our problem that will be implemented if the participants in the political process cannot negotiate an agreement. We then define $\mathbf{W}(\mathbf{z})$ to be the set of alternatives that Pareto dominate this outcome when the model is parameterized by $\mathbf{z}$. Pareto dominance is a necessary condition for a large class of political economic solution concepts; in any model requiring consensus among some set of players, a necessary condition for a policy to be a solution is that it yields each player in that set an expected utility as least as high as the outcome that would occur without agreement.

The default outcome in our model is denoted by $\mathbf{y}^d\left(s;\mathbf{z}^d\right)$, where $\mathbf{z}^d$ is a vector of parameters that relate to this 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. The vector of expected default utilities is thus:

$$\mathbf{Eu}^d(\mathbf{z}) = \int \mathbf{u}\left(\mathbf{y}^d\left(s;\mathbf{z}^d\right);\mathbf{z}^u\right) f\left(s;\mathbf{z}^s\right) ds.$$

Note that by definition, this expected utility is independent of every non-default policy $\mathbf{x}$ in $\mathbb{X}$. Once default utilities have been defined, the Pareto dominance political prediction mapping is specified as:

$$\mathbf{W}(\mathbf{z}) = \left\{\mathbf{x} \in \mathbb{X} : Eu_i(\mathbf{x};\mathbf{z}) \geq Eu_i^d(\mathbf{z}) \text{ for all } i\right\}. \tag{2.1}$$

2.2. **Probabilistic Political Viability.** If the parameter vector $\mathbf{z}$ were known by the researcher with certainty, the previous section would fully describe the political prediction of the model. As we have noted above, however, the empirical information available to researchers in the field of predictive political economy is typically very imprecise. Consequently, the modeler can have very little confidence in the appropriateness of any particular assignment of values to the parameter vector $\mathbf{z}$. To incorporate this lack of knowledge into our methodology, we model the components of $\mathbf{z}$ as stochastic; we define a random vector $\tilde{\mathbf{z}} \in \mathbb{Z}$ with density function $h(\tilde{\mathbf{z}})$ that represents the modeler's epistemic uncertainty about the true value of $\mathbf{z}$.[3] To reiterate, the

[2]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, our notion of political viability 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. 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 we shall demonstrate in section 3 below, it can be especially instructive to learn that certain highly publicized possibilities fail to satisfy even this relatively modest selection criterion.

[3]This approach admits the possibility that the researcher knows some components of the parameter vector with certainty. Specifically, suppose that $\mathbf{z}$ is an N-component vector, but for $i < n < N$ the value of component i is known to be $\bar{z}_i$. In this case we would specify the support of $\tilde{\mathbf{z}}$ to be $\mathbb{Z} = \prod_{i=1}^{n}\{\bar{z}_i\} \times \prod_{i=1+1}^{N} Z_i$.

vector $\tilde{\mathbf{z}}$ is random only from the perspective of the researcher; the stakeholders in our model are assumed to know the realization of $\tilde{\mathbf{z}}$. That is, each realization $\mathbf{z}$ of $\tilde{\mathbf{z}}$ corresponds to a specific parameterization of the model, in which stakeholders have uncertainty only about the realized state of the world. For this parameterization, the prediction of the model is given by $\mathbf{W}(\mathbf{z})$. In short, stakeholders are modeled as participating in a standard multi-player decision problem with exogenous uncertainty generated by a known process; the modeler, however, does not know which problem it is that stakeholders confront. One benefit of this approach is that it allows us to study the sensitivity of our political prediction mapping, $\mathbf{W}(\cdot)$, to the particular parameterization of the problem. In the remainder of this paper, we reserve the term *modeling uncertainty* to refer to the researcher's lack of knowledge about the precise value of $\mathbf{z}$.[4]

To study the role of modeling uncertainty, we define a *probabilistic political viability function*, which is a mapping $V : \mathbb{X} \to [0,1]$ where

$$V(\mathbf{x}) = Pr_{\mathbb{Z}}(\mathbf{x} \in \mathbf{W}(\mathbf{z})) \tag{2.2}$$

is the probability, computed over possible realizations of modeling uncertainty that a policy satisfies our viability criterion. The viability of a policy is defined to be the probability that it Pareto dominates the default.

To summarize the information provided by our viability function, we partition the policy space into "more likely" and "less likely" regions. Specifically, for some K, we specify a K-vector ρ of probability threshholds, 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 $\mathbb{C}_k^- = \{\mathbf{x} \in \mathbb{X} : V(\mathbf{x}) \leq \rho_k\}$. Formally, $\mathbb{C}_k^+$ and $\mathbb{C}_k^-$ are, respectively, the upper- and lower-contour sets of V corresponding to ρ_k. Under Pareto dominance, $\mathbb{C}_k^+$ 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 modeling choices. Conversely, a policy in the "lowest" lower-contour set $\mathbb{C}_1^-$ will be called *never politically viable;* for a policy in this category, we can be highly confident that it will not survive the political process, regardless of specific modeling choices.

Our approach is closely related to the "robust decisionmaking" approach developed in order to evaluate problems characterized by "deep uncertainty." Deep uncertainty refers to situations where the researcher

[4] The role that modeling uncertainty plays in our paper is different in an important way from the role that model uncertainty plays in Hansen & Sargent (2001) and the literature that this paper spawned.See footnote ???

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)[5] 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 statistics to report, for example, the mean value of a specific outcome variable over the entire set of simulations, as one would in a Monte Carlo setting, the results are interpreted as representing modeling uncertainty. If a potential solution (in our case a policy) 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.[6] Our probabilistic political viability approach follows the same logic. In our political economic context, just as in a decision-theoretic context, the value of a single optimal solution based on a single model specification becomes less useful the more sensitive the model outcome becomes to uncertainty regarding the model specification (Lempert *et al.* (2006)). In complicated problems, 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 in both approaches. The first role is the conventional one of representing the likelihood of realizations of states of the world, or known uncertainty. The second role is 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).

Our approach is related as well to both the “robust control” and the “info-gap” literatures, although less closely. Robust control is a means of modeling ambiguity-averse preferences (Hansen & Sargent, 2001). Due to the limits of knowledge regarding the factors driving species survivial, among other considerations, robust

[5]Deep uncertainty is closely related to the distinction introduced to economists by Frank Knight. As Gilboa *et al.* (2008) note “Since the early days of probability theory, there has been a distinction between probabilities that are given, as in a game of chance, and probabilities that are not given, but reflect a subjective degree of belief... Knight (1921) is typically credited with the distinction between situations of "risk" and of "uncertainty". The former designates situations in which probabilities are known, or knowable in the sense that they can be estimated from past data, calculated using the laws of probability, etc. By contrast, "uncertainty" refers to situations in which probabilities are neither known, nor can they be deduced, calculated, or estimated in an objective way (p. 173).”

[6]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)).

control is a natural choice for modeling many natural resource problems, including extractive fisheries and water allocation (Shaw & Woodward, 2008). Info-gap theory is designed to identify policies that decision-makers can be confident will meet an acceptability criterion (Ben-Haim, 2006). In both literatures, the goal is to single out a policy that meets an optimality criterion designed to address the well-known problems associated with decisionmaking under Knightian uncertainty (See Ellsberg, 1961; Gilboa & Schmeidler, 1989). Since we do not seek to identify a unique optimal solution, difficulties associated with ambiguity aversion do not arise in our context.

2.3. **Modeler and/or stakeholder uncertainty.** The model presented above makes a sharp distinction between parameter vectors and states of the world. We reserve the term "parameter" for variables whose values are known by stakeholders. The symbol $\mathbf{z}$ denotes the vector of all such variables. On the other hand, stakeholders are assumed to be uncertain about the realized value of s, the state of the world. Conventionally, a "state of the world" refers to a "move by nature" (Rasmusen, 2007, p.54). In our context, a classic example of such a move would be an event such as an earthquake.[7] In this paper, we use the term "state of the world" very broadly, to encompass any component of our model about which stakeholders are uncertain, including components that are not usually thought of as being determined by nature. These components include, in particular, certain random aspects of the mapping from policies to outcomes, and of the default outcome. Thus in our framework, uncertainty over states of the world includes stakeholder uncertainty about certain model coefficients.[8] For every model variable that we classify as state dependent, we must specify a distribution over the states of the world that represent stakeholders' uncertainty about it. The parameters of these distributions will be components of the vector $\mathbf{z}^s$ of distribution parameters. In many instances, we have no stronger grounds for confidence in any one parameterization of stakeholder uncertainty than we has confidence in any one parameterization of the variables which from the stakeholders' perspective are deterministic. In such cases, it is appropriate to model the relevant distribution parameters as stochastic from the perspective of the modeler, i.e., subject to modeling uncertainty.

For each exogenous variable in their model, researchers applying our methodology must decide whether to classify that variable as a constant (known both to stakeholders and the modeler), as a component of modeling uncertainty (known to stakeholders but unknown to the modeler), or as state-dependent (unknown

[7]From the perspective of a physicist, the forces determining whether or not an earthquake occurs may be considered deterministic. From the perspective of a stakeholder, however, the occurrence of an earthquake is unquestionably a move by nature.

[8]We model stakeholders in our framework as facing 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.

to both stakeholders and the modeler). The choice will depend on the characteristics of the particular application. In many instances, it would radically distort the problem being studied to assume that some variable were known by stakeholders, when in reality this knowledge could not possibly be available to them. For example, in many problems we study, the default outcome, implemented if a political solution cannot be reached, will be determined by a litigation process whose end result cannot possibly be predicted with certainty. More generally, however, the choice could be quite a subtle one and will depend on the specific question the researcher is interested in asking.

2.4. **Our Simulation Approach.** As discussed in the introduction, our methodology is designed to analyze one-of-a-kind, complex real-world policy problems. In these settings, it will be virtually impossible to express in tractable analytical form the key elements of our framework, in particular, $\mathbf{y}(\cdot)$, $\mathbf{W}(\cdot)$ and $V(\cdot)$. Accordingly, we implement our methodology using simulation methods: we assign specific functional forms to $\mathbf{y}(\cdot)$ and $\mathbf{u}(\cdot)$ and specify distributions over the sets of states of the world, $\mathbb{S}$, and realizations of modeling uncertainty $\mathbb{Z}$. We define the parameter space $\mathbb{Z}$ to be a hypercube. Lacking any basis on which to rank the relative likelihoods of alternative parameterizations, we invoke the principle of insufficient reason (Sinn, 1980) and assume that the elements of the random parameter vector $\tilde{\mathbf{z}}$ are independently and uniformly distributed. That is, for each dimension of $\tilde{\mathbf{z}}$ we specify an interval wide enough to include all reasonable values of the component and assume that each value in that interval is equally likely to be realized. Let $f(\cdot)$ denote the (constant) density defined on $\mathbb{Z}$. For each realization of $\tilde{\mathbf{z}}$ with distribution parameter subvector $\mathbf{z}^s$, the distribution over states of the world has density $h(\cdot;\mathbf{z}^s)$. Once again, we assume that $h(\cdot;\mathbf{z}^s)$ is a constant; that is, the subvector $\mathbf{z}^s$ specifies the supports of the various random variables. Now, for each $\mathbf{z} \in \mathbb{Z}$ (the "outer loop"), we compute, for each realization $s \in \mathbb{S}$ (the "inner loop"), players' payoffs for each policy in $\mathbb{X}$ and for the default outcome. We then take expectations over $\mathbb{S}$ to identify the PD set for the realization $\mathbf{z}$. This approach can be interpreted as providing, through the probabilistic viability function $V(\cdot)$ and its associated upper and lower contour sets $\mathbb{C}_k^+$ and $\mathbb{C}_k^-$, a comprehensive picture of political viability across the entire spectrum of plausible parameter configurations. In the following two sections, we apply this methodology to a specific policy problem, and study $\mathbf{W}(\cdot)$ and $V(\cdot)$ in that context.

3. The Delta Application

We illustrate the methodology proposed in Section 2 with a case study of the debate over the future of the Sacramento-San Joaquin Delta. This debate is a typical example of the kind of political-economic problems our methodology was designed to analyze. Data are scarce and their applicability is controversial. The problem is exceedingly complex and multi-faceted. The issues that arise involve market and non-market goods, privately owned and common-pool resources, and an intricate mix of economic, environmental and engineering objectives. The scientific relationships between key variables are imperfectly understood. There is a great deal of exogenous uncertainty. There is a diverse set of stakeholders whose conflicting, non-comparable interests cannot be balanced against each other using conventional utilitarian principles. While the economic and environmental problems are hard to analyze, the political issues are even more soIn short, the Delta conflict is a highly ideosyncratic, one-of-a-kind event and so is not a natural candidate for econometric analysis. Quoting Lempert *et al.* (2003) (p. xiii) "the long-term...presents a vast multiplicity of plausible (scenarios). Any one or small number of stories about the future is bound to be wrong." For these reasons, it is appropriate to search among potential resolutions for ones whose political viability is robust with respect to a wide array of possible characterizations of the political situation.

3.1. **The Sacramento-San Joaquin Delta.** The Sacramento-San Joaquin Delta is the largest estuary on the west coast and is home to a number of threatened or endangered species. It also serves as the hub of California's intricate water supply system. Two large river systems draining California's Central Valley flow into the Delta before emptying into San Francisco Bay, with the Sacramento River entering in the north and the San Joaquin River entering in the south.

Prior to large scale settlement of the region, the Delta was a marshy region of shifting channels and salinity. During low runoff periods, saline water from the San Francisco Bay would reach the western parts of the Delta; during high runoff periods, large quantities of fresh water in the river systems would push the salinity barrier further west into the Bay. Today, the region is a series of levee-protected islands surrounded by fixed channels. The vast majority of the San Joaquin River water is diverted upstream; the flow reaching the Delta through the river is primarily agricultural runoff. Large quantities of Sacramento River water are also diverted upstream. A substantial portion of what does reach the Delta is then pulled south, against natural flow patterns, to the large pumping plants at the southern end of the Delta. From these plants, water is exported south to agricultural users in the San Joaquin Valley and urban users in both Southern California and the Bay Area.[9] To protect the quality of water exported south, the salinity of the Delta is carefully regulated using timed releases of water to keep the boundary between fresh and saline water in a relatively fixed location.

[9]Some Bay Area communities draw water from the Delta in other locations.

Today, the Delta is widely acknowledged to be in crisis. The region serves two critical needs for California: ecosystem services and water infrastructure. While there has always been some tension between these goals, the conflict between them has intensified in recent years. Fish populations have crashed, leading to lawsuits filed under the Endangered Species Act (ESA). Five species are listed as threatened or endangered. In response, federal Judge Wanger imposed dramatic cuts in water exports in an effort to boost fish populations (United States District Court, 2007). The resulting reductions in water availability have contributed to rising unemployment rates in many agricultural regions reliant on the Delta for water.

