Market-Based Mechanisms*

Quitzé Valenzuela-Stookey† Francisco Poggi‡

August 25, 2021

Click here for the latest version

Abstract

This paper studies the general problem of a principal who conditions their actions on aggregate market outcomes as a proxy for an unobserved payoff-relevant state. Agents in the market have private information about the state, and their choices reflect both their beliefs about the state and their expectations of the principal's actions. We fully characterize the set of joint distributions of market outcomes, principal actions, and states that can be implemented in equilibrium. We focus in particular on implementation under constraints imposed by concerns about manipulation and equilibrium multiplicity. This characterization of the feasible set admits a tractable representation, and significantly simplifies the principal's design problem. We apply our results to study carbon credits, corporate bailouts, and monetary policy.

Keywords: Feedback effect, Market manipulation, Equilibrium multiplicity, Price informativeness, Rational expectations equilibrium.

*We are grateful to Marciano Siniscalchi, Jeff Ely, Eddie Dekel, Piotr Dworczak, Alessandro Pavan, Ludvig Sinander, and Gabriel Ziegler for discussion and feedback. We also thank seminar participants at Northwestern University, EWMES 2020, Bocconi University, California Institute of Technology, Cornell University, Carnegie Mellon, LSE, Microsoft Research, NYU Stern, Penn State, Toulouse School of Economics, UC Berkeley, UC Davis, University of Manchester, University of Mannheim, University of Pennsylvania, Vanderbilt University, and 2021 MaCCI Summer Institute in Competition Policy for helpful comments.

†Department of Economics, Duke University, Durham NC 27708, U.S.A., quitze.valenzuelastookey@duke.edu.

‡Department of Economics, University of Mannheim, Mannheim 68161, Germany.

1 Introduction

One of the fundamental insights of information economics, going back at least to Hayek (1945), is that market outcomes can aggregate dispersed information. As a result, policy makers facing uncertainty often use, or are encouraged to use, market outcomes, such as prices in financial markets, to inform their decisions. However, the use of market outcomes to inform policy making is complicated by the fact that the policy maker's own action may have a significant impact on the market in question. Market participants anticipate the policy maker's action, and this influences the market outcome. This creates a feedback loop between actions and market outcomes, which constrains the policy maker's ability to learn from the market. Market-based policies may also be vulnerable to manipulation by market participants. Moreover, the dependence of market outcomes on expectations of the policy maker's endogenously determined action can lead to multiplicity of equilibria, and potentially non-fundamental market volatility (Woodford, 1994).

This paper studies the general problem of using market outcomes to inform decision making in settings with feedback effects. To fix ideas, consider a regulator who wishes to limit firms' carbon emissions. The regulator wants to reduce emissions more aggressively if the cost to firms of reducing emissions, known as their abatement cost, is low. Abatement costs, however, are private information of the firms.

There is a long running debate, going back to Weitzman (1974), over whether it is more efficient to reduce emissions by setting quantities via caps, or via prices, for example by taxing emissions.[1] Cap-and-trade policies, employed in many countries, are a way of setting the overall quantity of emissions and allowing the market to allocate the volume across firms. Under such a policy, the regulator issues a fixed number of carbon credits, which are traded by firms in a credits market. Each credit entitles the holder to emit one ton of carbon.

Under such a policy, the credit price reveals important information about abatement costs.[2] It is natural to seek policies which respond to this information. One way to incorporate such information directly into emissions policy is through a "variable-volume credit policy".[3] Under such a policy, the regulator issues a fixed number of credits. In contrast to the standard cap-and-trade policy however, each credit does not correspond to a fixed emissions-volume allowance. Instead, the regulator announces a rule specifying the per-credit emissions-volume allowance as a

[1] These two methods are equivalent if the marginal costs and benefits of reducing emissions are known by the regulator, but are generally not otherwise. See Weitzman (1974).

[2] In particular, low credit prices indicate low abatement costs. As the socially optimal level of emissions is increasing in the abatement cost, the regulator who observes lower-than-anticipated prices will thus learn ex-post that the number of credits issued was too high. Consistent with such inferences, it has been observed that lower than expected credit prices create political pressure for regulators to reduce the size of future credit issuance (Flachsland et al., 2020). Indeed, some systems, such as the EU Emissions Trading Scheme, have mechanisms in place to try to control prices or adjust issuance in response to market conditions, and there is a movement towards expanding these mechanisms.

[3] Karp and Traeger (2021) also propose such a policy, which they refer to as a "smart cap".

function of the credit price. As before, credits then trade in a credits market. After the market closes, the per-credit volume is determined by the credit price and the rule announced by the regulator. In this way the regulator is able to control both prices *and* quantities, and, hopefully, make the emissions cap responsive to the unknown abatement costs.

The difficulty, from the regulator's perspective, is that the information about abatement costs revealed by credit prices depends on the joint distribution of prices and abatement costs, which is an equilibrium object. In particular, firms' value for credits, and thus the credit price, will depend on the anticipated per-credit volume, as well as their private information. As a result, the regulator's choice of decision rule mapping credit prices to the per-credit volume will itself shape the information revealed by credit prices. This feedback effect must be accounted for by the regulator when choosing their decision rule.

The results of this paper allow us to fully characterize the set of joint distributions of states, credit prices, and emissions quantities that the lender can implement by using a variable-volumes credit policy. More importantly, we characterize what joint distributions can be implemented as the unique equilibrium outcome while also preventing manipulation by firms. We return to the emissions-regulation example in detail in Section 6.1, where we show that the regulator can in fact use such a policy to achieve the first-best outcome.

Model

The example of the emissions regulator illustrates the key components of the general model that we study in this paper. A principal (the regulator) needs to choose an action (the emissions volume), but faces uncertainty regarding a payoff-relevant state of nature (firms' abatement costs). There is a market populated by agents (firms) who may have some private information about the state. The behavior of agents in the market depends on their beliefs about the state, as well as the anticipated action of the principal. The joint behavior of agents determines an aggregate market outcome (the price of emissions credits). For clarity, we refer to the market outcome as a price.

To exploit information aggregated by the market, the principal publicly commits in advance to a decision rule mapping the price to their action. This decision rule is the principal's design instrument.[4]

The general model of market-based policy making described above nests many of the instances considered in the literature. Our conceptual innovation, relative to this literature, is to bring a design and implementation perspective to the general problem of market-based policy making. From a design perspective, it is natural to fist ask what exactly the principal achieve by using such a market-based decision rule. In particular, what joint distributions of states, actions, and prices can the principal induce, or implement, in equilibrium?

There are, however, additional practical concerns that need be accounted for when designing

[4] We focus primarily on a principal who has commitment power, but briefly discuss the implications of our results for a model without commitment power in Appendix E

policy. We focus on what we believe are the two most salient issues: equilibrium multiplicity and market manipulation. If the principal's decision rule induces multiple equilibria then outcomes may be indeterminate, and this can lead to non-fundamental market volatility. Indeed, the endogeneity of the principal's action can exacerbate issues of equilibrium multiplicity (Bernanke and Woodford, 1997). As a result, there is great interest in designing policies for which a unique equilibrium outcome exists (Woodford, 1994).

The dependence of the principal's action on the price also opens the door for manipulation of the principal's decision via the market. For example, firms may distort their trade of carbon credits in order to induce the regulator to raise the per-credit volume. Market manipulation will be a concern, even in large markets, if agents in the market can significantly impact the principal's action by triggering small perturbations to the price.

Our goal is to understand how do deal with these concerns when designing policy, and in particular, identify what constraints they impose on the set of equilibria that the principal can implement. These are the central objectives that motivate this paper.

Contribution

The current paper makes four major contributions relative to the existing literature. First, we provide a general framework for studying market-based interventions in environments with feedback effects. We show that in a wide range of markets, equilibrium outcomes can be succinctly summarized via a reduced-form representation. By working directly with such a representation, we are able to derive general results which are applicable in diverse settings.

Second, we use this framework to fully characterize the feasible set in outcome space. In particular, we characterize the set of implementable joint distributions of states, principal actions, and prices. More importantly, given the concerns about manipulation and equilibrium multiplicity, we characterize what joint distributions can be induced as the *unique* equilibrium using a decision rule that satisfies a notion of *robustness to manipulation*. We refer to such equilibria as *continuously uniquely implementable* (CUI). We view this characterization as our primary contribution. The characterization admits a tractable representation, which significantly simplifies the analysis of optimal policy in applications (see Section 6.2 for a simple illustration).

Existing analyses of market-based policy design optimize over the space of decision rules. Generally, this approach requires one to impose restrictions on the environment and/or the admissible decision rules which make it possible to solve for equilibrium in closed form. In contrast, optimizing in the space of implementable joint distributions (described by maps from states to actions and prices) increases tractability. We are able to identify qualitative features of optimal policy even when a closed form solution is not available. We show that in many applications, it is sufficient for the principal to choose the map from states to actions to maximize their objective, subject only to the constraint that the induced map from states to prices be monotone. Moreover, we are able to highlight the cost imposed by restrictions on the decision rule which are sometimes used in the literature for tractability purposes (Section 6.3).

On a conceptual level, this characterization also reveals a surprising interaction between unique implementation and the constraints imposed by concerns about market manipulation. Singly, neither set of constraints imposes a substantive restriction on the implementable set. However jointly they have important implications for what the principal can achieve. To our knowledge, we are the first to consider these constraints jointly.

Third, we show that the constraints of unique implementation and robustness to manipulation imply a natural notion of robustness to model misspecification (Section 5.1). This means the principal's payoff is not highly sensitive to their potentially limited understanding of market fundamentals. Finally, the results also allow us to analyze optimal policy when the requirement of unique implementation is relaxed. In particular, we use our characterization of the implementable set to show that if the principal takes a worst-case approach to equilibrium multiplicity then the restriction to unique implementation is generally without loss of optimality. We also discuss optimal policy under alternative criteria for evaluating multiple equilibria.

Applications

The use of market outcomes to inform policy decisions occurs in many settings. Bank regulators may use market prices of bank securities to inform an intervention decision (Greenspan, 2001). Central banks condition monetary policy on macroeconomic indicators such as the unemployment rate or the rate of inflation. A fall in a company's stock price can prompt shareholder action to replace top management (Warner et al., 1988). Moreover, there is growing interest in "rules-based" policy, in which policy is conditioned on measurable outcomes in a pre-determined way, prompted in part by the slow and disjointed response to the current COVID-19 crisis. For example, it is argued that state-contingent debt instruments, in which payments are conditioned on variables such as GDP or commodities prices, should be used to reduce the need for protracted and costly sovereign debt restructurings (Cohen et al., 2020).

In Section 6 we study three applications in detail. In Section 6.1 we study the variable-volume credits policy for regulating emissions, described above. We show that such a policy can uniquely implement the regulator's first-best outcome in a way that is robust to market manipulation.

In Section 6.2, we apply our results to the problem of a government considering a bailout of a firm or industry. The government uses the firm's stock price to inform its decision.[5] We show that in this setting, the government's first best policy is CUI if and only if the positive social externalities from bailing out the company are high. In this case, the optimal decision rule involves a gradual transition from a large to a small bailout as the firm's stock price increases. We also characterize the optimal CUI policy when first-best is not feasible, which in this case involves a rapid reduction of the level of support as a function of the stock price, and show which policies are optimal when

[5]A related application is performance pricing in debt contracts, whereby the interest rate is conditioned on the borrower's financial ratios, e.g. interest coverage, or credit ratings (Grochulski and Wong, 2018).

the uniqueness requirement is relaxed.

In Section 6.3, we discuss the distinctive features of settings in which the principal attempts to "move against the market". For example, central banks often use open market operations (the principal's action) to reduce the interest rate (the price) during severe crises, while, absent interventions, interest rates would be increasing in the severity of the crisis (the state). We show that in such settings, it is necessary for the principal to use a non-monotone decision rule in order to avoid equilibrium multiplicity. This highlights the cost of placing ex-ante restrictions on the decision rule, for example restricting attention to linear decision rules, as is common in the literature for reasons of tractability. In the central bank example, restricting attention to monotone decision rules implies that the bank cannot induce lower interest rates when the crisis is more severe without also being vulnerable to non-fundamental volatility. This restriction comes from a surprising interaction between equilibrium multiplicity, monotonicity of the principal's decision rule, and bounds on the set of actions available to the principal (e.g. the size of asset purchases/sales). By allowing for more general decision rules we show that the central bank can uniquely implement essentially any decreasing map from the state to the interest rate. This application demonstrates the value of our characterization of the entire feasible set, beyond simply facilitating the search for optimal policies.

Related literature

The current paper is closely related to the literature on 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 documented the possibility of multiplicity of equilibria (see, among others Dow and Gorton (1997), Bernanke and Woodford (1997), and Angeletos and Werning (2006)). Closely related is the literature on prediction markets and conditional decision markets, e.g. Teschner et al. (2017), in which a principal conditions their actions on a market outcome. The current paper brings a design approach to policy making in these settings, formalizing the problem of policy design under commitment in a general setting and providing a full characterization of feasible policy outcomes while accounting for manipulation and equilibrium multiplicity concerns. Such a characterization is absent from the literature.

This paper is also related to the literature studying market-based intervention in the presence of feedback effects *without* commitment. Bond et al. (2010) study a problem similar to the emergency lending example of Section 2.4, but where the principal does not have commitment power. They identify that there cannot be an equilibrium in which the principal's first-best is achieved exactly in the situation in Figure 1b, when the induced price function would be non-monotone. In the language of the current paper, this is because the induced price and action functions violate the necessary measurability condition for implementability; the action must be measurable with respect to the price.[6] However, we show that if the principal is concerned with equilibrium multiplicity

[6] Bond et al. (2010) then observe that if the principal has access to a signal with a sufficiently narrow bounded support around the true state, they will be able to differentiate between high and low states which induce the same

and manipulation, then non-monotonicity of the price is problematic even if there is no violation of measurability, for example as depicted in Figure 2a. Identifying that monotonicity of the price is necessary for unique implementation under robustness to manipulation is one of our primary contributions.

