Market-Based Mechanisms *

Quitzé Valenzuela-Stookey Francisco Poggi †

October 9, 2020

Click here for the latest version

Abstract

Decision makers frequently condition their actions on economic outcomes, e.g. asset prices, that they believe convey information about an unknown state. However the decision maker's action, or expectations thereof, may also influence the outcome. In this paper we study the general problem of choosing decision rules mapping outcomes to actions in the presence of such feedback effects. We characterize the set of joint distributions of outcomes, actions, and states that can be implemented as the *unique* equilibrium by decision rules which satisfy a minimal notion of robustness to *manipulation*. Moreover, we show that all such equilibria are robust to model misspecification. This characterization of the feasible set greatly simplifies the problem of choosing decision rules. A simple graphical technique allows us to identify qualitative features of optimal policies. We illustrate the power of this approach with an application to corporate bailouts. The results are also useful for characterizing optimal decision rules when the requirement of unique implementation is relaxed.

1 Introduction

This paper is motivated by two observations. First, many economic agents base their decisions on aggregate outcomes such as prices, the unemployment rate, or the rate of inflation. Second, the aggregate outcomes used in decision making may in turn be affected by the actions taken,

*We are grateful to Marciano Siniscalchi, Jeff Ely, Alessandro Pavan, Eddie Dekel, Piotr Dworczak, and Ludvig Sinander for discussion and feedback, and to seminar participants at Northwestern University for helpful comments.

†Department of Economics, Northwestern University.

or expectations thereof. Together, these two conditions give rise to a feedback loop: outcomes affect actions, expectations of which in turn affect outcomes. We are interested in the use of aggregate outcomes to make decisions in the presence of such feedback effects. We refer to these as market-based decision problems.

Decision making in the presence of feedback effects appears in wide range of economic environments. Central banks often adjust the nominal interest rate in response to changes in inflation rate, or expert forecasts thereof, and a large literature exists discussing the efficacy of such policies given the feedback loop that they create.[1] Governments frequently decide on fiscal policy based on macroeconomic indicators, which are impacted by expectations of policy decisions. The debate in the U.S. surrounding the Paycheck Protection Program (PPP) serves as a recent example. Members of Congress cited rising unemployment as evidence of the need for robust interventions, while business based their personnel decisions in part on expectations of future PPP support.[2] Feedback effects are also present in corporate decision making and financial markets. Inclusion in the S&P 500 is based in part on market capitalization, which creates feedback between share prices and the inclusion decision (the decision maker here being the manager of the index). Part of the recent rise in the share price of Tesla may have been driven by speculation that the company would soon be included in the S&P 500, which would trigger a jump in demand from passive funds that track the index.[3] The fact that Tesla's shares fell 21% following its omission from the index supports this hypothesis. Alternatively, the managers of a publicly traded firm may abandon a new project if the share price drops following its announcement. In each of these situations, the action taken, or the expectation of which action will be taken, will in turn affect the variable on which decision making is conditioned.

In our model, a principal commits to a decision rule mapping some outcome, which we will refer to here as the price, to actions. There is a payoff relevant state of nature which is unknown to the principal. The price is the equilibrium outcome of some underlying price-formation game played by the agents. Our model accommodates a broad class of price-formation games which share the feature that the equilibrium price can be summarized by a function of agents' expectations of the principal's action and the state. We require, moreover, that agents' beliefs about the principal's action be consistent with the principal's announced decision rule, given the realized equilibrium price. This gives rise to the feedback effect.

The principal's payoff depends on the joint distribution of the action, state, and price. Therefore, rather than directly studying the choice of decision rules, we instead focus on the

[1] See for example Bernanke and Woodford (1997).

[2] See Thomas and Cutter (2020) and Hughes and Morath (2020) for contemporaneous media coverage.

[3] Stevens (2020).

induced mappings from states to actions and prices, referred to as action and price functions respectively. We then ask which action and price functions are *implementable*, i.e. can be induced as equilibrium outcomes by some decision rule. An action and price function pair will be *virtually implementable* if there are implementable functions arbitrarily nearby. Focusing on implementability, rather than on the decision rule directly, greatly simplifies the study of optimal policies.

This problem can also be viewed as an instance of constrained mechanism design. The principal is limited in that they can respond to the agent's actions/messages only through the market outcome (or price) which aggregates them. Our problem is further distinguished from other multi-agent mechanism design problems by two key features. First, we generally (although not always) think of the market as being large. When there are many agents in the market, their incentive compatibility constraints are relaxed; each agent's actions can have at most a small effect on the price. This may enable the principal to elicit information even in environments in which the usual single-crossing conditions on agent payoffs are violated. Second, there is a feedback effect. The price is determined by the state and agents' beliefs about the principal's action. We study problems in which each agent's beliefs about the principal's action are also shaped by the price. The feedback effect limits what the principal can achieve.

Practical concerns often constrain which decision rules the principal can use. We investigate three prominent concerns which are present in most environments. First, a salient feature of nearly all market-based decision problems is the potential for manipulation; interested agents may try to influence the principal's decision by manipulating the price. This is particularly easy to do when the principal's decision rule is a discontinuous function of the price. We say that a decision rule mapping prices to actions is *robust to manipulation* if it is continuous. This should be interpreted as a minimal criterion; it is necessary to prevent manipulation, but may not be sufficient depending on the application.

Second, the principal may be concerned about indeterminacy of outcomes when there are multiple equilibria under the announced decision rule. Non-fundamental volatility is well documented in environments such as asset markets where expectations play an important role in determining outcomes. There is therefore great interest in designing policies for which a unique equilibrium outcome exists, as discussed in Woodford (1994). This is especially true in problems, such as managing inflation, in which stability is a paramount concern. Moreover, conditioning policy decisions on prices often exacerbates equilibrium multiplicity Bernanke and Woodford (1997). We say that a decision rule is *robust to multiplicity* if it induces a unique outcome in almost all states.

Finally, the principal in general has limited information about the fundamentals of the economy. In particular, the precise map from states and expected actions to prices may be unknown. It is therefore desirable for the principal to use a decision rule that is robust to such uncertainty; small perturbations to the fundamentals should not lead to drastic changes in the joint distribution of states, prices, and actions. When a decision rule induces a map from fundamentals to outcomes that is suitably continuous, we say that the decision rule is *robust to structural uncertainty.*

We say that an action and price function pair is (*virtually*) *robustly implementable* if it is (virtually) implementable by a decision rule that is robust to manipulation and multiplicity. Our first major contribution is a characterization of the set of robustly implementable action functions. We show first that all decision rules that are robust to manipulation and multiplicity induce a price that is monotone in the state. The basic intuition for this property is the following. First, for a price function to be implementable and non-monotone it must be discontinuous; it cannot be that different actions are taken in different states yet induce the same price, since the decision rule is measurable with respect to the price. However, when there are discontinuities and non-monotonicities in the price functions there will be multiple equilibria. This is due to the fact that while the price function may be discontinuous, the decision rule must be continuous. Continuously "bridging the gaps" where the price function is discontinuous creates multiplicity.

The monotonicity of the price function, moreover, essentially characterizes robustly implementable action functions. Under some conditions, an action function is robustly implementable if and only if the associated price function is monotone. More generally, monotonicity plus another easily interpretable condition characterizes robust implementability. Additionally, we show that any decision rule that is robust to manipulation and multiplicity is robust to structural uncertainty. The characterizations of robustly implementable price functions greatly simplifies the problem of finding optimal decision rules.

When non-fundamental volatility of market outcomes is not a primary concern, the principal may be willing to tolerate indeterminacy of equilibrium, provided all equilibria give the principal a high payoff. We therefore consider decision making when the requirement of unique implementation is dropped. Nonetheless, we show that *any* continuous decision rule will induce at least one equilibrium that could in fact be implemented uniquely by an appropriate modification of the decision rule. This result has a number of important implications. First, if the principal takes a strict worst-case view of multiplicity, evaluation decision rules based only on the worst equilibrium that they could induce, then it is without loss to restrict attention to uniquely implementable outcomes. More generally, consider a principal who takes a lexicographic approach to multiple equilibria: the principal first evaluates a decision rule according to the worst equilibrium that it induces. Among those decision rules with the same worst-case equilibrium payoff, the principal chooses based on the best equilibrium (or some other function of the remaining equilibria).[4] The result mentioned earlier in this paragraph implies that the best "worst-case guarantee" can be found by optimizing over the set of virtually robustly implementable mechanisms. Suppose there is a unique decision rule M that achieves this worst case guarantee. Then the need to satisfy the worst-case guarantee implies that the principal must use a decision rule that coincides with M for all prices that can arise in the equilibrium under M. This observation can greatly simplify the problem of solving for optimal decision rules when the principal takes a lexicographic approach to multiplicity.

We illustrate the power of our results with an application to government bailouts of a company or industry. We show under which qualitative features of the environment the government's first best policy will be robustly implementable. We characterize the optimal robust policy when first-best is not feasible. We also show which policies will be optimal when the uniqueness requirement is relaxed.

This paper is part of a large literature related to the two-way feedback between financial markets and the real economy, beginning with Baumol (1965). For a survey of this literature see Bond et al. (2012). Among other contributions, this literature identified multiplicity of equilibria as a fundamental feature of feedback environments. Multiplicity is discussed in Dow and Gorton (1997), Bernanke and Woodford (1997), and Angeletos and Werning (2006). The current paper contributes to this literature by characterizing the set of policies under which multiplicity arises.

More specifically, this paper relates to decision making under commitment in the presence of two-way feedback. Important contributions include Bernanke and Woodford (1997), Bond et al. (2010), and Bond and Goldstein (2015). Bernanke and Woodford (1997) shows how the use of inflation forecasts to inform monetary policy can reduce the informativeness of forecasts. In the language of our paper, this occurs when induced market-outcome function (in this case the inflation forecast) violates the necessary monotonicity condition. Bernanke and Woodford (1997) restrict attention to linear decision rules, and show that equilibrium multiplicity can arise. Our analysis show that non-monotone decision rules may in fact be *necessary* to prevent multiplicity. Bond and Goldstein (2015) focuses on the how market-based interventions affect the efficiency of information aggregation by prices, when there is

[4]Such preferences are similar in spirit to these studied in the context of robust mechanism design (Börgers, 2017) and information design (Dworczak and Pavan, 2020).

aggregate uncertainty in the market. Unlike in our baseline model, in which there is no aggregate uncertainty, the price in Bond and Goldstein (2015) is a stochastic function of the state and expected government action. Nonetheless, our results on robustness to structural uncertainty show that robustly implementable policies will perform well even in the presence of aggregate uncertainty, provided the degree of uncertainty is not too large.

Other papers have noted that decisions based on market outcomes may be vulnerable to manipulation. Goldstein and Guembel (2008) studies manipulation by strategic traders when firms use share prices in secondary financial markets to guide investment decisions. In Lee (2019) a regulator uses stock-price movements of affected firms to determine whether or not to move forward with new regulation. In this model the discontinuous nature of the policy considered opens the door to manipulation. Relative to these papers, which explicitly model manipulation in specific settings, we study only a general necessary condition for preventing manipulation: continuity of the decision rule. Our characterization results, however, provide an important starting place for more application-specific models of manipulation.

In general, the current paper makes four major contributions to the existing literature. First, we provide a general framework for studying market-based mechanism design in this context. By focusing on implementable price and action functions, rather than directly on the decision rule, we are able to shed new light on the general structure of the problem. Second, and most importantly, we characterize the feasible set of implementable outcomes, accounting for the issues of manipulation and multiplicity that appear throughout the literature. This characterization greatly simplifies the analysis of optimal policy in applications. To our knowledge, we are the first to provide a tight characterization of uniquely implementable outcomes. Third, we show that robustness to manipulation and multiplicity implies robustness to structural uncertainty. This means the decision maker's payoff will not be overly sensitive to their limited understanding of market fundamentals. Robustness to structural uncertainty also allows us to apply our results to settings in which there is aggregate uncertainty, provided the degree of uncertainty is not too large. The latter is an important concern from a practical perspective. Finally, our results also allow us to analyse optimal policy when the principal is willing to tolerate multiplicity. As the issue of equilibrium indeterminacy is central to the literature, as noted by Woodford (1994) among others, we see this as an important new development.

The remainder of the paper is organized as follows. Section 2 introduces the model, and discusses the various robustness notions considered. Section 3 presents the main characterization results. Section 4 discusses optimal policy when the unique implementation restriction is relaxed. Section 5 explores the application to bailouts.

2 The model

The model consists of the following primitive objects.

i. The state space, denoted by Θ, which is a closed interval in $\mathbb{R}$.

ii. A convex set $\mathcal{A}$ of principal actions, which is a subset of a Banach space.

iii. A convex set $\mathcal{P} \subseteq \mathbb{R}$ of aggregate outcomes.

For clarity, we will refer to the aggregate outcome as the price, although the model applies to many situations in which the aggregate outcome is not a price, as will be discussed below. There are three periods; $0, 1, 2$. The timing of interaction is as follows.

0. The principal publicly commits to a decision rule $M : \mathcal{P} \mapsto \mathcal{A}$ specifying an action for each price.

1. The price is determined.

2. If the price is p, the principal takes the action $M(p)$.

It only remains to describe how the price is determined in period 1. This is covered in the following section.

2.1 Price formation

We first present a general reduced form representation of price formation. This will form the basis for most of the analysis of the paper. The reduced form can be thought of as summarizing the equilibrium outcomes of some game through which the price is determined. It is general enough to capture many of the leading applications of decision making under feedback. We discuss various micro-foundations for our reduced form approach.

2.1.1 Price formation: reduced form

For the reduced form, we assume that the price is a function of the state and the anticipated action of the principal. Thus the price is given by a function $R : \mathcal{A} \times \Theta \mapsto \mathcal{P}$. The precise interpretation of R is that if the state is θ and all agents believe that the principal will take action a then the price will be $R(a, \theta)$. Throughout, we maintain the assumption that R is continuous.

The defining feature of environments with feedback is that if the principal has announced decision rule M and the price is p then all agents anticipate that the principal will take action

$M(p)$. This situation arises naturally in many applications, and will be discussed further in the micro-foundations. The meaning of the function R will be further clarified when we discuss equilibrium. Before turning to equilibrium analysis we will discuss how equilibrium outcomes in various games can be summarised by a function $R : \mathcal{A} \times \Theta \mapsto \mathcal{P}$.

2.1.2 Price formation: micro-foundations

A formal treatment of various micro-foundations is presented in Appendix C. Here we will simply discuss informally two micro-foundations.

Asset market. This is one of the leading examples of decision making under feedback. Consider an environment in which there is fixed supply of a single asset and a continuum of traders. The asset pays a dividend that is a function of the state and the principal's action. Each trader receives a private signal that is partially informative about the state. Traders base their demand on a) the market price, b) the anticipated action of the principal, and c) their belief about the state. The latter is a function of both their private signal and information conveyed by the asset price. Since the principal's action is a function of the price, there is no ambiguity about the action given the observed price. A rational expectations equilibrium (REE) in this environment consists of a price function $\tilde{P} : \Theta \mapsto \mathcal{P}$ such that markets clear in each state θ given

- The anticipated action $M(\tilde{P}(\theta))$,
- The inferences made from the price given the function $\tilde{P}$.

In Appendix C we show that in such an environment, under some assumptions on information and payoffs, there exists a function $R : \mathcal{A} \times \Theta \mapsto \mathcal{P}$ that gives the REE clearing price given any decision rule M.

Expert forecasts.

In many situations agents may not observe the aggregate outcome when making the decisions that will, taken together, determine the aggregate outcome. For example, the unemployment rate in a given month is the result of the decisions of firms and workers who act without observing the realized unemployment rate. If, in such a situation, the principal makes a decision that is relevant for agents, based on the aggregate outcome then agents will need to predict the action that the principal will take. In many such settings, expert forecasts play an important role in agent decision making.

Suppose an economist receives a signal $\theta \in \Theta$ about the underlying state of the economy ω, and reports publicly their expectation $\hat{p}$ of the unemployment rate p. At the end of the month, the government observes p and chooses $a \in \mathcal{A}$ according to $M(p)$. The action here could be, for example, the amount of money to put into an employment subsidy program. The realized unemployment rate will depend on firm's expectations about a and the underlying state θ. Assume that firms trust the economist's forecast; they take it as an accurate prediction of the unemployment rate. Firm's them make their personnel decisions. The realized unemployment rate will given by $J(M(\hat{p}), \theta)$.

The economist recognizes the effect that their forecast has on firm behavior, and thus on the realized unemployment rate. The economist will take this into account when making their prediction. Thus their expectation of the unemployment rate will be given by

$$\hat{p} = \mathbb{E}[J(M(\hat{p}), \omega)|\theta] \equiv R(M(\hat{p}), \theta).$$

Such a fixed point exists when $\mathcal{A}$ is compact and M continuous. Note that R here is a function of the economist's signal, rather than the underlying state.

2.2 Implementation

We have not yet discussed the preferences of the principal. We will not make assumptions on these, other than that they do not depend directly on the announced decision rule M. Rather, the principal cares only about the joint distribution of states, actions, and prices. In other words, the principal cares about the equilibrium maps from states to actions and prices induced by their announced decision rule. A rational expectations equilibrium (REE) in this context consists of a price function $P : \Theta \mapsto \mathcal{P}$ such that $P(\theta) = R(M \circ P(\theta), \theta)$ for all $\theta \in \Theta$.

Let $Q : \Theta \mapsto \mathcal{A}$ be an *action function* and $P : \Theta \mapsto \mathcal{P}$ a *price function.* One question of interest is which pairs of price and action functions can be implemented as a rational expectations equilibrium.[5]

Definition. (Q, P) *is* ***implementable*** *if there exists* $M : \mathcal{P} \mapsto \mathcal{A}$ *such that*

1. $P(\theta) = R(M \circ P(\theta), \theta) \quad \forall\, \theta \in \Theta$ *(RE)*
2. $Q = M \circ P$. *(commitment)*

The RE (rational expectations) condition requires that the realized price be consistent with the anticipated action given decision rule M. The commitment condition simply says that the principal is in fact using decision rule M. Implementability can be equivalently defined without making explicit reference to the implementing decision rule M.

[5]If $\mathcal{A}$ is a compact convex subset of a Euclidean space then a rational expectations equilibrium will exist for any continuous M (see Lemma 8). Continuity of M is discussed in the next section.