In addition to the ecosystem concerns, the Delta faces a substantial risk of levee failure. The aging levees protecting Delta islands are widely acknowledged to provide insufficient protection. Due to land subsidence and sea level rise, many Delta island lie well below the water level. The levees protecting these islands are at risk from isolated failures, like the Jones Tract failure in 2004, and catastrophic simultaneous failures of many levees. There is particular concern that anticipated earthquakes on the region's faults could trigger many simultaneous failures. In the event of massive failure, water would rush in to fill the levee lined islands. As water filled the islands, saline water from San Francisco Bay would be drawn into the Delta, contaminating Delta water and making it unfit for drinking, agricultural production, and fish. The consequences to California's residents would be enormous; nearly two-thirds of the state's residents rely on the Delta for drinking water. It is anticipated that a major levee breach would cost between $8 and $15 billion (Lund *et al.*, 2008).

3.2. **Proposed solutions.** Several independent studies have studied how the state should respond to the impending crisis outlined above. (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). The Lund *et al.* (2008) report has been particularly influential and our analysis draws heavily on it. 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 and continue exporting water through it, build a canal or other conveyance 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 in a canal.

The first strategy, stopping all exports, would have sweeping consequences. Agricultural and urban interests currently reliant on the Delta for water would need to reduce their water use, find alternate sources of supply, or do some of both. Water conservation, land fallowing, wastewater recovery, and desalination would all likely play major roles in the adaptation. Each of these responses would be extremely costly; Lund *et al.*

(2008) estimate that stopping all water exports through the Delta would cost between 1.5 and 2.5 billion dollars per year. While this strategy would likely be the best option from an ecosystem prospective, it is important to recognize that even stopping all exports would not **guarantee** the recovery of endangered fish populations.

The second strategy is to reinforce the levee structure in the Delta while continuing to export water through it. This strategy is appealing in that it does not require the construction of a new canal or other conveyance infrastructure, whose eventual performance would be unknown and whose price tag would be high. Water managers would continue policies designed to keep the Delta's salinity below specified targets. Given the current risk of levee failure, a through-Delta strategy would require substantial investments in levee upgrades. However, most engineers believe that it would be impossible to eliminate the risk of catastrophic levee failure. As a result, choosing a through-Delta strategy implies accepting some degree of failure risk. Moreover, most ecologists believe that such a system is likely to be the worst of the four alternatives from an ecosystem perspective.

The third strategy is to build a canal, tunnel or other conveyance system around the Delta.[10] Today, all the water that reaches the lower end of the Sacramento River flows into the Delta. By exporting water around the Delta instead of through it, only water not destined for export would flow into the Delta. Still, several ecosystem impacts would be reduced under this system. In particular, the canal would eliminate the "flow reversals" that occur when the export pumps draw water against its natural flow patterns. It would also eliminate the need to regulate the salinity of the Delta as water would enter the canal upstream from the Delta. Finally, a canal would insulate the state's water supply from the risk of levee failure. A canal would be expensive to construct, but water users have pledged to pay for it in exchange for the security it offers. Although controversial, many biologists believe that such a system would be better for the region's ecosystem than the status quo, despite the lower quantities of fresh water that would flow into the Delta. However, in-Delta interests are worried about the impact of these lower inflows on Delta water quality. Moreover, many groups have expressed concern that the institutions charged with managing the canal might eventually bow to political pressure to renege on agreed upon limits on canal usage.

The final strategy is to combine the previous two: export some water through a canal and some through the Delta. This strategy is known as "dual-conveyance." Many believe that a dual conveyance strategy

[10]The idea of a tunnel (or pair of tunnels) rather than a surface canal is a relatively recent one *(Bay Delta Conservation Plan, 2010, p. 35)*. While there are important differences between the two conveyance systems, we view them as second order relative to the ones discussed in this paper. In what follows, we use the word "canal" as shorthand for "some surface or underground conveyence system that is isolated from the Delta itself."

could represent the best of both words: maintaining inflows to the Delta and therefore maintaining its water quality, providing the flexibility to route exports in the least harmful fashion at any particular point in time, and providing a secure export option in the event of levee failure. Others fear that it could be the worst of both worlds: expensive to construct and not guaranteeing either ecosystem health or water supply reliability by failing to separate these functions.

Opinions about how to proceed are numerous.Supporters of the canal include Arnold Schwarzenegger and Jerry Brown, California's past and present governors, an independent group of experts, many water export users, the Bay-Delta Conservation Plan committee, and the Nature Conservancy. On the other side, farmers, local residents and recreational users within the Delta (henceforth referred to collectively as "in-Delta interests") and some Northern California residents have declared their opposition. Many environmental groups are withholding judgment but have indicated that if appropriate safeguards were guaranteed, a canal might be part of a workable solution. The main source of concern among environmental groups is that the canal will not be operated in accordance with environmental protection laws and the Endangered Species Act (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. The situation is evolving rapidly; at the time of this writing, a forthcoming report from the National Academies of Science (National Research Council (2011)) regarding California's Draft Bay Delta Conservation Plan is anticipated to shift the terms of the discussion. In order to assess the political viability of the various options on the table, we construct a probabilistic political viability model of the political debate over the future of the Delta.

3.3. **A Probabilistic Political Viability Model of the Delta.** In this subsection we embed a computable model of the policy options facing the government into the framework introduced in Section 2. The model has several basic components: policy choices, stakeholders, uncertainty, outcomes, and utilities. Given the complexity of the choices facing the state, the full specification is quite detailed. This section provides a sketch of the entire model and focuses on the details that are critical for the results. The technical appendix provides a complete description of the model's functional form and parameter specifications.

Following Lund *et al.* (2008), we focus on the choice of how much water to export and in what fashion. 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. We chose these stakeholders because we believe that it would be politically very difficult to impose a solution to the Delta's problems over the vigorous objections of any one of them.[11] To the extent that each of these groups possesses some degree of veto power over the final solution, our political viability criterion—Pareto dominance—is a necessary condition for a solution to be sustainable.[12] Indeed, in his report on a decision to delay voting on a water bond that earmarked initial funding to pave the way for a canal, *Sacramento Bee* political reporter Dan Walters observed that the bond violated an "unwritten rule" 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). In their analysis of the Delta, Madani & Lund (2011a) consider four game-theoretic solution concepts but a necessary condition for a policy to solve any one of their games is that it (weakly) Pareto dominates the status quo.

3.3.1. *Policy Choices.* Our model focuses on the government's decision about how much water to export and the manner in which it will be exported. A policy in our model is represented by a pair $(x_{\text{ex}}, x_{\text{shr}}) \in \mathbb{X} \subset \mathbb{R}^2_+$, where x_{ex} is the total amount of water exported from the Delta region and x_{shr} is the share of exports routed through a canal or some other conveyance. We let x_{ex} vary from zero to 7.5 million acre feet (maf)[13] and let x_s vary from zero to one. Each of the solutions identified in Lund *et al.* (2008) can be represented as a specific point in $\mathbb{X}$. Our 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 canal and through-Delta pumping, we represent the dual conveyance alternative by dividing exports evenly between the two; alternate values of x_{shr} would correspond to different dual conveyance alternatives.

3.3.2. *Outcomes.* As noted in subsection 2.1, each policy vector is mapped to a stochastic outcome vector which represents the payoff-relevant consequences of implementing that vector. Many of these outcomes are financial. In particular, different export regimes impose different types of costs that are shared among three of our stakeholder groups: agricultural users and urban users that rely on Delta exports for some of their water, and the taxpayers. We include five specific costs in our model: costs due to reduced water exports, water treatment costs, levee maintenance costs, repair costs following a major collapse, and costs associated

[11] There is another interest group that has some degree of veto power: agricultural users upstream of the Delta. We have omitted them because of lack of data.
[12] It is less clear that the in-Delta interests have real veto power over Delta solutions, but they are certainly a vocal interest group exerting substantial influence over the process.
[13] 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

with a major collapse of the levee system.[14] Table 1 summarizes the allocation of these costs and the key pathways through which the policy vector $\mathbf{x}$ influences them; the technical appendix provides specific details. Several of the costs are borne only in the event that a canal either is or is not constructed. As a result, our mapping from policies to outcomes is discontinuous as we move from $x_{\text{shr}} = 0$ to $x_{\text{shr}} > 0$. This discontinuity induces discontinuities in stakeholder preferences that are discussed in section 4.

In addition to financial impacts, there are other outcomes that affect stakeholder utilities, including agricultural employment in the San Joaquin Valley, inflows to the Delta, and the possibility of fish extinction. The technical appendix includes details about the specific functions we used to calculate these outcomes conditional on the export regime; the role they play in stakeholder utility is discussed in subsection 3.3.5.

3.3.3. *Stakeholder Uncertainty.* All of the outcomes defined above are contingent on the state of the world s. The future state of the world incorporates all of the uncertainty that stakeholders confront about how policy choices translate into specific realizations of the random outcome vector.The uncertainty reflects scientific controversy among experts, as well as the inherently stochastic nature of the linkages between policy decisions and their ultimate impacts. Scientific disagreement about how fish populations will respond to various changes in water export regimes is a notable example. Another source of uncertainty is that certain payoff-relevant events will not occur until the future: these include the occurrence and timing of a major levee collapse, whether fish species will recover, and whether exports will be cut at some future time to aid the species' recovery. The technical appendix presents a detailed description of each of the components of our model that vary across states of the world and their assumed distributions.[15]

[14]The costs included in this category do not cover all potential consequences of a collapse. For instance, they do not include 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 in subsection 3.3.5.

[15]Our approach is related to the one taken by Madani & Lund (2011a). Their Monte Carlo analysis draws from a distribution similar to our distribution over states of the world. A key distinction between their analysis and ours is that every draw from

There is one last kind of uncertainty that plays a critical role both in the real-world policy debate and in our model: how, exactly, will future export regimes be implemented? Here we distinguish between the negotiated policy choice $\mathbf{x}$ that emerges from the political process—we call this the *declared policy*—and the *actual export regime in a particular state of the world.* The actual export regimes may in some states deviate from the declared policy for one of three reasons. The first is that reductions in exports may be required in order to comply with the ESA. For instance, if the political process agreed upon a Delta solution in which exports reverted to pre-Wanger levels (around 6 maf), fish populations would quite likely continue their rapid decline. In some states of the world, the judicial system would then intervene, mandating a significant cut in exports. In these states, actual exports would diverge from the declared policy. The second reason is that, for reasons that will be discussed in subsection 3.4, the institutions responsible for water management may fail to implement the declared export policy. The final reason is that if a major levee collapse occurs, the post-disaster export regime may differ from the pre-disaster export regime.

To model this relationship between declared export policy and the realized export regime, we define the function $\mathbf{g}(\mathbf{x}; s, \mathbf{z}_y)$, representing the actual export regime that results in state s from a policy choice $\mathbf{x}$, given the outcome parameter vector $\mathbf{z}^y$. As noted above, stakeholders' utilities depend on the outcome vector $\mathbf{y}$, which depends in turn on the *realized export regime*, so we write $\mathbf{y}(\mathbf{g}(\mathbf{x}; s, \mathbf{z}^y); s, \mathbf{z}^y)$, which is the vector of outcomes associated with the actual export level $\mathbf{g}$ induced by policy $\mathbf{x}$ in state s. Thus s can affect the outcome vector both directly, through state-contingent values such as the relationship between a particular export level and water supply costs, and indirectly, through its influence on the actual export regime. This functional form is a special case of the general form for the outcome function $\mathbf{y}(\mathbf{x}; s, \mathbf{z}^y)$ introduced in Section 2.

3.3.4. *The Default Outcome.* The default outcome in our model has both deterministic and random components. We assume that if stakeholders cannot agree on a policy alternative in $\mathbb{X}$, then no major policy change will be implemented. As a result, no money will be spent on maintaining the levees and a peripheral canal will not be built. Because of this neglect, the probability of a massive levee failure will increase. 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 default exports will depend on whether or not fish populations show signs of recovery and whether export reductions are imposed if not.

their distribution resolves all payoff-relevant uncertainty: their stakeholders face no uncertainty in the games that they play. By contrast, stakeholders in our model compare policy alternatives to the default outcome on the basis of expected utility computations across possible states of the world.

A modeling dilemma arises relating to the level of default exports. This level could be treated either as a parameter of our model, known to stakeholders but not to the modeler, or as dependent on the future state of the world, and hence unknown to both. The latter approach better reflects the true political situation but has a cost. If stakeholders take expectations over the many possible levels of default exports, its influence on the stakeholder's choice of action will be difficult to isolate. In particular, if the distribution describing stakeholder uncertainty about default exports is held constant across modeling uncertainty, then stakeholders face exactly the same kind of uncertainty regarding default exports in every realization of modeling uncertainty. If the distribution, which is known to stakeholders, is subject to modeling uncertainty—e.g., if the support of the default export distribution is unknown to the modeler—then the information our simulations provide about the influence of default exports is limited to how the model's predictions depend on the distribution, which is known ex ante by stakeholders. By contrast, if we treat the realization of default exports as known ex ante to stakeholders but a component of modeling uncertainty with a specified distribution, then when we run simulations over a wide range of model parameterizations, we generate helpful information about how robust are our political predictions to variations in the value of default exports. We are, in effect, running a sensitivity analysis on default exports and many other variables simultaneously.

Susan: We need to sell the importance of this experiment here.

Our simulation results illustrate the different consequences of these two approaches. In our primary simulation runs, we model default exports as state-contingent and treat the parameters governing the distribution of the relevant state as components of modeling uncertainty. To reflect stakeholder uncertainty, we write the default export regime as $\mathbf{g}^d\left(s, \mathbf{z}^d\right)$ and the default outcome as $\mathbf{y}^d\left(\mathbf{g}^d\left(s, \mathbf{z}^d\right); s, \mathbf{z}^d\right)$. As always, we assume that stakeholders know the value of $\mathbf{z}^d$ but not the realization of the state of the world s; as usual, $\mathbf{z}^d$ is a component of modeling uncertainty. In an auxilary simulation, discussed in subsection 4.3 below, we assume that the level of default exports is not state-contingent, and that stakeholders (but not the modeler) know the level of default exports, and there is modeling uncertainty about this level. For this case, the symbolic representation of the default outcome will be $\mathbf{y}^d\left(\mathbf{g}^d\left(\mathbf{z}^d\right); s, \mathbf{z}^d\right)$; that is, default exports are no longer state dependent. Of course, the components of the default parameter subvector $\mathbf{z}^d$ will be different in the primary and auxiliary simulations.