The most closely related work in the literature without commitment is Siemroth (2019), which studies a noisy REE market with a principal who learns from the asset price, similar to the setting in Section 4.7, and identifies conditions under which a rational expectations equilibrium exists when the principal lacks commitment. In contrast, we fully characterize unique implementability and solve for optimal policies under commitment.[7] Moreover, Siemroth (2019) restricts attention to equilibria in which the price function is continuous (not to be confused with continuity of the principal's decision rule). This is a substantive assumption, as it implies that the equilibrium, when it exists, is unique. In a noisy REE model *without* feedback effects, Pálvölgyi and Venter (2015) and Breon-Drish (2015) show that in general multiple equilibria are possible. Uniqueness holds only within the class of equilibria with continuous price functions. Multiplicity that arises even without feedback effects, for example if the principal does not condition on the price, can be called *fundamental multiplicity*. However, in settings with feedback effects there may also be equilibrium multiplicity caused by the endogeneity of the principal's action. Eliminating fundamental multiplicity by imposing continuity of the (endogenous) price function also eliminates multiplicity caused by action endogeneity. Moreover, it does so by imposing exogenous restrictions on the principal's policy.[8] The present paper's contribution is in characterizing the set of implementable outcomes; we do not restrict this set ex-ante by imposing continuity of the price function. Instead, we characterize the restrictions on the set of implementable outcomes imposed by robustness to multiplicity and manipulation.

Other papers have noted that policy based on market outcomes may be vulnerable to manipulation. Goldstein and Guembel (2008) study 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 Lee (2019), the discontinuous nature of the policy considered opens the door to manipulation. Motivated by these concerns, we characterize robustness to manipulation in the limit as agents in the market becomes small, and consider policies that are robust to manipulation in this sense.

This paper relates most directly to the literature on policy making under commitment in the

price, and thus overcome the measurability problem.

[7]Other important differences between the current paper and Siemroth (2019) are discussed in Section 4.7.

[8]Without commitment, one could argue that if the principal's best response is suitably continuous, it is natural to focus on equilibria with continuous price functions. This is not the case with commitment however; the principal may wish to commit to a policy that induces a discontinuous price function, even if the first-best action function is continuous.

presence of feedback effects. Important contributions include Bernanke and Woodford (1997), Ozdenoren and Yuan (2008), Bond and Goldstein (2015), Glasserman and Nouri (2016), Boleslavsky et al. (2017), and Hauk et al. (2020). Bernanke and Woodford (1997) show 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 the 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 shows that non-monotone decision rules may in fact be *necessary* to prevent multiplicity (Section 6.3). Bond and Goldstein (2015) focus on the how market-based interventions affect the efficiency of information aggregation by prices. In contrast to the current paper, traders in Bond and Goldstein (2015) care about the state only insofar as it allows them to predict the government's action. As a result, information aggregation is highly dependent on the decision rule.

Glasserman and Nouri (2016) show how equilibrium multiplicity issues that arise in a static setting may not be present in a dynamic trading model. In a static problem nearly identical to that depicted in Figure 1a, they show that equilibrium multiplicity will arise if the principal uses a discontinuous threshold rule. Restricting attention to such rules, they show that in a dynamic version of the model there may be a unique equilibrium. We observe that in this type of problem, in which the price function is monotone, the multiplicity issue can also be resolved by allowing for gradual adjustment of the principal's action. Our main focus however is on identifying what conditions are *necessary* for unique implementation.

Hauk et al. (2020) develop variational techniques for identifying optimal decision rules in settings with feedback effects. These techniques complement our results, which simplify the problem of identifying optimal policies by characterizing the feasible set in the space of action and price functions, rather than the space of decision rules.

In general, our primary contribution relative to this literature is the complete characterization of implementable outcomes, taking into account the practical concerns of equilibrium multiplicity and manipulation. Moreover, by providing a tractable framework for studying flexible market-based policy design in a general setting, we avoid the artificial restrictions imposed by some of the simplifying assumptions used in the literature.

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, both when the state space is one-dimensional and for multi-dimensional state spaces. Section 4 discusses different microfundations for the type of markets studied in the paper, namely those that admit a reduced-form representation. Section 5 analyses properties of the solution concept and extensions, such as what happens to optimal policy when the unique implementation restriction is relaxed. Finally, Section 6 explores the applications to carbon markets, bailouts, and monetary policy. In Section 7 we briefly discuss directions for future work.

2 Model

The baseline model consists of the following primitive objects.

i. The convex state space, denoted by $\Theta \subseteq \mathbb{R}^N$.

ii. A convex and compact 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 market outcomes.

For clarity, we 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 is discussed below. The state may contain dimensions that are not directly payoff relevant for the principal.[9]

The principal chooses a *decision rule* $M : \mathcal{P} \to \mathcal{A}$, which specifies the action taken by the principal as a function of the market outcome. The principal publicly commits to their decision rule.

The final primitive feature of our model is a market. In general, by market we mean a game played by a set of agents, or market participants, and a solution concept with the defining feature that in an equilibrium there is a price $p \in \mathcal{P}$ associated with each state. The market game is played following the principal's announcement of the decision rule M. Formally, this means that, given a decision rule M, for any equilibrium in the market there exists a *price function* $P : \Theta \to \mathcal{P}$ describing the price which realizes in each state. The specific types of markets that are accommodated by our analysis are discussed in subsequent sections.

In summary, the timing of interaction is as follows.

1. The principal publicly commits to a decision rule $M : \mathcal{P} \to \mathcal{A}$ specifying an action for each price.[10]

2. The market game is played and a price is determined.

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

2.1 Analysis

Our focus is on characterizing the set of equilibria which the principal can induce using a (market-based) decision rule. For the majority of this paper, we do not make explicit reference to the principal's preferences over equilibria. These we discuss only in the context of the applications studied in Section 6. Thus our characterization results can be viewed as the first step in solving the principal's design problem.

[9] For example, in a noisy REE model of an asset market, as in Grossman and Stiglitz (1980), the state will include the supply shock, in addition to the payoff relevant state. More generally, the state can represent the entire profile of agent's private signals, as in Jordan (1982).

[10] In Appendix E we briefly discuss the version of the model in which the principal cannot commit.

Our analysis can be divided into three steps. Conceptually, although not formally, these steps are analogous to those of Myerson (1981). First, we redefine the problem in outcome space, focusing directly on the set of implementable equilibria. This is analogous to recasting the problem in terms of direct revelation mechanisms in the classical mechanism design setting. Second, we show how to represent equilibrium outcomes in a way that facilitates a state-by-state analysis of the principal's problem. We do this be deriving a reduced-form representation of equilibrium in the original market. In Myerson (1981), state-by-state analysis is facilitated by writing the principal's payoff in terms of virtual values, which subsume the global IC constraints. Finally, we characterize the set of implementable outcomes under various policy desiderata. This produces a set of constraints on the principal's problem which significantly increase tractability.

Implementation in outcome space

Rather than study explicitly the principal's choice of decision rule, we instead focus directly on the objects of interest: the induced equilibrium outcomes. For any market equilibrium induced by a decision rule M, there exists a price function $P : \Theta \to \mathcal{P}$ describing the price which realizes in each state. Since the principal commits to their decision rule, an equilibrium in the market also induces an *action function* $Q : \{\Theta \mapsto \mathcal{A}\} \equiv M \circ P$. Our interest is in the set of implementable action and price functions, which also describe the set of joint distributions over states, actions, and prices.

Definition. *A price function P is implemented by a decision rule M if it is an equilibrium price function given M. A pair of action and price functions (Q, P) is **implemented** by M if P is implemented by M and $Q = M \circ P$.*

We say that a pair (Q, P) is *implementable* if it is there exists a decision rule that implements it. The condition $Q(\theta) = M(P(\theta))$ makes backing out a decision that implements (Q, P) straightforward.[11] The notion of equilibrium which determines which (Q, P) are implementable of course depends on the specific market in question.

There are a number of benefits of working directly in the space of action and price functions, rather than the space of decision rules, for the purposes of designing policy. For one, it is convenient to think about equilibrium multiplicity in this space. As discussed in the introduction, we are interested in unique implementation, which means characterizing the set of action and price functions that can be implemented as the unique equilibrium actions for some decision rule. Additionally, given principal preferences over equilibria, it is easier to optimize in the space of action and price

[11] In general, this condition only defines M on $P(\Theta)$. However, in the types of markets on which we focus (those which admit a reduced-form representation, as defined below) if (Q, P) is implementable then they are implemented by M as long as $Q(\theta) = M(P(\theta))$ on $P(\Theta)$. When additional desiderata are imposed on the decision rule, for example that it induce a unique equilibrium, then some care may be needed when specifying M on $\mathcal{P} \setminus P(\Theta)$. See the proof of Theorem 2.

functions, rather than directly in the space of decision rules. Finally, working in this space makes it easy to impose state-contingent constraints on the principal's actions. While we do not focus on this in the current paper, our results extend immediately to settings in which the principal can only commit to policies in which their equilibrium action in state θ belongs to some subset $a(\theta) \subsetneq \mathcal{A}$. Such constraints may be needed to accommodate settings in which the principal cannot realistically commit to take certain actions in some states.[12]

Reduction

As described above, one of the primitive features of the model is a market in which the price is determined. For the purposes of characterizing implementable outcomes and designing policy, it is helpful to impose additional structure on the market. Consider the variable-volume credits policy example. In this market, each firm's demand for credits depends only on their own abatement cost and the anticipated per-credit allowance (as well as the credit price). Thus, the equilibrium price the market can be described as a function of the anticipated per-credit volume and the vector of abatement costs (which is the state in this setting). Our approach is to show that in a wide range of markets, equilibrium outcomes can be summarized in a concise way via a reduced-form representation along these lines. Once this representation has been identified, we can use it to characterize the feasible set and design policy.

Definition. *The market* ***admits a reduced-form representation*** *if* $\exists$ *a function* $R : \mathcal{A} \times \Theta \to \mathcal{P}$ *such that for any* Q, P, M*, the pair* (Q, P) *is implemented by* M *iff for all* $\theta \in \Theta$

i. $Q(\theta) = M(P(\theta))$ *(commitment)*

ii. $P(\theta) = R(Q(\theta), \theta)$ *(market clearing)*

The commitment condition is clearly necessary in any market for Q and P to be equilibrium outcomes given M. The additional structure comes from the market clearing condition. This condition describes the effect that the principal's actions have on the market. The economic content of this condition can be separated into two parts. First, R is a function from $\mathcal{A} \times \Theta$ to $\mathcal{P}$, not from $\mathcal{A}^{\Theta} \times \Theta$ to $\mathcal{P}$. In other words, the equilibrium price at state θ depends only on the equilibrium action in that state, and not on global features of the equilibrium action function. This greatly simplifies the problem of characterizing the set of implementable outcomes, as it reduces the number of global constraints on the action and price functions. Second, the market-clearing function does not vary with M. In other words, the relationship between Q and P, the equilibrium outcomes, does not depend explicitly on the decision rule M. Thus, we can determine whether Q and P can be equilibrium outcomes for some decision rule, without making explicit reference to

[12]We could also easily impose price-contingent constraints on the principal's action. However these constraints could also be easily imposed when working directly in the space of decision rules.

which decision rule M will implement them.[13]

For the purpose of characterizing implementable equilibria, we will restrict attention to markets that admit a reduced-form representation. In Section 4, we discuss our approach to modeling the market in detail. We derive conditions under which a reduced-form representation can be identified in a wide range of markets, and clarify what types of markets will not admit such a representation.

Characterization

Using the reduced-form representation, we turn to characterizing the implementable set. Without further constraints, this set is easily identified.

Observation 1. *If the market admits a reduced for representation with market-clearing function R, (Q, P) is implementable iff*

1. $Q(\theta) \neq Q(\theta') \quad \Rightarrow \quad P(\theta) \neq P(\theta')$. *(measurability)*
2. $P(\theta) = R(Q(\theta), \theta) \quad \forall\, \theta \in \Theta$ *(market clearing)*

The measurability condition says that if the action in two states is to differ, then so must the price. This condition guarantees that there exists a P-measurable function M that induces action function Q. Clearly if this condition is violated there can not exist such M. Given an implementable (Q, P), the implementing decision rule can be easily identified. Measurability implies that the set $Q\left(P^{-1}(p)\right)$ is either empty or singleton; this defines M on $P(\Theta)$.

Observation 1 fully characterizes the set of implementable price and action functions. However, this characterization ignores manipulation and multiplicity considerations which are central to the policy design problem in many settings. When such constraints are taken into account a more meaningful characterization of the set of implementable equilibria can be given. In the following section we introduce what we see as the most important practical considerations for policy design. We provide a characterization of implementablity under the constraints imposed by these concerns. This characterization yields a tractable set of feasible policies, and is therefore useful for designing policy. We demonstrate the utility of this approach in the applications of Section 6.

2.2 Practical constraints

We focus on what we view as the two most important constraints which a policy maker might consider.

Manipulation

A salient concern in many market-based policy-making environments is that agents in the market may attempt to to manipulate the price in order to influence the principal's action. An

[13]This is not to say of course that different decision rules are not needed to implement different action and price function.

agent may manipulate the price by buying/selling an asset, releasing false information, or other means.[14] While agents are generally assumed to behave as price takers in the market models we consider (those that admit a reduced-form representation, as discussed in Section 4), we view the price-taking assumption as an idealization of a world in which agents are small, but may have some non-zero market power. The ability of a small (but not infinitesimal) agent to manipulate the principal depends on the sensitivity of the principal's decision rule mapping prices to actions. If, for example, the decision rule is discontinuous, then an agent will be able to induce a significant change in the principal's action by manipulating the price, even if their individual price impact is small.