Observation 1. (Q, P) *is implementable iff*

1. $P(\theta) = R(Q(\theta), \theta) \quad \forall\, \theta \in \Theta$ *(RE)*

2. $Q(\theta) \neq Q(\theta') \qquad \Rightarrow \qquad P(\theta) \neq P(\theta').$ *(measurability)*

Here the measurability condition guarantees that there exists a P measurable function M that induces action function Q. Clearly if this condition is violated there can exist no such M.

Observation 1 gives a characterization of the set of implementable (Q, P). However it is not, on its own, a very useful characterization for two reasons. First, it does not point to any general qualitative features of implementable mechanisms. Second, it ignores many important practical concerns that the principal may consider when choosing a decision rule. These practical concerns are the subject of the next section. It turns out that when these are taken into account a more meaningful characterization of the set of implementable mechanisms can be given.

2.3 Robustness

A number of practical concerns naturally arise when contemplating the use of market based decision rules. We focus on two factors to which the principal may wish their mechanism to be robust.

Manipulation. In most applications, a principal using a market-based decision rule should be concerned that interested parties may manipulate the price in order to influence the principal's action. For example, a judge conditioning their decision in a merger case on the stock price of a competitor should be concerned that both the competitor and the merging firms have incentives to manipulate the stock price. An agent may manipulate the asset price by buying/selling the asset, releasing false information, or other means.[6]

There are many potential models of manipulation, depending on the specific application considered. Rather than attempt to encompass all possible models of manipulation, we will instead propose a minimal definition of robustness.

Definition. *Say that a decision rule M is* ***robust to manipulation*** *if it is continuous.*

Clearly a discontinuous decision rule will be vulnerable to manipulation; an agent can induce a significant change in the principal's action by making an arbitrarily small change to the price. As a result the principal should, at a minimum, use a continuous decision rule

[6] Goldstein and Guembel (2008) discusses manipulation of this sort.

when there are concerns about manipulation. Additional restrictions may be required in specific settings, and will imply refinements of the set of decision rules that are robust to manipulation.

Continuity of the decision rule does capture a limiting notion of robustness to manipulation when moving the price is very costly. Assume that the cost of price manipulation is proportional to the magnitude of the induced change. For fixed costs of price manipulation and fixed private benefits, manipulation will not be profitable iff the decision rule is Lipschitz continuous, where the Lipshitz constant is determined by the costs and benefits of manipulation. In the limit we simply require continuity. Alternatively, one could define robustness to manipulation as Lipschitz continuity. This would not substantively change the analysis.

Multiple equilibria. The approach to multiple equilibria depends on the type of analysis being conducted. From an implementation perspective, the question is how to induce a given (Q, P) as equilibrium outcomes. In the implementation literature, this means that (Q, P) should be the unique equilibrium outcomes induced by some decision rule. The mechanism design perspective, on the other hand, is that the principal can choose from any of the equilibria induced by a given decision rule M. From this perspective, the goal of the principal is simply to induce (Q, P) as *an* equilibrium outcome.

We will consider both perspectives in this paper. To begin, we will take the implementation perspective that implementation must be done uniquely. We will then show how the results obtained can be related to a more permissive attitude towards multiplicity when the principal chooses which (Q, P) to implement.

Definition. *A decision rule M is* ***robust to multiplicity*** *if $\{p \in \mathcal{P} : p = R(M(p), \theta)\}$ is singleton for almost all θ.*

The requirement that there is a unique rational expectations price for almost all, rather than all, states is not important. Requiring uniqueness everywhere would not substantively change the results. This definition of robustness makes most sense when the principal maximizes expected utility and has an absolutely continuous prior H. If instead H has atoms then uniqueness should hold almost everywhere under H. There is not difficulty in accommodating this modification, although it requires rewording some of the results.

2.4 Robust implementation

Robust implementation accounts for manipulation and uniqueness concerns. We will refer to a decision rule that is robust to manipulation and structural uncertainty as *robust.*

Definition. *(Q, P) is* ***robustly implementable*** *if it is implementable by an M that is robust to manipulation and multiplicity.*

In other words, (Q, P) is *robustly implementable* if there exists $M : \mathbb{R} \mapsto \mathcal{A}$ such that:

1. $Q = M \circ P$

2. $P(\theta)$ is the solution to
$$p = R(M(p), \theta)$$
for all θ, and the unique solution for almost all θ.

3.
$$Q(\theta) \neq Q(\theta') \qquad \Rightarrow \qquad P(\theta) \neq P(\theta')$$

4. M is continuous.

There are two differences between implementability and robust implementability; the uniqueness requirement in condition 2 and the continuity requirement of condition 4. Continuity, as discussed before, is or definition of robustness to manipulation. Uniqueness is required almost everywhere. Nothing substantive changes if we require uniqueness everywhere. We will discuss non-uniqueness on zero measure sets when we give the characterization of robustly implementable mechanisms.

We will sometimes refer to an action function Q as robustly implementable, by which we mean that there exists a P such that the pair (Q, P) is robustly implementable, in similarly for price functions P.

At times, it will be convenient to discuss approximate, rather than exact, implementation. As is standard, we say that (P, Q) is virtually implementable if it can be approximated arbitrarily well by some implementable $(\hat{P}, \hat{Q})$. Say that Q' is an ε-approximation of Q if the set $\{\theta : Q(\theta) \neq Q'(\theta)\}$ has measure less than ε.

Definition. *(P, Q) is* ***virtually robustly implementable*** *if for any $\varepsilon > 0$ there exists an ε-approximation of Q that is robustly implementable.*

The characterization of robustly implementable (and virtually robustly implementable) outcomes will be one of the main results of this paper. It turns out that this characterization is central to understanding optimal decision rules even when the uniqueness requirement is relaxed.

3 Main results

We turn now to the main results of the paper. Some preliminary definitions and results are first needed. The following assumption on price formation will be maintained for most results.

Definition. *R is **weakly increasing in** θ if $\theta \mapsto R(a,\theta)$ is weakly increasing for all $a \in \mathcal{A}$*

We say that R is **strictly increasing** in θ if $\theta \mapsto R(a,\theta)$ is strictly increasing for all $a \in \mathcal{A}$. Note that the order used on Θ is irrelevant. All results that assume that R is weakly increasing continue to hold under the weaker assumption that R is weakly increasing in θ; $R(a,\theta'') \geq R(a,\theta')$ implies $R(a',\theta'') \geq R(a',\theta'')$ for all a', a''. Similarly results that assume that R is strictly increasing continue to hold as long as there exists some order on Θ such that $\theta \mapsto R(a,\theta)$ is strictly increasing for all a. Both strictly and weakly increasing R can be justified by natural assumptions on primitives in many micro-foundations, and is satisfied in all applications that we have come across.

For any decision rule M, define $\theta_M(p) = \{\theta \in \Theta : R(M(p),\theta) = p\}$. When R is strictly increasing in θ, $\theta_M(p)$ will be a function from $\mathcal{P} \mapsto \Theta \cup \varnothing$. When R is only weakly increasing $\theta_M(p)$ may be set valued.

The defining feature of robustly implementable outcomes is a monotone price. This condition is necessary under weakly increasing R.

Theorem 1. *Assume R is weakly increasing in θ. If M is robust to manipulation and multiplicity then it induces a monotone price function.*

Proof. In Appendix A.1 □

In other words, Proposition 1 says that if M is continuous and it induces a price function P that is non-monotone then there will be multiple equilibria. Monotonicity essentially characterizes robustly implementable outcomes, as the following two sections will show. Two features of Theorem 1 are worth emphasising. First, the induced equilibrium price function P need not be increasing; it may be monotonically decreasing. Second, monotonicity of the induced P is not simply a consequence of measurability. This would be the case if we required Q to be continuous. However Q need not be continuous. The assumption that M is continuous in no way implies continuity of Q.

3.1 Implementable action functions

In most situations the principal cares about the actions that they take. They may also care about the price, but since the price is determined in equilibrium by the action function it is sufficient to look only at actions.[7] In such cases the relevant question is what action functions Q are robustly implementable.

To understand the sufficient conditions for robust implementability, assume first that $\theta \mapsto R(a, \theta)$ is strictly monotone for all a. If Q is a continuous action function and $\theta \mapsto R(Q(\theta), \theta)$ is strictly monotone then clearly Q will be robustly implementable: define $M(p)$ as the unique function satisfying $M(R(Q(\theta), \theta)) = Q(\theta)$. This is well defined when $\theta \mapsto R(Q(\theta), \theta)$ is strictly monotone. M is continuous since Q is continuous, and induces a unique equilibrium since $|\{\theta \in \Theta : R(a, \theta) = p\}| \leq 1$ for all p under strict monotonicity of $\theta \mapsto R(a, \theta)$. This shows that strict monotonicity of the induced price function and continuity of Q are sufficient for robust implementation. Both can be relaxed only slightly to yield the necessary and sufficient conditions.

First, since we only require uniqueness *almost everywhere* Q can have countably many discontinuities, provided the discontinuities satisfy a certain condition.[8] Roughly, this condition says that Q can be well approximated by a continuous Q'. For any two actions $a', a'' \in \mathcal{A}$, a *path* from a' to a'' is a continuous function $\gamma : [0, 1] \mapsto \mathcal{A}$ such that $\gamma(0) = a', \gamma(1) = a''$. Say that there exists a *monotone path* from a' to a'' at θ if there exists a path γ from a' to a'' such that $x \mapsto R(\gamma(x), \theta)$ is strictly monotone.

Definition. *A discontinuity in Q at θ' is* ***bridgeable*** *if there exists a monotone path from $\lim_{\theta \nearrow \theta'} Q(\theta)$ to $\lim_{\theta \searrow \theta'} Q(\theta)$ at θ'.*

The following are weaker notions of bridgeability, which it will only be necessary to define on the extreme states $\underline{\theta}, \bar{\theta}$.

Definition. *A discontinuity in Q at $\bar{\theta}$ is* ***upper bridgeable*** *if there exists a path γ from $\lim_{\theta \nearrow \bar{\theta}} Q(\theta)$ to $Q(\bar{\theta})$ such that $R(\gamma(x), \bar{\theta}) \leq \max\{\lim_{\theta \nearrow \bar{\theta}} R(Q(\theta), \theta), R(Q(\bar{\theta}), \bar{\theta})\}$ for all $x \in [0, 1]$, with equality iff*

$$\gamma(x) = \underset{a \in \{\lim_{\theta \nearrow \bar{\theta}} Q(\theta), Q(\bar{\theta})\}}{\arg\max} R(a, \bar{\theta}).$$

Definition. *A discontinuity in Q at $\underline{\theta}$ is* ***lower bridgeable*** *if there exists a path γ from $\lim_{\theta \searrow \underline{\theta}} Q(\theta)$ to $Q(\underline{\theta})$ such that $R(\gamma(x), \underline{\theta}) \geq \min\{\lim_{\theta \searrow \underline{\theta}} R(Q(\theta), \theta), R(Q(\underline{\theta}), \underline{\theta})\}$ for all $x \in [0, 1]$, with equality iff*

$$\gamma(x) = \underset{a \in \{\lim_{\theta \searrow \underline{\theta}} Q(\theta), Q(\bar{\theta})\}}{\arg\min} R(a, \underline{\theta}).$$

[7] In Appendix A.7 we explore the case in which the principal only cares about the price function.

[8] If we require uniqueness everywhere then we can show that Q must be continuous. In this case conditions i and ii of Theorem 2 can be dropped: Q is robustly implementable iff it is continuous and $P(\theta)$ is strictly monotone.

Notice that a necessary condition for a discontinuity at θ to be bridgeable (or upper/lower bridgeable) is $\lim_{\theta\nearrow\theta'} R(Q(\theta),\theta) \neq \lim_{\theta\searrow\theta'} R(Q(\theta),\theta)$. We say that the environment is *fully bridgeable* if for every θ, this condition is also sufficient for bridgeability. Finally, say that the environment is *continuously bridgeable* if for any $\theta^* \in \Theta$ there exists $\varepsilon > 0$ such that if a', a'' is bridgeable at θ^* and $R(a'',\theta) \neq R(a',\theta)$ for all $\theta \in [\theta^*, \theta^* + \varepsilon]$ then there exists a sup-norm continuous function $\sigma(\cdot|a', a'') : [\theta^*, \theta^* + \varepsilon] \mapsto \mathcal{A}^{[0,1]}$ such that $\sigma(\theta|a', a'')$ is a monotone path from a' to a'' for all $\theta \in [\theta^*, \theta^* + \varepsilon]$. Say that the environment is *continuously fully bridgeable* if it is full bridgeable and continuously bridgeable. Bridgeability, and the related notions, will be discussed further following the statement of the results.

Theorem 2. *Assume $\theta \mapsto R(a,\theta)$ is strictly increasing for all $a \in \mathcal{A}$. Then Q is robustly implementable iff*

i. $P(\theta) := R(Q(\theta),\theta)$ is strictly monotone.

ii. Any discontinuity in Q on the interior of Θ is bridgeable.

iii. A discontinuity in Q at $\underline{\theta}$ is lower bridgeable, and at $\bar{\theta}$ is upper bridgeable.

Proof. In Appendix A.2. □

If the principal's payoffs are invariant to changes on zero measure sets then condition *iii.* can be ignored for the purposes of choosing optimal policies; we can restrict attention to Q that are continuous at the endpoints. When the environment is fully bridgeable the type of discontinuities in Q that are allowed can be more easily characterized.

Corollary 1. *Assume $\theta \mapsto R(a,\theta)$ is strictly monotone for all $a \in \mathcal{A}$, and the environment is fully bridgeable. Then Q is robustly implementable iff*

i. $P(\theta) := R(Q(\theta),\theta)$ is strictly monotone.

ii. If Q is discontinuous at θ then so is P.

Finally, it will be useful to know when condition *ii* in Corollary 1 is redundant. This will be the case when any discontinuities that violate this condition can be well approximated. Say that Q has a *degenerate discontinuity* at θ if Q is discontinuous at θ and P is not. The environment is *correctable* if for and $\varepsilon > 0$, any strictly monotone Q, and any θ at which Q has a degenerate discontinuity, there exists a monotone Q' that has no degenerate discontinuities in $(\theta - \varepsilon, \theta + \varepsilon)$ and such $Q' = Q$ on $\Theta \setminus (\theta - \varepsilon, \theta + \varepsilon)$. Sufficient conditions for correctability are discussed in Appendix B.

Corollary 2. *Assume $\theta \mapsto R(a, \theta)$ is strictly monotone for all $a \in \mathcal{A}$, and the environment is fully bridgeable and correctable. Then Q is virtually robustly implementable iff*

i. $P(\theta) := R(Q(\theta), \theta)$ is strictly monotone.

ii. The set of states at which Q is discontinuous has zero measure.

Proof. In Appendix A.3. □

Bridgeability of discontinuities and correctability of the environment are less transparent conditions than monotonicity of the price, and so it will be useful to know general conditions under which they satisfied. If $\mathcal{A}$ is a subset of $\mathbb{R}$ then clearly a discontinuity at θ with left limit $\underline{a}$ and right limit $\bar{a}$ is bridgeable iff $a \mapsto R(a, \theta)$ is strictly monotone on $[\min\{\underline{a}, \bar{a}\}, \max\{\underline{a}, \bar{a}\}]$.[9] When the action space is multi-dimensional the condition becomes more difficult to check, but also easier to satisfy. For example, Proposition 13 shows that full bridgeability is satisfied when the action space is multi-dimensional and R satisfies a weak monotonicity condition.

A full discussion of bridgeability, corrrectability, and related notions is contained in Appendix B. This section gives general conditions under which every discontinuity is bridgeable. In most applications encountered in the literature it is easy to verify that the environment is continuously fully bridgeable and correctable. Even when it is not, the states at which these conditions fail are readily identifiable.

An alternative way to understand the conditions of Theorem 2 is in terms of approximations to Q. Say that Q' is a continuous ε-approximation of Q if Q is continuous and $\lambda(\{\theta \in \Theta : Q(\theta) \neq Q'(\theta)\}) < \varepsilon$, where λ is Lebesgue measure.[10]

Proposition 1. *Assume $\theta \mapsto R(a, \theta)$ is strictly increasing for all $a \in \mathcal{A}$ and the environment is continuously fully bridgeable. Then if Q is robustly implementable there exists a continuous ε-approximation Q' that is robustly implementable, for any $\varepsilon > 0$.*

Proof. In Appendix A.4. □

In other words, Proposition 1 says that the space $\mathcal{Q}$ of continuous Q which induce a strictly monotone price is dense in the space of robustly implementable Q. Proposition 1 can help simplify the problem of solving for an optimal policy. Any such $Q \in \mathcal{Q}$ will be robustly implementable, and by Proposition 1 there is no loss of optimality, provided the principal's payoffs are continuous.

[9] This does not mean that $a \mapsto (a, \theta)$ is monotone in the same direction in every state; it could be increasing in some states and decreasing in others.

[10] Since Θ and $\mathcal{A}$ are compact, if there is an ε-approximation for any ε then there is a sequence that approaches Q in the L^p norm, for any $p < \infty$.

Relaxing strict monotonicity to weak monotonicity we can obtain a similar characterization to Theorem 2. It is necessary however to add an additional condition to account for actions for which the induced price is constant over an interval of states. Let $r(a,p) = \{\theta \in \Theta : R(a,\theta) = p\}$. Under strict monotonicity $r(a,p)$ contains at most one state for all $a \in \mathcal{A}, p \in \mathcal{P}$. Under weak monotonicity however $r(a,p)$ may be a non-degenerate interval.

Let $P(\theta) = R(Q(\theta),\theta)$, and suppose $r(Q(\theta'), P(\theta')$ is non-degenerate. If $Q(\theta'') \neq Q(\theta')$ for some $\theta'' \in r(Q(\theta'), P(\theta')$ then clearly there will be multiplicity, since $R(Q(\theta'),\theta')$ is an REE price in state θ''. The only modifications needed to extend Theorem 2 are those that rule out such instances of multiplicity. Say γ is a *proper monotone path* from a' to a'' at θ if it is a monotone path, and moreover $r(\gamma(x), R(\gamma(x),\theta)) = \theta$ for all $x \in [0,1]$. A discontinuity in Q at θ is *properly bridgeable* if there exists a proper monotone path from $\lim_{\theta' \nearrow \theta} Q(\theta)$ to $\lim_{\theta' \searrow \theta} Q(\theta)$ at θ. The environment is fully properly bridgeable if all non-degenerate discontinuities are properly bridgeable. Note that if $\theta \mapsto R(a,\theta)$ is strictly increasing for all a then proper bridgeability is equivalent to bridgeability.