3.3.5. *Stakeholder Utility.* With the exception of environmentalists, each stakeholder group has a CES utility function defined over the components of the outcome vector that we assume affects its utility. Many of these outcomes are financial and some change discontinously at the boundaries of the policy space. Each group's preferences over outcomes induce preferences over policy variables, although the linkages are not immediately transparent. We discuss the critical linkages briefly below and in more detail in the results section. Environmentalists are a special case. For tractability, they are concerned exclusively about the survival of two fish species, Delta smelt and salmon. Because fish survival is a binary variable, the CES specification we use for other groups is inappropriate for environmentalists. The following list introduces the arguments of each stakeholder group's utility function and, in simplistic terms, the relationship between these arguments and the two policies being negotiated. The two letter code after each group's title will be used as shorthand to identify groups when we present our results. Further discussion of our utility specifications is provided in subsection 4.1.1 below.

State taxpayers (Tx): We assume that 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 a canal does not exist. Thus Tp's utility is increasing in x_{shr}, decreasing in x_{ex}, and jumps up as we move from $x_{\text{shr}} = 0$ to $x_{\text{shr}} > 0$.

Urban users (Ur): This group, an aggregate of urban interests in Southern California and the San Francisco Bay Area, 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 ecosystem cutbacks increase as water exports are shifted from a canal to the Delta, so urban user utility also increases with x_{shr} .

Agricultural users (Ag): This group includes farmers in the the agricultural regions of the San Joaquin Valley. The two arguments in Ag's utility function are farming profits and the level of agricultural employment. Ag's preferences over $\mathbb{X}$ are very similar to those of Ur, although Ag's utility decreases faster than Ur's as x_{ex} falls since both Ag's profits and agricultural employment 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 is a proxy for the quality of water in the Delta, which we do not model explicitly, but which is highly correlated with Delta inflows. In the absence of a canal, Delta inflows are determined by factors exogenous to our model—hydrological variables and upstream diversions. If a canal were built, then any water exported through it would reduce inflows into the Delta. The second argument, levee maintenance, is a function of the amount of water exported through the Delta: we assume that 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}. The dominant effect of increasing exports is, at high values of x_{shr}, to reduce 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): Environmentalists are concerned exclusively with the survival of two fish species: Delta smelt and salmon. We define four state-dependent utility levels for Ev. If both species survive, its utility is 1; if neither survive it is zero. If only one of the two species survive, its utility is some number between 0 and 1, depending on which species survives.[16] Ev's expected utility thus increases as the probability of survival increases and hence decreases with x_{ex}. The impact of x_{shr} is more complex and is discussed in detail in section 4.

Following the approach in Section 2, we write the vector of expected utilities resulting from policy $\mathbf{x}$ as:

$$\mathbf{Eu}\left(\mathbf{x};\mathbf{z}\right) = \int \mathbf{u}\left(\mathbf{y}\left(\mathbf{g}\left(\mathbf{x};s,\mathbf{z}^{y}\right);s,\mathbf{z}^{y}\right);\mathbf{z}^{u}\right) h\left(s;\mathbf{z}^{s}\right) ds$$

and the expected default utility vector as:

$$\mathbf{Eu}^{d}\left(\mathbf{z}\right) = \int \mathbf{u}\left(\mathbf{y}^{d}\left(\mathbf{g}^{d}\left(s,\mathbf{z}^{d}\right);s,\mathbf{z}^{d}\right);\mathbf{z}^{u}\right) h\left(s;\mathbf{z}^{s}\right) ds.$$

3.3.6. *Modeling Uncertainty.* In preceding subsections, we identified several components of the parameter space $\mathbb{Z}$ that are components of *modeling uncertainty*. In particular, we have essentially no empirical basis for assigning specific values either to the coefficients of stakeholders' utilities or to the paramerization of stakeholders uncertainty about the default level of exports (see subsection 3.3.4). The first set of parameters includes the environmentalists' payoffs when only one species survives and, for each of the other groups, its level of risk aversion, its elasticity of substitution, and the relative weights that it assigns to different objectives. The second set are coefficients specifying the distribution over the states of the world that relate to the default outcome. The remaining components of $\mathbb{Z}$ are modeled as degenerate random variables—i.e., ones with singleton support (see footnote 3)—representing parameters known to both stakeholders and the researcher; a list of these is provided in the technical appendix. Table 2 lists each non-degenerate element of $\mathbb{Z}$, along with upper and lower bounds on the interval of values that we consider plausible for that random variable. Since we have no basis for specifying an informative prior over these intervals, the principle of

[16] Woodward & Shaw (2008) specify environmentalists' utility in the same way in their robust control model of water management in the presence of an endangered species.

TABLE 2. Elements and Distribution of Modeling Uncertainty

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 substiution	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.

insufficient reason dictates that each element should be independently and uniformly distributed over the interval we specify as its support.

3.3.7. *Probabilistic Viability of Delta Solutions.* In Section 2 (equations 2.1 and 2.2 respectively) we defined $\mathbf{W}(\mathbf{z})$, the Pareto dominant set (PD) given a particular realization $\mathbf{z}$ of $\tilde{\mathbf{z}}$, and $V(\mathbf{x})$, the probability with respect to modeling uncertainty that the policy $\mathbf{x}$ belongs to the Pareto dominant set. Note that for any given parameterization $\mathbf{z}$, the payoff-relevant characteristics of each policy alternative $\mathbf{x}$ is a random vector because the state of the world is unknown. There are two sources of randomness. First, as discussed in subsection 3.3.3, the total level of exports and the share flowing through the canal are state-dependent. Second, for any realized level of exports, there is uncertainty about how this level will map into payoff relevant outcomes.For any given realization $\mathbf{z}$ of $\tilde{\mathbf{z}}$, however, the PD set is deterministic—a policy is either dominated in expectation by the default for at least one player or it is not.

3.4. **Institutional Mistrust.** In our discussion of the peripheral canal option in subsection 3.2 above, we noted that a major obstacle to the adoption of this alternative is widespread mistrust of the institutions that would be charged with managing the canal. Many groups, including in-Delta interests and some environmentalists, have repeatedly expressed concerns that once a canal is built, commitments to keep exports low enough to maintain Delta water quality will not be honored. Essentially, they fear that if the total capacity available for exports were increased by building a canal, this entire capacity would be maximizally utilized, regardless of the declared level of exports.[17] This concern has been exacerbated by

[17]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

calls from engineers to size the canal or conveyance to match engineering constraints rather than to implement any particular export level. The engineering reasons for this approach are compelling: a large canal 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 unauthorized exports.

We model the impact of institutional mistrust on the political process in a stylized way by introducing an additional state-dependent variable: with probability $\lambda \in [0,1]$, water users will convince regulators to fill the canal to its capacity level at all times; with probability $1-\lambda$, the canal will be operated in accordance with the declared policy $\mathbf{x}$. As discussed in subsection 3.3.3, the realized export regime may diverge from the declared policy—i.e. $\mathbf{g}(\mathbf{x}; s, \mathbf{z}^y) \neq \mathbf{x}$—due to court-imposed restrictions related to the ESA. The possibility of unauthorized exports introduces an additional reason for divergence: if the canal were kept full, both the total level of exports and the share of exports routed through the canal would be higher than the declared policy. As with any other component of s, stakeholders take expectations over the possibility that export level commitments are not honored; the magnitude of λ is another component of the distribution parameter $\mathbf{z}^s$, which is assumed to be known by stakeholders. In section 4, we consider the comparative statics effect on political viability of increasing the value of λ. More precisely, we examine how the upper and lower contour sets of $V(\mathbf{z};\cdot)$ vary as the degree of mistrust increases.

4. Results

Our first set of results focuses on the set of policies that Pareto dominate the default outcome given one particular parameterization of our model, i.e., one particular realization of modeling uncertainty. A natural realization to consider is the expected value $\bar{\mathbf{z}} = \int_{\mathbb{Z}} \mathbf{z} f(\mathbf{z}) d\mathbf{z}$ of the random variable $\tilde{\mathbf{z}}$ with respect to modeling uncertainty.[18] Next, we increase mistrust λ, and examine the changes in size and location of the Pareto dominant (PD) set — $\mathbf{W}(\bar{\mathbf{z}};\lambda)$ — for this parameterization (see eq. 2.1). We then introduce modeling uncertainty into the analysis, and examine the probabilistic properties of the PD set over a wide range of alternative parameterizations. Specifically, we classify policies according to the probability with respect to modeling uncertainty that they dominate the default outcome. We use graphical analysis to illuminate the relationship between our probabilitistic set measures and institutional mistrust. Our final set of results illustrates our discussion in subsection 2.3 about how in our framework insights may be obtained

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)

[18] Since the elements of the modeling uncertainty vector $\mathbf{z}$ are independently distributed, $\bar{z}_i = E(\mathbf{z}_i)$, for each i .

by assuming that stakeholders know the value of a particular parameter, even though this parameter is in reality unknowable.

4.1. **One realization of modeling uncertainty.** To identify the PD set when the modeling uncertainty vector $\mathbf{z}$ is set equal to its expected value, $\bar{\mathbf{z}}$, we use Monte Carlo methods. We draw a large sample from the space $\mathbb{S}$ of states of the world in accordance with the distribution $h(\cdot;\mathbf{z}^s)$. For each policy vector $\mathbf{x}$ in a fine grid of points in the policy space $\mathbb{X}$ and each realized state s in our sample, we compute stakeholder group k's utility $u_k(\mathbf{x};s,\bar{\mathbf{z}})$ from $\mathbf{x}$ given s, and conditional on the parameter vector $\bar{\mathbf{z}}$. We then average $u_k(\mathbf{x};\cdot,\bar{\mathbf{z}})$ over our sample to obtain the expected utility $Eu_k(\mathbf{x};\bar{\mathbf{z}})$ that k derives from $\mathbf{x}$. We repeat this approach to compute k's expected utility, $Eu_k^d(\bar{\mathbf{z}})$ from the default outcome. This exercise identifies the PD set $\mathbf{W}(\bar{\mathbf{z}})$.

4.1.1. *Perfect Trust.* Initially, we analyze the PD set assuming institutions are perfectly trusted by stakeholders. That is, recalling that λ is the probability that water institutions honor their commitments, we identify $\mathbf{W}(\bar{\mathbf{z}};\lambda=0)$, which is the shaded area in Figure 4.1. We describe below the components of this figure in some detail; this explanation will facilitate the discussion which follows about $\mathbf{W}(\bar{\mathbf{z}};\lambda)$, for other values of λ, and about our probabilistic political viability function $V(\cdot)$. The boundaries of Figure 4.1 coincide with the boundaries of our 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 canal, x_{shr}, increases.[19] The filled circles in the figure are stylized depictions of the four policy alternatives discussed in detail in the PPIC report.[20]

Each contour line depicted in Figure 4.1 represents the participation constraint for one of the stakeholder groups, i.e., the set of policy options which yield that group an expected utility level 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. The label attached to each constraint is the abbreviation for the corresponding stakeholder group, defined in subsection 3.3.5.

The shapes of the participation constraints in Figure 4.1 reveal a great deal about the preferences of the various stakeholder groups. Both Ag and Ur will veto policies that involve low levels of exports. Both groups are concerned primarily with achieving a high base level of exports but will trade slightly smaller total exports for larger shares through the canal; consequently, their participation constraints slope steeply

[19]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.

[20]As noted above, the vertical location of 0.5 assigned to the dual conveyance option is arbitrary; many other locations with the same horizontal coordinate as this circle would be equally plausible representations of the PPIC's notion of dual conveyence.

FIGURE 4.1. The PD set with perfect trust conditional on $\bar{\mathbf{z}}$

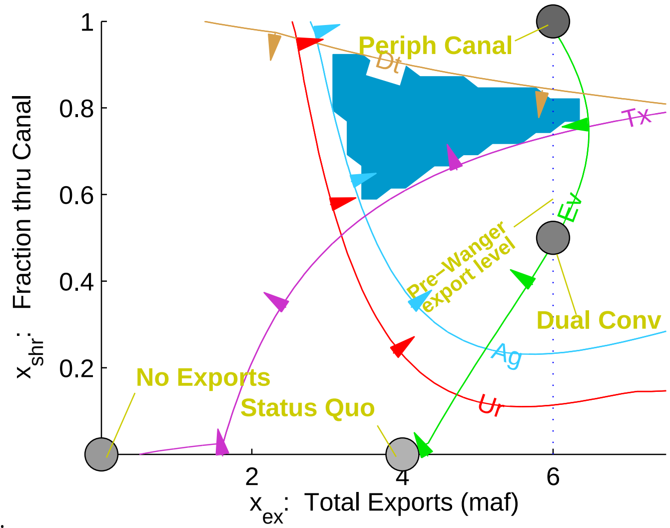

downward. Tp will veto policy vectors in the lower right corner of the space, since it is in this region that the levels of exports through the Delta, and hence expenditures on levee maintenance, will be the highest. Ev will veto policies that involve high levels of exports, although this group is more willing to accept exports if they are routed at least partially through the canal. This reflects the conclusion in Lund *et al.* (2008) that fish populations are more likely to recover under either a dual conveyance or pure peripheral canal option than if exports are pumped exclusively through the Delta. The precise shape of Ev's participation constraint is an artifact of our modeling choices; it is a consequence of our quadratic specification of the fish survival probabilities and our decision to represent the dual conveyence option as splitting exports equally between the canal and the Delta. For Dt, the two outputs which matter—freshwater inflows to the Delta and expenditures on levee maintenance—both decrease with exports through a canal; hence this group will veto policies in the uppermost region of the policy space. Because levee maintenance expenditures increase with total exports, Dt will accept even very high levels of x_{ex}, provided x_{shr} is sufficiently low. As x_{ex} increases, there is a decline in the maximum level of x_{shr} that is acceptable to Dt. It is suprising that Dt is willing to accept such a large fraction of the available alternatives, since Delta interests have always been vociferious in its opposition to a peripheral canal. The reason is that, in our model, these alternatives are being compared to a default outcome that is extremely unsatisfactory in expectation: we assume that 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.