In order to maintain consistency between the idealized model in which agents are price takers and one in which agents are small, but may have a non-zero price impact, it seems natural restrict the principal to use a continuous decision rule. However the restriction to everywhere-continuous decision rules is stronger than is needed to address these concerns. As Theorem 3 shows, it is enough to have continuity in the neighborhood of any equilibrium price to guarantee robustness to small perturbations to market fundamentals. Similarly, if a discontinuity in M occurs at a price which is far from any which could arise in equilibrium then manipulation via a small price impact will not be possible. We therefore allow for discontinuities in the decision rule, provided they do not occur near equilibrium prices.

Formally, fix a market and let the principal's decision rule be M. For each state θ, let $P_M(\theta)$ be the set of prices that are equilibrium outcomes in this market when the state is θ, given decision rule M. Let $\bar{P}_M := P_M(\Theta)$ be the set of all possible equilibrium prices given M, and let $cl(\bar{P}_M)$ be the closure of this set.[15]

Definition. *A function $M : \mathcal{P} \to \mathcal{A}$ is **essentially continuous** if it is continuous on an open set containing $cl(\bar{P}_M)$.*

In other words, an essentially continuous decision rule can have discontinuities only where there are no nearby equilibrium prices. Let $\mathcal{M}$ be the set of essentially continuous decision rules. Throughout, we will restrict attention to decision rules in $\mathcal{M}$. We will at times refer to this as a continuity requirement; although it does not imply that M must be everywhere continuous, it has the same intuitive content. Discontinuities in M are only needed when the principal attempts to "move against the market", as discussed in Section 6.3, and then are only needed above the highest equilibrium price and below the lowest equilibrium price.

A related concern to that of manipulation is that if M is discontinuous then the set of equilibrium outcomes may be overly sensitive to the model fundamentals, in particular to the function R, about which the principal may well have imperfect knowledge. Indeed, Lemma 6 shows that if M has a discontinuity at at some price which could occur in equilibrium then the equilibrium outcomes

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

[15] If the market admits a reduced-form representation then $\bar{P}_M = \cup_{\theta \in \Theta} \{p \in \mathcal{P} : R(M(p), \theta) = p\}$.

will respond discontinuously to changes in R. Decision rules for which the equilibrium outcomes respond continuously to perturbations of R, which we refer to as *robust to structural uncertainty*, are discussed in Section 5.1. These results can also be used to model manipulation which translates into perturbations to R

The restriction to essentially continuous decision rules ensures consistency between the model with small, but not atomistic, agents, and the model with infinitesimal agents. For tractability we generally want to work in the limiting model in which agent's are infinitesimal, but we do not wish to artificially disregard any manipulation concerns by doing so.

Equilibrium multiplicity

The second critical concern is equilibrium multiplicity. The dependence of the principal's action on the endogenously determined price can lead to multiple equilibria, since there may be multiple self-fulfilling beliefs that agents in the market can hold about what action the principal will take (Bernanke and Woodford, 1997). This type of multiplicity is pervasive in market-based policy problems. In reality, the principal is often unable to select which equilibrium will be played. Moreover, the fact that there are multiple equilibria could lead to non-fundamental volatility in the market, as agents coordinate on one or another belief about what action the principal will take. This type of volatility is a first order concern in many settings in which market-based policies are used, such as monetary policy (Woodford, 1994).

Because of these concerns one of our central objectives is to characterize what price and action functions the principal can implement uniquely. We say that an action an price function (Q, P) are *uniquely implementable* if they are the unique equilibrium outcomes given some decision rule M. Alternatively, we can state this uniqueness property in terms of the decision rule M. Since the function M maps prices to principal actions, it is sufficient to specify that there is a unique equilibrium price in every state.

Definition. *M is **robust to multiplicity** if $P_M(\theta)$ is singleton for all θ.*[16]

It will also be useful to consider a slightly weaker notion of robustness to multiplicity. A decision rule M is *weakly robust to multiplicity* if $P_M(\theta)$ is singleton for almost all θ.[17]

Our primary focus is on unique (or almost-everywhere unique) implementation. As discussed above, unique implementation is desirable in many settings, especially those in which non-fundamental volatility is a first-order concern. However we show in Section 5.2 how our results extend to settings in which the principal is willing to tolerate multiple equilibria.

[16] Recall that when the market admits a reduced-form representation, $P_M(\theta) = \{p : p = R(M(p), \theta)\}$

[17] This definition of robustness is natural when the principal maximizes expected utility and has an absolutely continuous prior H. If instead H has atoms then the definition should be modified so that the requirement of a unique price holds almost everywhere under H. There is no difficulty in accommodating this modification into the analysis, although it requires rewording some of the results.

2.3 Constrained implementation

We first analyse the problem of implementation subject to the constraints imposed by concerns about market manipulation and equilibrium multiplicity.

Definition. (Q, P) *is* ***continuously uniquely implementable (CUI)*** *if it is implementable by an essentially continuous* M *that is robust to multiplicity.*

When the market admits a reduced-form representation, (Q, P) is continuously uniquely implementable if there exists $M \in \mathcal{M}$ such that:

1. $Q = M \circ P$
2. For all θ, $P(\theta)$ is the unique solution to $p = R(M(p), \theta)$.
3. $Q(\theta) \neq Q(\theta') \quad \Rightarrow \quad P(\theta) \neq P(\theta')$

There are two differences between implementability and CUI; the uniqueness requirement in Condition 2 and the continuity requirement that $M \in \mathcal{M}$. Continuity, as discussed above, reflects manipulation concerns. If Condition 2 holds for almost all θ, rather than all θ, then we say that (Q, P) is *continuously weakly uniquely implementable (CWUI).* There is no substantive difference between the two notions, but in some cases the weaker notion yields a more transparent characterization.

We sometimes refer to an action function Q as CUI, by which we mean that there exists a P such that the pair (Q, P) is CUI, in similarly for price functions P. Moreover, at times, it is 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 CUI*** *if for any* $\varepsilon > 0$ *there exists an* ε*-approximation of* Q *that is CUI.*

The characterization of CUI and CWUI outcomes is one of the main results of this paper. It turns out that this characterization is also central to understanding optimal decision rules even when the uniqueness constraint is relaxed.

2.4 A brief illustration: emergency lending

To illustrate the key results, consider the problem of an international lender such as the IMF or World Bank deciding on the size of an emergency loan to extend to a country experiencing a crisis. The lender is unaware of the precise severity of the crisis, which is represented by an unknown state $\theta \in [\underline{\theta}, \bar{\theta}]$; lower states represent greater severity. Dispersed information regarding the state may be at least partially reflected in the price of government bonds. For simplicity, imagine that all traders in the bond market know the true state (this assumption is purely for illustrative purposes; it does not affect the results discussed here and is not required in the general model).

Let $\pi(a,\theta)$ be the value of government bonds, i.e. their ex-post payout, if the lender extends a loan of size $a \in [0,\bar{a}]$ and the state is θ. For any loan amount a, bond values are increasing in θ. A large emergency loan leads to higher bond prices when the crisis is severe, as it reduces the probability of default in the short term. However bondholders may also worry that the increase in the country's debt burden could have adverse long-term affects. For example, the increase in the debt burden may lead to debt overhang and push the country down the back side of the debt Laffer curve, as investors worry that long-term growth will be negatively affected by the higher taxes needed to service the increased debt burden (Cordella et al., 2010). If the current crisis is mild, this effect may dominate, in which case bond prices will react negatively to the lender's intervention.[18] These considerations are captured by the following two assumptions on bond values:

1. There exists $\theta^* \in [\underline{\theta},\bar{\theta}]$ such that $\pi(\cdot,\theta)$ is increasing for $\theta \leq \theta^*$ and decreasing for $\theta > \theta^*$.

2. $\pi_2(a,\theta)$ is decreasing in a.

The lender would like to extend emergency relief only when the crisis is severe.[19] For simplicity, assume there exists a state $\theta^\bullet$ such that the the lender's payoff is increasing in a when $\theta \leq \theta^\bullet$, and decreasing in a when $\theta > \theta^\bullet$. As a result, the lender would ideally like to extend the maximal loan amount $\bar{a}$ if and only if $\theta \leq \theta^\bullet$, and otherwise extend no loan. We refer to this policy as the first-best action function. A higher $\theta^\bullet$ corresponds to a more interventionist policy on the part of the lender. The lender is likely to be interventionist if the country is very poor, in which case the short-run welfare losses from government austerity are large, or if the country is central to the global economy, because in this case a recession there will have large spillover effects on other countries. Figure 1a illustrates an interventionist first-best action function in which $\theta^\bullet > \theta^*$. The solid lines denote the bond values as a function of the state under the two extreme actions 0 and $\bar{a}$. The dashed blue line is the price function P^* induced by the first-best action function. Note that for each price p there is at most a single state θ such that $P^*(\theta) = p$. It is therefore possible to choose a decision rule mapping prices to actions that implements the first-best action function. In fact, Proposition 2 implies that this first best in Figure 1a will be continuously weakly uniquely implementable (there will only be multiple equilibrium actions in state $\theta^\bullet$). In this case, the decision rule which uniquely implements the first-best involves a gradual reduction in the level of support as the bond price increases over an intermediate range.

[18]Indeed, Cordella et al. (2010) find that the strongest empirical evidence of a negative relationship between debt and growth is for countries with relatively good policies and institutions.

[19]It could be that the lender does not wish to make a loan if the crisis is too severe, and the loan is unlikely to be repaid. Preferences of this sort are covered in Section 6.2.

Figure 1b illustrates a conservative first-best policy. In this case the lender is unwilling to intervene in some states in which bondholders would like the government to receive a large emergency loan. This is likely the most realistic scenario for middle-income countries. In this case the first-best action function cannot be implemented by a market-based decision rule. This is due to the fact that for prices in (p', p'') the price function does not reveal enough information: upon observing such a price the decision maker cannot tell if the state is below $\theta^\bullet$, in which case the action $\bar{a}$ should be taken, or above $\theta^\bullet$, in which case the action should be 0. In other words, the action is not measurable with respect to the induced price.

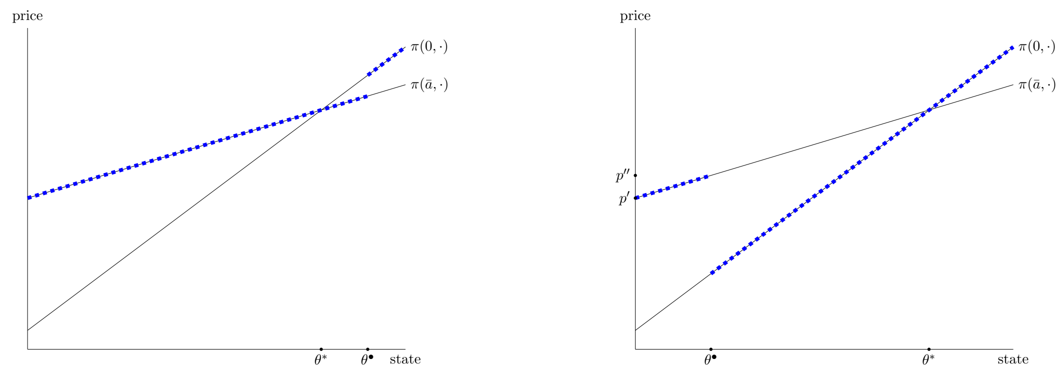

(a) First-best is CWUI

(b) First-best not implementable

Figure 1: First-best

Consider the modification of the conservative first-best action function illustrated in Figure 2a, which is a natural way to eliminate the measurability problem discussed above. This requires making an intermediate loan for states in (θ', θ''), where the lender would prefer not to intervene at all. Given this modification, for any price p there is a unique state θ such that $P^*(\theta) = p$, and so this action function is implementable. In fact, it is implementable with a continuous decision rule.

Unfortunately, is not possible to continuously and *uniquely* implement a policy resembling that of Figure 2a. In fact, any continuous decision rule M that implements this action function induces at least one equilibrium in which large loans are made for all states in $(\theta^\bullet, \theta')$. Our main results illustrate why this is the case. As a modification of the first best along the lines of Figure 2a is a natural way to deal with measurability, this example also illustrates how our results, in identifying the equilibrium-multiplicity implications of such a policy, are useful from an applied perspective.

3 Main characterization results

In this section, we argue that the defining feature of CUI outcomes is a monotone price. Specifically, we show that a monotone price function is necessary when the market-clearing function R is weakly increasing in the state (Theorem 1), and essentially sufficient under additional mild conditions (Theorem 2).

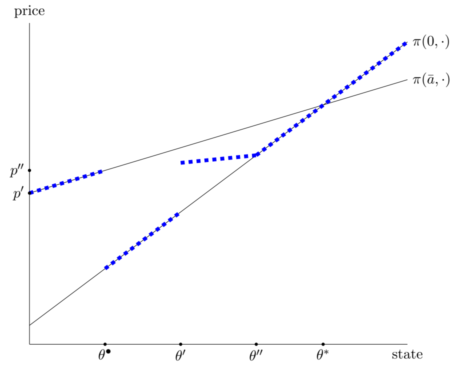

(a) Implementable, not CWUI

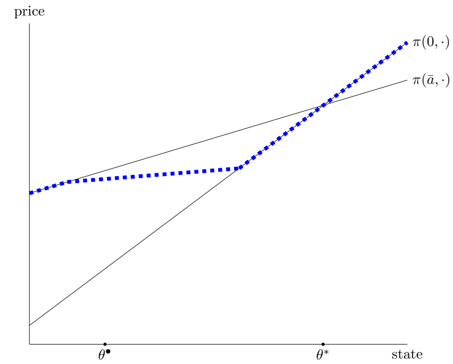

(b) (virtually) optimal CWUI policy

Figure 2: Implementable policies

To present these results, we assume a state space that is an open bounded interval $\Theta = (\underline{\theta}, \bar{\theta})$. These assumption are maintained for most of the paper, however we also extend the results to settings with closed and multidimensional Θ.[20] We also assume that the market clearing function R is continuous. Finally, we restrict attention to market-clearing functions R that are (weakly) increasing in the state θ for all actions. Both monotonicity and continuity of R can be justified by natural assumptions on primitives in many micro-foundations, as discussed in Section 4.