Proposition 2. *Assume $\theta \mapsto R(a,\theta)$ is weakly increasing for all a. Then Q is robustly implementable iff*

i. $P := R(Q(\theta),\theta)$ is weakly monotone.

ii. Any discontinuity in Q on the interior of Θ is properly bridgeable.

iii. A discontinuity in Q at $\underline{\theta}$ is lower bridgeable, and at $\bar{\theta}$ is upper bridgeable.

iv. $Q(\theta) = Q(\theta')$ for all $\theta' \in R(Q(\theta), P(\theta))$ and all θ.

Proof. In Appendix A.6. □

The first condition is necessary by Theorem 1. It is sufficient given the other two conditions. Under monotonicity, condition *iii.* guarantees that the measurability restriction is satisfied, so an implementing M can be found. Condition *ii.* guarantees that a continuous M can be found that implements Q.

3.2 Structural uncertainty

Another practical concern of the principal is that the price may be influenced by uncertain factors other than the state in which the principal is interested. For example, the presence of noise/liquidity traders in an asset market could introduce aggregate uncertainty. As a consequence, the price may not be a deterministic function of the state and anticipated action. Additionally, the principal may simply have limited information about market fundamentals, which within the model translates into uncertainty about the function R. The principal will want to choose a decision rule that is robust to these types of uncertainty.

Endow the space of market clearing functions $R : \mathcal{A} \times \Theta \mapsto \mathbb{R}$ with the sup-norm. Let $\mathcal{C}$ be the set of continuous functions on $\mathcal{A} \times \Theta$. For $S \subseteq \Theta$, define the *S-supnorm* distance between $f : \Theta \mapsto \mathcal{A}$ and $g : \Theta \mapsto \mathcal{A}$ by

$$\sup_{\theta \in S} |f(\theta) - g(\theta)|.$$

For a given decision rule M and market clearing function R, let $\tilde{Q}_R(\theta|M) := \{a \in \mathcal{A} : M(R(a,\theta)) = a\}$. In words, $\tilde{Q}_R(\theta|M)$ is the set of actions that are consistent with rational expectations in state θ. Say that $R \rightrightarrows \tilde{Q}_R$ is *almost continuous* at R if $\forall\ \varepsilon > 0\ \exists\ S \subseteq \Theta$ with $F(S) > 1 - \varepsilon$ such that $R \rightrightarrows \tilde{Q}_R(\theta)$ is upper and lower hemicontinuous at R, with the S-supnorm (meaning when the range of $R \rightrightarrows \tilde{Q}_R(\theta)$ is endowed with the S-supnorm).

Definition. *A decision rule M is **robust to structural uncertainty** if $R \rightrightarrows \tilde{Q}_R$ is almost continuous at R.*

The interpretation of this definition is that the decision rule should induce almost the same joint distribution of states and actions for small perturbations to the market clearing function. This in turn implies that the principal's expected payoff will be continuous in the function R. It turns out that decision rules that are robust to manipulation and multiplicity are robust to structural uncertainty.

Theorem 3. *If a decision rule M is robust to manipulation and multiplicity then it is robust to structural uncertainty.*

Proof. Proof in Appendix A.8.1 □

The important implication of Theorem 3 is that small changes in R lead to small changes in the principal's expected payoff. Formally, for any selection from $\theta \mapsto \tilde{Q}(\theta, R)$, i.e. any function $Q : \Theta \mapsto \mathcal{A}$, such that $Q(\theta) \in \{\tilde{Q}(\theta, R)\}$ for all θ, abuse notation and write $Q \in \tilde{Q}(\cdot, R)$. Let the principal's expected payoff for a Q be given by

$$U(Q) = \int_{\Theta} u(\theta, Q(\theta)) dH(\theta)$$

where $u : \Theta \times \mathcal{A} \mapsto \mathbb{R}$ is continuous and H is absolutely continuous with respect to Lebesgue measure.[11] Let $\mathcal{U}(R) = \{v \in \mathbb{R} : \exists\ Q \in \tilde{Q}(\cdot, R) \text{ with } U(Q) = v\}$ be the set of payoffs consistent with equilibria induced by M, given market clearing function R.

[11] Alternatively, we could dispense with absolute continuity and define robustness to multiplicity in terms of H.

Proposition 3. *Let M be a decision rule that is robust to manipulation and multiplicity for market clearing function R (i.e. M is continuous and $\{p : R(M(p), \theta) = p\}$ is singleton for almost all θ). Then $\mathcal{U}(R)$ is upper and lower hemicontinuous at R on $\mathcal{C}$.*

Proof. in the Appendix A.8.2. □

4 Beyond uniqueness

The following two propositions are extremely useful when relaxing the requirement of unique implementation. They allow us to use the previous characterization the study this new problem.

Proposition 4. *Assume R is weakly increasing in θ. Any continuous decision rule admits an equilibrium with a monotone price function.*

Proof. In Appendix A.7.4 □

Proposition 5. *Assume R is weakly increasing in θ and that the environment is fully properly bridgeable. Then any continuous decision rule admits an equilibrium with outcomes (Q, P) that are robustly implementable.*

Proof. In Appendix A.7.5 □

Consider now a principal who takes a lexicographic approach to multiplicity. The principal first ranks decision rules according to their worst case outcomes. Among those decision rules with the same worst-case payoff, the principal then chooses the one with the highest best-case payoff. By Proposition 4 we know that the highest worst-case guarantee is exactly the payoff of the best robust decision rule. Once this value has been determined, the goal of the principal is to choose the decision rule with the best equilibrium outcome, subject to the worst-case bound.

Assume first that the principal's payoffs do not depend directly on the price; the principal cares only about the joint distribution of states and actions (similar discussion will apply to other preferences). Assume that there is a unique optimal robustly implementable action function Q^* (similar discussion applies to virtual implementation), implemented uniquely by decision rule M^*. If this is the case then, by Proposition 4, the principal needs to choose a decision rule that implements Q^* as one of its equilibrium outcomes. This pins down the decision rule for all prices in the range $\{R(Q^*(\theta), \theta) : \theta \in \Theta\}$; any optimal decision rule must coincide with M^* for such prices. Moreover, Q^* will be an equilibrium outcome of any such decision rule. Thus the problem of a principal who wishes only to satisfy robustness to manipulation can be stated as follows: choose M' subject to

1. $M'(p) = M^*(p) \; \forall \; p \in P^*(\Theta)$,

2. M' is continuous and satisfies measurability.

5 Bailouts

The government is considering a bailout for a publicly traded company, which it considers strategically important.[12] The company's business prospects $\theta \in \Theta$, representing the demand environment, competition, future costs, etc., are unknown. The government chooses a level of support $a \in \mathcal{A} = [0, \bar{a}]$. For each level of support the share price is a strictly increasing function of the state. We make two additional assumptions regarding the share price.[13]

1. The slope of $\theta \mapsto R(a, \theta)$ is decreasing in a.

2. There exists a state θ^* such that $a \mapsto R(a, \theta)$ is strictly increasing for $\theta < \theta^*$ and strictly decreasing for $\theta > \theta^*$.

The first assumption represents the belief on the part of investors that government involvement in the firm will reduce upside when business prospects are good. This could be because the bailout involves the government taking a role in management, for example by gaining seats on the board. An alternative interpretation is that the bailout takes the form of forgivable loans, such that the amount owed is increasing in the state. The second assumption captures the fact that when business prospects are sufficiently bad, the bailout is necessary to sustain the operations of the business. When business prospects are sufficiently good however, the adverse effects of government intervention dominate. These features are derived from the discussion around recent bailouts, for example that of Lufthansa by the German government. [14]

[12] Alternatively, the bailout could be for an entire industry, in which many of the firms are publicly traded.

[13] These assumptions can be directly related to the asset dividends, as discussed in Appendix C.

[14] In the Lufthansa case, one large shareholder, Heinz Hermann Thiele, threatened to veto the proposed bailout, which involved the government taking a 20% stake in the company and receiving seats on the board. Thiele was reportedly concerned that the government stake would make it harder to restructure and cut jobs. On the other hand, supervisory board chairman Karl-Ludwig Kley emphasised Lufthansa's dire prospects: "We don't have any cash left. Without support, we are threatened with insolvency in the coming days." Lufthansa shares rose 20% when Thiele announced that he would support the deal (Wissenbach and Taylor, 2020).

The government does not wish to give any support to the company if the state is below some threshold θ'. In such cases the business is not considered viable, and the government prefers to let it fail. On the other hand, if the state is above some threshold $\theta'' > \theta'$, the government would also like to offer no support. In this case the government believes that the business can survive without intervention. Its payoff $u(a, \theta)$ is therefor decreasing in a. The government would like to intervene when the state is in $[\theta', \theta'']$. In these states the government's payoff $u(a, \theta)$ is increasing in a. The principal maximizes expected utility, and has an absolutely continuous prior H.

Figure 1 illustrates the situation in which $\theta^* \in [\theta', \theta'']$. The blue lines correspond to the price function P^* induced by the first-best action function Q^*. Since the price function is strictly increasing and Q is continuous, the first-best is robustly implementable by Theorem 2.

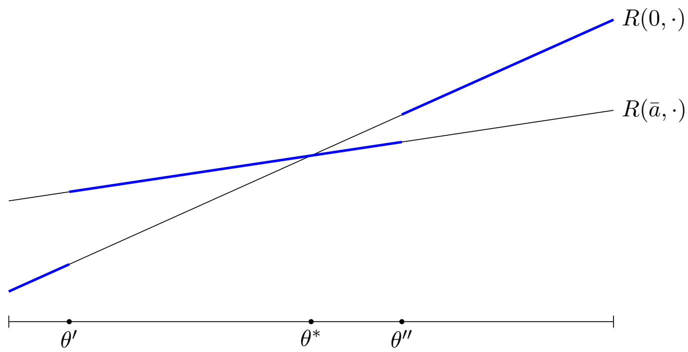

Figure 1: First-best is robustly implementable

A relatively interventionist government is represented by $\theta^* < \theta''$. In this case the government would like to intervene even in states in which investors would prefer no bailout. This will be the case when the strategic importance of the company is high, for example when the company is involved in national security, employs a large number of workers, or engages in production which has large technological spillovers.

Although the first-best is robustly implementable when $\theta^* < \theta''$ must take care in choosing the appropriate implementing decision rule, so as to avoid multiplicity. There are a continuum of decision rules that implement the first-best as an equilibrium outcome. The decision rule for prices in $P^*(\Theta)$ is clearly determined by the desired action function. However the action function alone does not pin down the decision rule for prices in $\tilde{P} \setminus P^*(\Theta)$. Consider the prices in the range $(R(0, \theta'), R(\bar{a}, \theta'))$. For such prices M must satisfy $p = R(M(p), \theta')$. If the government responds to much to price changes in this range, meaning that M increases faster than what this condition implies, then there will equilibria in which action $a > 0$ is taken for states below θ'. Similarly if the government under-responds then there will be equilibria in which action $a < \bar{a}$ is taken for states above θ'. A similar analysis applies to the discontinuity in P^* at θ''.

Suppose instead that $\theta^* > \theta''$. In this case the government is lassiez faire; it does not wish to intervene in states (θ'', θ^*) in which investors would welcome a bailout. The price function associated with the first-best outcome is depicted in Figure 2. In this case the price is non-monotone, and is therefore neither robustly implementable nor virtually robustly implementable. In fact, in this case it is not even implementable, as it violates measurability. The optimal virtually robustly implementable outcome is found by ironing the price function to eliminate non-monotonicity.

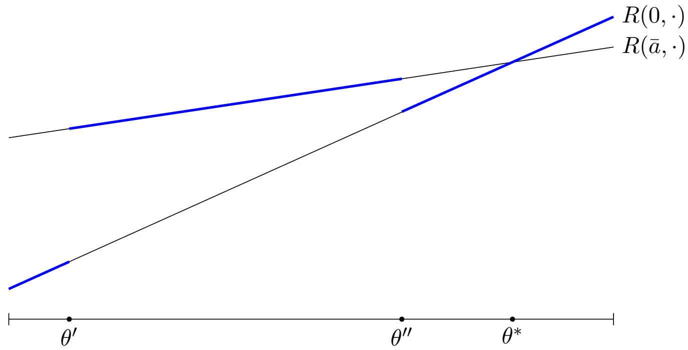

Figure 2: First-best not implementable

The price function for the virtually optimal decision rule is pictured in Figure 3. It is characterized by a state $\hat{\theta}$ at which the ironing begins. For any $\hat{\theta} \in [\theta', \theta'']$ the government's payoff is given by

$$\int_{\underline{\theta}}^{\theta'} u(0,\theta)dH(\theta) + \int_{\theta'}^{\hat{\theta}} u(\bar{a},\theta)dH(\theta) + \int_{\hat{\theta}}^{t(\hat{\theta})} u(\alpha(\theta,\hat{\theta}),\theta)dH(\theta) + \int_{t(\hat{\theta})}^{\bar{\theta}} u(0,\theta)dH(\theta),$$

where $\alpha(\theta, \hat{\theta})$ is defined by $R(\alpha(\theta, \hat{\theta}), \theta) = R(\bar{a}, \hat{\theta})$, and $t(\hat{\theta})$ by $R(0, t(\hat{\theta}) = R(\bar{a}, \hat{\theta})$. Here $\alpha(\theta, \hat{\theta})$ is decreasing in its first argument and increasing in the second, and $t(\hat{\theta})$ is decreasing. Assuming R and u are differentiable, the optimal $\hat{\theta}$ can be identified via the first order condition.

Suppose that in the case of $\theta^* > \theta''$ the government is willing to tolerate some multiplicity, and takes the lexicographic approach described in Section 4. The question is whether or not the government can improve their upside while still guaranteeing the payoff given by the virtually optimal decision rule. Assume for simplicity that there is a unique $\hat{\theta}$ that defines the virtually optimal price function (if there are multiple such $\hat{\theta}$ the same analysis applies to any selection). Then, as shown in Section 4, the virtually optimal decision rule is pined down on

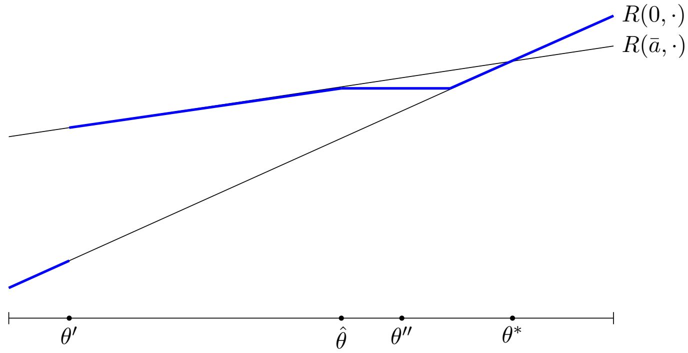

Figure 3: virtually optimal decision rule

$P^*(\Theta)$. The only potential changes that could be made to the decision rule when allowing for multiplicity are on $(R(0,\theta'), R(\bar{a},\theta'))$. It is easy to see from Figure 3 however, that changing the decision rule on this range will can only induce equilibria in which lower actions are taken on (θ',θ'') or higher actions are taken on $[\underline{\theta},\theta')$. Neither of these modifications benefits the principal. Thus relaxing the unique implementation requirement does not change the optimal decision rule.

References

G.-M. Angeletos and I. Werning. Crises and prices: Information aggregation, multiplicity, and volatility. *american economic review*, 96(5):1720–1736, 2006.

S. Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69(4):861–889, 2001.

W. J. Baumol. *The stock market and economic efficiency*. Fordham University Press, 1965.

B. S. Bernanke and M. Woodford. Inflation forecasts and monetary policy. *Journal of Money, Credit, and Banking*, pages 653–684, 1997.

P. Bond and I. Goldstein. Government intervention and information aggregation by prices. *The Journal of Finance*, 70(6):2777–2812, 2015.

P. Bond, I. Goldstein, and E. S. Prescott. Market-based corrective actions. *The Review of Financial Studies*, 23(2):781–820, 2010.

P. Bond, A. Edmans, and I. Goldstein. The real effects of financial markets. *Annu. Rev. Financ. Econ.*, 4(1):339–360, 2012.

T. Börgers. (no) foundations of dominant-strategy mechanisms: a comment on Chung and Ely (2007). *Review of Economic Design*, 21(2):73–82, 2017.

J. Dow and G. Gorton. Stock market efficiency and economic efficiency: is there a connection? *The Journal of Finance*, 52(3):1087–1129, 1997.

P. Dworczak and A. Pavan. Preparing for the worst but hoping for the best: Robust (bayesian) persuasion. mimeo, Northwestern University, 2020.

I. Goldstein and A. Guembel. Manipulation and the allocational role of prices. *The Review of Economic Studies*, 75(1):133–164, 2008.

S. Hughes and E. Morath. Lawmakers seek new ways to cover workers' paychecks during coronavirus. *The Wall Street Journal*, 2020. URL `https://www.wsj.com/articles/lawmakers-seek-new-ways-to-cover-workers-paychecks-during-coronavirus-11586692801`.

Y.-H. A. Lee. A model of stock market-based rulemaking. *Available at SSRN 3440321*, 2019.

P. Stevens. Tesla could soon join the s&p 500 — but inclusion isn't automatic, even with a full year of profitability. *CNBC*, 2020. URL `https://www.cnbc.com/2020/07/21/tesla-isnt-a-gurantee-for-the-sp-500-even-with-year-of-profits.html`.

P. Thomas and C. Cutter. Companies cite new government benefits in cutting workers. *The Wall Street Journal*, 2020. URL `https://www.wsj.com/articles/companies-cite-new-government-benefits-in-cutting-workers-11586264075`.

I. Wissenbach and E. Taylor. Lufthansa soars after top shareholder backs bailout. *Reuters*, 2020. URL `https://www.reuters.com/article/us-health-coronavirus-lufthansa-rescue/lufthansa-soars-after-top-shareholder-backs-bailout-idUSKBN23W0SE`.