As noted in subsection 3.3.2, there is a discontinuity in the mapping from policies to expected utilities when $x_{\text{shr}} = 0$. This reflects an important discontinuity in the real-world political-economic landscape: if a canal is built, it will have a very high option value, even if $x_{\text{shr}} \approx 0$, i.e., even if under the negotiated agreement, it would hardly ever be used. Moreover, we assume that the size of the canal that will be constructed is independent of x_{shr}, provided that $x_{\text{shr}} > 0$. Thus the state-contingent costs and benefits of a canal both change discontinuously at $x_{\text{shr}} = 0$. Specifically, we assume that provided a canal exists, all exports will be routed through it in the event of catastrophic damage to Delta levees. Thus, both Ag's and Ur's potential exposure to environmental disaster jumps up when $x_{\text{shr}} = 0$, and with it, the minimal level of x_{ex} that either group will accept. Moreover, in the absence of an alternative conveyance option, a major levee collapse will lead to one of two outcomes: either a canal will be built or extensive levee repairs will be undertaken. The added cost of these repairs implies that the maximum level of x_{ex} (and thus the maximum level of 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 level of x_{ex}, fish survival probabilities are lowest when the realized share through the canal is zero (i.e., all exports flow through the Delta). We assume that if $x_{\text{shr}} > 0$, all exports would with probability one be shifted exclusively to a canal following a disaster; in constrast, if $x_{\text{shr}} = 0$, exports will, with positive probability, continue to be routed exclusively through the Delta, with strong negative implications for the fish. In short, even if a canal would be used only in the event of a distaster, its existence would contribute in expectation to the survival chances of the fish. For this reason, the maximum level of x_{ex} that Ev will accept falls discontinuously when $x_{\text{shr}} = 0$.

The shaded region in Figure 4.1 is the PD set at the mean level $\bar{\mathbf{z}}$ of modeling uncertainty, i.e., the set of policies that are not vetoed by any stakeholder group under this parameterization of the model. Note first that this set is nonempty, i.e., there do exist policies that Pareto dominate the default. We can conclude from this that if our model, when parameterized by $\bar{\mathbf{z}}$, is a reasonable stylization of the actual political process, then we cannot rule out the possibility that *some* negotiated solution will emerge from the political process. Moreover, for this parameterization, two necessary conditions for a policy to be acceptable to all stakeholders are that total exports will not exceed the pre-Wanger level of 6 maf, and no more than half of all exports will flow through the Delta. Finally, under this parameterization of our model, at least three of the four options identified by the PPIC report 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: if exports

FIGURE 4.2. Impact of mistrust on the location of PD set with parameterization $\bar{\mathbf{z}}$

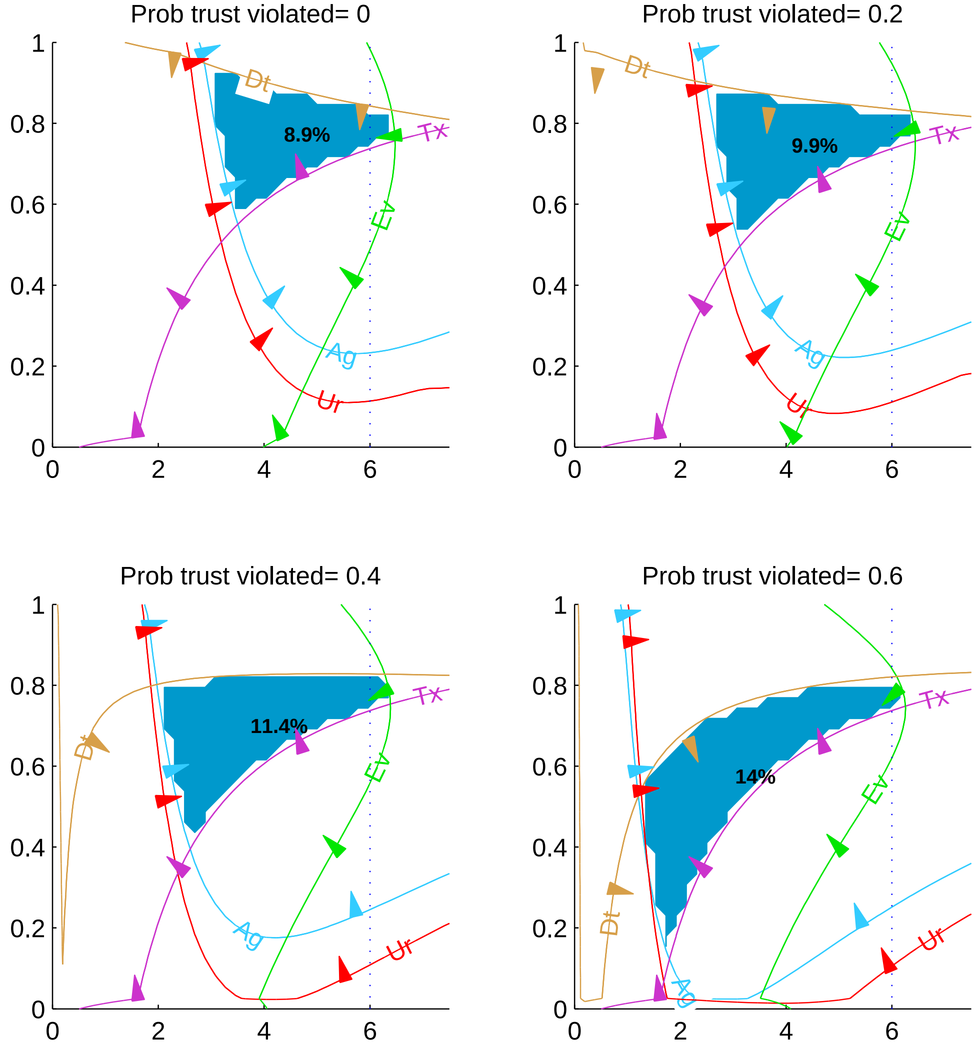

were split evenly between canal and Delta, it would be vetoed; however, if the share through the canal were higher (roughly between 60% and 80%), then the dual conveyance option might survive the political process.

4.1.2. *Impact of Mistrust.* In subsection 3.4 we introduced our approach of modeling mistrust: we assume that with probability $\lambda \in \{0, 0.2, 0.4, 0.6\}$, the canal will be filled to capacity. In this subsection we compare the PD sets, i.e. the sets $\mathbf{W}(\bar{\mathbf{z}}; \lambda)$, for these four levels of λ. That is, we restrict our attention to just one parameterization of the model, given by $\bar{\mathbf{z}}$. Figure 4.2 is the analog of Figure 4.1, but for all four levels of mistrust. The first panel replicates Figure 4.1. 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. The location of each number roughly indicates the center of the corresponding set.

The successive increases in mistrust induce shifts in stakeholders' participation constraints, and hence the location and size of the PD set. At higher levels of mistrust, both Ag and Ur are willing to accept lower levels of declared total exports, because the actual levels of total exports (and shares through the canal) will exceed declared policy levels. As mistrust increases, Ev's participation constraint also moves to the left and its curvature increases.[21] The shift in Dt's participation constraint as mistrust increases is particularly dramatic. The change in its curvature between $\lambda = 0.2$ and $\lambda = 0.4$ can be traced to the change in the shape of the contours of the expected actual share routed through a canal. Canal exports lower Dt's expected payoff because the water is diverted upstream of the Delta, thus lowering Delta water quality by reducing freshwater inflows. Tp's participation constraint is independent of the level of mistrust, since all of the costs for which this group is responsible are independent of the level of mistrust.

The combined impacts of these shifts determine how the shape and location of the PD set changes as mistrust increases. Several changes are immediately apparent from Figure 4.2. First note that the *size* of the set increases monotonically with mistrust, from 8% to 14% of the entire space. This result seems counter-intuitive at first glance: one could interpret the size of the PD set as a summary measure of the prospects for a successful negotiation, and one would certainly expect that these prospects would be diminished in the presence of mistrust. But this intuition does not take into account that there are gainers as well as losers from mistrust. As mistrust increases, the constraints of the mistrusted groups slacken, while those of the mistrusters tighten; the former are more willing to come to an agreement upon which they may be able to renege; the latter are less willing to agree, because the agreement may be reneged upon. Specifically, as noted above, the participation constraints for Ag, Ur and Ev all move to the left. The net impact of these leftward shifts on the size of the PD set can be decomposed into two factors. The first is the relative rates at which these shifts occur. These rate changes depend, in turn, on considerations such as the relative risk aversion of the different stakeholders: as the figures indicate, Ag's and Ur's constraints are shifting left at a significantly faster rate than Ev's. It is the second factor, however, that plays the more significant role. Because of the interaction between Tp's and Dt's participation constraints, the PD set is much narrower at its right-hand edge than at its left-hand edge: that is, 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

[21] The upper bound on declared export levels to which Ev will agree shrinks as the probability of export levels in excess of declared levels increases, while there is a contraction in the maximum fraction through the Delta that Ev will accept. The downward shift is driven largely by that fact that under our specification of the fish survival probability functions, survival rates are highest when roughly 75% of exports flow through a canal. Moreover, as Figure 4.3 indicates, declared shares through the canal correspond to higher and higher actual shares as mistrust increases. Hence declared share levels that were initially near the upper boundary of Ev's acceptable set become unacceptable as mistrust increases.

FIGURE 4.3. Expected export configurations for parameterization $\bar{\mathbf{z}}$
Insert new 2x4 figure here.

and Ag/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.

It is striking in Figure 4.2 that Dt's constraint increasingly limits the set of possible agreements. As we noted earlier, 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 clearly higher than that of any other group. But once the probability of a trust violation reaches 0.6, the fraction of alternatives that Dt will accept is particularly small. Our model thus suggests that a critical factor driving Dt's highly publicized opposition to the canal is its strong belief that agreed upon restrictions on exports are unlikely to be honored in practice.

It is also clear from the figure that as mistrust increases, the PD set moves down and to the left; the acceptable policies are characterized by fewer total exports and, of these, less through the canal. The downward shift is driven primarily by the shift in Dt's constraint; the leftward shift by the shifts in Ag's, Ur's and Ev's constraints. The shift toward lower exports is not surprising, but this figure tells only part of the story, i.e., what happens to the levels *of declared* water exports in the PD set. As mistrust increases, so does the probability that water exports will fill the canal to capacity. Hence the reduction in *expected actual* exports associated with alternatives in the PD set is less than the figure would suggest, indeed, expected actual exports could in fact actually increase.

Figure 4.3 illustrates how mistrust affects the mapping from negotiated policies to expected actual export configurations.[22] Each panel in Figure 4.3 is a contour plot: each solid line is a locus of policy vectors for which the expected level of actual exports or actual share through the canal is the labeled amount. Included as a reference, the solid area in each panel is the PD set corresponding to λ, i.e. the set $\mathbf{W}(\bar{\mathbf{z}};\lambda)$. With perfect trust, the contours for expected actual exports in the top row of Figure 4.3 are nearly vertical lines; however, each *declared* level implies a lower *expected actual* level because of the possibility of fish-mandated cutbacks. For any declared policy such that $x_{\text{shr}} > 0$ and $x_{\text{ex}} > 0$, the corresponding expected level of actual exports increases with mistrust; in the last column, expected actual exports are expected to be nearly half the pre-Wanger level even when the declared export level is almost zero.[23] Now consider the second row of Figure 4.3. With perfect trust, the contours are perfectly horizontal and the declared and actual expected

[22]Throughout this discussion of the mapping, we take expectations over only those states of the world for which a major levee collapse does *not* occur.
[23]If $x_{\text{ex}} = 0$, no canal is built and the possibility of additional exports is foreclosed.

FIGURE 4.4. Probability of Pareto dominance with perfect trust

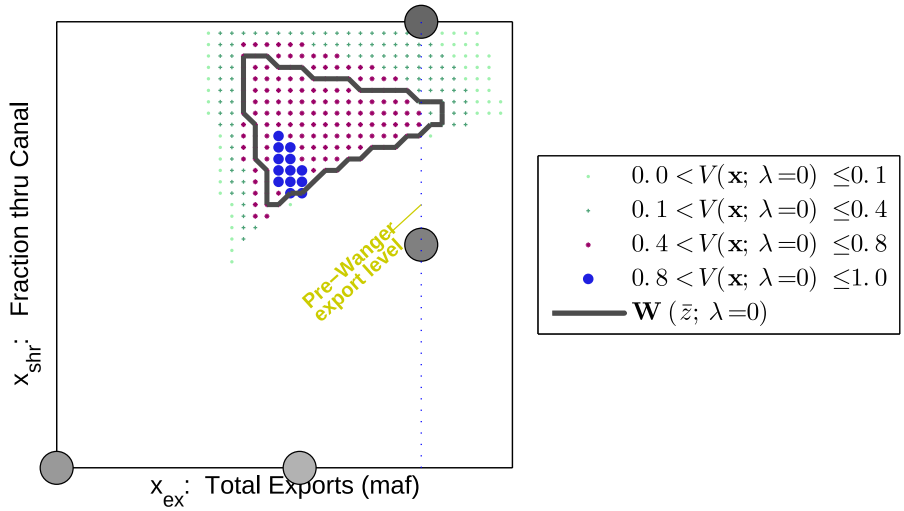

shares coincide. However, since any unauthorized excess exports would flow through a canal, the actual share through the canal increases considerably as mistrust increases. To see this, note that at low levels of declared total exports, very small amounts would flow through the Delta, whether or not commitments are honored; on the other hand, if water users succeed in lobbying regulators to increase actual exports, large quantities of water will flow through the originally nearly empty canal, dramatically increasing the actual share that flows through it.

4.2. **Probabilistic Pareto dominance.** In the preceding subsection, we examined the properties of our model under one particular parameterization of modeling uncertainty. Many of these properties, however, are determined by the interplay among a large number of parameters. To illustrate, recall how in Figure 4.2 the size of the PD set increases monotonically with mistrust. Trends such as this one depend on the interaction between factors such as the parameters defining stakeholders' utility functions. But very little information is available about these coefficients; we have no reason to believe that the parameterization on which Figure 4.2 is based is a particularly accurate reflection of reality. A natural next step, then, is to compile statistical information about the impact of mistrust, based on a large sample of possible model parameterizations. Accordingly, we summarize in this subsection the data generated by repeating the analysis in subsection 4.1 for 1000 draws from the space of modeling uncertainty, $\mathbb{Z}$. We first set mistrust $\lambda = 0$ and evaluate our probabilistic political viability function, $V(\cdot; \lambda = 0)$, at each element of the policy space, $\mathbb{X}$. (Recall that $V(\mathbf{x}; \lambda)$ is the probability with respect to modeling uncertainty that $\mathbf{x}$ belongs to the PD set.) We then repeat this process for the other three levels of mistrust.

4.2.1. *Perfect Trust.* Figure 4.4 partitions our policy space $\mathbb{X}$ into regions depending on the probability that each policy belongs to the PD set. The legend of Figure 4.4 indicates how points are classified. We will say that a policy as *robustly politically viable* (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 solid circles. Policies that are politically viable for at least one parameterization will be referred to as *possibly politically viable* (PPV). The PPV set includes all of the points marked with some symbol in the figure. Finally, policies in the white (unshaded) region are said to be *never politically viable* (NPV).[24] The solid line is added as a point of reference: it is the boundary of the shaded area in Figure 4.1, denoting the PD set given $\bar{\mathbf{z}}$. The figure is in some sense similar to Figure 4.1; both suggest that most of the policies that have some chance of emerging from the political process involve export levels less than pre-Wanger levels, which are routed primarily but not exclusively through a peripheral canal or other conveyance. Yet Figure 4.4 contains far more information than Figure 4.1.