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

Note that the order used on Θ is irrelevant, provided continuity is satisfied.[21] Finally, we add two technical conditions on the market-clearing function for the extreme states. First, we assume that the function converges uniformly to the extreme states. In other words, $R(\cdot, \theta_n)$ converges uniformly as $\theta_n \to \underline{\theta}$ and $\bar{\theta}$. This guarantees that continuity is preserved for the limit functions $\underline{R}(a) := \inf_{\theta\in\Theta} R(a,\theta)$ and $\bar{R}(a) := \sup_{\theta\in\Theta} R(a,\theta)$. Second, we assume that for all $p \in \mathcal{P}$, $\underline{R}^{-1}(p)$ and $\bar{R}^{-1}(p)$ are the union of finitely many connected subsets of $\mathcal{A}$. Given continuity, this assumption means that at the extreme states the market-clearing price does not oscillate too frequently (as a function of the action).

Theorem 1 (Necessity)**.** *Assume R is weakly increasing in θ. If P is implemented by a decision rule $M \in \mathcal{M}$ that is weakly robust to multiplicity, then P is monotone.*

Proof. In Appendix A.2. □

[20] For closed interval states, see Appendix A.4. The results can be immediately extended to multidimensional state spaces, as long as there is a complete order on Θ and R is increasing with respect to that order. We extend the results to multidimensional state spaces without this specific feature in Section 3.1.

[21] All results that assume that R is strictly (weakly) increasing continue to hold under the weaker assumption of comonotonicity: there exists a complete order on Θ with respect to which R is strictly (weakly) increasing.

In other words, if $M \in \mathcal{M}$ induces a price function P that is non-monotone then there will be multiple equilibria. A feature of Theorem 1 that is worth emphasising is that the induced equilibrium price function P need not be increasing; it may be monotonically decreasing, even when R is strictly increasing in θ.

A monotone price function, together with the market-clearing condition, is not sufficient for CUI. The following theorem characterizes the set of continuously uniquely implementable pairs (P, Q) when R is strictly increasing in θ. First, lets name a specific set of actions.

Definition. *$a \in \mathcal{A}$ is* maximal at the bottom *iff $\underline{R}(\cdot)$ has a local maximum at a. a is* minimal at the top *iff $\bar{R}(\cdot)$ has a local minimum at a.*

The importance of these actions for the characterization lies in the possibility of assigning actions to prices outside of $P(\Theta)$ in a way that there is no state for which these prices and actions are compatible with an equilibrium. If an action a is maximal at the bottom, there is no way to construct a continuous relationship between actions close to a and prices close to $\underline{R}(a)$ such that $p > R(a, \theta)$ for all $\theta \in \Theta$.

Theorem 2. *Assume R is strictly increasing in θ. Then (Q, P) is CUI iff*

1. *$P(\theta) = R(Q(\theta), \theta)$ for all θ,*

2. *P is strictly monotone.*

3. *Q is continuous and $\bar{Q} := \lim_{\theta \to \bar{\theta}} Q(\theta)$ and $\underline{Q} := \lim_{\theta \to \underline{\theta}} Q(\theta)$ exist. Moreover, if P is decreasing, then $\underline{Q}$ is not maximal at the bottom and $\bar{Q}$ is not minimal at the top.*

Proof. In Appendix A.3. □

The first point in Theorem 2 is simply the market clearing condition that was already necessary for implementation (Observation 1). It is worth noting that continuity of Q is not implied by the continuity of M, but is instead a consequence of the robustness to multiplicity. [22]

Notice that for any (Q, P) that is CUI, the continuity of Q implies continuity of P and thus $P(\Theta)$ must be convex. Given (Q, P) satisfying condition 1 of Theorem 2, and with P increasing, it is straightforward to construct an M that continuously uniquely implements it: for prices in $P(\Theta)$ simply choose the action that is consistent $M(p) = Q \circ P^{-1}(p)$, and then use $\bar{Q}$ for prices above $\sup P(\Theta)$ and $\underline{Q}$ for prices below $\inf P(\Theta)$. Moreover, this implies that if (Q, P) is CUI and P is increasing then it can be implemented by a continuous M.

[22] In Section 5.2 we show that by slightly relaxing the concept of robustness to multiplicity (allowing multiple equilibria on a zero-measure set of states) we get a characterization that allows for discontinuous Q, so we don't see this as a critical characteristic of implementable pairs. The monotonicity of P, on the other hand, is the essential point.

When P is decreasing, the construction of M leaving actions constant for prices outside of $P(\Theta)$ does not work. The last part of condition 3 of Theorem 2 guarantees that there exists a continuous path of actions for prices slightly above $\sup P(\Theta)$ and slightly below $\inf P(\Theta)$ such that there is no state compatible with an equilibrium.

Next, we provide two other important characterizations: first, we generalize Theorem 2 to R weakly increasing. Second, we characterize the CWUI set.

CUI with weakly increasing R. Relaxing the assumption of strictly increasing $\theta \mapsto R(a, \theta)$ to weakly increasing, we 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. When R is strictly increasing in θ, $R(a, \theta) = R(a, \theta')$ implies that $\theta = \theta'$. Under weakly increasing R however, this might not be the case.

Consider (P, Q) implementable. (P, Q) satisfies market-clearing and measurability by Observation 1. Moreover suppose that $R(Q(\theta), \theta) = R(Q(\theta), \theta')$ with $\theta \neq \theta'$. If $Q(\theta') \neq Q(\theta)$ then there will be multiplicity, since by measurability $P(\theta') \neq P(\theta)$ but $P(\theta) = R(Q(\theta), \theta')$ is a market clearing price in state θ'. The only modifications needed to extend Theorem 2 are those that rule out such instances of multiplicity.

Proposition 1. *Assume R is weakly increasing in θ. Then (Q, P) is CUI iff*

1. *$P(\theta) = R(Q(\theta), \theta)$ for all θ,*
2. *P is weakly monotone.*
3. *Q is continuous and $\bar{Q} := \lim_{\theta \to \bar{\theta}} Q(\theta)$ and $\underline{Q} := \lim_{\theta \to \underline{\theta}} Q(\theta)$ exist. Moreover, if P is decreasing, then $\underline{Q}$ is not maximal at the bottom and $\bar{Q}$ is not minimal at the top.*
4. *$P(\theta) = P(\theta')$ OR $P(\theta) = R(Q(\theta), \theta')$ implies $Q(\theta') = Q(\theta)$ for all θ, θ'.*

Proof. In Appendix A.5. □

CWUI with strictly increasing R. The only substantive difference between CUI and CWUI outcomes is that the action function need not be continuous. This illustrates that the essential feature of manipulation and multiplicity-proof implementation is monotonicity of the price function. To establish this, we begin with some preliminary observations. First, if (Q, P) are CWUI, then, since P must be monotone by Theorem 1, any discontinuity in P must be a jump discontinuity, and P can have at most countably many discontinuities. Moreover, Q can be discontinuous at θ only if P is as well: otherwise it would not be possible for Q to be implemented by an M that is continuous at $P(\theta)$. Thus Q can also have at most countably many discontinuities. Finally, recall that $M \in \mathcal{M}$ must be continuous on $cl(P(\Theta))$. This implies the following.

Lemma 1. *The one-sided limits of any CWUI Q, denoted by $\lim_{\theta \nearrow \theta'} Q(\theta)$ and $\lim_{\theta \searrow \theta'} Q(\theta)$, must exist for all θ'.*

Proof. Proof in Appendix C.2.2. □

Suppose P has a discontinuity at θ^*, and let $\underline{p} = \lim_{\theta \nearrow \theta^*} P(\theta)$ and $\bar{p} = \lim_{\theta \searrow \theta'} P(\theta)$. Say that this discontinuity in P at θ^* is *bridgeable* given Q if there exists a continuous function $\gamma : [\min\{\underline{p}, \bar{p}\}, \max\{\underline{p}, \bar{p}\}] \to \mathcal{A}$ such that *i)* $\gamma(\underline{p}) = \lim_{\theta \nearrow \theta^*} Q(\theta)$, *ii)* $\gamma(\bar{p}) = \lim_{\theta \searrow \theta^*} Q(\theta)$, and *iii)* $p = R(\gamma(p), \theta^*)$ for all $p \in [\min\{\underline{p}, \bar{p}\}, \max\{\underline{p}, \bar{p}\}]$. We say that the environment is *fully bridgeable* if for any (Q, P), all discontinuities in P are bridgeable.

Observation 2. *A discontinuity in P at θ^* is bridgeable iff there exists a continuous function $\gamma : [0,1] \to \mathcal{A}$ such that i) $\gamma(0) = \lim_{\theta \nearrow \theta^*} Q(\theta)$, ii) $\gamma(1) = \lim_{\theta \searrow \theta^*} Q(\theta)$, and iii) $x \mapsto R(\gamma(x), \theta)$ is strictly monotone.*

Observation 2 is useful because the condition that $x \mapsto R(\gamma(x), \theta)$ is strictly monotone is easier to check than the fixed-point condition in the definition of bridgeability.

Proposition 2. *Assume R is strictly increasing in θ. Then (Q, P) is CWUI iff*

1. *$P(\theta) = R(Q(\theta), \theta)$ for all θ.*
2. *P is strictly monotone.*
3. *If Q is discontinuous at θ^* then P has a bridgeable discontinuity at θ^*.*
4. *$\bar{Q} := \lim_{\theta \to \bar{\theta}} Q(\theta)$ and $\underline{Q} := \lim_{\theta \to \underline{\theta}} Q(\theta)$ exist. Moreover, if P is decreasing, then $\underline{Q}$ is not maximal at the bottom and $\bar{Q}$ is not minimal at the top.*

Proof. Proof in Appendix C.2.3. □

The bridgeability condition in Proposition 2 is the most difficult to verify in practice. Fortunately, it can be ignored in most relevant environments, as these satisfy full bridgeability. In such settings, Q can be discontinuous at θ^* iff P is as well, or equivalently: $\lim_{\theta \nearrow \theta^*} R(Q(\theta), \theta) \neq \lim_{\theta \searrow \theta^*} R(Q(\theta), \theta)$. Thus, from a practical perspective, many applied problems can be solved simply by optimizing over the action function Q subject to the constraint that $R(Q(\theta), \theta)$ be strictly monotone.

Full bridgeability is satisfied in many natural settings. For example, suppose the principal's action consists of mixtures over a set of consequences, i.e. $\mathcal{A} = \Delta(Z)$ for some finite set Z, where each consequence is associated with a value $\pi(z, \theta)$. Fixing the state θ, any action $a \in \Delta(Z)$ induces a distribution over values via the function $\pi(\cdot, \theta)$. If $R(a'', \theta) > R(a', \theta)$ whenever the distribution induced by a'' first-order stochastically dominates that induced by a', then the environment is fully bridgeable. In other words, a weak monotonicity notion suffices for bridgeability. This result, along with more general sufficient conditions for bridgeability, is discussed in Appendix C.3.

3.1 Multidimensional state space

Suppose that Θ is an open subset of $\mathbb{R}^N$, endowed with the usual product partial order. Assume that $R : \mathcal{A} \times \Theta \to R$ is continuous, and is increasing with respect to the partial order on Θ, i.e.

Definition. *Say that R is strictly increasing if $\theta'' > \theta'$ implies $R(a, \theta'') > R(a, \theta')$ for all a.*[23]

Define $\bar{R}(a, \theta) = \{\theta' : R(a, \theta') = R(a, \theta)\}$. That is, $\bar{R}(a, \theta)$ is the level set of $R(a, \cdot)$ corresponding to the price $R(a, \theta)$. Under the assumptions of continuous and strictly increasing R, $\bar{R}(a, \theta)$ is a one-dimensional curve in Θ for any a, θ. The problem can be reduced to one with a uni-dimensional state space if and only if $\bar{R}(a, \theta) = \bar{R}(a', \theta)$ for all a, a' and θ; if this condition does not hold then there is no complete order on Θ with respect to which R is monotone for any a. Nonetheless, we are able to characterize CUI in this setting.

Clearly if there is an equilibrium under M in which the principal takes action a in state θ then there is an equilibrium in which the principal takes action a for all states in $\bar{R}(a, \theta)$. Therefore robustness to multiplicity implies that for all θ, $Q(\theta') = Q(\theta)$ for all $\theta' \in \bar{R}(Q(\theta), \theta)$. Robustness to multiplicity also implies that $Q(\theta) \neq Q(\theta') \Rightarrow \bar{R}(Q(\theta), \theta) \cap \bar{R}(Q(\theta'), \theta') = \varnothing$; otherwise there would be multiple equilibrium actions for any states in the intersection. As a result of these two observations, the previous characterizations of CUI action and price functions can be extended without much difficulty. The following result is analogous to Theorem 2 in the uni-dimensional case.

Proposition 3. *Assume strictly increasing R. Then if (Q, P) is CUI*

i. $P(\theta) = R(Q(\theta), \theta)$.

ii. *P is strictly monotone (in the product partial order on Θ).*

iii. *Q is continuous.*

iv. *For all θ, $Q(\theta') = Q(\theta)$ for all $\theta' \in \bar{R}(Q(\theta), \theta)$.*

v. $Q(\theta) \neq Q(\theta') \Rightarrow \bar{R}(Q(\theta), \theta) \cap \bar{R}(Q(\theta'), \theta') = \varnothing$.