M. Woodford. Determinacy of equilibrium under alternative policy regimes. *Economic Theory*, 4(3):323–326, 1994.

Appendix

A Proofs of Section 3

It will be useful to establish some properties of θ. Let $\bar{P}$ be the set of p for which $\theta(p) \neq \varnothing$. Corollary 3 below implies that the set of p for which $\theta(p) = \varnothing$ is open.

Lemma 1. *If M is continuous then $p \mapsto \theta(p)$ is compact-valued, and it is upper hemicontinuous on $\bar{P}$.*

Proof. Compact valued is easy: if $R(M(p), \theta) - p \neq 0$ then by continuity of R this holds for all θ' in a neighborhood of θ.

Now upper hemicontinuity. Let V be an open neighborhood of $\theta(p)$. Then $\Theta \setminus V$ is compact, so there exists $\kappa > 0$ such that $R(M(p), \theta) - p > \kappa$ for all $\theta \in \Theta \setminus V$. Then by continuity of R, M there exists an open neighborhood U of p such that $R(M(p'), \theta) - p' > \kappa$, and thus $\theta(p') \in V$, for all $p' \in U \cap \bar{P}$. Thus $p \mapsto \theta(p)$ is upper hemicontinuous. $\square$

Lemma 2. *If R is weakly increasing in θ then $\theta_M(p)$ is convex valued.*

Proof. $\theta_M(p) = \{\theta \in \Theta : R(M(p), \theta) = p\}$. If $R(M(p), \cdot)$ is monotone, $R(M(p), \theta') = R(M(p), \theta'') = p$ implies $R(M(p), \theta) = p$ for all $\theta \in (\theta', \theta'')$. $\square$

A.1 Proof of Theorem 1

Proof. Assume without loss of generality that that $R(a, \cdot)$ is increasing for all a. Lemmas 5 and 6 and Corollary 3 apply.

Let $\theta_1 < \theta_2 < \theta_3$ be interior, and suppose $P(\theta_1) > P(\theta_2)$ and $P(\theta_3) > P(\theta_2)$ (the other type of non-monotonicity is dealt with symmetrically). We first want to show that $[P(\theta_2), \max\{P(\theta_1), P(\theta_3)\}] \subseteq \bar{P}$. Suppose not, so there is some $p' \in [P(\theta_2), \max\{P(\theta_1), P(\theta_3)\}]$ such that $\theta(p) = \varnothing$. Then p' is either Type L or Type H. Assume it is Type H (symmetric argument using Lemma 5 if it is Type L). First, assume $p' \in [P(\theta_2), \min\{P(\theta_1), P(\theta_3)\}]$. By Lemma 6, part (i) there is a $p > p'$ such that $\theta_2 \in \theta(p)$. Thus there is multiplicity in state θ_2. Then by continuity of M and R, and R weakly increasing in θ, there is multiplicity for all θ in $[\theta_2, \theta_2 + \varepsilon)$ and/or $[\theta_2, \theta_2 - \varepsilon)$ for some $\varepsilon > 0$, violating multiplicity. Thus $[P(\theta_2), \min\{P(\theta_1), P(\theta_3)\}] \subseteq \bar{P}$. Suppose instead that $p' \in [\min\{P(\theta_1), P(\theta_3)\}, \max\{P(\theta_1), P(\theta_3)\}]$. The by Lemma 5, either part (i) or part *(ii)*, there is multiplicity in one of θ_1, θ_3. Then there is multiplicity on a positive measure set, since these are interior.

Assume that $P(\theta_3) \geq P(\theta_1)$ (symmetric argument for reverse inequality). Suppose there exists $\theta' \geq \theta_2$ such that $\theta' \in \theta(P(\theta_1))$. Note that R weakly increasing in θ implies that

$\{\theta \in \Theta : R(a, \theta) = p\}$ is convex for all a, so $\theta_2 \in \theta(P(\theta_1))$. Thus if such a θ' exists there will be multiplicity in state θ_2, and, by the same argument as above, there will be multiplicity for a positive measure of states.

It remains to show that the existence of such a θ' is implied by our assumptions. Suppose instead that $\theta(P(\theta_1)) \subseteq [\underline{\theta}, \theta_2)$. We will show that this implies that there exists $p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, so there is multiplicity in θ_2, and as before this will imply multiplicity for a positive measure of states. Suppose no such p' exists. Let $\tilde{p} = \sup\{p \in [P(\theta_1), P(\theta_3)] : \max \theta(p) < \theta_2\}$, which is well defined by Lemma 1. Since $\theta(\tilde{p})$ is convex, the assumption that no such p' exists implies that either $\max \theta(\tilde{p}) < \theta_2$ or $\min \theta(\tilde{p}) > \theta_2$ Then we have a violation of upper hemicontinuity at $\tilde{p}$. Thus there exists $p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, as desired.

Now, for the extreme states we want to see that the monotonicity is maintained. Let θ be interior, then $p(\bar{\theta}) = R(M(p(\bar{\theta})), \bar{\theta}) > R(M(p(\bar{\theta})), \theta)$ and $R(M(p(\bar{\theta})), \theta) > \underline{p} := \inf\{\bar{P}\}$. So there is an equilibrium price $\tilde{p} \in (p(\bar{\theta}), \underline{p})$. If $p(\bar{\theta}) > \bar{p} := \sup\{\bar{P}\}$, then the previous does not imply multiplicity, but if $p(\bar{\theta}) < \bar{p}$, it does for all θ with associated prices in $(p(\bar{\theta}), \bar{p})$. The same argument holds to prove that $p(\underline{\theta}) \leq \underline{p}$. □

A.2 Theorem 2

Proof. First for necessity. Theorem 1 implies that P must be weakly monotone. If it is not strictly monotone then it will violate measurability, given that $R(a, \cdot)$ is strictly monotone. This proves necessity of *i.*

To show necessity of *ii.*, suppose Q has a discontinuity at an interior state θ' that is not bridgeable. Given that we have established *i*, assume without loss or generality that P is strictly increasing. Under strict monotonicity of $\theta \mapsto R(a, \theta)$, we have $|\theta_M(p)| \leq 1$. Thus Lemma 3 implies that $(\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta)) \subseteq \bar{P}$. If $\theta_M(p) \neq \theta'$ for some $p \in (\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta))$ then there will by multiplicity by Lemma 7. But if the discontinuity at θ' is not bridgeable then there is no continuous M such that the following three conditions hold: *i)* $\lim_{\theta \nearrow \theta'} M(P(\theta)) = \lim_{\theta \nearrow \theta'} Q(\theta)$, *ii)* $\lim_{\theta \searrow \theta'} M(P(\theta)) = \lim_{\theta \searrow \theta'} Q(\theta)$ and, *iii)* $\theta_M(p) = \theta'$ for all $p \in (\lim_{\theta \nearrow \theta'} P(\theta), \lim_{\theta \searrow \theta'} P(\theta))$. To see this, notice that any such M would constitute a monotone path from $\lim_{\theta \nearrow \theta'} Q(\theta)$ to $\lim_{\theta \searrow \theta'} Q(\theta)$.

Now for necessity of *iii.* Continue to assume without loss that P is strictly increasing. We will prove this for $\bar{\theta}$; $\underline{\theta}$ is symmetric. If there is a discontinuity at $\bar{\theta}$ that is not upper bridgeable then there must exist $p \in (\lim_{\theta \nearrow \bar{\theta}} P(\theta), P(\bar{\theta}))$ such that $p < R(M(p), \bar{\theta})$. Let $p_1 = \lim_{\theta \nearrow \bar{\theta}} P(\theta)$. Then $p_1 = R(M(p_1), \bar{\theta}) > R(M(p_1), \underline{\theta})$, by the assumption of strict monotonicity of R. Then there exists $p' \in (p_1, p]$ such that $\theta_M(p') = \theta' < \bar{\theta}$. But then Lemma 7 implies that there is multiplicity on $(\theta', \bar{\theta})$, which is a contradiction.

Now for sufficiency. Assume without loss that P is strictly increasing. Define $M(p) = Q(P^{-1}(p))$, which is well defined on $P(\Theta)$ by *i.* Moreover M is continuous on $P(\Theta)$ since under *ii.* Q is continuous at any interior state θ at which P is continuous. It remains to define M on $\mathcal{P} \setminus P(\Theta)$.

First, suppose Q (and therefore P, under *ii*) has a discontinuity at an interior θ, and let p', p'' and a', a'' be the left and right limits of P and Q respectively. By *ii* there exists a monotone path γ from a' to a''. Then for each $p \in (p', p'')$ there is a unique $x(p) \in [0,1]$ such that $R(\gamma(x(p)), \theta) = p$. Define $M(p) = \gamma(x(p))$ on (p', p'').

Finally for discontinuities at $\bar{\theta}$ (the argument for $\underline{\theta}$ is symmetric). Let γ be a path satisfying the conditions for an upper bridge. Then there exists a continuous and strictly monotone function $b : \gamma([0,1]) \mapsto (p', p'')$ such that $b(\gamma(0)) = p'$, $b(\gamma(1)) = p''$, and $b(a) > R(a, \bar{\theta})$ for all $a \in \gamma([0,1]) \backslash \{a', a''\}$. Then let $M(p) = b^{-1}(p)$. By construction $p > R(M(p), \bar{\theta})$ for all $p \in (p', p'')$, so such p are type H and there is no multiplicity. □

A.3 Corollary 2

Proof. First necessity. If Q is not strictly monotone then for ε small enough there will be no ε-approximation that is strictly monotone. Thus by Theorem 1 there are no robustly implementable ε-approximations. Suppose *ii* is violated. Since P is strictly monotone it can have at most countably many discontinuities. Thus Q must have a positive measure δ of degenerate discontinuities. If $\varepsilon < \delta$ then for any ε-approximation of Q there will be a degenerate discontinuity that is outside of the set of states for which $Q' \neq Q$. But then Q' has a degenerate discontinuity, and so is not robustly implementable.

Now for sufficiency. Given Corollary 1 we need only show that Q can be approximated around all degenerate discontinuities. This follows immediately from the definition of correctable. □

A.4 Proposition 1

Proof. Corollary 1 implies that P is strictly monotone and that whenever Q is discontinuous so is P. Thus P (and Q) can have at most countably many discontinuities. The proposition will follow if we can show that for any $\varepsilon > 0$ and any θ^* at which Q is discontinuous, we can continuously approximate Q around θ^* without changing Q outside of $(\theta^* - \varepsilon, \theta^* + \varepsilon)$.

Since R is continuous and P is discontinuous at θ^*, there exists $\delta < \varepsilon$ and $\theta' \in (\theta^*, \theta^* + \delta)$ such that $Q(\theta^*), Q(\theta')$ and δ satisfy the conditions of continuous bridgeability. Thus there exists a continuous Q' on $[\theta^*, \theta^* + \delta]$ such that $Q'(\theta^*) = \lim_{\theta \nearrow \theta^*} Q(\theta)$, $Q'(\theta^* + \delta) = Q(\theta^* + \delta)$, and $R(Q'(\theta),\theta)$ is strictly increasing on $[\theta^*,\theta^*+\delta]$. Since ε was arbitrary, this gives the desired approximation. $\square$

A.5 Lemma 3

Lemma 3. *Assume R is weakly increasing in θ. For any M that is robust to manipulation and multiplicity, let p_1, p_2 be prices such that $\theta_M(p_1)$ and $\theta_M(p_2)$ are contained in the interior of θ. Then*

$$[\min\{p_1,p_2\},\max\{p_1,p_2\}] \in \bar{P}.$$

Proof. By Theorem 1, the price function P is monotone, so without loss of generality assume that it is increasing, and let $p_2 > p_1$. Assume towards a contradiction that there exists $p \in (p_1,p_2)$ such that $\theta_M(p) = \varnothing$. By Lemma 4 p is either type H or type L. Suppose it is type L, i.e. $R(M(p),\underline{\theta})-p > 0$. Since $\theta_M(p_1) \neq \varnothing$, it must be that $R(M(p_1),\underline{\theta})-p_1 \leq 0$. Moreover, since $\underline{\theta} \notin \theta_M(p_1)$ by assumption, the inequality is strict: $R(M(p_1),\underline{\theta}) - p_1 < 0$. Then by continuity there exists $p' \in (p_1,p)$ such that $R(M(p'),\underline{\theta}) - p' = 0$. Let $\theta_1 = \min\theta_M(p_1)$, which exists by Lemma 1 (by assumption $\theta_1 > \underline{\theta}$). Since P is increasing, $p' > p_1 > P(\theta)$ for all $\theta \in [\underline{\theta},\theta_1)$. Then by Lemma 7 there is multiplicity for all states in $\theta \in [\underline{\theta},\theta_1)$, which is a contradiction. If p is type H then the proof is symmetric, using p_2 rather than p_1. $\square$

A.6 Proof of Proposition 2

Proof. This essentially follows from Theorem 2. The only modifications are the following. Condition *iv* is clearly necessary and sufficient to for there to be no monotonicity involving actions in $Q(\Theta)$. The modification of *ii* from bridgeable to properly bridgeable is necessary and sufficient for there to be no multiplicity involving actions not in $Q(\Theta)$. There is no need to modify condition *iii* since it guarantees existence of that are all type H (at $\bar{\theta}$) or type L (at $\underline{\theta}$), and thus involve no multiplicity.

$\square$

A.7 Implementable price functions

In some cases the principal may not care directly about the actions they take, only about the price that they induce. In this section we ask the following question: for which price functions there exists an action function such that (Q,P) is (robustly) implementable. We call such a P (robustly) implementable.

Definition. *A price function $P : \Theta \to \mathbb{R}$ is in range if for each $\theta \in \Theta$, $P(\theta) \in R(\mathcal{A},\theta)$.*

We say that *identification* holds if for each p, a there is at most one state θ such that $R(a, \theta) = p$.

Proposition 6. *Under identification, a price function is implementable if and only if it is in range and an injection.*

Proof. in the Appendix A.7.1 $\square$

We will call intersection states the ones where there is at least two different actions a_1, a_2 with $R(a_1, \theta) = R(a_2, \theta)$. Let Θ_I be the set of such states. We will make the following extra assumptions on R:

Monotonicity. For any θ, any a'', a' such that $R(a'', \theta) > R(a', \theta)$, and any $p \in [R(a', \theta), R(a'', \theta)]$, there is a unique α such that $R(\alpha a' + (1 - \alpha)a'', \theta) = p$.

Isolated intersections. For every $\theta \in \Theta_I$, there exists an $\epsilon > 0$ such that $B_\epsilon(\theta) \cap \Theta_I = \{\theta\}$.

Intersection smoothness. $R_2(a, \theta)$ exists for every intersection state θ and a that puts weight only on intersecting actions for that state.

Definition. *A price function $P : \Theta \to \mathbb{R}$ satisfies the* ***kink's condition*** *iff there exist C^1 functions $\bar{P}$ and $\underline{P}$ in range and such that $\bar{P}(\theta) \geq P(\theta) \geq \underline{P}(\theta)$.*

The kink's condition effectively means that every kink of P in the upper envelope of $R(a, \theta)$ is concave, and every kink in the lower envelope is convex. Moreover, the kink's condition implies that if there is a θ such that $R(\mathcal{A}, \theta)$ is a singleton, the price function has to be differentiable at θ.

Proposition 7. *Under identification, a price function is robustly implementable if and only if it is in range, strictly monotone, and satisfies the kink's condition.*

Proof. in the Appendix A.7.2. $\square$

There are two primary components of the proof of Proposition 7. The first is that $\bar{P}$ is convex. The second is to show that if $\theta(p)$ is non-monotone then there will be multiplicity. Identification is used to prove both parts of the proposition, but it is not necessary for either. One simple relaxation under which the result is preserved is to allow for actions with constant payoffs.

Weak identification. R is weakly increasing in θ. Moreover, if $R(a, \cdot)$ is not strictly increasing then it is constant.

Proposition 8. *Under weak identification, a pair (Q, P) is implementable if and only if $Q(P^{-1}(p))$ is a singleton for all $p \in P(\Theta)$.*

Proposition 9. *Under R weakly increasing in θ, a price function is robustly implementable if and only if it is in range and weakly monotone and whenever it is flat at a price p, it is so for the whole set $\theta_M(p)$.*

Proof. in the Appendix A.7.3. □

A.7.1 Proof of Proposition 6

Proof. ($\Rightarrow$): suppose not an injection. There are θ and θ' with $P(\theta) = P(\theta')$. By identification, $R(Q(\theta), \theta) \neq R(Q(\theta'), \theta')$, which by rational expectations means that $P(\theta) \neq P(\theta')$, a contradiction. If not in range, then there exist a $\theta \in \Theta$ such that $P(\theta) \notin R(\mathcal{A}, \theta)$, i.e. there is no $a \in \mathcal{A}$ such that $R(a, \theta) = P(\theta)$, so $R(Q(\theta), \theta) \neq P(\theta)$, violating rational expectations.

($\Leftarrow$): Since $P(\theta)$ is in range, for each $\theta \in \Theta$ there exists a a with $R(a, \theta) = P(\theta)$. let's define $Q(\theta)$ by a selection in the rational expectations condition: $R(Q(\theta), \theta) = P(\theta)$. Measurability is satisfied trivially since $P(\theta) \neq P(\theta')$ for all $\theta \neq \theta'$. □

A.7.2 Proof of Proposition 7

Proof. ($\Rightarrow$): Take M that implements P. For all $p \in R(\mathcal{A}, \Theta)$ there is at most a unique $\theta \in \Theta$ that satisfies $R(M(p), \theta) = p$. Otherwise identification would be violated. This defines a function $\theta(p)$.

Let $\bar{P}$ be the set of all prices for which there is an interior solution. We want to show that $\bar{P}$ is convex. Pick $p, p' \in \bar{P}$ and $\alpha \in (0, 1)$ we want to see that $p_\alpha := \alpha p + (1 - \alpha)p' \in \bar{P}$. Let θ and θ' the associated states of p and p'. Continuity of R plus identification imply strict monotonicity of R in θ and for all a. Assume without loss that $\theta' > \theta$.