One critical difference between the two figures is the interpretation of the unshaded region of the policy space. Policies in the unshaded region of Figure 4.1 are Pareto dominated by the default for the single parameterization $\bar{\mathbf{z}}$ of our model. By contrast, NPV policies in Figure 4.4 are Pareto dominated by the default, for *virtually all* reasonable realizations of modeling uncertainty. a striking property of the figure is that all four policies singled out by Lund *et al.* (2008) are NPV. Indeed, *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. Lund *et al.* (2008) were similarly skeptical of all no-canal alternatives, noting that despite the environmental benefits, stopping all exports is simply too expensive for the state, while continuing to rely exclusively on through Delta exports carries unacceptable risks to both water supply reliability and the ecosystem. Moreover, our analysis suggests all points on the top edge of the graph (corresponding to pure canal alternatives) are *also* NPV. These alternatives are always vetoed by at least one stakeholder. Water users are unwilling to pay for a canal if export levels are too low to justify the cost of construction. On the other hand, Dt will veto any pure canal alternative if export levels are too high, because of the combined impact of two negative consequences: water quality in the Delta will decline relative to the default and expenditures on levee maintenance will remain at zero. Finally, our analysis also suggests that a dual-conveyance alternative with pre-Wanger exports evenly split between a canal and the Delta is NPV, although other dual-conveyance configurations are PPV.

[24]Using the terminology defined in subsection 2.2, our vector of probability thresholds is given by $\rho = (0.0, 0.1, 0.4, 0.8)$.The RPV set corresponds to $\mathbb{C}_4^+$; the PPV set corresponds to $\mathbb{C}_1^+$, and the NPV set corresponds to $\mathbb{C}_1^-$.

At the other end of the likelihood spectrum, the set of robustly politically viable policies in Figure 4.4 is considerably smaller than the shaded set $\mathbf{W}(\bar{\mathbf{z}};\lambda)$ depicted in Figure 4.1, and bounded by the solid line in this figure. While roughly 8% of the policy space belongs to $\mathbf{W}(\bar{\mathbf{z}};\lambda)$, less than 1% of the policy space is robustly politically viable. Put another way, only 8% of the policies inside the solid line are robustly politically viable, although almost all of them satisfy the Pareto criterion with probability at least 40%. There is **no** policy that Pareto dominates the default for more than 85% of the realizations of $\mathbf{z}$. These data illustrate the obvious point that inclusion in the PD set for the mean realization of modeling uncertainty is no guarantee of robust political viability.

4.2.2. *Impact of Mistrust.* Equipped with our probabilistic notion of political viability, we now examine further the role that mistrust plays in our model of the political negotiation process. The first panel of Figure 4.5 replicates Figure 4.4; the remaining panels show the impact of increasing mistrust. Legends for these figures are the same as for Figure 4.4. In each panel, we again overlay for reference the boundaries of the shaded set $\mathbf{W}(\bar{\mathbf{z}};\lambda)$ in the corresponding panel of Figure 4.2. As mistrust increases, the set of PPV policies increases; this effect is driven primarily by slackenings, at least for some parameterizations, in the participation constraints for Ag and Ur: as discussed in subsection 4.1.2, both Ag and Ur will be more willing to accept somewhat lower values of both x_{ex} and x_{shr}, the more likely it is that the canal will be utilized to capacity. Just how much more accommodating Ag and Ur become depends on how their utility functions are parameterized. On the other hand, the set of robustly politically viable policies shrinks dramatically, virtually disappearing even for $\lambda = 0.2$. While the first effect is consistent with the information provided by Figure 4.2, the latter effect seems inconsistent with the increasing size of the PD sets in this figure.

Why does mistrust so dramatically reduce the set of RPV policies? We identify three possible explanatory factors based on possible changes in the distribution of the size of the PD sets across model parameterizations. The actual answer to our question will be some combination of these components. First, there could be a decrease in the first moment of the distribution of the size of the PD sets, i.e., a decrease in the number of policy alternatives per parameterization that dominate the default. Second, there could be a *higher moment size* effect, i.e., an increase in the variability of PD set sizes. Specifically, if these sets shrink for a significant fraction of draws from $\mathbb{Z}$, our robustness criterion–satisfying the PD criterion with high probability—will become more stringent. Our third explanatory factor is a *higher moment location* effect: as the "center

FIGURE 4.5. Impact of mistrust on the probability of Pareto dominance

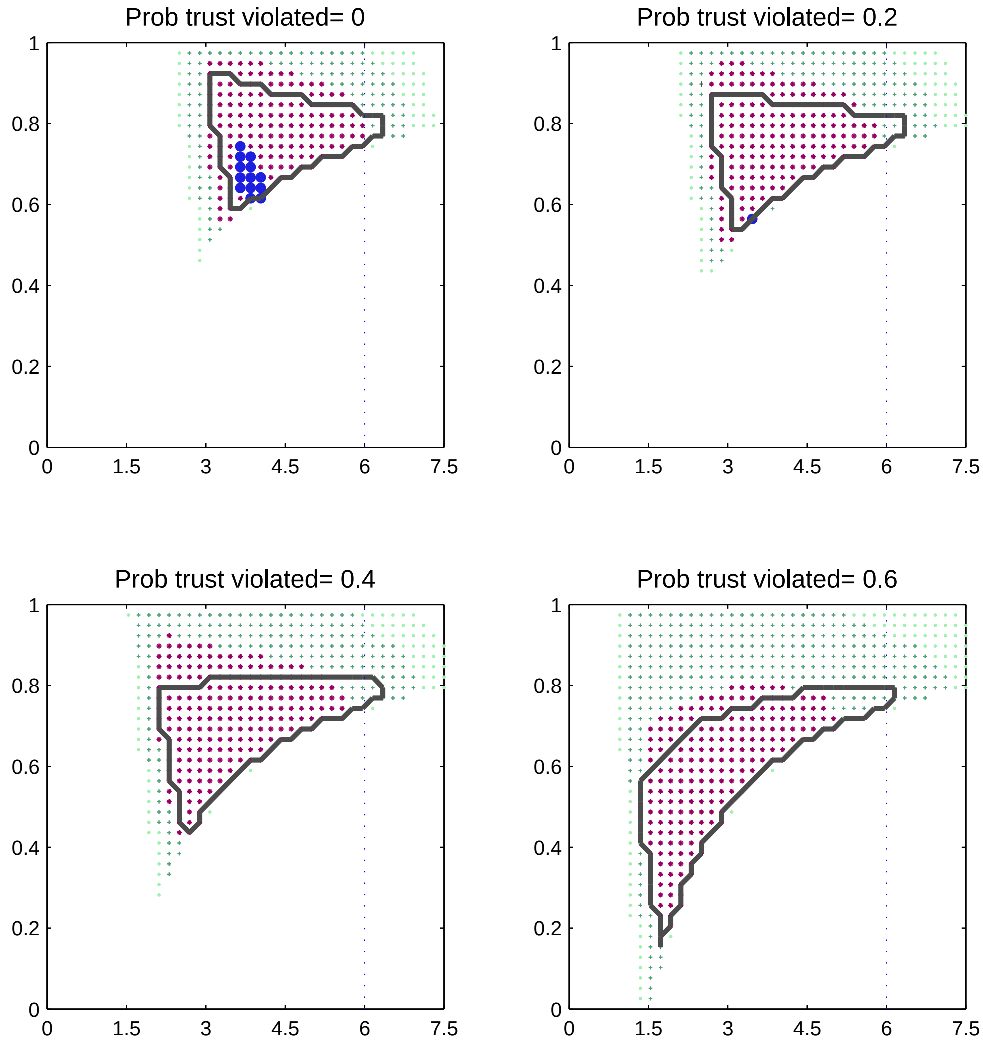

of gravity" in policy space of the PD sets varies more with model parameterizations, the criterion that a particular policy option is PD with high probability becomes more demanding.[25]

Figure 4.6 demonstrates that in our case, both of the higher moment effects play a role, but the first moment size effect does not. Our data strongly indicate that the first moment size effect does not contribute to the shrinking of the RPV sets. Figures 4.2 and 4.5 hinted that the first moment size effect was unlikely to play a role. In Figure 4.2, the PD sets for parameterization $\bar{\mathbf{z}}$ increase with mistrust, as do the PPV sets in Figure4.5 Figure 4.6 confirms this and elucidates the roles of the higher moment effects. While Figure 4.5 indicates the probability with respect to modeling uncertainty that any given policy vector will Pareto dominate the

[25]Note that a *first moment location* effect would not cause the RPV to shrink. If the sole impact of mistrust were to shift the center of gravity of the PD set on average, then the location of the RPV set would change but ~~not~~ its sizewould not.

FIGURE 4.6. Distribution of set measures across modeling uncertainty

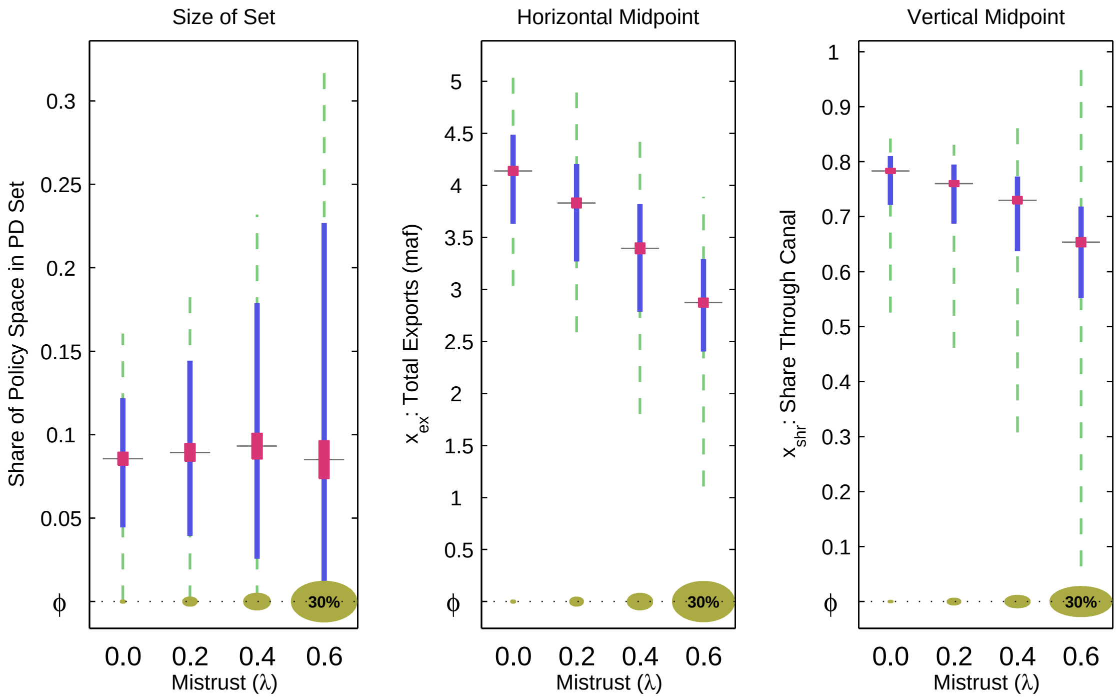

default, Figure 4.6 provides summary information about the stochastic properties of the PD set as a whole. The reported information is obtained by computing three functions that map each parameterization vector $\mathbf{z}$ to certain characteristics of the PD set. The three properties are: the size of the set and the locations of its horizontal and vertical midpoints. The latter two measures are, for each vector $\mathbf{z}$, the means of the *declared* levels of exports and shares through the canal that satisfy the Pareto criterion for that vector $\mathbf{z}$. Using modified box-and-whisker plots, Figure 4.6 summarizes for different levels of mistrust the distribution of these measures across modeling uncertainty. In each panel, and for each value of λ, the solid horizontal line indicates the median value of our sample from $\mathbb{Z}$ for the measure being plotted, while the dotted horizontal line denotes the corresponding number for the set $\mathbf{W}(\bar{\mathbf{z}};\lambda)$ depicted in Figure 4.2 (e.g., in the left panel, the dotted lines are the sizes of the sets for the mean parameterization $\bar{\mathbf{z}}$). The thick, squat rectangles denote 95% confidence intervals for the population medians; a shift in medians from one level of mistrust to the next will be significant if and only if the thick rectangles encasing the dotted lines for these levels have no intersection. The thin, elongated rectangles denote the interquartile ranges of the sample data; the whiskers (thin, dashed lines) indicate the support of the sample data. At the bottom of each panel, the filled ovals indicate the probability that the PD set is empty for the corresponding level of mistrust: for each λ, the width of the corresponding oval is proportional to the percentage of parameterizations for which the PD sets are empty, given that level of mistrust.

As the left panel of Figure 4.6 illustrates, the average size of the Pareto set $\mathbf{W}(\cdot;\lambda)$ varies insignificantly with λ, another indicator that the RPV shrinkage is not due to the first moment size effect. By contrast, both the inter-quartile range and the support of the entire sample increase dramatically. Obviously, the size variable is censored at zero; reflecting this, there is a striking increase in the percentage of parameterizations for which the PD set is empty. These fractions are insignificant for low levels of λ, but increase dramatically for higher levels, reaching 30% when $\lambda = 0.6$. At this level of mistrust, a policy would fail our RPV criterion, even if it belonged to the PD set for *every* realization of modeling uncertainty such that the PD set were nonempty. Given this degree of variation, the RPV set is *necessarily* empty when λ is high. To summarize, the data strongly indicate that the second moment size effect contributes significantly to the evaporation of the RPV set seen in Figure 4.5.

Figure 4.6 confirms that the second moment location effect also plays a role. The trends in each locational median are large, negative and significant, reflecting the shift down and to the left of the PPV sets in Figure 4.5. In the middle panel, we see 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. However, in the right panel, we see that the vertical dispersion of the PD sets increases with mistrust, especially at high levels, confirming that the second moment location effect also plays a role in reducing the size of the RPV set.