Proof. in Appendix A.6. □

The conditions of Proposition 3 are also sufficient, except that it may be possible for Q to have discontinuities at the states associated with the highest prices (see Section 3 for discussion), and additional conditions are required on the limit actions, analogous to Condition 4 in Theorem 2.

[23] This can be easily relaxed to weakly increasing.

4 Understanding price formation

In order to characterize implementable and robustly implementable price and action functions in Section 3, we made use of the assumption that the market admitted a reduced-form representation. In this section we discuss how the existence of a reduced-form representation simplifies the analysis, what types of markets have this property, and how our results can be extended to markets which fail to admit such a representation.

4.1 Benefits of working with reduced-form representation

Identifying a reduced-form representation of equilibrium in a given market facilitates state-by-state analysis of the principal's policy design problem, and makes it possible for us to work effectively in the space of action and price functions. As discussed in Section 2.1, there are two key features of the reduced-form representation. First, R is a function from $\mathcal{A} \times \Theta$ to $\mathcal{P}$, not from $\mathcal{A}^{\Theta} \times \Theta$ to $\mathcal{P}$. In other words, the equilibrium price at state θ depends only on the equilibrium action in that state, and not on global features of the equilibrium action function. This greatly simplifies the problem of characterizing the set of implementable outcomes, as it reduces the number of global constraints the action and price functions. Second, the market-clearing function $R : \mathcal{A} \times \Theta \to \mathcal{P}$ does not depend on M. As Observation 1 makes clear, in order to determine if Q and P can be equilibrium outcomes for some M, we therefore do not need to know what precise M implement them.

An additional benefit of being able to summarize the equilibrium price via the market-clearing function R is that the principal does not need to know the details of the market micro-structure in order to design policy. The market-clearing function represents the equilibrium relationship between the principal's action, state, and price. Since it does not depend on the decision rule M, it can be estimated using data from a market in which the principal's action is not conditioned on the price, or in which some other decision rule was used. Thus a principal contemplating the introduction of a market-based decision rule can use historical aggregate data to estimate the function R and design the decision rule, without being subject to the Lucas critique that a change in the policy regime will change the relationship between the fundamentals (the state and anticipated action) and aggregate outcomes (price)(Lucas et al., 1976).

4.2 Deriving a reduced-form representation

For a market to admit a reduced-form representation means that the equilibrium price in a given state θ is uniquely determined by the action $Q(\theta)$ taken in that state. Conversely, the market fails to admit a reduced-form representation if and only if the following holds

1. Given some decision rules M and M' there exist equilibria with action and price functions (Q, P) and (Q', P') respectively and a state θ, such that $Q(\theta) = Q'(\theta)$ and $P(\theta) \neq P'(\theta)$.

In particular, letting $M = M'$ and $Q = Q'$, one possible reason for the failure of a reduced-form representation is the following

2. Given some decision rule M, there exist multiple equilibria with the same action function, but different price functions.

If the only reasons for the failure of the market to admit a reduced-form are of the second type, then the market admits a representation via a market-clearing correspondence $R : \mathcal{A} \times \Theta \mapsto 2^{\mathcal{P}}$. It is relatively straightforward to extend our analysis to this type of market. The more interesting and challenging scenario is when there are when there are failures that are not of this type. Such failures of existence of the reduced-form representation occur because of *global effects*: it is not sufficient to know the equilibrium action in state θ in order to determine the equilibrium price in that state (or even the set of equilibrium prices in state θ).

We generally focus on markets in which market participants are price takers. If each agent in the market understands that their individual action may have a significant impact on the price, then clearly their behavior will depend directly on global features of the principal's decision rule. Thus it is generally not possible to derive a reduced-form representation of equilibrium outcomes, which does not depend explicitly on the decision rule used. However, we understand the price-taking assumption to be an approximation of behavior in a market in which agents are small, but may have some non-zero price impact. This is the reason for considering decision rules that are robust to manipulation, as previously discussed. By studying decision rules that are robust to manipulation, we ensure that the predictions of the model in which agents are assumed to be price takers are good approximations to the outcomes in a model in which agents may have some small degree of market power.

4.3 Reduced form: private values

In is easy to see that the market admits a reduced-form representation when the behavior of market participants depends only on their private information, the price, and the anticipated principal action. Assume that the market consists of a set I of agents. Each agent $i \in I$ receives a private signal s_i. Agents are price takers and have private values: their action depends only on their own signal, price, and the anticipated principal action.[24] Let $x_i(s_i, p, a)$ be the action of agent i. For example, in carbon credits market example in the introduction, each agent is a firm, which learns it's own marginal abatement cost s_i and decides on a quantity x of credits to purchase, anticipating that the principal will set the per-credit volume at a. A firm's payoff here does not depend directly on the abatement costs of other firms.

[24]While we make no assumption about the size of the market here, the price taking assumption of course generally fits best in markets with many small agents.

The price $F(\{x_i\}_{i \in I})$ is determined by the actions of all agents. Let $\theta = \{s_i\}_{i \in I}$ be the profile of signal realizations.[25] In equilibrium it must be that

$$P(\theta) = F(\{x_i(s_i, P(\theta), Q(\theta))\}_{i \in I}).$$

If $p \mapsto F(\{x_i(s_i, p, a)\}_{i \in I})$ has a unique fixed point for every action a and signal profile θ (for example, if F is strictly monotone and $p \mapsto x_i(s_i, p, a)$ is strictly monotone for all s_i, a) then the market admits a reduced-form representation.[26] Moreover, the market clearing function defined by this condition is continuous provided F and x_i are continuous (by Berge's maximum theorem).

4.4 Reduced form: equilibrium inferences and global effects

To illustrate the challenges involved in modeling markets in which global effects are present, and to understand the reasons such effects might arise, consider a simple rational expectations equilibrium model of an asset market. There is a single asset, which pays an ex-post dividend $\pi(a, \theta)$ that depends on the principal's action and the state. Assume that the aggregate supply of the asset is fixed, and normalize this to zero. In Section 4.7 we consider the important extension to the noisy REE model in which there are stochastic shocks to the aggregate supply.

There are a continuum of investors $i \in [0, 1]$, each of whom observes a private signal s_i that is informative about the state. An investor's payoff $u_i\big(x \cdot (\pi(a, \theta) - p)\big)$ depends on the quantity x that they purchase, the price p of the asset, and the asset dividend (where u_i is strictly increasing). After observing their private signal, each investor submits a demand schedules to a market maker, which specifies their quantity demanded for every price. The market maker then chooses a price to clear the market.

The key feature of this environment is that in addition to their private signals, investors learn about the state from the price. In contrast to the private-values setting discussed in Section 4.3, other investors' signals are informative about a payoff-relevant state, and thus investors draw inferences from the price. When formulating their demand given a price p, investors condition on the set of states at which p is the equilibrium price. Formally, fix the principal's decision rule $M : \mathcal{P} \mapsto \mathcal{A}$ mapping prices to actions. A rational expectations equilibrium (REE) consists of a price function P_M mapping states to prices such that two conditions hold.

i. Investors optimize, conditioning on signal and price:

$$X_i(p, s_i) = \arg\max_x \mathbb{E}\left[u_i\big(x \cdot (\pi(M(p), \theta) - p)\big) \mid s_i, P_M(\theta) = p\right] \tag{1}$$

[25] In some settings, such as the model of Section 4.4, we can define the state to be a lower-dimensional sufficient statistic for the profile of signal realizations.

[26] If there are multiple fixed points then the market admits a representation via a market-clearing correspondence, as discussed in Section 4.2.

ii. Markets clear in all states:[27]

$$\int X_i\left(P_M(\theta), s_i\right) di = 0 \quad \forall \quad \theta \in \Theta.$$

The difficulty with analysing market-based policy in this environment can be seen by examining (1). The principal's decision rule affects investors in two ways. The first is a direct *forward guidance* effect: the decision rule determines what action investors anticipate, conditional on the price, and thus affects the anticipated dividend $\pi(M(p), \theta)$. However there's also and indirect *informational effect*, arising from the fact that when formulating their demand for the price of p, investors condition on the event $\{\theta \in \Theta : P_M(\theta) = p\}$. The decision rule shapes the entire equilibrium price function P_M, and thus determines what information investors infer about the state from the price. The subtlety of this informational effect is that investor beliefs in a give state will depend on the equilibrium price and principal actions in other states. Thus global properties of the decision rule and the equilibrium price and action functions will matter for determining the price in a given state. Such global dependence makes it more difficult to analyse the principal's problem in outcome space (the space of price and action functions); modifying the action and price function for some states may necessitate modifications elsewhere. This introduces global constraints into the principal's problem.

To understand this difficulty, consider the REE asset market model described above, and let Q_1, P_1 a implementable action and price function. The price function, depicted in Figure 3a, is constant over the interval $[\theta_1, \theta_3]$. Let Q_2 be another action function, such that $Q_2(\theta) = Q_1(\theta)$ for $\theta \leq \theta_2$ and $Q_2 \neq Q_1$ elsewhere. We want to know if Q_2 is implementable, and if so, what the corresponding price function will look like. It is natural to expect that if Q_2 is implementable, the corresponding price function P_2 will differ from P_1 for states above θ_2. Suppose that $P_2 > P_1$ above θ_2. However, can it be the case that the price functions also differ below θ_2, where the action functions are the same? Suppose that this is not the case; $P_2 = P_1$ below θ_2. Let $\theta^* \in (\theta_1, \theta_2)$ be a state in which Q_1 and Q_2 coincide, so $Q_1(\theta^*) = Q_2(\theta^*) = a^*$. In the Q_1 equilibrium, the information revealed by a price of $P_1(\theta^*)$ is $\{\theta : P_1(\theta) = P_1(\theta^*)\} = [\theta_1, \theta_3]$. Therefore, in state θ^* investor i's demand is give by

$$X_i(P_1(\theta^*), s_i) = \arg\max_x \mathbb{E}\left[u_i\big(x \cdot (\pi(a^*, \theta) - p)\big) \mid s_i, \theta \in [\theta_1, \theta_3]\right]$$

Similarly, in the Q_2 equilibrium, a price of $P_2(\theta)$ reveals that $\theta \in [\theta_1, \theta_2]$, so i's demand in state θ is given by

[27] By specifying the market-clearing condition state-by-state, we are assuming that the state captures all uncertainty in the market. In general, we can define the state to mean the entire profile of signal realizations for each investor in the market. However in this model with a continuum of investors and conditionally independent signals, it follows from the usual "continuum law of large numbers" convention that the distribution of signal realizations in the population is identified by θ.

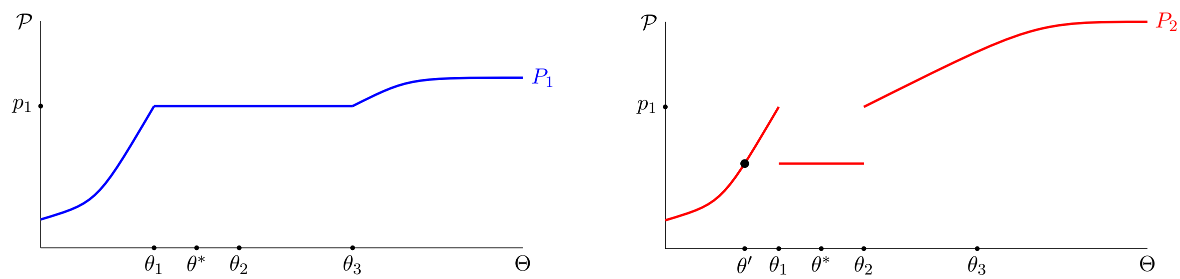

(a) Q_1 equilibrium

(b) Q_2 not implementable

Figure 3: Information effects and global dependence

$$X_i(P_2(\theta^*), s_i) = \arg\max_x \mathbb{E}\left[u_i\big(x \cdot (\pi(a^*, \theta) - p)\big) \mid s_i, \theta \in [\theta_1, \theta_2]\right]$$

Notice that the beliefs of investor i in this case are first-order stochastically dominated by those in the Q_1 equilibrium. If $\theta \mapsto \pi(a^*, \theta)$ is strictly increasing then the quantity demanded by every investor will by higher under the FOSD dominant beliefs. This means that in order for markets to clear at state θ^* in the Q_2 equilibrium, the price must be lower than in the Q_1 equilibrium. Thus it cannot be that $P_1 = P_2$ for all states below θ_2. However, if in the Q_2 equilibrium the price must be lower for states in $[\theta_1, \theta_2]$, as depicted in Figure 3b, then it may be that (Q_2, P_2) is not even implementable. This will be the case if there is some state $\theta' < \theta_1$ such that $P_2(\theta') = P_2(\theta^*)$, but $Q_2(\theta') \neq Q_2(\theta^*)$, as the principal's action must be measurable with respect to the price.

4.5 Reduced-form representation in REE

The reason for the global dependence illustrated in this example is that a given state may belong to different public information sets (the information revealed by the price) in different equilibria. Informational effects arising from inferences drawn by market participants from the price function are one potential source of global dependence which leads to failure of the market to admit a reduced-form representation. Indeed, it is difficult to conceive of a market without such informational effects which fails to admit a reduced-form representation, provided the state is appropriately defined to capture all the relevant uncertainty in the market. However, rational expectations models, in which such informational effects do arise, are central to the analysis of asymmetric information in markets. We therefore study under what conditions such markets admit a reduced-form representation. Surprisingly, reduced-form representations can be derived under fairly weak conditions on the market. These conditions do not imply that there are no informational effects: the principal's decision rule may indeed determine what information is reveled by the price to market participants. We show that nonetheless, these informational effects can be summarized via a reduced-form representation.