We will prove that

$$R(M(p_\alpha), \theta) \leq p_\alpha \leq R(M(p_\alpha), \theta') \tag{1}$$

Consider a violation of the second inequality. If $p_\alpha > R(M(p_\alpha), \theta')$ notice that also, $p = R(M(p), \theta) < R(M(p), \theta')$. Therefore, we have

$$p_\alpha - R(M(p_\alpha), \theta') > 0 \qquad \text{and} \qquad p - R(M(p), \theta') < 0$$

By continuity and since $\theta' \in \Theta^o$, there exists an $\bar{\varepsilon} > 0$ such that for all $\tilde{\theta} \in B_{\bar{\varepsilon}}(\theta')$

$$p_\alpha - R(M(p_\alpha), \tilde{\theta}) > 0 \qquad \text{and} \qquad p - R(M(p), \tilde{\theta}) < 0$$

By continuity there is a $p_1 \in (p_\alpha, p)$ with $p_1 - R(M(p_1), \theta') = 0$. But for $0 < \epsilon < \bar{\varepsilon}$ we have $p' = R(M(p'), \theta') < R(M(p'), \theta' - \epsilon)$. There exists a $p_2 \in (p_\alpha, p')$ such that $R(M(p_2), \theta' - \epsilon) = p_2$. $p_1 \neq p_2$, so there is multiplicity in a set of states $[\theta', \theta' + \epsilon)$. With a similar logic we can rule out $p_\alpha \leq R(M(p_\alpha), \theta)$.

Finally, by continuity of π in θ and using Equation (1), there is a $\hat{\theta}$ in (θ, θ') such that $p_\alpha - R(M(p_\alpha), \hat{\theta}) = 0$, therefore $p_\alpha \in \bar{P}$.

The function $\theta(p)$ is continuous in the set $\bar{P}$. However we could have discontinuities for the two prices that are associated with the extreme states $\bar{\theta}$ and $\underline{\theta}$.

We show now that $\theta(p)$ is monotone in $\bar{P}$. Suppose that is not, i.e. there are prices $p_l < p_m < p_h$ such that either $\theta(p_m) < \min\{\theta(p_l), \theta(p_h)\}$ or $\theta(p_m) > \max\{\theta(p_l), \theta(p_h)\}$. Suppose the first (the symmetric argument holds for the other case). Then for all $\theta \in (\theta(p_m), \min\{\theta(p_l), \theta(p_h)\})$ and by continuity there are prices $p_\theta^1 \in (p_l, p_m)$ and $p_\theta^2 \in (p_m, p_h)$ with $\theta(p_\theta^1) = \theta(p_\theta^2)$. This violates multiplicity.

We can invert $\theta(p)$ in $\bar{P}$. The only problem is when $\theta(p)$ is flat, but any selection would give us that the inverse is strictly monotone.

Now, for the extreme states we want to see that the monotonicity is maintained. Let θ be interior, then $p(\bar{\theta}) = R(M(p(\bar{\theta})), \bar{\theta}) > R(M(p(\bar{\theta})), \theta)$ and $R(M(p(\bar{\theta})), \theta) > \underline{p} := \inf\{\bar{P}\}$. So there is an equilibrium price $\tilde{p} \in (p(\bar{\theta}), \underline{p})$. If $p(\bar{\theta}) > \bar{p} := \sup\{\bar{P}\}$, then the previous does not imply multiplicity, but if $p(\bar{\theta}) < \bar{p}$ It does for all θ with associated prices in $(p(\bar{\theta}), \bar{p})$. The same argument holds to prove that $p(\underline{\theta}) \leq \underline{p}$.

($\Leftarrow$): P is strictly monotone and bounded so there is a countable number of discontinuities. Fill those to get a continuous and monotone $\theta(p) := \sup\{\theta : P(\theta) < p\}$.

Let $\bar{M} : P(\Theta) \rightrightarrows \mathcal{A}$ be the set of actions that give price p at the corresponding state i.e. $a \in \bar{M}(p)$ if and only if $R(a, \theta(p)) = p$.

If p is not an intersection price, then $\bar{M}$ is LHC at p. Therefore in a ball around p there is a continuous selection. If p is a interior intersection, then we can consider the set of actions that are not involved in the intersection and select a continuous M. $\square$

A.7.3 Proof of Proposition 9

Proof. Assume without loss of generality that that $R(a, \cdot)$ is increasing for all a. Lemmas 5 and 6 and Corollary 3 apply.

Let $\theta_1 < \theta_2 < \theta_3$ be interior, and suppose $P(\theta_1) > P(\theta_2)$ and $P(\theta_3) > P(\theta_2)$ (the other type of non-monotonicity is dealt with symmetrically). We first want to show that $[P(\theta_2), \max\{P(\theta_1), P(\theta_3)\}] \subseteq \bar{P}$. Suppose not, so there is some $p' \in [P(\theta_2), \max\{P(\theta_1, P(\theta_3)\}]$

such that $\theta(p) = \varnothing$. Then p' is either Type L or Type H. Assume it is Type H (symmetric argument using Lemma 5 if it is Type L). First, assume $p' \in [P(\theta_2), \min\{P(\theta_1, P(\theta_3)\}]$. By Lemma 6, part (i) there is a $p > p'$ such that $\theta_2 \in \theta(p')$. Thus there is multiplicity in state θ_2. Then by continuity of M and R, and R weakly increasing in θ, there is multiplicity for all θ in $[\theta_2, \theta_2 + \varepsilon)$ and/or $[\theta_2, \theta_2 - \varepsilon)$ for some $\varepsilon > 0$, violating multiplicity. Thus $[P(\theta_2), \min\{P(\theta_1, P(\theta_3)\}] \subseteq \bar{P}$. Suppose instead that $p' \in [\min\{P(\theta_1, P(\theta_3), \max\{P(\theta_1, P(\theta_3)\}]$. The by Lemma 5, either part (i) or part *(ii)*, there is multiplicity in one of θ_1, θ_3. Then there is multiplicity on a positive measure set, since these are interior.

Assume that $P(\theta_3) \geq P(\theta_1)$ (symmetric argument for reverse inequality). Suppose there exists $\theta' \geq \theta_2$ such that $\theta' \in \theta(P(\theta_1))$. Note that R weakly increasing in θ implies that $\{\theta \in \Theta : R(a, \theta) = p\}$ is convex for all a, so $\theta_2 \in \theta(P(\theta_1))$. Thus if such a θ' exists there will be multiplicity in state θ_2, and, by the same argument as above, there will be multiplicity for a positive measure of states.

It remains to show that the existence of such a θ' is implied by our assumptions. Suppose instead that $\theta(P(\theta_1)) \subseteq [\underline{\theta}, \theta_2)$. We will show that this implies that there exists $p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, so there is multiplicity in θ_2, and as before this will imply multiplicity for a positive measure of states. Suppose no such p' exists. Let $\tilde{p} = \sup\{p \in [P(\theta_1), P(\theta_3)] : \max \theta(p) < \theta_2\}$, which is well defined by Lemma 1. Since $\theta(\tilde{p})$ is convex, the assumption that no such p' exists implies that either $\max \theta(\tilde{p}) < \theta_2$ or $\min \theta(\tilde{p}) > \theta_2$ Then we have a violation of upper hemicontinuity at $\tilde{p}$. Thus there exists $p > P(\theta_1)$ such that $\theta_2 \in \theta(p')$, as desired.

Now, for the extreme states we want to see that the monotonicity is maintained. Let θ be interior, then $p(\bar{\theta}) = R(M(p(\bar{\theta})), \bar{\theta}) > R(M(p(\bar{\theta})), \theta)$ and $R(M(p(\bar{\theta})), \theta) > \underline{p} := \inf\{\bar{P}\}$. So there is an equilibrium price $\tilde{p} \in (p(\bar{\theta}), \underline{p})$. If $p(\bar{\theta}) > \bar{p} := \sup\{\bar{P}\}$, then the previous does not imply multiplicity, but if $p(\bar{\theta}) < \bar{p}$, it does for all θ with associated prices in $(p(\bar{\theta}), \bar{p})$. The same argument holds to prove that $p(\underline{\theta}) \leq \underline{p}$. □

Strictly monotone price functions may have jump discontinuities. However such price functions can always be approximated arbitrarily well by continuous and strictly increasing functions.

The principal might be willing to accept multiplicity if all the equilibria induced by a mechanism are good. In particular, it is reasonable to assume that if all equilbria induced by a given mechanism M are better then the best robustly implementable equilibrium then the principal will prefer M to any robustly implementable equilibrium. We first need some intermediate results.

Lemma 4. *Fix a continuous M. Assume $R(a, \cdot)$ is (weakly) increasing for all a (the same* *holds if decreasing, with $\underline{\theta}$ and $\bar{\theta}$ switched). Then each p such that $\theta(p) = \varnothing$ is of one and only one of the following two types:*

- *Type L: $R(M(p), \underline{\theta}) > p$.*
- *Type H: $R(M(p), \bar{\theta}) < p$.*

Proof. Since $\theta \mapsto R(M(p), \theta)$ is increasing p cannot be of both types. If p is of neither then by continuity there exists a $\theta \in [\underline{\theta}, \bar{\theta}]$ such that $R(M(p), \theta) = p$. But then $\theta(p)$ is not empty. □

Corollary 3. *The set of prices $\{p : \theta(p) = \varnothing\}$ is open.*

Lemma 5. *Assume $R(a, \cdot)$ is (weakly) increasing for all a (the same holds if decreasing, with $\underline{\theta}$ and $\bar{\theta}$ switched) and M is continuous. Let p be Type L and $\theta'' > \theta'$.*

i. If there exists $p'' > p$ such that $\theta'' \in \theta(p'')$ then there exists $p' \in (p, p'']$ such that $\theta' \in \theta(p')$.

ii. If there exists $p'' < p$ such that $\theta'' \in \theta(p'')$ then there exists $p' \in [p'', p)$ such that $\theta' \in \theta(p')$.

Proof. We will prove (i), the proof for (ii) is symmetric. $R(M(p), \underline{\theta}) > p$ since p is type L. Moreover, under monotonicity

$$p'' = R(M(p''), \theta'') \geq R(M(p''), \theta') \geq R(M(p''), \underline{\theta}).$$

Then by continuity of R and M, there exists $\underline{p} \in (p, p'']$ such that $R(M(\underline{p}), \underline{\theta}) = \underline{p}$. By monotonicity we have $R(M(\underline{p}), \theta') \geq R(M(\underline{p}), \underline{\theta}) = \underline{p}$ and $p'' = R(M(p''), \theta'') \geq R(M(p''), \theta')$. Then by continuity of R, M there exists $p' \in [\underline{p}, p'']$ such that $R(M(p'), \theta') = p'$, so $\theta' \in \theta(p')$ as desired. □

Lemma 6. *Assume $R(a, \cdot)$ is (weakly) increasing for all a (the same holds if decreasing, with $\underline{\theta}$ and $\bar{\theta}$ switched) and M is continuous. Let p be type H and $\theta'' > \theta'$.*

i. If there exists $p' > p$ such that $\theta' \in \theta(p')$ then there exists $p'' \in (p, p']$ such that $\theta'' \in \theta(p'')$.

ii. If there exists $p' < p$ such that $\theta' \in \theta(p')$ then there exists $p'' \in [p', p)$ such that $\theta'' \in \theta(p'')$.

Proof. Analogous to that of Lemma 5. □

Lemma 7. *(Generalized intermediate value theorem). Let $F : [0,1] \mapsto [0,1]$ be a compact and convex valued, upper hemicontinuous correspondence. Let $p_1 < p_2$. Let $y_1 \in F(p_1)$ and $y_2 \in F(p_2)$. Then for any $\tilde{y} \in (\min\{y_1, y_2\}, \max\{y_1, y_2\})$ there exists $p \in [p_1, p_2]$ such that $\tilde{y} \in F(p)$.*

Proof. Define $p^* := \sup\{p \in [p_1, p_2) : \max F(p) < \tilde{y}\}$. If $p^* = p_1$ then $\max F(p) \geq \tilde{y}$ for all $p \in (p_1, p_2)$. Assume none of these hold with equality (otherwise we are done). Then if $\min F(p) \leq \tilde{y}$ for some $p \in (p_1, p_2]$ then we are done, by convexity of F. So suppose $\min F(p) > \tilde{y}$ for all $p \in (p_1, p_2]$. Then $\tilde{y} \in F(p_1)$: otherwise, by convexity of $F(p_1)$, we have $\max F(p_1) < \tilde{y}$, which violates upper hemicontinuity. Thus we are done if $p* = p_1$.

Suppose instead that $p^* = p_2$. If $\min F(p_2) \leq \tilde{y}$ then we are done, by convexity of $F(p)$. Suppose $\min F(p_2) > \tilde{y}$. Then by the definition of p^*, it must be that for any $\varepsilon > 0$ there exists $p \in (p_2 - \varepsilon, p_2)$ such that $\max F(p) < \tilde{y}$. But this violates upper hemicontinuity of F at p_2. Thus we are done if $p^* = p_2$.

It only remains to address the case of $p^* \in (p_1, p_2)$. It must be that $\max F(p^*) \geq \tilde{y}$: if not then by upper hemicontinuity there exists $\varepsilon > 0$ such that $\max F(p) < \tilde{y}$ for all $p \in [p^*, p^* + \varepsilon)$, but this would contradict the definition of p^*. If $\min F(p^*) \leq \tilde{y}$ then we are done, by convexity. So suppose $\min F(p^*) > \tilde{y}$. Then by upper hemicontinuity there exists $\varepsilon > 0$ such that $\min F(p) > \tilde{y}$ for all $p \in (p^* - \varepsilon, p^*]$. But this contradicts the definition of p^*. □

A.7.4 Proof of Proposition 4

Proof. Assume for the proof that $R(a, \cdot)$ is increasing for all $a \in \mathcal{A}$. The proof if decreasing is symmetric (using the corresponding versions of Lemmas 5 and 6).

Fix a decision rule M. Consider the mapping $\theta(p)$ on $\bar{P}$. Since we are not assuming uniqueness $\bar{P}$ may not be convex. For any θ', the set $\theta^{-1}(\theta')$ is compact: if $R(M(p), \theta') \neq p$ then this holds for all p' in a neighborhood p by continuity of M and p.

Take any $\underline{p} \in \theta^{-1}(\underline{\theta})$ and $\bar{p} \in \theta^{-1}(\bar{\theta})$. If $\underline{p} = \bar{p}$ then we are done: convexity of $\theta_M(p)$ (Lemma 2) implies that there is a constant, and thus monotone, equilibrium price function. Assume instead that $\underline{p} > \bar{p}$. Since $\theta^{-1}(\underline{\theta})$ is compact, $\underline{p}' = \min\{p \in \theta^{-1}(\underline{\theta}) : p \geq \bar{p}\}$ exists. Notice that there are no Type L price in $[\bar{p}, \underline{p}']$; if there were then by Lemma 5 there would be a $\bar{p} < p' < \underline{p}'$, $p' \in \theta^{-1}(\underline{p})$, which would contradict the definition of $\underline{p}'$. Now let $\bar{p}' = \max\{p \in \theta^{-1}(\bar{\theta}) : p \leq \underline{p}'\}$. Then by Lemma 6 there are no Type H prices in $[\bar{p}', \underline{p}']$. Thus there are no Type L or Type H prices in $[\bar{p}', \underline{p}']$, i.e. $[\bar{p}', \underline{p}'] \subseteq \bar{P}$. By a symmetric argument, if we start from a $\underline{p} < \bar{p}$ then we can construct $\underline{p}', \bar{p}'$ such that $[\underline{p}', \bar{p}'] \subseteq \bar{P}$. Again, if $\bar{p}' = \underline{p}'$ we are done.

Assume for the remainder of the proof that $\underline{p} \in \theta^{-1}(\underline{\theta})$, $\bar{p} \in \theta^{-1}(\bar{\theta})$, $\underline{p} < \bar{p}$, and $[\underline{p}, \bar{p}] \subseteq \bar{P}$. The proof for the reverse inequality is symmetric.

Claim 1. $\theta^{-1}(\theta) \cap [\underline{p}, \bar{p}] \neq \varnothing$ for all $\theta \in \Theta$.

To show this, suppose as a contradiction that $p' = \min\{p \in \theta^{-1}(\theta) : p \geq \underline{p}\} > \bar{p}$ for some θ' (the argument for the other type of violation is symmetric). By a similar argument to that given above (using Lemmas 5 and 6) we can show that $[\underline{p}, p'] \subseteq \bar{P}$. Note that if there exists $\theta \in [\underline{\theta}, \theta']$ such that $\theta^{-1}(\theta) = \bar{p}$, then by convexity of $\theta(p)$, we have that $\bar{p} \in \theta^{-1}(\theta')$, contradicting the assumption that $p' > \bar{p}$. It remains to show that such a θ exists. This follows immediately form Lemma 7. This proves Claim 1.

Finally, we need to construct an increasing equilibrium price function (in this case it must be increasing since $\underline{p} < \bar{p}$). Consider an arbitrary price function $\tilde{P}$ such that $\tilde{P}(\theta) \in \theta^{-1}(\theta) \cap [\underline{p}, \bar{p}]$ for all θ, $\tilde{P}(\underline{\theta}) = \underline{p}$, and $\tilde{P}(\bar{\theta}) = \bar{p}$. We will show that any violations of monotonicity can be ironed without leading to further violations.

Claim 2. Suppose $\tilde{P}(\theta'') < \tilde{P}(\theta') < \tilde{P}(\theta''')$ for $\theta''' > \theta'' > \theta'$. Then there exists $p \in \theta^{-1}(\theta'') \cap [\tilde{P}(\theta'), \tilde{P}(\theta''')]$.

This claim follows immediately from Lemma 7.

The existence of a monotone price function follows. □

A.7.5 Proof of Proposition 5

Proof. By Proposition 4, M admits an equilibrium with a monotone price function P. Let Q be the associated action function. For any state θ such that $r(Q(\theta), P(\theta))$ is non-degenerate, let $\hat{Q}(\theta') = R(Q(\theta), \theta)$ for all $\theta' \in r(Q(\theta), P(\theta))$. Clearly $\hat{P}(\theta) := R(\hat{Q}(\theta), \theta)$ will also be monotone, and $(\hat{Q}, \hat{P})$ is also implemented by M. It remains to show that M can be modified on $\mathcal{P} \setminus \hat{P}(\Theta)$ in order to implement $(\hat{Q}, \hat{P})$ uniquely. This follows from Proposition 2. Note that $\hat{Q}$ will have no degenerate discontinuities since M was assumed to be continuous. □

A.8 Proof of Section 3.2

Lemma 8. *Given a continuous function $F : X \times (0,1) \mapsto X$ on a compact, convex subset X of a Euclidean space, define the function*

$$G(t) = \{x \in X : F(x,t) = x\}.$$

Then $G(t)$ is non-empty for all t (Brouwer's fixed point theorem). Moreover, if $G(t)$ is single valued then G is upper and lower hemicontinuous at t.