Figure 4.7 provides information about the types of changes we observe in the PD set for various parameterizations. While Figure 4.6 relates to the distribution of certain set measures for different levels of mistrust, Figure 4.7 concerns the distribution of *changes* in these set measures, providing additional insight into the impact of mistrust. We first compare the left panels of the two figures, focusing in particular on high levels of mistrust. In Figure 4.6, the median size of the PD set declines (although not significantly) as λ increases from 0.4 to 0.6; moreover, as we have seen, the percentage of parameterizations for which the PD set is empty increases dramatically over this range. These statistics contrast sharply with what we observe in Figure 4.7: the size of the PD set *increases* for roughly 60% of the parameterizations; the median change in the size of the PD set size is positive and significantly different than zero; and there is a statistically significant increase in the size of the median change as mistrust increase. Taken together these statistics suggest that starting from a high level of mistrust, the impact of a further increase in mistrust can be quite different, depending on the particular parameterization of the model.[26]On the other hand, the data summarized by

[26]One interesting property of the data is not evident from our figures: there is a positive correlation between the size of the PD set with perfect trust and the change in its size due to mistrust. That is, for a $\mathbf{z}$ such that $\mathbf{W}(\mathbf{z};0)$ is relatively large (relatively small), $\mathbf{W}(\mathbf{z};\lambda)$ will tend to increase (decrease) with λ.

Figure 4.7. Distribution of Changes in Set Measures Across Modeling Uncertainty

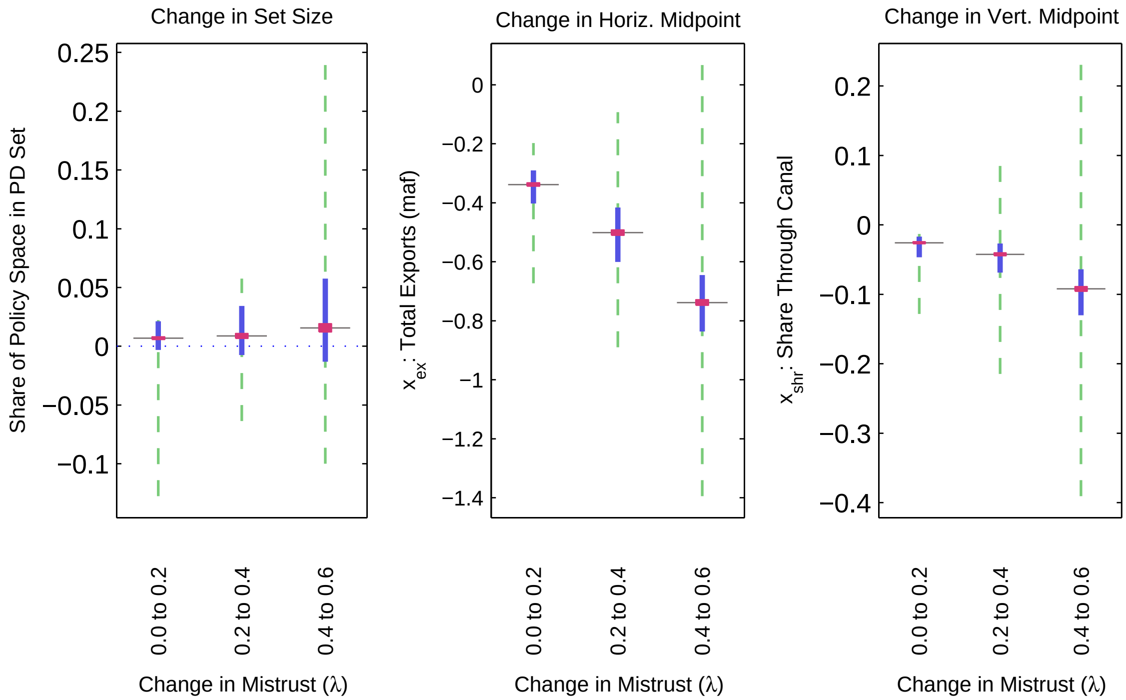

the middle and right panels of Figure 4.7 is entirely consistent with the corresponding data in Figure 4.6: as mistrust increases, the horizontal midpoint of the PD set moves left in virtually every realization, and down in over 75% of the realizations. Finally, perhaps the most striking aspect of Figure 4.7 is the increase in the variability of all three statistics as the level of mistrust increases.

4.3. **Modeling Uncertainty versus States of the World.** In subsection 3.3.4, we discussed two possible approaches to modeling default exports, denoted by $\mathbf{g}^d$. In this subsection, we report on a simulation illustrating the second approach. We shall use the terms "base" and "alternate" to distinguish between our original simulation, discussed in previous subsections, and the one discussed in this one. The two simulations are identical except for the ways in which we treat $\mathbf{g}^d$.In our base simulation, this variable is state dependent; the support of its distribution is a component of modeling uncertainty, i.e., it is a component of the default parameterization subvector $\mathbf{z}^d$. To address the question discussed above, our alternative simulation assumes that stakeholders know $\mathbf{g}^d$, although we, the researchers, do not; more precisely, we replace the component of $\mathbf{z}^d$ that defines $\mathbf{g}^d$'s support with another parameter representing the actual realization of $\mathbf{g}^d$. The distribution over realizations of $\mathbf{g}^d$ is chosen in the alternate case so that for any set G, the probability that the realized export level $\mathbf{g}^d$ belongs to G is identical in both simulations.[27]

Figure 4.8 is the analogue of Figure 4.2. Both depict the PD sets for the mean model parameterization $\bar{\mathbf{z}}$. The two figures are quite similar except that, necessarily, the PD sets in each panel of the current figure are smaller their analogs in Figure 4.2. This is because for any given model parameterization, the only difference between the alternate and the base simulation is that the default outcome in the alternate is less unsatisfactory: for each realization of modeling uncertainty, a random component of the default is replaced by its mean value; because each stakeholder group is risk averse, its expected payoff from the default outcome is unambiguously higher than in the base simulation. The figures also indicate that the higher is the level of mistrust, the greater is the difference between the sizes of the corresponding Pareto sets.

The change in our probabilistic classification is more dramatic. Figure 4.9 is the modified version of Figure 4.5. Figure 4.9 overlays the classification in 4.5 with the analogous classification for the alternate simulation. Each row of the figure represents a different level of mistrust; each column corresponds to one element of the partition we introduced in Figure 4.1. For example, in the first (second) column the points marked with small dots are politically viable in the base simulation for at least one (at least 10%) of the realizations of modeling uncertainty; and so on. The solid lines are the boundaries of the corresponding sets for the alternate simulation.

We derive two key insights from Figure 4.9. First, when stakeholders know the realization of default exports, the set of RPV policies disappears entirely, even under perfect trust. This is a special case of a general result: when stakeholders know the level of default exports prior to evaluating policy options, the set of PPV policy options is larger, while the set of RPV policies is smaller. These general trends make sense intuitively: compared to the base simulation, there is in the alternate an additional element of variability ranging across model parameterizations: previously, different draws from $\mathbb{Z}$ corresponded to different degrees of uncertainty about default exports, but the expected level of this variable remained constant; now, each draw from $\mathbb{Z}$ corresponds to a distinct realization of this variable. Because each parameterization in the alternate simulation is more distinct in this one respect, RPV is a more demanding requirement in this case, while PPV is less so. Second, our qualitative results on the impact of mistrust are more or less unaffected by

[27] In our base case, the ultimate realization of $\mathbf{g}^d$ was determined as a composition of two distributions, one defined on $\mathbb{S}$, the other on $\mathbb{Z}$; the parameters defining the distribution over states of the world affecting $\mathbf{g}^d$ were themselves random variables, depending on the realization of modeling uncertainty. For consistency, we define in the current simulation a single distribution on (one dimension of) $\mathbb{Z}$ that is equal to the original composition. For this reason, realizations of $\mathbf{g}^d$ are not all equally probable. This distinguishes $\mathbf{g}^d$ from all other components of modeling uncertainty; the distribution over every other dimension of $\mathbb{Z}$ is uniform.

FIGURE 4.8. PD Set at $\bar{\mathbf{z}}$; default exports are a component of modeling uncertainty.

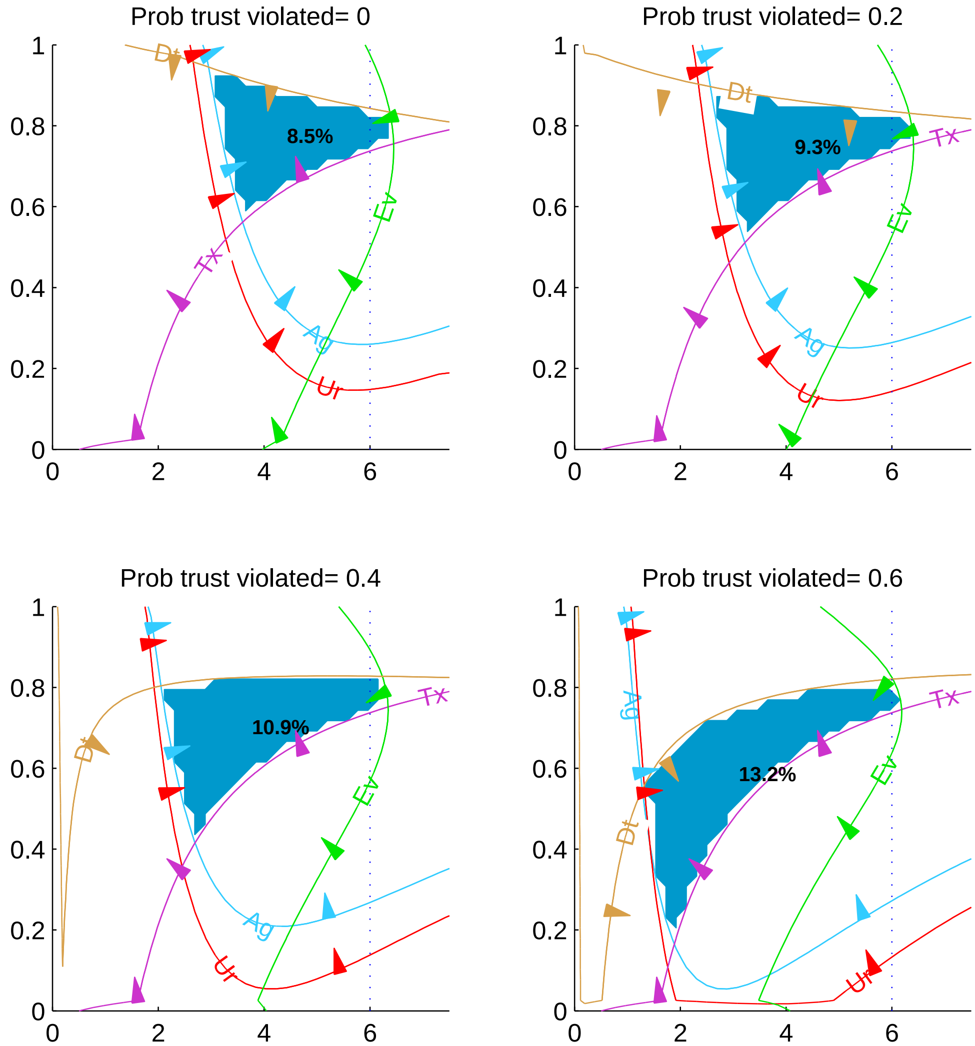

our treatment of default exports. For the three categories of policies that are non-empty in both simulations (i.e. the first three columns of the figure), the changes with mistrust in both size and location of the other categories (the first the columns of the figure) are qualitatively similar. Moreover, the four policies singled out by Lund *et al.* (2008) are NPV in both simulations.

5. Conclusions

In this paper we develop a new methodology for predictive political economy modeling. The approach is designed for the purpose of analyzing complex real-world, one-of-a-kind, political negotiations, typically involving tradeoffs between economic vs environmental objectives, market vs non-market valuations, and

FIGURE 4.9. Probabilistic political viability: default exports are a component of modeling uncertainty

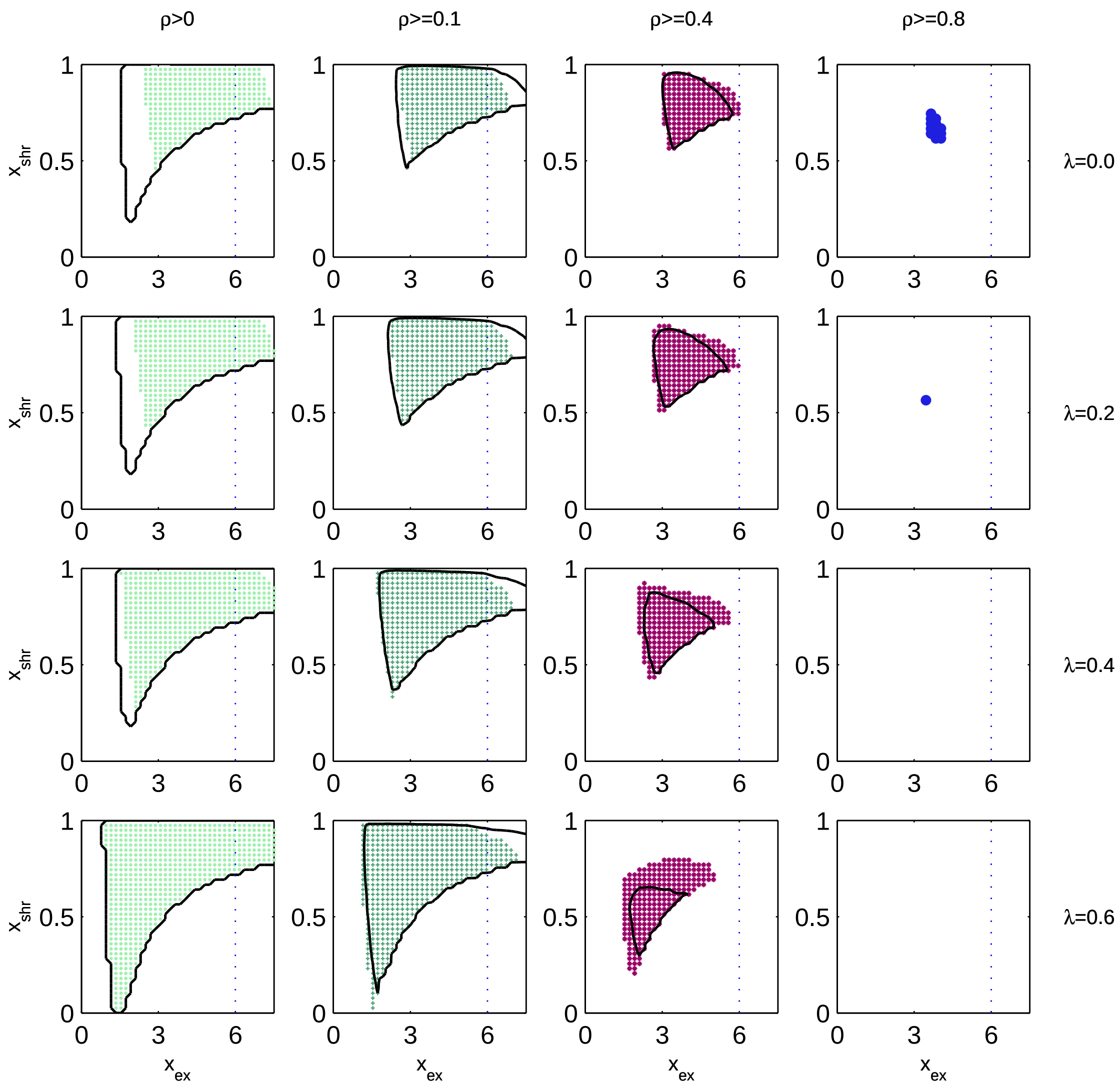

private vs public goods. The end goal of our analysis is to make assessments about which policy outcomes might emerge from the political process. In this context, we as modelers face deep or Knightian uncertainty about many critical components of the analysis, and have no clear basis for selecting one particular model of the decision-making process. To confront these difficulties, we analyze a large, parameterized family of models, and seek to identify outcomes which are robustly political viable over a wide variety of possible specifications. Lacking sufficient knowledge to specify which game-form best captures the real-world political process we are modeling, we adopt a minimal predictive criterion—Pareto dominance—rather than impose any one specific game-theoretic solution concept.