Consider a more general version of the asset market model described above. There are a unit mass of investors. Investors receive conditionally independent signals s_i about the state, with conditional distribution $h(\cdot|\theta)$ on $[\underline{s}, \bar{s}]$. The ex-post payoff to investor 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, strictly concave in x (to guarantee a unique solution), and continuous in x, θ.[28] For a fixed action a the demand of investor i who observes signal s and knows that the state is in $\mathcal{I} \subseteq \Theta$ is given by

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

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

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

Since there is a continuum of investors and a finite number investor 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 $\tilde{P} : \Theta \to \mathbb{R}$, let $\mathcal{I}_{\tilde{P}} : \Theta \to 2^{\Theta}$ 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.

There is one technical complication which must be confronted when working with a continuous state space. This is the fact that zero-measure perturbations to public information sets do not affect conditional beliefs. Thus it is possible to create equilibira by perturbing the price function on a zero measure set, which violate the conditions for the market to admit a reduced form. However, we can show that in such cases the market admits an almost everywhere (AE) reduced form. Let (Ω, Σ, μ) be a measurable space.

Definition. *The market* ***admits an AE reduced-form representation*** *if* $\exists$ *a function* R : $\mathcal{A} \times \Theta \to \mathcal{P}$ *such that for any* Q, P, M

1. *If the pair* (Q, P) *are equilibrium outcomes given* M *then*

   i. $Q(\theta) = M(P(\theta))$ *for all* $\theta \in \Theta$ *(commitment)*

   ii. $P(\theta) = R(Q(\theta), \theta)$ *for almost all* $\theta \in \Theta$ *(AE market clearing)*

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

2. The pair (Q, P) are equilibrium outcomes given M if for all $\theta \in \Theta$

i. $Q(\theta) = M(P(\theta))$ *(commitment)*

ii. $P(\theta) = R(Q(\theta), \theta)$ *(market clearing)*

If the market admits an AE reduced-form representation, we can use this representation when designing policy and characterizing the implementable set. The only difference is that statements about equilibrium uniqueness must be qualified: there will always be additional equilibria, but these differ from the equilibrium associated with the AE reduced-form representation only on a set of measure zero. In other words, we are only be able to identify M which are *weakly* robust to multiplicity. Otherwise, there is no difference between working with a reduced-form representation and an AE reduced-form representation.

In many settings, there is a monotone relationship between investors' private signals and their actions. It turns out that this is sufficient to guarantee existence of an AE reduced-form representation. Let $\geq$ be a complete order on the state space. Define the level set of $\geq$ as $L_\theta \equiv \{\theta' \in \Theta : \theta' \geq \theta\} \cap \{\theta' \in \Theta : \theta \geq \theta'\}$, and let the upper-set be $U_\theta = \{\theta' \in \Theta : \theta' \geq \theta\}$.

Increasing Differences. $V_i(a, \theta, x, p)$ satisfies increasing differences in x, θ.

Belief Monotonicity. $h(\cdot|\theta'')$ strictly MLRP dominates $h(\cdot|\theta')$ for $\theta'' > \theta'$.

Increasing Differences implies in particular that if $\theta' \in L_\theta$ then $V_i(a, \theta', x, p) - V_i(a, \theta', x', p) = V_i(a, \theta, x, p) - V_i(a, \theta, x', p)$. In other words, i has the same preferences over quantities in states θ and θ', conditional on a, p.

To see how these two assumptions imply that the market admit a reduced form, consider the example illustrated in Figure 3. The issue encountered there is that since state θ^* belonged to different public information sets in the Q_1 and Q_2 equilibria, i.e. different level sets of the equilibrium price function, the demands in state θ^* could also differ. In particular, we posited that if higher states are associated with higher aggregate beliefs in the population (Belief Monotonicity) then demand would be higher in state θ^* when this state belongs to the public information set $[\theta_1, \theta_3]$ then when it belongs to the public information set $[\theta_1, \theta_2]$. This conclusion holds when higher beliefs are associated with higher demands (an implication of strictly Increasing Differences). The flaw with the above of reasoning is that if demands are strictly increasing in private signals conditional on the public information set $[\theta_1, \theta_3]$ then we cannot have a constant price over this interval to begin with: demand would be higher at higher states within this interval. Thus it must be that demand is constant as a function of private signals, which in turn implies that aggregate demand will be the same whether the public information is $[\theta_1, \theta_3]$ or $[\theta_1, \theta_2]$.

The key observations that we make use of in order to show prove that the market admits a reduced form are 1) that the principal's action is measurable with respect to the price, and 2) that public information sets revealed to investors are exactly the level sets the price function. The following proposition formalizes the above argument.

Proposition 4. *Assume there is a complete order on Θ such that Increasing Differences and Belief Monotonicity are satisfied. Then the market admits an AE reduced-form representation. In particular, $R(a, \theta) = P^*(a, L_\theta, \theta)$.*

Proof. Proof in Appendix B.1. □

Continuity of the market clearing function $R(a, \theta) = P^*(a, L_\theta, \theta)$ is guaranteed by continuity of $\theta \mapsto h(\cdot|\theta)$ and continuity of V_i.

4.6 REE without a reduced-form representation

The results of Section 4.5 gave general conditions under which a rational expectations market admits an AE reduced-form representation. It was essential that there existed a *complete* order on the state space such that Increasing Differences and Belief Monotonicity held. While both conditions apply also if there is only a partial order on Θ, this would not be sufficient to guarantee our result. There are some markets, especially those for which the state space is multi-dimensional, in which such a complete order does not exist.

One approach to this problem is to impose additional refinements to the solution concept, which allow one to derive an AE reduced-form representation. An alternative approach, which we focus on here, is to anticipate that our objective will in the end be *unique* implementation, and derive a reduced-form representation under this restriction.

Definition. *The market* ***admits a reduced-form representation under uniqueness*** *if $\exists$ a function $R : \mathcal{A} \times \Theta \to \mathcal{P}$ such that for any Q, P, M, the pair (Q, P) are the unique equilibrium outcomes given M iff for all θ*

i. $Q(\theta) = M(P(\theta))$ *(commitment)*

ii. $P(\theta) = R(Q(\theta), \theta)$ *(market clearing)*

iii. $\{p : p = R(M(p), \theta)\}$ *is singleton* *(uniqueness)*

The only difference between this definition and that of the reduced form is that we require that (Q, P) are the unique equilibrium outcomes given M, and impose the uniqueness condition. As Section 4.7 shows, there are important models in which the additional structure imposed by the uniquness requirement is crucial for deriving a reduced-form representation.

4.7 Noisy REE in asset markets

Asset markets are an important setting in which decision making under feedback effects occurs. Since Grossman and Stiglitz (1980) and Hellwig (1980), the noisy rational expectations model has been a workhorse model for studying asymmetric information in asset markets. This model adds shocks to aggregate supply, interpreted as noise or liquidity traders, to a rational expectations model of the asset market. The standard approach, without feedback effects, is to assume joint normality of asset returns and aggregate demand shocks, and look for equilibria in which the price is linear in trader's private signals. Breon-Drish (2015) generalizes the noisy REE model to allow for non-normal distributions of states and supply shocks. This section extends results from Breon-Drish (2015) to a setting with feedback effects.

The setting is as follows. There is a single asset that pays an ex-post dividend of $\pi(a, \omega)$, where $\omega \in \Omega$ is referred to as the payoff-relevant state. We assume that π is continuous and is affine in θ for all a; $\pi(a, \omega) = \beta_0^a + \beta_1^a \omega$. Each investor observes an additive signal $s_i = \omega + \varepsilon_i$, where $\varepsilon_i \sim N(0, \sigma_i^2)$, where σ_i^2 lies in a bounded set. The supply shock is a random variable z taking values in $\mathcal{Z}$. We assume that z has a truncated normal distribution. That is, z is the restriction of a normal random variable $\hat{z} \sim N(0, \sigma_Z^2)$ to the interval $[b_1, b_2]$, with $-\infty \leq b_1 \leq 0 \leq b_2 \leq \infty$ (note that this assumption accommodates un-truncated supply shocks as well). For simplicity, let $b_1 = -b_2$; this does not affect the results. The state θ consists of both the payoff-relevant state ω and the supply shock z.

There are are a continuum of investors $i \in [0, 1]$, each with CARA utility $u(w) = -\exp\left\{-\frac{1}{\tau_i} w\right\}$. The ex-post payoff to an investor who purchases x units of the asset at price p when the principal takes action a is given by $-\exp\left\{-\frac{1}{\tau_i} x(\pi(a, \theta) - p)\right\}$, where τ_i lies in some bounded set. We assume that the distribution of private signals in the population is uniquely determined by the state ω (this is the usual "continuum law of large numbers" convention). Let $x_i(p|a, \mathcal{I}, s_i)$ be the demand of investor i when the price is p, the anticipated principal action is a, and the public information revealed by the price is that $(\omega, z) \in \mathcal{I}$, and i's private signal is s_i. Aggregate demand can be written as $X(p|a, \mathcal{I}, \omega)$.

As in the one-dimensional case, equilibrium consists of a price function as well as a specification of the public information for off-path prices. To be precise, fixing a decision rule M, an equilibrium is characterized by a price function $P : \Omega \times \mathcal{Z} \to \mathcal{P}$ and an off-path inference function $\lambda : \mathcal{P} \backslash P(\Omega, \mathcal{Z}) \to 2^{(\Omega, \mathcal{Z})}$. The price function P is such that for all (ω, a) market's clear given the anticipated action and the information revealed by the price, that is: $X(P(\omega, z)|M(P(\omega, z), \mathcal{I}(\omega, z), \omega) = z$, where $\mathcal{I}(\omega, z) = \{(\omega', z') : P(\omega', z') = P(\omega, z)\}$. For the off-path information, we assume only that it is consistent with market clearing (when possible), that is: $\lambda(p) \subseteq \{(\omega, z) : X(p|M(p), \lambda(p), \omega) = z\}$.

The approach in this setting is somewhat distinct from the previous analysis, in that we do not begin by deriving a reduced-form representation directly from the primitives of the model. Rather, we show that there exists a reduced-form representation that can be used to design policy *under* *uniqueness*. However, the search for truly unique implementation is hopeless in the noisy REE model studied here, since there are multiple (meaningfully different) equilibria even when there is no policy feedback, that is, fixing the principal's action (Pálvölgyi and Venter, 2015). We therefore focus here on a more limited, but still meaningful, notion of uniqueness. What we really want to rule out is multiplicity arising from the endogeneity of the principal's action. Therefore, in the context of the noisy REE model, we require a weaker notion of robustness to multiplicity: for this model we say that M is robust to multiplicity if there is a unique market clearing price in every state, fixing the inferences draw from prices, that is, fixing the public information sets associated with each price both on and off path. Another interpretation of this requirement is that there must be a unique market clearing price fixing the demand schedules submitted to the market maker by each agent. If there are multiple equilibria in this sense, then moving between them requires no change in the behavior of market participants, simply a change in the selection of the market clearing price by the market maker.

Relative to the model with a one-dimensional state space, the complication in this setting in which the state space $\Theta = \Omega \times \mathcal{Z}$ is two dimensional, is that there is no easy way to narrow down the space of possible public information sets that can be revealed by the price. This makes it difficult to derive a reduced-form representation of the market ex-ante, without strong restrictions on the set of possible equilibria. To deal with this difficulty, we instead analyse directly the problem of characterizing what equilibria can be induced with a decision rule $M \in \mathcal{M}$ that is robust to multiplicity. Under these restrictions, we show that the market admits a reduced-form representation under uniqueness. To do this, we first need some preliminary results. The first concerns the continuity of aggregate demand.

Lemma 2. *For and $\mathcal{I} \subseteq \Omega \times \mathcal{Z}$, $p \in \mathcal{P}$, and $a \in \mathcal{A}$, the function $\omega \mapsto X(p|a, \mathcal{I}, \omega)$ is Lipschitz continuous.*

Proof. Proof in Appendix B.1.1. □

Note that since the distribution of signals in the population is uniquely determined by ω (following the usual continuum law of large numbers convention) it cannot be that any public information set $\mathcal{I}$ contains states (ω', z') and (ω', z'') with $z'' \neq z'$, since the aggregate demand would not be the same in both cases. Therefore, the distribution of ω conditional on $\mathcal{I}$ cannot have atoms. The following lemma strengthens this observation slightly, by showing that in fact, given Lipschitz continuity of aggregate demand, the distribution of ω conditional on $\mathcal{I}$ will be absolutely continuous.

Lemma 3. *For any $p \in \mathcal{P}$ and $a \in \mathcal{A}$, let $\mathcal{I}$ be a set satisfying $\mathcal{I} \subseteq \{(\omega, z) : X(p|a, \mathcal{I}, \omega) = z\}$. Then the distribution of ω conditional on $\mathcal{I}$ is absolutely continuous.*

Proof. Proof in Appendix B.1.2. □

Lemmas 2 and 3 did not make use of much of the structure that we have assumed; for example, CARA utility and truncated-normal noise distributions. Using these properties, we can establish further characteristics of public information sets. The following lemma says that any public information set, either one revealed on-path by the price or by the off-path inference function, must lie in a linear subset of $\Omega \times \mathcal{Z}$. In other words, and such $\mathcal{I}$ must be a subset of some set of the form $\{(\omega, z) : k \cdot \omega - z = \ell\}$ for some $k > 0$ and ℓ.

Lemma 4. *Let $\mathcal{I}$ satisfy $\mathcal{I} \subseteq \{(\omega, z) : X(p|a, \mathcal{I}, \omega) = z\}$ for some p, a. Then there exists $k > 0$ and ℓ such that $\mathcal{I} \subseteq \{(\omega, z) : k \cdot \omega - z = \ell\}$*

Proof. Proof in Appendix B.1.3. □

The following proposition identifies exactly which hyperplanes the public information sets can lie in.