Proof. Since $G(t)$ is single valued upper hemicontinuity implies lower hemicontinuity. We want to show that for any open neighborhood V of $G(t)$ there exists a neighborhood U of t such that $G(t') \subseteq V$ for all $t' \in U$.

Claim 1. For any open neighborhood V of $G(t)$ there exists a $\kappa > 0$ such that

$$|F(x,t) - x| > \kappa \quad \forall \;\; x \in X \setminus V.$$

The proof of claim 1 is as follows. $X \setminus V$ is a closed subset of a compact set, and thus compact. The function $x \mapsto F(x,t) - x$ is continuous, so it attains its minimum on $X \setminus V$. Since $G(t)$ is unique and $G(t) \notin X \setminus V$, this minimum is strictly greater then zero, so the desired κ exists.

To complete the proof of Lemma 8, we need to show that there exists an open neighborhood U of t such

$$|F(x,t') - x| > \kappa \quad \forall \;\; x \in X \setminus V, \;\; t' \in U.$$

By continuity of $t' \mapsto F(x,t') - x$, for each x there exists a ε_x such that $|t' - t| < \varepsilon_x$ implies $|F(x,t') - x| > \kappa$. For each x, define $\ell(x,\varepsilon) = \min\{|F(x,t') - x| : |t' - t| \leq \varepsilon/2\}$, which exists by continuity of F and compactness of $|t' - t| \leq \varepsilon/2$. Define

$$B(x) = \{x' \in X : \ell(x', \varepsilon_x) > \kappa\}.$$

By continuity of $x \mapsto F(x,t') - x$, $B(x)$ contains an open neighborhood of x (Berge's maximum theorem). Let $\tilde{B}(x)$ be this open neighborhood. The set $\cup_{x \in X \setminus V} \tilde{B}(x)$ covers $X \setminus V$. Then by compactness of $X \setminus V$ there exists a finite sub-cover. Let u be the smallest ε_x corresponding to an x such that $\tilde{B}(x)$ is in the finite sub-cover. Then $U = \{t' \in (0,1) : |t' - t| < u\}$. □

Proposition 10. *Given a continuous function $F : X \times \Theta \times (0,1) \mapsto X$ on a compact, convex subset X of a Euclidean space, define the function*

$$G(t,\theta) = \{x \in X : F(x,\theta,t) = x\}.$$

Then $G(t,\theta)$ is non-empty for all t,θ (Brouwer's fixed point theorem). Moreover, let S be any compact subset of Θ such that $G(t,\theta)$ is single valued for all $\theta \in S$. Then $G(t,\theta)$ is upper and lower hemicontinuous in its first argument at t, uniformly over S.

Proof. Since $G(t,\theta)$ is unique on S it suffices to show upper hemicontinuity. Let $V(\theta)$ be an open neighborhood of $\theta \mapsto G(t,\theta)$ on S. Without loss of generality (since Θ is compact and $G(t,\theta)$ single valued on S), let $V(\theta) = \{x \in X : |G(t,\theta) - x| < \delta\}$ for some $\delta > 0$, or equivalently, $V(\theta) = \cup_{x \in G(t,\theta)} N_\delta(x)$. We want to show that there exists an open neighborhood U of t such that $t' \in U$ implies $G(t',\theta) \subseteq V(\theta)$ for all $\theta \in S$.

Claim 1. $X \setminus V(\theta)$ is upper and lower hemicontinuous on S.

The proof of Claim 1 is as follows. Since $G(t, \theta)$ is single valued,

$$X \setminus V(\theta) = X \setminus N_\delta(G(t, \theta))$$

where $N_\delta(x)$ is the open ball around x with radius δ. We first show upper hemicontinuity. Let W be an open set containing $X \setminus V(\theta)$. Without loss of generality, let

$$W = X \setminus \bar{N}_{\delta-\rho}(G(t, \theta))$$

for some $\rho \in (0, \delta)$ where $\bar{N}_{\delta-\rho}(x)$ is the closed ball around x with radius $\delta - \rho$.[15] By Lemma 8, we know that $\theta \mapsto G(t, \theta)$ is upper and lower hemicontinuous at all $\theta \in S$. By upper hemicontinuity of $\theta \mapsto G(t, \theta)$ at θ, there exists an open neighborhood B of θ such that $\theta' \in B$ implies $|x - G(\theta, t)| < (\delta - \rho)/2$ for all $x \in G(\theta', t)$. Then $\bar{N}_{\delta-\rho}(G(t, \theta)) \subset \cup_{x \in G(t,\theta')} N_\delta(x) = V(\theta')$ for all $\theta' \in B$. Thus $V(\theta') \subset W$ for all $\theta' \in B$, which shows upper hemicontinuity.

For lower hemicontinuity, let $W \subset X$ be an open set intersecting $X \setminus V(\theta)$. This holds if and only if there exists $x' \in W$ such that $|x' - G(t, \theta)| > \delta$. By upper hemicontinuity of $\theta \mapsto G(t, \theta)$ at θ, there exists an open neighborhood B of θ such that $\theta' \in B$ implies $|x - G(\theta, t)| < (|x' - G(t, \theta)| - \delta)/2$ for all $x \in G(\theta', t)$. Then $\theta' \in B$ implies $|x' - x| > \delta$ for all $x \in G(t, \theta')$. Thus $x' \notin \cup_{x \in G(t,\theta')} N_\delta(x) = V(\theta')$, so $W \cap X \setminus V(\theta') \neq \varnothing$ for all $\theta' \in B$, which shows lower hemicontinuity. This completes the proof of Claim 1.

We know from Lemma 8 that for each $\theta \in S$ there exists $\varepsilon_\theta, \kappa_\theta > 0$ such that

$$|t' - t| < \varepsilon_\theta \Longrightarrow |F(x, \theta, t') - x| > \kappa_\theta \;\; \forall\, x \in X \setminus V(\theta). \tag{2}$$

Claim 2. For each $\theta \in S$ there exists an open neighborhood $B(\theta)$ of θ such that

$$\theta' \in B(\theta) \text{ and } |t' - t| < \varepsilon_\theta \Longrightarrow |F(x, \theta, t') - x| > \kappa_\theta \;\; \forall\, x \in X \setminus V(\theta'),$$

where $\varepsilon_\theta, \kappa_\theta$ satisfy (2).

The proof of this claim is as follows. Define

$$z(\theta, \varepsilon) := \min\{|F(x, \theta, t') - x| : |t' - t| \leq \varepsilon/2,\; x \in X \setminus V(\theta)\},$$

which is well defined by compactness of $X \setminus V(\theta)$. By Berge's maximum theorem and Claim 1, $\theta \mapsto z(\theta, \varepsilon)$ is continuous at any $\theta \in S$. By (2) we know that $z(\theta, \varepsilon_\theta) > \kappa_\theta$ for all $\theta \in S$. Then for any $\theta \in S$ there exists an open neighborhood $B(\theta)$ of θ such that $\theta' \in B(\theta)$ implies $z(\theta', \varepsilon_\theta) > \kappa_\theta$. This proves Claim 2.

[15] W so defined is open in X, but not in the space of which X is a subset.

To complete the proof of Proposition 10, note that $\cup_{\theta \in S} B(\theta)$ is an open cover of S. By compactness of S there exists a finite sub-cover. Let I be the set of $\theta \in S$ that index this sub-cover. Let $\varepsilon = \min\{\varepsilon_\theta : \theta \in I\}/2$. Then

$$|t' - t| < \varepsilon \Longrightarrow |F(x, \theta, t') - x| > 0 \ \ \forall\, x \in X \setminus V(\theta) \text{ and } \theta \in S.$$

Since $G(t', \theta)$ is non-empty for all t', θ we have that $|t' - t| < \varepsilon$ implies that for all θ, $G(t', \theta) \subseteq V(\theta)$, which shows upper hemicontinuity as desired. $\square$

A.8.1 Proof of Theorem 3

Proof. Let $F(a, \theta, t) = M(R(a, \theta, t))$, where t continuously parameterizes the function R. Then F is continuous since M is continuous by robustness to manipulation. Moreover, $G(t, \theta) = \tilde{Q}(\theta, t)$ will be single valued on all but a zero-measure set of states, by robustness to multiplicity. Therefore for any $\varepsilon > 0$ we can find a compact set S such that $G(t, \theta)$ is single valued for all $\theta \in S$. Then Proposition 10 applies, which gives the result. $\square$

A.8.2 Proof of Proposition 3

Proof. First, note that $|\mathcal{U}(R)| = 1$. It is non-empty since Lemma 8 and robustness to multiplicity imply that $\theta \mapsto \tilde{Q}(\theta, R)$ is a continuous function on all but a zero measure set of states, and is thus $\theta \mapsto u(\theta, Q(\theta)$ integrable for all $Q \in \tilde{Q}(\cdot, R)$. It is single valued since all $Q \in \tilde{Q}(\cdot, R)$ are the same on all but a zero measure set of states.

Since $|\mathcal{U}(R)| = 1$, upper hemicontinuity implies lower hemicontinuity, so it suffices to show the former. Thus we want to show that for any $\delta > 0$ there exists an open neighborhood $B \subseteq \mathcal{C}$ of R such that $R' \in B$ implies $\mathcal{U}(R') \subseteq (\mathcal{U}(R) - \delta, \mathcal{U}(R) + \delta)$. Since the set of states at which $\tilde{Q}(\theta, R)$ is not single valued has zero measure, for any $\varepsilon > 0$ there exists a compact set S such that $\{\theta \in \Theta : |\tilde{Q}(\theta, R)| \neq 1\} \subset \Theta \setminus S$ and $\lambda(\Theta \setminus S) < \varepsilon$ (where λ is Lebesgue measure). For any such S, there exists a neighborhood $B_S \subset C$ of R such that $R' \in B_S$ implies

$$\int_S u(\theta, Q(\theta)) dH(\theta) - \mathcal{U}(R)| < \delta/2$$

Taking ε small enough gives implies $\mathcal{U}(R') \subseteq (\mathcal{U}(R) - \delta, \mathcal{U}(R) + \delta)$ as desired. $\square$

A.8.3 Structural uncertainty with multiplicity

This section explores robustness to structural uncertainty without robustness to multiplicity. The following condition will be useful to show lower hemicontinuity of the set of equilibrium actions.

Local monotonicity. *Local monotonicity* holds at θ, p if the following holds: $R(M(p), \theta) = p$ and for all $\varepsilon > 0$ there exist $p'', p' \in N_\varepsilon(p)$ such that

$$R(M(p''), \theta, t) - p'' > 0 > R(M(p'), \theta, t) - p'.$$

We will also say that *local monotonicity* holds in state θ if it holds for all p such that $R(M(p), \theta) = p$.

Lemma 9. *Assume that M is continuous, and let $R(M(p), \theta, t)$ be continuous. Then $t \mapsto \tilde{Q}(\theta, t)$ is upper hemicontinuous at t. Moreover $t \mapsto \tilde{Q}(\theta, t)$ is lower hemicontinuous at t if and only if local monotonicity holds at θ.*

Proof. The proof of upper hemicontinuity is the same as in the proof Lemma 8. It remains to show lower hemicontinuity.

Assume local monotonicity holds. Let $a \in \tilde{Q}(\theta, t)$, and let V be an open neighborhood of a. We want to show that there exists an open neighborhood U of t such that $\tilde{Q}(\theta, t') \cap V \neq \varnothing$ for all $t' \in U$. Let $p = R(a, \theta)$. Since M is continuous, there exists a $\varepsilon > 0$ such that $M(p') \in V$ for all $p' \in N_\varepsilon(p)$. Thus lower hemicontinuity of $t \mapsto \tilde{Q}(\theta, t)$ will follow from lower hemicontinuity of $t \mapsto \{p' \in \mathbb{R} : R(M(p'), \theta, t) = p'\}$.

By assumption, there exist p', p'' in $N_\varepsilon(p)$ such that

$$R(M(p''), \theta, t) - p'' > 0 > R(M(p'), \theta, t) - p'.$$

Then by continuity of R in t, there exists an open neighborhood U of t such that

$$R(M(p''), \theta, t') - p'' > 0 > R(M(p'), \theta, t') - p'$$

for all $t' \in U$. Then by continuity of M, for all such t' there exists $p_{t'} \in (p', p'')$ such that $R(M(p_{t'}), \theta, t') = p_{t'}$, which is what we wanted to show.

The converse direction for lower hemicontinuity is easy. Suppose local monotonicity does not hold at θ. Then let p and $\varepsilon > 0$ be such that $R(M(p), \theta) = p$ and $R(M(p'), \theta) \leq p'$ for all $p' \in N_\varepsilon(p)$ (symmetric proof if $\geq$). Then it is easy to see that perturbing R to $R - \sigma$ for some $\sigma > 0$ will lead to a violation of lower hemicontinuity. $\square$

Lemma 9 made use of Local Monotonicity. In order to show that the set of principal payoffs is upper hemicontinuous we want to say that this holds on all but a zero measure set of states.

Let $R : \mathcal{A} \times \Theta \times (0, 1) \mapsto \mathbb{R}$ be continuous. For a fixed M, define

$$Y(t, \theta) = \{p : R(M(p), \theta, t) = p\}.$$

$(Y(t, \theta) = \theta^{-1}(\theta)$ for fixed t. The distinction between the two is made for notational convenience.) As preliminaries, observe the following properties:

1. Since Θ and $\mathcal{A}$ are compact. We will consider small neighborhoods of t, so for some closed neighborhood N of t define $\underline{p} := \min_{a,\theta,t'} R(a, \theta)\ s.t.\ t' \in N$ and $\bar{p} := \max_{a,\theta,t'} R(a, \theta)\ s.t.\ t' \in N$.

2. For any θ, t, the set $Y(t, \theta)$ is compact.

The following assumption will, essentially, characterize robustness to manipulation and structural uncertainty.

Local monotonicity when flat. If $\theta \mapsto R(M(p), \theta)$ is constant in a neighborhood of θ and $R(M(p), \theta) = p$ then local monotonicity holds at θ, p

Lemma 10. *Assume that M is continuous, that $\{\theta \mapsto R(a, \theta, t)\}$ are weakly increasing for all a, t, and that $Y(t, \theta)$ is infinite for no more than a zero measure set of states.*[16] *Then then the set of states for which local monotonicity does not hold has zero measure if and only if M satisfies local monotonicity when flat.*

Proof. First, suppose local monotonicity when flat holds. We wish to show that the set of states for which $|Y(t, \theta)|$ local monotonicity fails is countable. Let θ be such that $|Y(\theta, t)| < \infty$. By continuity and local monotonicity when flat, there exists an $\varepsilon > 0$ such that $|\theta' - \theta| < \varepsilon$ implies that θ' satisfies local monotonicity (regardless of whether or not θ did). To see this, consider fist p such that p, θ satisfy local monotonicity. For each such $p \in Y(\theta, t)$ there exist p', p'' such that

$$R(M(p''), \theta, t) - p'' > 0 > R(M(p'), \theta, t) - p'.$$

Then for each such p there exists $\varepsilon_p > 0$ such that

$$R(M(p''), \theta', t) - p'' > 0 > R(M(p'), \theta', t) - p'$$

for all θ' such that $|\theta' - \theta| < \varepsilon_p$. Let ε be the minimum over ε_p for all such p.

Now consider p such that $R(M(p), \theta) = p$ and local monotonicity is not satisfied; assume there exists $\sigma > 0$ such that $R(M(\tilde{p}), \theta) < \tilde{p}$ for all $\tilde{p} \in N_\sigma(p) \setminus p$ (symmetric proof for $>$; the case of weak inequality is addressed below). Local monotonicity when flat implies that $\theta \mapsto R(M(p), \theta)$ cannot be flat in a neighborhood of θ. By R weakly increasing in θ, local

[16] Alternatively, could assume that the set of states such that $Y(t, \theta)$ is finite is dense in Θ, and derive this condition.

monotonicity must be satisfied at p for all these states $\theta' \in (\theta, \theta + \delta)$ for some $\delta > 0$. If $\theta \mapsto R(M(p), \theta)$ is strictly monotone at θ' (assume without loss that it is increasing) then there exists $\kappa > 0$ such that $R(M(\tilde{p}), \theta') - \tilde{p} < 0$ for all $\tilde{p} \in N_\sigma(p)$ and $\theta' \in (\theta - \kappa, \theta)$. Then $N_\sigma(p) \cap G(t, \theta') = \varnothing$, for all $\theta' \in (\theta - \kappa, \theta)$. Moreover for $\theta\ \in (\theta, \theta + \kappa)$ there are two prices in $N_\sigma(p)$ that satisfy local monotonicity.

The final case is of weak inequality: consider p such that $R(M(p), \theta) = p$ and assume there exists $\sigma > 0$ such that $R(M(\tilde{p}), \theta) = \tilde{p}$ for all $\tilde{p} \in N_\sigma(p)$ (if p is the endpoint of the set of such prices it is dealt with in the same way as the previous cases). Then by R weakly increasing in θ and local monotonicity when flat, $\theta' \mapsto R(M(p'), \theta')$ is strictly increasing for all these prices. Thus there exists $\kappa > 0$ such that $N_\sigma(p) \cap G(t, \theta') = \varnothing$, for all $\theta' \in N_\kappa(\theta)$.

Since $\theta \mapsto G(t, \theta)$ is upper hemicontinuous at θ (essentially the same proof as Lemma 9), the above argument shows that there exists a neighborhood of θ such that local monotonicity is satisfied for all θ' in this neighborhood. Then there can be only countably many θ that violate local monotonicity and have $|\theta_M^{-1}(\theta)| < \infty$ (since each is containted in an open interval of states in which it is the unique state violating local monotonicity).