We illustrate our approach by applying it to the current political-economic conflict over the future of California's Sacramento-San Joaquin Delta. This application is particularly well suited to our mode of analysis. It is characterized by a variety of the problems that plague political economists, including public goods, missing markets, and tradeoffs between noncomparable objectives. Commentators agree that a political compromise cannot be implemented without the approval of a broad spectrum of stakeholders; moreover, it is relatively straightforward to identify the range of possible ways that the future may unfold if a compromise cannot be reached. For these reasons, Pareto Dominance is a well-defined concept, and is indeed a necessary condition for a solution to the problem. Finally, a small number of potential solutions that have been identified as focal points of the policy debate. It is instructive to use our methodology to assess the viability of these candidate solutions.

Our specific results regarding the policy debate are consistent with the conventional wisdom. Many experts agree that there is no hope that a consensus solution can be reached. Our analysis strongly supports this wisdom: we should that none of the most widely discussed options Pareto dominate the default, for *any* of the model specifications that we consider. Moreover, only a very small number of policy options that meet our criterion of robust political viaibility. In contrast to the conventional wisdom, our analysis identifies a broad range of options that do meet our standard for political viability, for at least a minority of the model specifications that we consider. Under the conventional wisdom, an important inhibitor of a consensus solution is stakeholders' mutual mistrust. Our analysis illustrates that mistrust dramatically reduces the set of robustly politically viable policy options, although it also increases the set of options which are viability with low probability. In particular, the impact of mistrust is to reduce the political viability of solutions involving high levels of water exports, transported primarily through a new conveyance that bypasses the Delta. Thus, our analysis highlights the importance of designing Delta governance institutions, which could potentially improve welfare by reducing or limiting the extent of mistrust.

Our analysis also contributes to the broader political economy methodology literature. Researchers in a wide range of fields are concerned with how to address Knightian uncertainty. In contexts characterized by this kind of uncertainty, there is little benefit to be gained from seeking to identify the "right" model or the "right" solution. As we illustrate, it is potentially more productive to identify a reasonable family of models and a set of candidate solutions that are robust to a wide range of model specifications. Moreover, our approach illustrates the usefulness of complex, yet still stylized, computable political economy models, demonstrating that they can aid in identifying which model components are critical drivers of the model's results, and which ones leave the model conclusions relatively intact.

References

Bacher, Dan. 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.

Bay Delta Conservation Plan. 2010. *Highlights of the BDCP*. Bay Delta Conservation Plan.

Ben-Haim, Yakov. 2006. *Info-gap decision theory: Decisions under severe uncertainty*. 2nd edition edn. London: Academic Press.

Blue Ribbon Task Force. 2007. *Delta Vision: Our Vision for the California Delta*. Tech. rept.

Blue Ribbon Task Force. 2008. *Our Vision for the California Delta*. 2 edn.

Cooley, Heather, Christian-Smith, Juliet, & Gleick, Peter H. 2008. *More with Less: Agricultural Water Conservation and Efficiency in California, A Special Focus on the Delta*. Oakland, CA: Pacific Insitute.

Delta Vision Committee. 2008. *Delta Vision Committe Implementation Report*. Tech. rept.

Dessai, S., Hulme, M., Lempert, R., & Pielke Jr, R. 2009. Climate Prediction: a Limit to Adaptation. *Pages 49–57 of:* Adger, W.N., Lorenzoni, I., O'Brien, K.L., & Adgar, WN (eds), *Living with Climate Change: Are There Limits to Adaptation?* Cambridge University Press.

Ellsberg, D. 1961. Risk, ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, **75**(4), 643–669.

Gilboa, I., & Schmeidler, D. 1989. Maximin Utility with Non-unique Prior. *Journal of Mathematical Economics*, **18**(2), 141–153.

Gilboa, I., Postlewaite, A.W., & Schmeidler, D. 2008. Probability and Uncertainty in Economic Modeling. *The Journal of Economic Perspectives*, **22**(3), 173–188.

Hanemann, M., & Dyckman, C. 2009. The San Francisco Bay-Delta: A Failure of Decision-making Capacity. *Environmental Science and Policy*, **12**(6), 710–725.

Hansen, L.P., & Sargent, T.J. 2001. Robust Control and Model Uncertainty. *The American Economic Review*, **91**(2), 60–66.

Knight, Frank H. 1921. *Risk, Unvertainty and Profit*. Boston: Houghton Mifflin.

Lempert, R., Nakicenovic, N., Sarewitz, D., & Schlesinger, M. 2004. Characterizing Climate-Change Uncertainties for Decision-Makers. An Editorial Essay. *Climatic Change*, **65**(1), 1–9.

Lempert, R.J. 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.J., Popper, S.W., & Bankes, S.C. 2003. *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-term Policy Analysis*. Rand Corp.

Lempert, R.J., Groves, D.G., Popper, S.W., & Bankes, S.C. 2006. A general, analytic method for generating robust strategies and narrative scenarios. *Management science*, **52**(4), 514.

Lund, Jay, Hanak, Ellen, Fleenor, William, Bennett, William, Howitt, Richard, Mount, Jeffrey, & Moyle, Peter. 2008. *Comparing Futures for the Sacramento-San Joaquin Delta*. San Francisco: Public Policy Institute of California.

Lund, Jay R, Hanak, Ellen, Fleenor, William, Howitt, Richard, Mount, Jeffrey, & Moyle, Peter. 2007. *Envisioning Futures for the Sacramento-San Joaquin Delta*. Public Policy Institute of CA.

Madani, K., & Lund, J.R. 2011a. A Monte-Carlo Game Theoretic approach for Multi-criteria Decision Making under uncertainty. *Advances in Water Resources*, **34**, 607–618.

Madani, K., & Lund, J.R. 2011b. California's Sacramento-San Joaquin Delta Conflict: from Cooperation to Chicken. *mimeo*.

Murphy, J.M., Sexton, D.M.H., Barnett, D.N., Jones, G.S., Webb, M.J., Collins, M., & Stainforth, D.A. 2004. Quantification of modelling uncertainties in a large ensemble of climate change simulations. *Nature*, **430**(7001), 768–772.

National Research Council. 2011. *A Review of the Use of Science and Adaptive Management in California's Draft Bay Delta Conservation Plan*. Washington, D.C.: National Academies Press.

Piani, C., Frame, DJ, Stainforth, DA, & Allen, MR. 2005. Constraints on Climate Change From a

Multi-thousand Member Ensemble of Simulations. *Geophys. res. lett*, **32**, 23.
Rasmusen, E. 2007. *Games and information: An Introduction to Game Theory*. Wiley-Blackwell.
Rougier, J. 2007. Probabilistic inference for future climate using an ensemble of climate model evaluations. *Climatic Change*, **81**(3), 247–264.
Shaw, W. Douglass, & Woodward, Richard T. 2008. Why Environmental and Resource Economists should care about Non-expected Utility Models. *Resource and Energy Economics*, **30**, 66–89.
Sinn, H.W. 1980. A Rehabilitation of the Principle of Insufficient Reason. *The Quarterly Journal of Economics*, **94**(3), 493.
Stainforth, D.A., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, DJ, Kettleborough, JA, Knight, S., Martin, A., Murphy, JM, *et al.* 2005. Uncertainty in Predictions of the Climate Response to Rising Levels of Greenhouse Gases. *Nature*, **433**(7024), 403–406.
United States District Court. 2007 (December). *OCAP Final Interim Relief Court Order*.
Walters, Dan. 2010 (August 11). *Water Bond Violated Unwritten Capitol Rule*. Sacramento Bee.
Woodward, R.T., & Shaw, W.D. 2008. Allocating resources in an uncertain world: Water management and endangered species. *American journal of agricultural economics*, **90**(3), 593–605.

Appendix A. Appendix: Model Details

Since outcomes depend on the actual export regime, we have defined each of the functions below using $\mathbf{g}$ as our argument instead of $\mathbf{x}$. The precise function for $\mathbf{g}(\mathbf{x})$ is described in subsection A.3 below.

A.1. Flow Variables.

We make frequent use of three flow variables that are derived from our 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 canal (χ_{pc}). To calculate the Delta inflow, we also use a constant giving the current inflow from the Sacramento River into the Delta (χ_{in_0}). Because the canal intakes would be upstream of the Delta, water that is exported through a canal would not flow into the Delta while water that is exported through the Delta would. Thus, we have:

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

A.2. Fish Survival Probabilites.

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:

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

where the constants a_i, b_i, c_i, and d_i were chosen to reproduce the survival probabilities in the PPIC report for the four policies they considered and the exponents α_i are allowed to vary across terms. The precise values of the parameters are part the state of the world. Each state of the world corresponds to a particular survival probability within the Lund *et al.* (2008) range for each of their policy options. The coefficients in the survival functions are calibrated independently for each draw from $h(\cdot)$. We interpret the probabilities as representing the stakeholders' a priori beliefs about the distribution of s_i^f.

A.3. Actual Export Regime.

As noted in subsection 3.3.3, there is a divergence between the declared export policy, $\mathbf{x}$, and the actual export regime $\mathbf{g}$. This subsection describes the function $\mathbf{g}(\mathbf{x}; s, \mathbf{z}_y)$ that 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 the impact of each.

A.3.1. *Mistrust.*

As described in subsection 3.4, λ gives 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 $(1-\lambda)\,x_{\text{ex}} + \lambda\chi_{max}$ and the actual share through the canal is given by $(1-\lambda)\,x_{\text{shr}} + \lambda\left(1 - \frac{\chi_{\delta}}{\chi_{max}}\right)$. We denote the resulting vector $\mathbf{g}^M(\mathbf{x};\lambda)$ with the M indicating an adjustment for mistrust.

A.3.2. *Ecosystem Driven Cutbacks.*

In addition to mistrust, the possibility that managers can alter the export policy after observing a signal about fish survival 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 occured. Therefore, in this subsection, we describe 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 drive cutbacks are realized.

We assume that decisions about whether to require cutbacks (i.e. to 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). However, we assume that before making a decision, managers will observe some information (a signal) that will allow them to update their prior belief about the distribution of s_i^f. Specifically, we define a second random variable ω_i for each species. This variable takes on the values $\{\mathsf{good}, \mathsf{bad}\}$ where *good* implies that managers have received positive news about fish populations leading them to believe the probability of survival is now higher than $\bar{fp}_i$ and bad implies that managers have received negative information about fish populations leading them to believe the probability of survival is now lower than $\bar{fp}_i$. Formally, we define the managers' updated probability distribution for s_i^f as $\hat{fp}_i$ and assume that its value is given by:

$$\hat{fp}_i\left(\mathbf{g}^E, \omega_i\right) = \begin{cases} \eta_i + (1-\eta_i)\,\bar{fp}_i\left(\mathbf{g}^E\right) & \text{if } \omega_i = \mathsf{good} \\ (1-\eta_i)\,\bar{fp}_i\left(\mathbf{g}^E\right) & \text{if } \omega_i = \mathsf{bad} \end{cases} \tag{A.1}$$

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 simplicitity, we assume that the distribution of ω_i, conditional on declared policy $\mathbf{x}$, is given by

$$\omega_i\left(\mathbf{g}^M\right) = \begin{cases} \mathsf{good} & \text{with probability } \bar{fp}_i\left(\mathbf{g}^M\right) \\ \mathsf{bad} & \text{with probability } 1 - \bar{fp}_i\left(\mathbf{g}^M\right). \end{cases} \tag{A.2}$$

The formulation embodied in Equations (A.1) and (A.2) guarantees that if managers decide not change the policy in response to the signal, there is consistency between the intial prior and the expected fish survival after observing the signal. In other words, these equations guarantee that:

$$\bar{fp}_i\left(\mathbf{g}^M\right) = Pr\left(\omega_i\left(\mathbf{g}^M\right) = \mathsf{good}\right)\hat{fp}_i\left(\mathbf{g}^M, \mathsf{good}\right) + Pr\left(\omega_i\left(\mathbf{g}^M\right) = \mathsf{bad}\right)\hat{fp}_i\left(\mathbf{g}^M, \mathsf{bad}\right).$$

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 redcuce 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 $R = 1$. We assume that if cutbacks occur, managers reduce exports of all types by a constant proportion μ. Therefore, we have:

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

A.3.3. *Post-Collapse Exports.* The final source of variation between the declared policy and the actual export regime has a different character. 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 follow Lund *et al.* (2008) in hypothesisizing that there is a constant annual probability of major levee collapse (p_{annFail}). This probability is calculated from a cummulative probability of failure over Yrs years given by p_{fail} according to the formula

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

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

We assume that following a major levee collapse, all exports will be shifted to the canal if one exists. If a canal has not been built, the state can either build one after the disaster or repair the Delta levees and continue pumping exclusively through the Delta.[28] We model the decision of whether to build a canal or repair as an element of the future state of the world given by the random variable ξ whose distribution is governed by p_ξ which is the probability a canal is built following a major levee collapse.

[28]A third option would be to cease exports in the wake of a major levee collapse. We agree with Lund *et al.* (2008) that such an outcome is unlikely since it would likely cost the state more than constructing a canal would.

This structure implies that the export regime after a disaster is given by

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

The export regime before a disaster is simply

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

A.4. **Agricultural Employment.** The level of agricultural employment in the San Joaquin Valley depends on total water exports and is given by

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

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-Wanger export levels.