Proposition 5. *Assume CARA utility, π affine in θ and continuous, additive normal signal structure and truncated-normally distributed supply shocks. Then there exists a unique (up to positive transformations) function $L^* : \Omega \times \mathcal{Z} \times \mathcal{A} \to \mathbb{R}$ defined by*

$$L^*(\omega, z|a) = \left(\frac{1}{\beta_1^a} \int_i \frac{\tau_i}{\sigma_i^2} di \right) \cdot \omega - z \tag{2}$$

such that for any M, if $\mathcal{I}$ is the public information revealed at price p (in which case $\mathcal{I} \subseteq \{(\omega, z) : X(p|a, \mathcal{I}, \omega) = z\}$) then $L^*(\omega'', z''|M(p)) = L^*(\omega', z'|M(p))$ for all $(\omega'', z''), (\omega', z') \in \mathcal{I}$

Proof. Proof in Appendix B.1.4. □

We now wish to use these properties, in particular Lemma 4, to identify features of equilibrium. Proposition 5 identifies the hyperplane to which each information set belongs. Following Breon-Drish (2015), we refer to these hyperplanes as *linear statistics*. So in other words, the public information will always reveal *at least* the associated linear statistic. The following result says that in fact, under robustness to multiplicity and $M \in \mathcal{M}$, the equilibrium price function will reveal *exactly* the linear statistic, and no more.

Proposition 6. *Maintain the assumptions of Proposition 5. If $M \in \mathcal{M}$ is robust to multiplicity then the level sets of the equilibrium price function $\tilde{P}$ are given by $\{(\omega, z) : L^*(\omega, z|M(p)) = \ell\}$ for some ℓ, where L^* is given by (2).*

Proof. Proof in Appendix B.1.5 □

The idea behind Proposition 6 is illustrated in Section 4.7. The left panel illustrates a situation in which the level set of the price function at $p = 4$ is a strict subset of the linear statistic $L^*(\omega, z|M(4))$. The dotted line is the segment of the linear statistic which is omitted from the level set. Since conditioning on the truncated level set induces higher posterior beliefs about ω than conditioning on the entire linear statistic, the price in these states would be lower than in an equilibrium in which the action was fixed at $M(4)$ for all prices. This would imply that there does not exist a reduced form representation. The representation is saved, however, by the uniqueness requirement. In the situation depicted in Figure 4a, we show that there are additional equilibria in which the action $M(4)$ is taken for states on the dotted line segment, violating the uniqueness requirement.

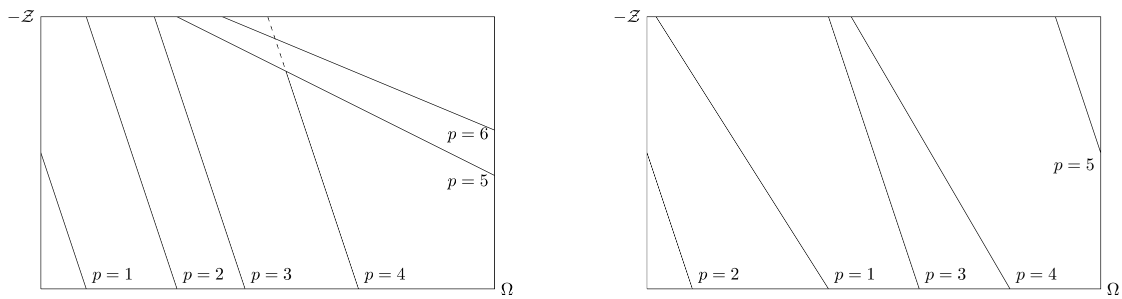

(a) Multiplicity at $p = 4$

(b) No multiplicity

Figure 4: Intersecting linear statistics imply multiplicity

From (2) we can see how the principal's action affects information aggregation; the higher is β_1^a, i.e. the more sensitive the asset value is to the state, the smaller the coefficient on θ in the equilibrium statistic. As a result, the price is less informative about the state. This is because when β_1^a is high, each trader's private signal is less informative about the asset value. As a result, traders place less weight on their private information relative to the information revealed by the price. The linear statistics for a fixed action $a \in \mathcal{A}$ are pictured in Figure 5 (in this figure the sign of the supply shock has been reversed, so that prices are increasing in the usual Euclidean product order). The slope of the linear statistics, is $-\frac{1}{\beta_1^a}\int_i \frac{\tau_i}{\sigma_i^2} di$, which again illustrates that the price reveals more precise information about ω the lower is β_1^a.

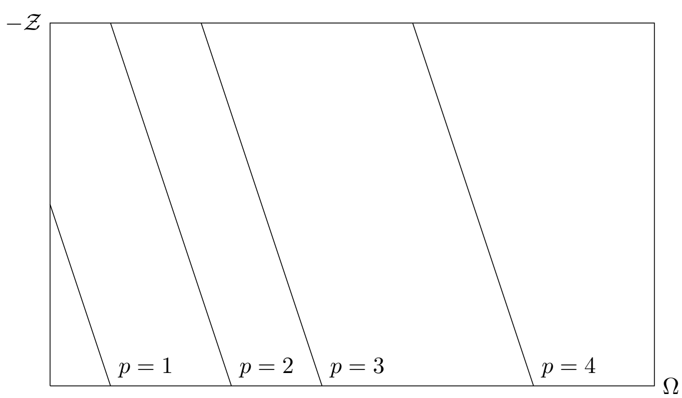

Figure 5: Linear statistics, fixed action $a \in \mathcal{A}$

The proof of Proposition 5 also yields an expression for $R(a, \theta, z)$, although for the current purposes it is sufficient to note simply that such a function exists and is strictly increasing (with the product partial order on $\Theta \times \mathcal{Z}$).

Corollary 1. *Assume CARA utility, π affine in θ, additive normal signal structure and truncated-normally distributed supply shocks. Then the market admits a reduced form representation under uniqueness, given by $R : \mathcal{A} \times \Theta \times \mathcal{Z} \to \mathcal{P}$. Moreover, $(\theta, z) \mapsto R(a, \theta, z)$ is strictly increasing for all a.*

There are three differences between the environment of Proposition 5 and that of Breon-Drish (2015) Proposition 2.1. First, the signal σ_i observed by each investor is given by the state plus noise, as opposed to the asset return plus noise as in Breon-Drish (2015). This is immediately handled by a suitable change of variables, given the assumption that $\theta \mapsto \pi(a, \theta)$ is affine for all a. Second, we allow here for the supply shock to follow a truncated normal distribution, where Breon-Drish (2015) considers only the un-truncated distribution. This requires generalizing Breon-Drish (2015) Proposition 2.1, which is relatively straightforward. Finally, and most importantly, the current setting features a feedback effect, whereas asset returns follow a fixed distribution in Breon-Drish (2015). We show how the results for the fixed-action case imply the desired result when there is feedback.

The approach here is similar to Siemroth (2019). However that paper assumes that the asset value is additively separable in the state and the principal's action. This is more than a technical assumption; it implies, as the author demonstrates, that the information revealed by the price is the same in all equilibria, regardless of the principal's actions. In contrast, we show precisely how the relationship between the principal's action and the asset value affects the information content of the price. This connection between the principal's action and the type of information revealed by the price has important implications for equilibrium multiplicity, as discussed below (e.g. Lemma 5). Siemroth (2019) also restricts attention to equilibria in which the price function is continuous, which has substantive implications, as discussed in the introduction.

Given Corollary 1, we can apply Proposition 3 to the noisy REE setting. The problem of finding optimal policies is generally complicated by the additional restrictions *iii.* and *iv.*, relative to the uni-dimensional case. In some cases however, these constraints simplify the problem. For example, if the supports of the noise term and the state are unbounded, these conditions and the expression for L^* in (2) have the following implication.

Lemma 5. *In the noisy REE model with normally distributed supply shocks and unbounded Θ, any CUI action function must be such that $\beta_1^{Q(\theta,z)} = \beta_1^{Q(\theta',z')}$ for all $(\theta, z), (\theta', z') \in \Theta \times \mathcal{Z}$.*

In other words, Lemma 5 says that any CUI action function can only use actions for which the slope of the asset payoff with respect to the state is the same. This will not be true when the supply shocks are bounded; in this case additional action functions will be CUI.

5 Properties and extensions

In this section we discuss properties of CUI policies, and study optimal policy when the unique implementation requirement is relaxed.

5.1 Structural uncertainty

Another practical concern of the principal, aside from manipulation and multiplicity, 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 translates into uncertainty about the function R. It is therefore desirable that the decision rule be robust to such perturbations, at least when the degree of uncertainty is small.[29]

Assume throughout this section that Θ is closed and bounded. Endow the space of market-clearing functions $R : \mathcal{A} \times \Theta \to \mathbb{R}$ with the sup-norm. 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 market clearing in state θ.

An open neighborhood of $\tilde{Q}_R(\cdot|M)$ is a set-valued and open-valued correspondence $U : \Theta \rightrightarrows 2^{\mathcal{A}}$ such that $\tilde{Q}_R(\theta|M) \subset U(\theta)$ for all θ. The map $R \rightrightarrows \tilde{Q}_R(\theta|M)$ is *uniformly continuous* at R if it is uniformly upper and lower hemicontinuous. That is, for any open neighborhood U of $\tilde{Q}_R(\cdot|M)$ and any open-valued correspondence $V : \Theta \rightrightarrows 2^{\mathcal{A}}$ such that $\tilde{Q}_R(\theta|M) \cap V(\theta) \neq \varnothing$ for all θ, there exists an open neighborhood N of R such that $\hat{R} \in N$ implies, for all $\theta \in \Theta$, *i)* $\tilde{Q}_{\hat{R}}(\theta|M) \subset U(\theta)$, and *ii)* $\tilde{Q}_{\hat{R}}(\theta|M) \cap V(\theta) \neq \varnothing$.

For any $S \subseteq \Theta$ let $\tilde{Q}_{R|S}$ be the restriction of $\tilde{Q}_R$ to S. Say that $R \rightrightarrows \tilde{Q}_R$ is *almost uniformly continuous* at R if $\forall\, \varepsilon > 0\ \exists\, S \subseteq \Theta$ with $\lambda(S) > 1 - \varepsilon$ such that $R \rightrightarrows \tilde{Q}_{R|S}(\theta|M)$ is uniformly continuous at R (where S replaces Θ in the definition of uniform continuity).

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

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

The interpretation of robustness to structural uncertainty 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 CUI (CWUI) implies implementability via a decision rule that is (weakly) robust to structural uncertainty.

[29] An alternative approach to additional dimensions of uncertainty is to model them explicitly. We show how this can be done in Section 3.1

Theorem 3. *If (Q, P) are CUI then they are implementable given market-clearing function R with an essentially continuous decision rule that is robust to multiplicity and structural uncertainty at R. If (Q, P) are CWUI then they are implementable with an essentially continuous decision rule that is weakly robust to multiplicity and weakly robust to structural uncertainty at R.*

Proof. Proof in Appendix C.1.1 □

The important implication of Theorem 3 is that small changes in R lead to small changes in the principal's expected payoff. Despite the fact that under the perturbed market-clearing function R' there may be multiple equilibria, the joint distribution of states, prices and actions associated with each one will be close to that of the original equilibrium under R.

If M is robust to multiplicity but has discontinuities on $\bar{P}_M$ then it will not be robust to structural uncertainty, at least when the discontinuity is not essential, i.e. when the left and right limits of M exist.[30] As discussed in Section 2, this further motivates the restriction to essentially continuous decision rules. Let $\theta_M(p|R) = \{\theta \in \Theta : R(M(p), \theta) = p\}$ be the set of states at which p could be an equilibrium price under M and R, and let $\bar{P}_M(R) := \{p \in \mathcal{P} : \theta_M(p) \neq \varnothing\}$ be the set of prices that could arise in equilibrium.

Lemma 6. *Assume that M satisfies robustness to multiplicity. If M has a non-essential discontinuity on $\bar{P}_M(R)$ then it is not robust to structural uncertainty at R.*

Proof. Proof in Appendix C.1.2. □

Lemma 6 shows that essential continuity is, to an extent, necessary for robustness to structural uncertainty.

5.2 Beyond uniqueness

When non-fundamental volatility is not a primary concern, the principal may be willing to tolerate the existence of multiple equilibria. This may be the case, even if the principal is not able to control which equilibrium will be played, provided all possible equilibria are "good" from their perspective. The following two results allow us to use the previous characterizations to study problems in which the strict uniqueness requirement is not imposed. They key insight is that even if a decision rule induces multiple equilibria, at least one of these will be weakly uniquely implementable. This is established via the following intermediate result.

Proposition 7. *Assume R is weakly increasing in θ. If $M \in \mathcal{M}$ induces multiple equilibria then at least one has a monotone price function (strictly monotone if R is strictly increasing in θ).*

[30] Given that $\mathcal{A}$ is compact, an essential discontinuity can be pictured as a point at which M oscillates with vanishing wavelength. The only potential benefit to the principal of using a discontinuous M is to avoid multiplicity, but an essential discontinuity is not useful in this regard.

Proof. In Appendix C.2.1 □

Theorem 1 says that monotonicity of the price function is a necessary condition for CWUI. Monotonicity is not in general sufficient. However, if we know that P is monotone *and* is induced by some $M \in \mathcal{M}$ then monotonicity of P suffices for CWUI in many settings. This is the case when the environment is fully bridgeable, as defined in Section 3. Under this assumption, any increasing selection from the price functions induced by M is CWUI.

Theorem 4. *Assume R is strictly increasing in θ and the environment is fully bridgeable. If $M \in \mathcal{M}$ induces multiple equilibria then at least one is characterized by (Q, P) that are CWUI.*

Proof. Proof in Appendix C.2.4. □

The important implication of Theorem 4 is that if the principal takes a strict worst case view of multiplicity then it is without loss of optimality to restrict attention to CWUI outcomes. That is, if the principal evaluates a decision rule M according to the worst equilibrium that it induces, then the principal may as well restrict attention to M that are weakly robust to multiplicity.