Suppose local monotonicity when flat does not hold at θ, i.e. there is a p such that $R(M(p), \theta) = p$, $\theta \mapsto R(M(p), \theta)$ is constant in a neighborhood of θ and local monotonicity does not hold at p, θ then there is a neighborhood of θ such that $R(M(p), \theta') = p$ and local monotonicity fails at θ', p for all θ' in the neighborhood. $\square$

Proposition 11. *Let $S \subseteq \Theta$ be a compact set such that $|Y(t, \theta)| < K < \infty$ for all $\theta \in S$. Then $t \mapsto Y(t, \theta)$ is upper hemicontinuous at t, uniformly over S. Moreover, $t \mapsto Y(t, \theta)$ is lower hemicontinuous uniformly over S if and only if each $\theta \in S$ satisfies local monotonicity.*

Proof. The proof for upper heimcontinuity is the same as in Proposition 10. The only difference is that $G(t, \theta)$ is finite valued on S, rather than single valued as in Proposition 10. However since $|g(t, \theta)|$ is uniformly bounded on S the proof is essentially unchanged.

It remains to prove lower hemicontinuity. Since M is continuous and $[\underline{p}, \bar{p}]$ is compact M is uniformly continuous on $[\underline{p}, \bar{p}]$. Thus it suffices to show that $Y(t, \theta)$ is lower hemicontinuous. Let $y(t, \theta)$ be any selection from $Y(t, \theta)$. For each $\theta \in S$ define $V(\theta) = \{p \in \mathbb{R} : |p - y(t, \theta)| < \delta\}$ for some $\delta > 0$. Since $Y(t, \theta)$ is finite for all $\theta \in S$ it suffices to show that there exists an open set U of t such that $V(\theta) \cap Y(t', \theta) \neq \varnothing$ for all $t \in U$.

Fix a $\theta \in S$. There are at most $K < \infty$ values $\{p_i\}_{i=1}^K$ such that $R(M(p_i), \theta) = p_i$. Local monotonicity is satisfied at each p_i, so there exist $p_i', p_i'' \in N_{\delta/2}(p_i)$ such that $R(M(p_i''), \theta, t) - p_i'' > 0 > R(M(p_i'), \theta, t) - p_i'$. Moreover, by continuity there exists $\varepsilon_i > 0$ such that $|\theta' - \theta| < \varepsilon_i$ implies $R(M(p_i''), \theta', t) - p_i'' > 0 > R(M(p_i'), \theta', t) - p_i'$. In particular, this implies that

$Y(t,\theta') \cap (\min\{p_i', p_i''\}, \max\{p_i', p_i''\}) \neq \varnothing$ for all such θ'. Define

$$\underline{x}_i'' := \min\{R(M(p_i''), \theta', t) - p_i'' : |\theta' - \theta| < \varepsilon_i/2\},$$

which is well defined and strictly positive since the objective is continuous and the constraint set compact. Similarly, define

$$\bar{x}_i' := \max\{R(M(p_i'), \theta', t) - p_i' : |\theta' - \theta| < \varepsilon_i/2\}.$$

We will have $\bar{x}_i' < 0$. Thus for all θ' such that $|\theta' - \theta| < \varepsilon_i/3$ we have

$$R(M(p_i''), \theta', t) - p_i'' > \underline{x}_i'' > 0 > \bar{x}_i' > R(M(p_i'), \theta', t) - p_i'.$$

Then by uniform continuity of $t' \mapsto R(a, \theta, t')$ on a closed neighborhood of t (Heine-Cantor theorem) there exists σ_i such that $|t' - t| < \sigma_i$ implies

$$R(M(p_i''), \theta', t') - p_i'' > 0 > R(M(p_i'), \theta', t') - p_i'$$

for all θ' such that $|\theta' - \theta| < \varepsilon_i/3$. This implies that $Y(t', \theta') \cap (\min\{p_i', p_i''\}, \max\{p_i', p_i''\}) \neq \varnothing$ for all such θ', t'.

Let $\sigma_\theta = \min_i \sigma_i$ and $\varepsilon_\theta = \min_i \varepsilon_i/3$. Note that $Y(t, \theta)$ is upper hemicontinuous in θ on S (similar to the proof of Lemma 9. This means that there exists $\tilde{\varepsilon}$ such that $\theta' \in S$, $|\theta' - \theta| < \tilde{\varepsilon}$ implies $Y(t, \theta') \subseteq \cup_{i=1}^K (\min\{p_i', p_i''\}, \max\{p_i', p_i''\})$. Without loss of generality, let $\varepsilon_\theta < \tilde{\varepsilon}$. Recall that $y(t, \theta)$ was an arbitrary selection from $Y(t, \theta)$. Then we can summarize the results thus far as follows: $\theta' \in S$, $|\theta' - \theta| < \varepsilon_\theta$, and $|t' - t| < \sigma_\theta$ implies $(y(t, \theta) - \delta, y(t, \theta) + \delta) \cap Y(t', \theta') \neq \varnothing$.

The open set $\cup_{\theta \in S}(\theta - \varepsilon_\theta, \theta + \varepsilon_\theta)$ is a cover of S. By compactness of S there exists a finite subcover. Let σ be the minimum over σ_θ for the θ represented in this subcover. Then $|t' - t| < \sigma$ implies $V(\theta) \cap Y(t', \theta) \neq \varnothing$ as desired.

□

Proposition 12. *Assume R is weakly increasing in θ. Let M be a decision rule that is robust to manipulation for market clearing function R. Assume that $\{p : R(M(p), \theta, t) = p\}$ is finite for almost all θ. Then $\mathcal{U}(t)$ is upper hemicontinuous at t on $\mathcal{C}$. Moreover, $\mathcal{U}(t)$ is lower hemicontinuous at t on $\mathcal{C}$ if and only if local monotonicity when flat holds for all θ such that $\{p : R(M(p), \theta, t) = p\}$ is finite.*

Proof. First, note that for any $\varepsilon > 0$ there exists a $K \in \mathbb{N}$ such that the measure of states for which $|Y(t, \theta)| > K$ is less than ε. This holds since the set of states for which $|Y(t, \theta)| = \infty$ has zero measure by assumption. Moreover, by Lemma 10 we have that the set of states such that $|Y(t,\theta)| \leq K$ which do not satisfy local monotonicity has zero measure. Thus for any $\varepsilon > 0$ there exists a compact set $S \subseteq \Theta$ with measure greater than $1-\varepsilon$ such that $|Y(t,\theta)| < K$ and local monotonicity is satisfied for all $\theta \in S$. Then we can apply Proposition 11 for arbitrarily small ε to obtain the result. $\square$

B Bridgeability

This section discusses bridgeability further. We provide sufficient conditions for the various notions of bridgeability, and show that they are satisfied in common settings.

Let $(\mathcal{A}, \succ)$ be a partially ordered set. Say $(\mathcal{A}, \succ)$ is *upward directed* if for any two $a'', a' \in \mathcal{A}$ there exists $c \in \mathcal{A}$ such that $c \succ a''$ and $c \succ a'$. Downward directed is defined analogously.[17] We use the notation $a''_\alpha a' \equiv \alpha a'' + (1-\alpha)a'$. Say that $\succ$ is preserved by mixtures if for any $a'' \succ a'$ and $\alpha \in (0,1)$, $a'' \succ a''_\alpha a' \succ a'$. Finally, say that $a \mapsto R(a,\theta)$ is *strongly monotone with respect to* $\succ$ if $a'' \succ a'$ and $a'' \neq a'$ implies $R(a'',\theta) > R(a',\theta)$. We use the notation $a''_\alpha a' \equiv \alpha a'' + (1-\alpha)a'$. The following proposition gives sufficient conditions for full bridgeability, but it is also useful because the proof of the existence of a monotone path is constructive. This construction could potentially be useful in applications.

Proposition 13. *Let $(\mathcal{A}, \succ)$ be a partially ordered set that is both upward and downward directed, and such that $\succ$ is preserved by mixtures. If $R(\cdot,\theta)$ is strongly monotone with respect to $\succ$ then there is a monotone path between a' and a'' at θ iff $R(a'',\theta) \neq R(a',\theta)$*

Proof. The condition $R(a',\theta) \neq R(a'',\theta)$ is obviously necessary. It remains to show that it is sufficient. That is, we want to show that there exists a monotone path between any $a'', a' \in \mathcal{A}$ such that $R(a',\theta) \neq R(a'',\theta)$. Assume without loss that $R(a'',\theta) > R(a',\theta)$. If $a'' \succ a'$ then the ray from a'' to a' is a monotone path. This follows since $\succ$ is preserved by mixtures and $R(\cdot,\theta)$ is strongly monotone.

Suppose a' and a'' are not ordered. Let $\bar{a}$ be an upper bound for a'', a', i.e. $\bar{a} \succ a''$ and $\bar{a} \succ a'$, and let $\underline{a}$ be a lower bound. Both exist since $(\mathcal{A}, \succ)$ is upward and downward directed. By continuity of R, there exists $\bar{\lambda} \in (0,1)$ such that $R(\bar{a}_{\bar{\lambda}} a', \theta) = R(a'',\theta)$. Similarly there exists $\underline{\lambda} \in (0,1)$ such that $R(a''_{\underline{\lambda}} \underline{a}), \theta) = R(a',\theta)$.

We will now construct one half of the monotone path from a' to a''. Let $t : [0,1] \mapsto [\bar{\lambda}, 1] \times [0,1]$ be a continuous and strictly monotone function, and let $t_i(x)$ be the i^{th} coordinate of $t(x)$. For each $x \in (0,1)$, we have $R(\bar{a}_{t_1(x)} a', \theta) > R(a'',\theta)$, $R(\underline{a}_{t_2(x)} a', \theta) < R(a',\theta)$, and $\bar{a}_{t_1(x)} a' \succ \bar{a}_{t_1(x)} a'$. These properties follow from strong monotonicity of R and the fact that $\succ$ is preserved under mixtures.

[17] A lattice is an upward and downward directed set, but the converse is not true.

For each $x \in (0,1)$, define $f(x)$ by $R((\bar{a}_{t_1(x)}a')_{f(x)}(\underline{a}_{t_2(x)}a'),\theta) = xR(a'',\theta)+(1-x)R(a',\theta)$. We claim that $x \mapsto (\bar{a}_{t_1(x)}a')_{f(x)}(\underline{a}_{t_2(x)}a')$ is a continuous function. It is a well defined function by strong monotonicity of R. It is continuous since R and t are continuous. Moreover, by construction $x \mapsto R((\bar{a}_{t_1(x)}a')_{f(x)}(\underline{a}_{t_2(x)}a'),\theta)$ is strictly increasing, and $(\bar{a}_{t_1(0)}a')_{f(0)}(\underline{a}_{t_2(0)}a') = a'$. Therefore $x \mapsto (\bar{a}_{t_1(x)}a')_{f(x)}(\underline{a}_{t_2(x)}a')$ forms one half of a monotone path from a' to a''. The other half of the monotone path is defined analogously, using a'' and $\underline{\lambda}$ in place of a' and $\bar{\lambda}$. $\square$

Proposition 13 makes it easy to identify when a discontinuity will be bridgeable. For example, it implies that when $\mathcal{A}$ is a chain a gap between a' and a'' will be bridgeable at θ iff $R(\cdot,\theta)$ is strictly monotone on (a',a'').

More importantly, Proposition 13 implies that every discontinuity will be bridgeable when $\mathcal{A} = \Delta(Z)$, i.e. the set of distributions on some set Z, under mild assumptions on R. Let $\pi(z,\theta)$ be a real valued function, with $\theta \mapsto \pi(z,\theta)$ continuous for all z. For example, $\pi(a,\theta)$ could represent a company's cash flow as a function of the state and government intervention $z \in Z$. In state θ, any $a \in \mathcal{A}$ in induces a distribution $F(a,\theta)$ on $\mathbb{R}$ via $\pi(\cdot,\theta)$. Let $\succ_{FOSD}$ be the first-order stochastic dominance order. This partial order on $\Delta(\mathbb{R})$ induces a preorder $\succeq$ on $\mathcal{A}$. Define $a'' \succ a'$ by $a'' \succeq a'$ and $\neg(a' \succeq a'')$ if $a'' \neq a'$, and $a' \succ a'$ for all a'. If $\pi(z',\theta) \neq \pi(z'',\theta)$ for all $z'' \neq z'$ then $\succeq = \succ$. Then $a \mapsto R(a,\theta)$ is strongly monotone if $F(a'',\theta) \succ F(a',\theta)$ implies $R(a'',\theta) \succ R(a',\theta)$. The partially ordered set $(\mathcal{A},\succ)$ satisfies the conditions of Proposition 13 (when $\pi(z',\theta) \neq \pi(z'',\theta)$ for all $z' \neq z''$ it is in fact a lattice).

Corollary 4. *If $\mathcal{A} = \Delta(Z)$ and for all θ $a \mapsto R(a,\theta)$ is strongly monotone with respect to the order induced by first-order stochastic dominance, then the environment is fully bridgeable.*

It will also be useful to establish a related notion of bridgeability. Say that there exists a monotone path from (a',θ') to (a'',θ'') if there exists a continuous function $\gamma : [0,1] \mapsto \mathcal{A} \times [\theta',\theta'']$ such that $\gamma(0) = (a',\theta')$, $\gamma(1) = (a'',\theta'')$, $x \mapsto \gamma_1(x)$ is weakly increasing and $R(\gamma_1(x),\gamma_2(x))$ is strictly increasing. The path is *strongly monotone* if moreover $\gamma_1(x)$ is strictly increasing.

Recall that the environment is *continuously bridgeable* if for any $\theta^* \in \Theta$ there exists $\varepsilon > 0$ such that if a', a'' is bridgeable at θ^* and $R(a'',\theta) \neq R(a',\theta)$ for all $\theta \in [\theta^*, \theta^*+\varepsilon]$ then there exists a sup-norm continuous function $\sigma(\cdot|a',a'') : [\theta^*,\theta^*+\varepsilon] \mapsto \mathcal{A}^{[0,1]}$ such that $\sigma(\theta|a',a'')$ is a monotone path from a' to a'' for all $\theta \in [\theta^*,\theta^*+\varepsilon]$. Say that the environment is *continuously fully bridgeable* if it is full bridgeable and continuously bridgeable.

Lemma 11. *Assume $\theta \mapsto R(a,\theta)$ is strictly monotone for all $a \in \mathcal{A}$ and the environment is continuously fully bridgeable. Then the environment is correctable if for all θ such that $R(a',\theta) = R(a'',\theta)$ for all $a \in \mathcal{A}$ the following holds: there exists $a_1 \in \mathcal{A}$ and $\delta > 0$ such that*

i. $R(a_1,\theta') > R(a,\theta')$ *for all* $a \neq a_1$ *and* $\theta' \in (\theta - \delta, \theta)$.

ii. $R(a,\theta') > R(a_1,\theta')$ *for all* $a \neq a_1$ *and* $\theta' \in (\theta, \theta + \delta)$.

Proof. If P is decreasing then the existence of an approximating Q' around any degenerate discontinuity follows immediately from continuous bridgeability. Assume therefore that P is increasing.

First suppose that there is a degenerate discontinuity at some θ such that there exist a', a'' with $R(a',\theta) \neq R(a'',\theta)$. Assume there exists $\bar{a}$ such that $R(\bar{a},\theta) > R(Q(a),\theta)$ (the argument for the reverse inequality is symmetric). Since P is increasing and continuous at θ, for any $\varepsilon > 0$ there exists $\theta'' \in (\theta, \theta + \varepsilon)$ and $\theta' \in (\theta - \varepsilon, \theta)$ such that $R(\bar{a},\theta) > R(Q(\theta''),\theta)$ for all $\theta \in (\theta', \theta'')$. Then since the environment is continuously bridgeable (in particular between $Q(\theta'')$ and $\bar{a}$) there exists a continuous Q' on $[\theta', \theta'']$, with the corresponding P' strictly increasing, such that $R(Q'(\theta'),\theta') \in (R(Q(\theta'),\theta'), R(Q(\theta''),\theta''))$ and $Q'(\theta'') = Q(\theta'')$. Thus the environment is correctable.

Now suppose there is a degenerate discontinuity at some θ such that $R(a'',\theta) = R(a',\theta)$ for all $a', a'' \in \mathcal{A}$. Then under conditions i and ii the following Q' satisfies the conditions for correcting the degenerate discontinuity: for any $\varepsilon < \delta$, $Q' = a_1$ on $(\theta - \varepsilon, \theta + \varepsilon)$ and equals Q elsewhere. $\square$

Note that Lemma 11 implies that the environment is correctable if for all θ there exist a', a'' such that $R(a',\theta) \neq R(a'',\theta)$. Except for unusual cases, the environment will be continuously fully bridgeable when it is fully bridgeable. For example, the environment of Corollary 4 is continuously fully bridgeable when Z is finite and $\theta \mapsto \pi(z,\theta)$ is differentiable for all z.

Lemma 12. *Assume $\mathcal{A} = \Delta(Z)$ for some finite Z, $\pi(z,\theta)$ is differentiable for all z, and for all θ, $a \mapsto R(a,\theta)$ is strongly monotone with respect to the first-order stochastic dominance induced order. Then the environment is continuously fully bridgeable.*

Proof. First suppose $\min_{z'',z' \in Z} |\pi(z'',\theta^*) - \pi(z',\theta^*)| > 0$. Then by continuity of $\theta \mapsto \pi(z,\theta)$, there exists $\varepsilon > 0$ such that $\pi(z'',\theta) > \pi(z',\theta) \Leftrightarrow \pi(z'',\theta^*) > \pi(z',\theta^*)$ for all $\theta \in [\theta^*, \theta^* + \varepsilon]$ and z', z''. Thus the partial order on $\mathcal{A}$ induced by first-order stochastic dominance is the same for all $\theta \in (\theta^* - \varepsilon, \theta^* + \varepsilon)$. This implies that the join and meet are the same for any

a', a'', and so the construction used in the proof of Proposition 13 can make use of the same join and meet. Then the conditions of continuous bridgeability are implied by continuity of R.