A.5. **Costs.**

Reduced water exports. Lund *et al.* (2008) provide estimates of the total costs to the state of reducing exports from pre-Wanger levels to three levels: no exports, 25% of pre-Wanger levels, and 50% of pre-Wanger levels. Their detailed results in Appendix J divided these costs into specific regions. We used that information to calibrate two functions of the form:

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

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

Water treatment. Water that flows through the Delta must be treated due to high salinity. The total treatment cost is proportional to the amount of water flowing through the Delta and is given by

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

where χ_0 is the pre-Wanger level of exports (all of which flow through the Delta) and C_0^{treat} is the Lund *et al.* (2008) estimate of the costs of treating this much water.

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

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

and the share of treatment costs paid by Ur are given by

$$\zeta_{\mathsf{Ur}}^{\mathrm{treat}}=\left(1-\zeta_{\mathsf{Ag}}^{\mathrm{treat}}\right).$$

Construction Costs. Since we assume that the size of the canal is independent of the planned level of exports through the canal, the annualized cost of constructing a canal is given by

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

$C_0^{\text{construct}}$ is the Lund *et al.* (2008) estimate of the total cost of constructing a canal.[29] These costs are allocated among Ag, Ur, and Tp according to the share vector ζ^{treat}.

Levee maintenance. The cost of maintaing the levees is given by

$$C^{\text{maintain}} = C_0^{\text{maintain}} \left(\frac{\chi_\delta}{\chi_0}\right)^{\vartheta^{\text{maintain}}}$$

where C_0^{maintain} is the estimate in Lund *et al.* (2008) of treatment with pre-Wanger exports all flowing through the Delta.

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

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

These collapse costs consider only the costs that could be avoided if exports were instead routed through a canal. Any costs of collapse that are unaffected by the export regime (e.g. ecosystem impacts or the inundation of Delta islands) are constant throughout our space and thus not included.

Repair following a major collapse. The cost of reparing after a major collapse depends on whether a canal is built or the levees are repaired and through Delta exports continue. The repair cost are thus given by

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

Stakeholder Net Benefits. Three of our stakeholders (Ag, Ur, and Tp) pay all of the costs in the model. For these three stakeholders their net financial benefits are given by

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

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

A.6. **Stakeholder Utility.**

Taxpayers. Taxpayers are assumed to be risk-neutral and have utility of the form

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

This utility function is not subject to modeling uncertainty outside of any effects on **g**.

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

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

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.

[29] We assume that if a canal is constructed following a disaster in the default, the cost is given by $(1 + P_d)\, C_0^{\text{construct}}$.

Agricultural Users. The agricultural interests in our model care about both agricultural employment (y_{employ})and their benefits net of all the costs they pay (B_{Ag}). They have CES utility given by

$$u_{\mathbf{u}(\mathsf{Ag})}(\mathbf{g}) = \left[w^{\text{employ}} y_{\text{employ}}(\mathbf{g})^{\text{eSubExp}_{\mathsf{Ag}}} + \left(1 - w^{\text{employ}}\right) B_{\mathsf{Ag}}(\mathbf{g})^{\text{eSubExp}_{\mathsf{Ag}}}\right]^{\frac{\gamma^{\mathsf{Ag}}}{\text{eSubExp}_{\mathsf{Ag}}}}$$

where $\text{eSubExp} = \frac{\text{eSub}-1}{\text{eSub}}$ 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. The in-Delta interests are concerned with maintenance expenditures (C^{maintain}) and Delta inflow (χ_{in}, a proxy for Delta water quality) and have CES utility. To avoid problems when one of their utility arguments is equal to zero, we add a constant $q_{\text{maintain}}$to the maintenance expenditures before calculating utility. Their utility function is

$$u_{\mathbf{u}(\mathsf{Dt})}(\mathbf{g}) = \left[w^{\text{maintain}}\left(C^{\text{maintain}}(\mathbf{g}) + q_{\text{maintain}}\right)^{\text{eSubExp}_{\mathsf{Dt}}} + \left(1 - w^{\text{maintain}}\right)\chi_{in}(\mathbf{g})^{\text{eSubExp}_{\mathsf{Dt}}}\right]^{\frac{\gamma^{\mathsf{Dt}}}{\text{eSubExp}_{\mathsf{Dt}}}}$$

$\text{eSubExp} = \frac{\text{eSub}-1}{\text{eSub}}$ governs the elasticity of substitution, w^{maintain} 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_{maintain}.

Environmentalists.

$$u_{\mathbf{u}(\mathsf{Ev})}(\mathbf{g}) = \begin{cases} v^{\mathsf{Ev}}_{\text{both}} & \text{if } \hat{f}p_i(\mathbf{g}) = 1 \text{ for } i = \{\text{smelt}, \text{salmon}\} \\ v^{\mathsf{Ev}}_{\text{salmon}} & \text{if } \hat{f}p_{\text{salmon}}(\mathbf{g}) = 1 \text{ and } \hat{f}p_{\text{smelt}}(\mathbf{g}) = 0 \\ v^{\mathsf{Ev}}_{\text{smelt}} & \text{if } \hat{f}p_{\text{smelt}}(\mathbf{g}) = 1 \text{ and } \hat{f}p_{\text{salmon}}(\mathbf{g}) = 0 \\ v^{\mathsf{Ev}}_{\text{none}} & \text{if } \hat{f}p_i(\mathbf{g}) = 0 \text{ for } i = \{\text{smelt}, \text{salmon}\} \end{cases}$$

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. We thus define their utility conditional on the state of the world as

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

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

A.7. **Default.** In both of our simulations, the level of default exports is influenced by both modeling uncertainty and the future state of the world. The distinction between the two experiments is the distribution over possible values. The variable x^d_{ex} gives the best case level of total exports in the default. We introduce a state-contigent random variable ϕ that is distributed uniformly on the interval $(0, 1)$. The best case level of default exports is given by

$$x^d_{\text{ex}} = \phi\left(\overline{x^d_{\text{ex}}} + v_{dx}\right) + \phi\left(\overline{x^d_{\text{ex}}} - v_{dx}\right)$$

In our main simulations, ϕ is state-contigent, $\overline{x^d_{\text{ex}}}$ is a constant and v_{dx} is an element of modeling uncertainty. This formulation implies that x^d_{ex} is distributed uniformly across an interval centered on $\overline{x^d_{\text{ex}}}$ whose width varies with modeling uncertainty. In the auxiliary simulation discussed in subsection 4.3, we assume that

TABLE 3. Elements of modeling uncertainty with mathematical symbols identified

	Variable	Lower Bound	Upper Bound
w^{employ}	Weight on jobs vs money in Ag utility	0.2	0.8
w^{maintain}	Weight on maintenance vs inflows in Dt utility	0.2	0.8
q_{maintain}	Constant in Dt utility	0.05	0.15
$\text{eSub}_{\mathsf{Ag}}$	Ag elasticity of substiution	0.5	1.5
$\text{eSub}_{\mathsf{Dt}}$	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
$v^{\mathsf{Ev}}_{\text{smelt}}$	Ev utility if only smelt survive*	0.25	0.75
$v^{\mathsf{Ev}}_{\text{salmon}}$	Ev utility if only salmon survive*	0.25	0.75
υ_{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.

the best case level of default exports is no longer state contigent. We treat both ϕ and υ_{dx} as elements of modeling uncertainty.

Since a canal will not be built in the default, $x^d_{\text{shr}} = 0$ and mistrust has no impact on default exports. However, the probabilities of both additional ecosystem driven cutbacks and future levee failures remain. This in both of our simulations, the actual default export regime is state contigent. We introduce a distinct variable R^d to indicate whether ecosystem cutbacks occur in the default; the probability that $R^d = 1$ is given by ν_0, which is set equal to 1 in our simulations as shown in Table 5. This implies that default exports after uncertainty about ecosystem cutbacks is resolved are given by

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

The mapping from the actual default export regime to payoff-relevant outcomes is the same as the mapping from the actual export regime induced by an agreement except for two specific differences. First, we assume that in the default $C^{\text{maintain}}\left(\mathbf{g}^d\right) = 0$ regardless of the level of default exports. Second, we assume that if a canal is constructed following a major levee collapse in the default, the canal will be more expensive to build due to pre-existing enmity between the stakeholders. Specifically, we assume that in the default, the construction cost is given by $(1 + P_d)\, C^{\text{construct}}\left(\mathbf{g}^d\right)$.

A.8. **Model Coefficients.** The coefficients of our model that are treated as modeling uncertainty are described in Table 2 in the main text. Table 3 replicates this information, adding an initial column identifying the mathematical symbol used for these coefficients in the appendix. We assume that each of these variables are independent and uniformly distributed on the given interval.

Table 4 lists the primitive variables governing our uncertainty about the future state of the world. Each of these variables is assumed to vary with the state of the world. Again, the variables are assumed to be independent and uniformly distributed on the given intervals.

TABLE 4. State dependent variables and their distributions

Symbol	Variable	Lower Bound	Upper Bound	Source
fp^{Smelt}_{PC}	Smelt surival prob with $\mathbf{x} = (6, 1)$	0.1	0.4	PPIC
fp^{Salmon}_{PC}	Salmon survival prob with $\mathbf{x} = (6, 1)$	0.2	0.5	PPIC
fp^{Smelt}_{Dual}	Smelt survival prob with $\mathbf{x} = (6, 0.5)$	0.1	0.4	PPIC+assumption
fp^{Salmon}_{Dual}	Salmon survival prob with $\mathbf{x} = (6, 0.5)$	0.2	0.5	PPIC+assumption
fp^{Smelt}_{Thru}	Smelt survival prob with $\mathbf{x} = (6, 0)$	0.05	0.3	PPIC
fp^{Salmon}_{Thru}	Salmon survival prob with $\mathbf{x} = (6, 0)$	0.1	0.3	PPIC
fp^{Smelt}_{No}	Smelt survival prob with $\mathbf{x} = (0, \cdot)$	0.3	0.6	PPIC
fp^{Salmon}_{No}	Salmon survival prob with $\mathbf{x} = (0, \cdot)$	0.4	0.8	PPIC
$\zeta^{\text{collapse}}_{\mathsf{Tp}}$	Tp share of collapse costs	0.2	1	Assumption
$\zeta^{\text{repair}}_{\mathsf{Tp}}$	Tp share of maintenance costs	0.5	1	Assumption/SU
$\vartheta^{\text{maintain}}$	Exponent in maintenance cost function	0.5	1.5	Assumption
$\vartheta^{\text{repair}}$	Exponent in repair cost function	0.1	0.2	Assumption
C^{collapse}_0	Total cost of collapse ($ billion)	7.8	15.7	PPIC
C^{repair}_0	Total cost of repairs ($ billion)	0.2	2.5	PPIC
$C^{\text{construct}}_0$	Total cost of canal construction ($ billion)	4.75	9.75	PPIC
C^{treat}_0	Annualized treatment cost for $\chi_\delta = 6$ ($ billion/yr)	0.3	1	PPIC
$\vartheta^{\text{rx}}_{\mathsf{Ag}}$	Exponent in Ag scarcity cost function	−3.62	−1.58	Derived from PPIC
$\vartheta^{\text{rx}}_{\mathsf{Ur}}$	Exponent in Ur scarcity cost function	−6.35	−1.97	Derived from PPIC
μ	Export reduction share if fish don't recover	0.25	0.4	PPIC
p_{fail}	Cummulative failure probability over Yrs	0.34	0.95	PPIC
$C^{\text{nx}}_{\mathsf{Ag}}$	Scarcity cost to ag with no exports ($ billion/yr)	0.49	0.96	Derived from PPIC
$C^{\text{nx}}_{\mathsf{Ur}}$	Scarcity cost to ur with no exports ($ billion/yr)	1.1	1.54	Derived from PPIC
C^{maintain}_0	Maintainence costs for $\mathbf{x} = (6, 0)$ ($ billion/yr)	1	2	PPIC
$\vartheta^{\mathsf{Ag}\text{treat}}$	Exponent in Ag treatment cost share	0.3	0.4	Derived from PPIC
ϕ	Location of x^d_{ex} in its interval	0	1	Bounds assumed
$\hat{p}_{\text{d}}$	% Increase in annual failure probability in default	0	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.)

The functions presented above rely on several variables that are treated as constants within our model. These are listed in Table 5

TABLE 5. Constants in the Model

	Variable	Value	Source
$\zeta^{\text{collapse}}_{\mathsf{Ur}}$	Share of the collapse cost paid by Ur	.2	Assumption
$\zeta^{\text{collapse}}_{\mathsf{Ag}}$	Share of the collapse cost paid by Ag	.2	Assumption
$\zeta^{\text{repair}}_{\mathsf{Tp}}$	Share of the repair cost paid by Tp	1	Assumption
$\zeta^{\text{construct}}_{\mathsf{Ur}}$	Share of the canal construction cost paid by Ur	.45	Assumption
$\zeta^{\text{construct}}_{\mathsf{Ag}}$	Share of the canal construction cost paid by Ag	.45	Assumption
$\zeta^{\text{construct}}_{\mathsf{Tp}}$	Share of the canal construction cost paid by Tp	.1	Assumption
$\zeta^{\text{maintain}}_{\mathsf{Ur}}$	Share of the costs of maintaining the Delta paid by Ur	.1	Assumption
$\zeta^{\text{maintain}}_{\mathsf{Ag}}$	Share of the costs of maintaining the Delta paid by Ag	.1	Assumption
B^0_{Tp}	Base benefit level for Tp ($ billion/yr)	100	2011-12 state budget (immaterial anyway)
B^0_{Ur}	Base benefit level for Ur ($ billion/yr)	5	Assumption
B^0_{Ag}	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	SJV Ag jobs w/no exports (millions)	.5	PPIC
ε_1	Increase in SJV 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
$\hat{\zeta}^{\text{treat}}_{\mathsf{Ag}}$	Share of treatment cost paid by Ag in base case	.4	Derived from PPIC
$\underline{x_{\text{ex}}}$	Minimum value of total exports in the policy space (maf)	0	Choice
$\overline{x_{\text{ex}}}$	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
$v^{\mathsf{Ev}}_{\text{both}}$	Ev utility if both species survive	1	Immaterial assumption
$v^{\mathsf{Ev}}_{\text{none}}$	Ev utility if neither species survives	0	Immaterial assumption
ν_0	Probability of ESA enforcement in the default	1	Assmption
χ_{max}	Maximum feasible export level (maf)	7.5	Downstream constraints per PPIC
p_ξ	Probability of peripheral canal following a disaster	.5	Assumption
$\overline{x^d_{\text{ex}}}$	Mean of the x^d_{ex} distribution	4	Assumption