Now suppose $\pi(z'', \theta^*) = \pi(z', \theta^*)$ for all $z'', z' \in B \subset Z$. Suppose that for any $\delta > 0$ there exists $\theta \in [\theta^*, \theta^* + \delta]$ and $z'', z'' \in B$ such that $\pi(z'', \theta) > \pi(z', \theta)$. Then by differentiability of π in θ, there exists a set $C \subset B$ and $\varepsilon > 0$ such that such that *i)* $\pi(z'', \theta) = \pi(z', \theta)$ for all $\theta \in [\theta^*, \theta^* + \delta]$ and all $z', z'' \in C$, and *ii)* $\pi(z'', \theta) > \pi(z', \theta) \Leftrightarrow \pi(z'', \theta') > \pi(z', \theta')$ for all $\theta, \theta' \in (\theta^*, \theta^* + \varepsilon]$ and all $z', z'' \in Z \setminus C$. Then the FOSD-induced order on $\mathcal{A}$ is the same for any $\theta', \theta'' \in [\theta^*, \theta^* + \delta]$. Moreover, this order is a superset of the FOSD-induced order at θ^*: if a'' first-order stochastically dominates a' at $\theta' \in (\theta^*, \theta^* + \delta]$ then it will also do so at θ^*. Thus for any a', a'' we can use the join and meet for the FOSD order induced by $\theta \in (\theta^*, \theta^* + \delta]$ to construct the monotone path θ^* as well. Then the conditions of continuous bridgeability are implied by continuity of R. □

C Micro-founding R

C.1 Asset market

We show here that summarizing the market through the function R is consistent with a model of information aggregation. Suppose there is a unit mass of traders. Traders receive conditionally independent signals σ_i about the state, with conditional distribution $h(\cdot|\theta)$. Assume that $h(\cdot|\theta) \neq h(\cdot|\theta')$ for all $\theta \neq \theta'$. Traders are expected utility maximizers. The payoff to trader i who purchases a quantity x of the asset when the principal takes action a, the state is θ, and the asset price is p is given by $V_i(a, \theta, x, p)$, which is assumed to be strictly decreasing in p.[18] For a fixed action a the demand of trader i who observes signal σ and knows that the state is in $\mathcal{I} \subseteq \Theta$ is given by

$$x_i(p|a, \sigma_i, \mathcal{I}) = \max_x E[V_i(a, \theta, x, p)|\sigma, \mathcal{I}].$$

Assume $p \mapsto x_i$ is strictly decreasing for all i (which is implied by assuming, for example, that $(x, p) \mapsto v_i(a, \theta, x, p)$ satisfies strict single crossing). Trader heterogeneity, both of utilities and beliefs, is allowed for, but for simplicity assume that there are are finitely many trader types, meaning finitely many distinct demand functions in the population. Normalizing the aggregate supply of the asset to zero, the market clearing condition is

$$\int_0^1 x_i(p|a, \sigma_i, \mathcal{I}) di = 0.$$

[18] For example, each trader has a strictly increasing Bernoulli utility function u_i and wealth w_i, and $V_i(a, \theta, x, p) \equiv u_i(x(\pi(a, \theta) - p) + w_i)$.

Since there is a continuum of traders and a finite number trader types aggregate demand is deterministic, conditional on the state and the principal action a. Thus we can write market clearing in state θ as

$$X(p|a,\mathcal{I},\theta)=0.$$

Let $P^*(a,\mathcal{I},\theta)$ be the unique price that clears the market.

Given any price function $\tilde{P}:\Theta\mapsto\mathbb{R}$, let $\mathcal{I}_{\tilde{P}}:\Theta\mapsto 2^{\Theta}$ be the coarsest partition with respect to which $\tilde{P}$ is measurable. We say that $\tilde{P}$ *induces* partition $\mathcal{I}_{\tilde{P}}$.

A *rational expectations equilibrium* (REE) given decision rule M consists of a price function $\tilde{P}$ such that $X(\tilde{P}(\theta)|M(\tilde{P}(\theta)),\mathcal{I}_{\tilde{P}}(\theta),\theta)=0$ for all θ. Let $\mathcal{M}$ be the set of decision rules for which there exists a REE. For any decision rule $M\in\mathcal{M}$, let $\tilde{P}_M$ be the associated REE price function.

The population distribution of signals is different for any distinct $\theta,\theta'\in\mathcal{I}$. It is therefore natural to assume that, unless all states in $\mathcal{I}$ are payoff equivalent, there will exist some pair of states $\theta,\theta'\in\mathcal{I}$ such that $P^*(a,\mathcal{I},\theta)\neq P^*(a,\mathcal{I},\theta')$.

A1. For any $a\in\mathcal{A}$ and $\mathcal{I}\subseteq\Theta$, if $P^*(a,\mathcal{I},\theta)=P^*(a,\mathcal{I},\theta')$ for all $\theta,\theta'\in\mathcal{I}$ then $P^*(a,\mathcal{I},\theta)=P^*(a,\theta,\theta)$ for all $\theta\in\mathcal{I}$.

This assumption is discussed further following the statement of the proposition.

We want to show the equivalence between implementable mechanisms and rational expectations equilbiria.

Proposition 14. *Under A1, there exists a function $R:\mathcal{A}\times\Theta\mapsto\mathbb{R}$ such that for any decision rule M there exists a rational expectations equilibrium with price function $\tilde{P}$ if and only if M implements $\tilde{P}$ given market clearing function R.*

Proof. First, we want to show that there exists an R such that for any decision rule M, if there exists a REE given M, with price function $\tilde{P}$, then M implements $\tilde{P}$ given market clearing function R. Suppose that for decision rules M_1,M_2 there exist REE, with price functions $\tilde{P}_1$ and $\tilde{P}_2$ respectively. Let $\mathcal{I}_{\tilde{P}_1}$ and $\mathcal{I}_{\tilde{P}_2}$ be the partitions of Θ induced by $\tilde{P}_1$ and $\tilde{P}_2$ respectively.

Define $R(a,\theta)=\{\tilde{P}_M(\theta):M\in\mathcal{M},\ M(\tilde{P}_M(\theta))=a\}$. That is $R(a,\theta)$ is the set of prices that can be supported as part of a REE for which the equilibrium action in state θ is a.

We want to show that R as defined above is a function. In other words, we want to show that if for some state θ, the equilibrium mixed is a under both M_1 and M_2 (that is, $M_1(\tilde{P}_1(\theta))=M_2(\tilde{P}_2(\theta))=a$), then $\tilde{P}_1(\theta)=\tilde{P}_2(\theta)$. Since $\tilde{P}_j$ induces $\mathcal{I}_{\tilde{P}_j}$, it must be that

$P^*(a, \mathcal{I}_{\tilde{P}_j}, \theta') = \tilde{P}_j(\theta)$ for all $\theta' \in \mathcal{I}_{\tilde{P}_j}$ for $j \in \{1,2\}$. Then A1 implies that $P^*(a, \mathcal{I}_{\tilde{P}_j}, \theta) = P^*(a, \theta, \theta)$ for $j \in \{1,2\}$, so $\tilde{P}_1(\theta) = \tilde{P}_2(\theta)$ as desired.

The other direction is straightforward. By the definition of implementation, if M implements $\tilde{P}$ given market clearing function R then $\tilde{P}$ is a REE price function given decision rule M. □

A1 is an assumption on the payoff structure and the information structure. It is satisfied in typical models of the asset market. For example, A1 will hold if the function $\theta \mapsto v_i(a, \theta, x, p)$ is strictly monotone for all a, p, $x > 0$, and all i; and the distribution of posteriors induced by $h(\cdot|\theta)$ is monotone (in an appropriate sense) in θ.[19]

The following are sufficient conditions for A1, along with a concrete example that satisfies these conditions. For the example, let $\mathcal{A} = [0,1]$, $v_i(a, \theta, x, p) = u(x \cdot (\pi(a, \theta) - p) + w_i)$ and assume that $\pi(a, \theta)$ is weakly increasing in θ.

1. *Ordered signals.* Assume that $h(\cdot|\theta'') >_{MLR} h(\cdot|\theta')$ for all $\theta'' > \theta'$. This implies that the posteriors induced by signals are also ordered by MLR; higher signals induce MLR higher posteriors over Θ.

   Example: $\sigma = \theta + \varepsilon$, where ε is zero-mean noise.

2. *Single-crossing between x, θ.* We want individuals to demand more of the asset when they get a high signal. Assume therefore that $V_i(a, \theta, x, p)$ satisfies single crossing between x and θ. Monotonicity of demand is implied by standard MCS results (see Athey (2001)).

   Example: $u(x(\pi(a, \theta) - p) + w_i)$ satisfies single crossing in x and θ when $\theta \mapsto \pi(a, \theta)$ is increasing.

3. *Payoff equivalence.* For any $\mathcal{I}$, we want demand to be strictly increasing in θ unless $V_i(a, \theta, x, p) = V_i(a, \theta', x, p)$ for all $\theta, \theta' \in \mathcal{I}$.

   Example: This holds given the assumptions made thus far (in particular, monotonicity of π).

[19] A sufficient condition for the monotonicity of $\theta \mapsto v_i(a, \theta, x, p)$ is co-monotonicity of $\theta \mapsto v_i(a, \theta, x, p)$ for all a (when $v_i(a, \theta, x, p) \equiv u_i(x \cdot (\pi(a, \theta) - p) + w_i)$ this is equivalent to co-monotonicity of $\pi(a, \cdot)$). Posterior monotonicity will hold, for example, if $\sigma = \theta + \delta$ for some continuously distributed zero mean random variable δ.

C.2 Forecasts and macro aggregates

Many policy decisions are made with reference to macroeconomic outcomes. For example, the government may decide to increase the amount of unemployment benefits or fund worker-retention programs depending on initial jobless claims or the unemployment rate. Many such problems also have a dynamic component. For example, businesses deciding whether or not to fire employees may care about the future unemployment rate both as a signal of demand and as a determinant of government worker-retention policies. In such settings, forward looking agents often make use of expert forecasts of the relevant macro variables, such as the unemployment rate.

C.2.1 Policy decision up-front

The policy maker may prioritize timeliness over accuracy when making certain policy decisions. In such cases it will be necessary for the policy maker to take an action before the relevant aggregate outcome has been realized. The policy maker will therefore make use of expert forecasts. For example, consider the problem of the government choosing the level of unemployment benefits. The policy maker may wish to act before relevant data, such as the unemployment rate in the coming month, has been collected. It must therefore relay on forecasts of the relevant variables. For simplicity, assume that the government conditions its benefits policy exclusively on expert forecasts of the unemployment rate for the coming month (it is straightforward to incorporate other sources of information).[20]

Forecasters wish to provide accurate estimates of the aggregate outcome (We will refer to this simply as the outcome from now on). If there are many forecasters, each individual expects their prediction to have only a small effect on overall expectations.[21] However they recognize that overall expectations will be used by the policy maker to take an action. These two factors imply that forecasters' private information will shape their expectations of policy decisions, which in turn will affect their forecasts.

This situation is easiest to model if we assume that forecasters observe each others' forecasts, and can make revisions based on what others say. We won't get into why might still "agree to disagree" even when they observe eachothers' forecasts. The consistency condition is that each forecaster doesn't want to change their forecast given those of the others, and the announced policy rule. Assume that for any fixed action by the policy makers such a fixed point exists. Then we are back to the original situation.

[20]Another example is inflation expectations. Forward guidance could destroy the informativeness of the signal.

[21]Bloomberg surveys around 80 economists for predictions on the monthly unemployment rate.

Formally, this model is very similar to the market price model. Assume there is a continuum of forecasters $\mathcal{F}$. Each forecaster $i \in \mathcal{F}$ receives a signal σ_i about the state. Signals are conditionally independent across forecasters. Forecasters make predictions about the value of some variable v, which will not be realized until after the principal has taken an action. Forecastes may have different models of the world, i.e. ways to map their information to a prediction, but assume for simplicity that there are only finitely many models in the population.

The principal bases their decision on some real-valued function of the profile of forecasts, the forecast aggregate. Forecasters iteratively revise their predictions based on their observations of the forecast aggregate. Thus we are looking for a rational expectations equilibrium conditional on the principal's announced decision rule M. In this context, assumption $A1$ can be restated as follows.

Fixed any principal action a, $\mathcal{I} \subseteq \Theta$ and $\theta \in \Theta$, and value of the forecast aggregate f. Assume all forecasters know that the principal will take action a, that the state is in $\mathcal{I}$, and that the value of the forecast aggregate is f (in addition the their private signals). Let the $X(f|a, \mathcal{I}, \theta)$ be the new value of the forecast aggregate after forecasters have a chance to revise their predictions. This is a deterministic function since there are a continuum of forecasters with i.i.d. signals. Then forecasts reach a fixed point when

$$X(f|a, \mathcal{I}, \theta) = f.$$

Assume that there is unique fixed point for any $a, \mathcal{I}, \theta$ (which will be the case, for example, when individual forecasts, as well as the aggregator, are monotone in f), and denote this fixed point by $F^*(a, \mathcal{I}, \theta)$.

A1'. For any $a \in \mathcal{A}$ and $\mathcal{I} \subseteq \Theta$, if $F^*(a, \mathcal{I}, \theta) = F^*(a, \mathcal{I}, \theta')$ for all $\theta, \theta' \in \mathcal{I}$ then $F^*(a, \mathcal{I}, \theta) = F^*(a, \theta, \theta)$ for all $\theta \in \mathcal{I}$.

Assumption A1' is satisfied, for example, when the distribution of beliefs induced in the population is monotone (in an FOSD sense, with an appropriate order on beliefs) in the state, individuals forecasts are monotone in their beliefs, and the forecast aggregate is monotone in individual forecasts (in an FOSD sense).

The existence of R in this setting follows from Proposition 14

C.2.2 Policy decision ex-post

Some policy decisions may condition on realized outcomes, rather than expectations. For example, Congress extended the time frame for spending PPP funds after observing that companies had difficulty re-hiring employees. Congress also approved a second tranche of PPP funds after the first was exhausted. In this case companies will condition their payroll decisions or loan applications on expectations of future aggregate outcomes. Forecasters make predictions knowing that *i*) expectations will shape business decisions, and *ii*) business decisions will shape the policy response. Again, assume forecasters observe each others' forecasts. Then we need a fixed point that takes into account the feedback of forecasts on policy through business decisions.

Formally this case is similar to that discussed above. There are two periods. Some variable v will realize in the second period, and the principal will take an action in the second period conditional on v. Assume that the principal commits to a rule M mapping v to an action (We will discuss why later).

A unit mass of economic agents, call them individuals, care about the principal's future action, as well as some underlying state θ. In order to predict what the principal's action will be, individuals rely on the predictions of a set $\mathcal{F}$ of forecasters. Individuals are fairly simplistic: they aggregate forecaster predictions in some way, and assume that this forecast aggregate f will be the true value. They choose their actions based on the action implied by the principal's decision rule, as well as their own private information. Assume that individuals do not infer anything about the state from the forecasters' predictions.[22] The actions off all individuals, along with the state, jointly determine the outcome v. When all individuals expect the principal to take action $a \in \mathcal{A}$ and the state is θ, the aggregate outcome in the second period will be given by $J(a, \theta)$.

As before, there is a unit mass of forecasters, each of whom receives a private signal about the state. Forecasters observe the current value of the forecast aggregate and revise their decisions. As before a fixed point is reached when $X(f|a, \mathcal{I}, \theta) = f$. This function incorporates the fact that a affects the aggregate outcome through $J(a, \theta)$. Assuming $A1'$ holds, we have an R function by Proposition 14.

The interesting part of the ex-post decision model is that principal is not intending to use the equilibrium variable, in this case the forecast, to make a decision. The principal may not even be able to commit to a mapping M from the aggregate outcome to an action. It could just be that agents anticipate the principal to behave in a certain way ex-post. Nonetheless, the forecast will be determined as a fixed point, and this will impact the aggregate outcome, and thus the principal's decision.

[22]**NOTE:** it should be possible to add this inference without too much complication, probably with a version of assumption A1.

C.2.3 Alternative model

Suppose that there is a single forecaster who gets a signal σ. The forecaster is aware of the effect that their prediction will have on individual behavior. The forecaster simply reports their expectation of the outcome v, when this is well defined. This will be well defined iff there is a fixed point to the function $f \mapsto \mathbb{E}[J(M(p),\theta)|\sigma]$. Let $R(a,\sigma) = \mathbb{E}[J(M(p),\theta)|\sigma]$. The analysis of the paper applies, with σ replacing R.

C.2.4 Adding forecast uncertainty

The fact that forecasters receive conditionally independent signals may seem unrealistic. It is straightforward to generalize to a situation in which signals are correlated. Assume that the state consists of a pair (κ,θ). As before, θ is the payoff-relevant state. κ simply determines the joint distribution of signals. Signals are conditionally independent given (κ,θ).

Let Σ be the space of population signal profiles $\{\sigma_i\}_{i\in\mathcal{F}}$. Assume that there is a complete order on the space of signal profiles, which can be represented by a bijection $b : \Sigma \mapsto [0,1]$.[23] Since b is a bijection, $b(\{\sigma_i\}_{i\in\mathcal{F}})$ contains the same information as $\{\sigma_i\}_{i\in\mathcal{F}}$. Then when all forecasters expect the principal to take action a and know that the current forecast aggregate is f, and know that $b(\{\sigma_i\}_{i\in\mathcal{F}}) \in \mathcal{I} \subseteq [0,1]$, then the updated forecast aggregate will be $X(f|a,\mathcal{I},b)$. Then the analysis proceeds as before, except that b replaces θ. The principal will have to account for the residual uncertainty when choosing a decision rule.

The discussion in this section applies when the set $\mathcal{F}$ of forecasters is finite. However the assumption of a continuum of forecasters remains convenient for two reasons. First, the assumption that such a bijection b exists makes more sense when there is a continuum of forecasters (see the example in the footnotes). Second, when there are finitely many forecasters they will behave strategically. For example, there in general no reason to expect that forecasts should reach a fixed point when forecasters take into the effect that their forecasts have on the forecast aggregate. For example, with a single forecaster trying to minimize the expected difference between their prediction and the actual outcome, unless they perfectly observe the state, it may be optimal for them to make a forecast that they know cannot be correct.

[23] For example, $\sigma_i = \varepsilon_i + \kappa + \theta$, where ε_i are i.i.d. random variables with a common bounded support distribution. In this case the order on the set of signal profiles is given by the population mean $\int_{i\in\mathcal{F}} \sigma_i di$.