The conclusion of Theorem 4 can be extended in two ways. First, the result extends to weakly increasing R, when the environment satisfies a slightly stronger notion of bridgeability. Second, since CWUI price and action functions can generally be very well approximated by CUI price and action functions, we can replace CWUI with virtually CUI in the conclusion of Theorem 4. This requires some mild additional conditions, which guarantee that any CWUI (Q, P) can be approximated arbitrarily well by some CUI $(\hat{Q}, \hat{P})$. These conditions are satisfied, for example, in the setting with $\mathcal{A} = \Delta(Z)$ discussed above. We omit the formal statement of these results in the interest of brevity.

Theorem 4 also simplifies the problem of a principal who takes a less extreme approach to multiplicity than the strict worst-case preferences described above. Consider a principal who lexicographicaly evaluates policies which induce 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 that each induces (or indeed some other function of the remaining equilibria).[31] By Theorem 4 we know that the highest worst-case guarantee is exactly the maximum payoff over the subset of decision rules in $\mathcal{M}$ that are weakly robust to multiplicity. Once this value has been determined, the goal of the principal is to choose the decision rule with the best equilibrium outcome, subject to not inducing any equilibrium with a payoff below this 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 (a similar discussion will apply to

[31] 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).

other preferences). Assume that there is a unique optimal CWUI action function Q^*, implemented uniquely by decision rule M^* (if there are multiple optimal CWUI action functions then Condition 1 in Proposition 8 below must hold for one of them). If this is the case then, by Theorem 4, the principal with lexicographic preferences wants to choose a decision rule that implements Q^* as one of its equilibrium outcomes; if Q^* is not one of the equilibrium outcomes then there will be some other CWUI action function induced by the decision rule, which will be worse than Q^* by definition. 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. This discussion implies the following.

Proposition 8. *Let $\mathcal{Q}^*$ be the set of optimal CWUI action functions. Then the optimization constraints of the principal with lexicographic multiplicity preferences can be stated as follows: choose $\hat{M}$ subject to*

1. *$\exists Q \in \mathcal{Q}^*$ such that $\hat{M}(R(Q(\theta), \theta)) = Q(\theta)$ for all $\theta \in \Theta$,*
2. *$\hat{M} \in \mathcal{M}$.*

As illustrated in the application of Section 6.2, these constraints can greatly simplify the problem of finding optimal policies for a principal with lexicographic preferences over multiple equilibria.

6 Applications

6.1 Variable-volume carbon credits

Consider the problem of the emissions regulator discussed in the introduction. Such a policy is referred to by Karp and Traeger (2021) as a "smart cap".[32] The socially optimal level of emissions is determined by the marginal cost to firms of reducing their emissions, know as the abatement cost, and the marginal social benefit of reducing emissions. Assume that the regulator knows the social benefit of reducing emissions, but does not know firms' abatement costs.[33] Firms have private information about these costs.

Let q be the quantity of "clean air" produced by society. The societal benefit of clean air is given by $B(q)$. The social cost of producing q units of clean air is unknown to the regulator. This cost depends on the cost to emissions-producing firms of reducing their emissions. We parameterize the cost by $C(q, \theta)$, where θ is unknown to the regulator.

Under a variable-volume credits policy the regulator issues a unit mass of credits. The regulator's action space $\mathcal{A} = [0, 1]$ is the per-credit emissions volume allowance. If the per-credit volume

[32]Karp and Traeger (2021) show that a smart cap can implement the regulator's first best, but do not consider uniqueness and manipulation constraints.

[33]In reality, there may also be uncertainty about the social benefit of reducing emissions.

allowance is a, the quantity of clean air is given by $1-a$. The regulator's decision rule specifies the per-credit volume as a function of the price for credits.

There are a continuum of firms $i \in [0,1]$. Each firm observes its own cost type s_i. The distribution of costs in population is F_θ, where $\theta \in [0,1]$ and $\theta \mapsto F_\theta$ is increasing in the FOSD order. A firm's payoff is given by $u(a \cdot x, s_i) - p \cdot x$, where a is per-credit volume and x is number of credits purchased. Assume u is continuous; strictly increasing and strictly concave in its first argument; and strictly decreasing and convex in its second argument. Notice that this is a private-values setting; a firm's payoff does not depend directly on the abatement costs of others. Credits are traded in a competitive market. Denote the firm's demand by

$$X(p, a, s_i) = \arg\max_x \quad u(a \cdot x, s_i) - p \cdot x.$$

Demands are unique under the maintained assumptions on u, and strictly decreasing in p. Given an action function $Q : \Theta \mapsto \mathcal{A}$, it must be that the equilibrium in state θ is the unique value satisfying

$$\int_\theta X(p, Q(\theta), s) dF_\theta(s) = Q(\theta) \tag{3}$$

Thus condition (3) implicitly defines a reduced-form representation for the credits market, where the market-clearing function $R(a,\theta)$ is continuous, strictly decreasing in its first argument, and strictly increasing in its second.

The regulator's first-best action function is given by

$$Q^*(\theta) = \arg\max_a B(1-a) - C(1-a, \theta).$$

Assume that $\theta \mapsto C_1(q,\theta)$ is continuous and strictly increasing. Then the first-best cannot be implemented by setting prices or quantities alone, since Q^* is strictly increasing. However since Q^* is continuous and the associated price function P^* is strictly increasing, the first-best can be implemented uniquely by a decision rule that is robust to manipulation, by Proposition 9. The implementing decision rule in this case is continuous and strictly increasing.

In fact, since the first best price function P^* fully reveals the state and the first-best is CUI, the first-best here can also be implemented even if the principal lacks commitment power. Formally, $M(p) = Q^*(P^{*-1}(p))$. In other words, in equilibrium the principal learns the state perfectly, and does not need to commit to take an ex-post sub-optimal action in order to induce this equilibrium.

6.2 Bailouts

Consider a government deciding on the size of a bailout for a publicly traded company.[34] We show here how features of the environment, such as the size of the strategic spillovers from the company

[34] Alternatively, the bailout could be for an entire industry, in which many of the firms are publicly traded. The analysis here is also very similar to the problem of an international lender using sovereign debt prices to inform emergency lending decisions, as discussed in Section 2.4.

to the rest of the country, determine whether the lender's first-best can be achieved, and shape how responsive the loan amount is to the bond price under the lender's optimal policy.

The government chooses a level of support $a \in \mathcal{A} = [0, \bar{a}]$. The company's business prospects $\theta \in \Theta$, representing the demand environment, competition, future costs, etc., are unknown. Higher states represent better prospects; for each level of support the share price is a strictly increasing function of the state. For exposition purposes, we work directly with the reduced-form representation, with the understanding that this can be derived from a market as described in Section 4. We make two assumptions regarding the share price.

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, or carries negative stigma (Che et al., 2018). An alternative interpretation is that the bailout takes the form of forgivable loans. As a result of ex-post loan forgiveness, the effective amount owed is increasing in the state (which will be revealed *ex post*). 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.[35]

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. Additionally, 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. The government's payoff $u(a, \theta)$ is therefore decreasing in a for $\theta \notin (\theta', \theta'')$. 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 government maximizes expected utility, and has an absolutely continuous prior H.

Figure 6 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^*. Assumption 2 on R (above) implies

[35] 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).

that the environment is fully bridgeable (see Appendix C.3). Thus the first best is CWUI by Proposition 2.

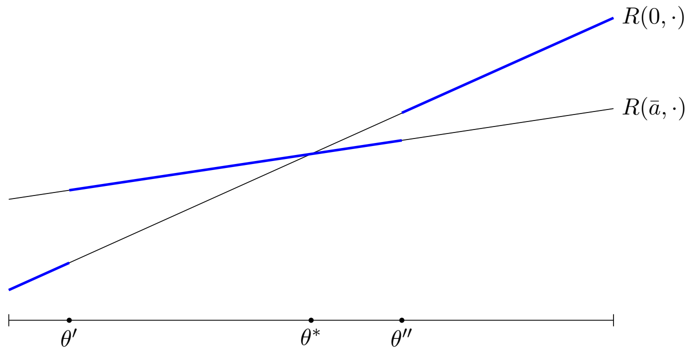

Figure 6: First-best is CWUI

If $\theta^* < \theta''$ the government is relatively interventionist. 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 company is considered highly strategically important, 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 CWUI when $\theta^* < \theta''$, the government 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 all such prices, M must satisfy $p = R(M(p),\theta')$, which implies that the government must gradually increase the size of the bailout as a function of the price. If the policy responds too rapidly price increases in this range, then there will equilibria in which action $a > 0$ is taken for states below θ'. On the other hand, if the government under-responds then there will be equilibria in which action $a < \bar{a}$ is taken for states above θ'. Similar restrictions apply 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 7. In this case the price is non-monotone, and is therefore not CWUI. In fact, in this case it is not even implementable, as it violates measurability. The optimal CWUI outcome is found by flattening the price function to eliminate non-monotonicity.

The price function for the virtually optimal decision rule is pictured in Figure 8 (it will only be virtually optimal since the price must be strictly increasing, but can have an arbitrarily small slope). It is fully characterized by a state $\hat{\theta}$ at which the flattening begins. In order to approximately implement such an equilibrium, the government uses a decision rule such that the size of the bailout

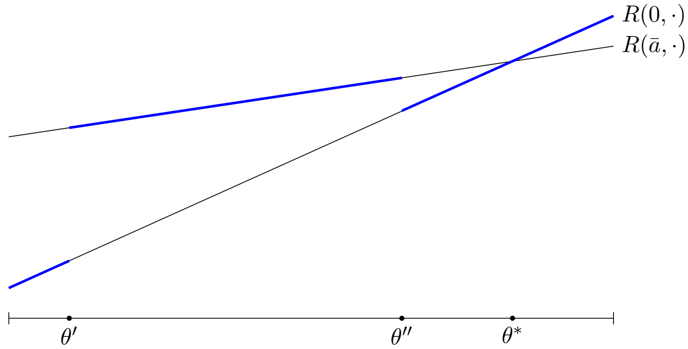

Figure 7: First-best not implementable

decreases rapidly from 0 to $\bar{a}$ over an interval $[R(\bar{a},\hat{\theta}), R(\bar{a},\hat{\theta})+\varepsilon]$, where ε can be made arbitrarily small. In other words, when the government is less inclined to intervene ($\theta^* > \theta''$) the optimal policy will be *more* responsive to the price, compared to the case in which the government is predisposed to intervene ($\theta^* < \theta''$).

The optimal policy is easily identified analytically via a first-order condition. 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.

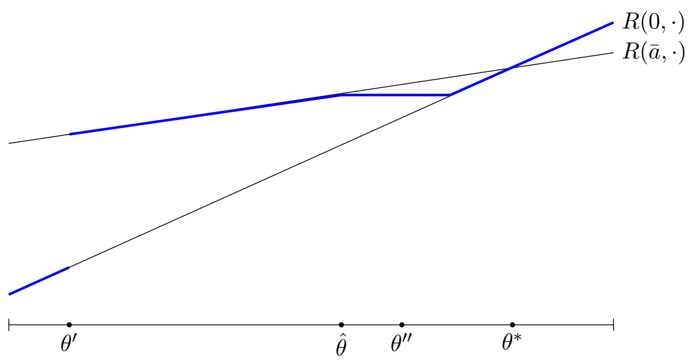

Figure 8: virtually optimal decision rule

The virtually optimal policy can involve two types of loss for the government: under-support for the company for states below θ'', or over-support for states above θ''. In fact, the optimal policy will entail both types of loss. To see this, suppose $\hat{\theta} = \theta''$. There is a first-order gain from lowering $\hat{\theta}$ by a small ε, since this means that less support needs to be offered on the entire interval

$(\theta'', t(\theta''))$. The loss, which results from less support being offered on $(\theta'' - \varepsilon, \theta'')$, is second order. An analogous argument applies to raising $\hat{\theta}$ when $t(\hat{\theta}) = \theta''$.

We can also use the results of Section 5.2 to study optimal policy if the government is willing to tolerate multiple equilibria. Suppose that in the case of $\theta^* > \theta''$ the government is willing to tolerate some multiplicity, and takes the lexicographic approach described in Section 5.2. 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 5.2, the virtually optimal decision rule is pined down on $P^*(\Theta)$. 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 8 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 virtually optimal policy.

6.3 Moving against the market

In this section, we explore the distinctive features of a set of applications in which the principal would like to induce a *decreasing* price. As before, $\theta \mapsto R(a, \theta)$ is increasing. These are therefore situations in which the principal is working to move prices against the market. The following are two such instances.

Monetary policy in a crisis

During the financial crisis of 2008 and the Covid-19 recession of 2020, central banks moved aggressively to lower interest rates. In this application the unknown state is the severity of the liquidity crisis faced by firms, and the market price is the interest rate. The action is the size of asset purchases made by the central bank through open market operations. The central bank's objective is to implement an interest rate that is decreasing in the state via their open market operations.

Grain reserves

Many developing countries manage grain reserves as a tool for stabilizing the grain price and responding to food shortages. The state here is the size of a demand or supply shock, the price is the grain price, and the action is the size of grain purchases/sales. Depending on the nature of the crisis and the structure of the grain market, the government may wish to implement a decreasing price. If the government has limited capacity to make direct transfers to households it may wish to implement transfers by lowering the grain price when there is a severe crisis. For example, suppose that grain is a Giffen good. If there is an employment crisis outside of agriculture the price of grain may rise, absent government intervention.[36] In this case the government may wish to subsidize non-agricultural households by lowering the grain price.

Throughout this section, we maintain the assumptions that $\mathcal{A} = [\underline{a}, \bar{a}] \in \mathbb{R}$ and that $\theta \mapsto R(a, \theta)$ is strictly increasing for all a, and that $a \mapsto R(a, \theta)$ is strictly decreasing for all θ (that this function is decreasing as opposed to increasing is simply a normalization). A deceasing price function is possible if and only if $R(\underline{a}, \underline{\theta}) > R(\bar{a}, \bar{\theta})$. Figure 9 depicts such an environment.

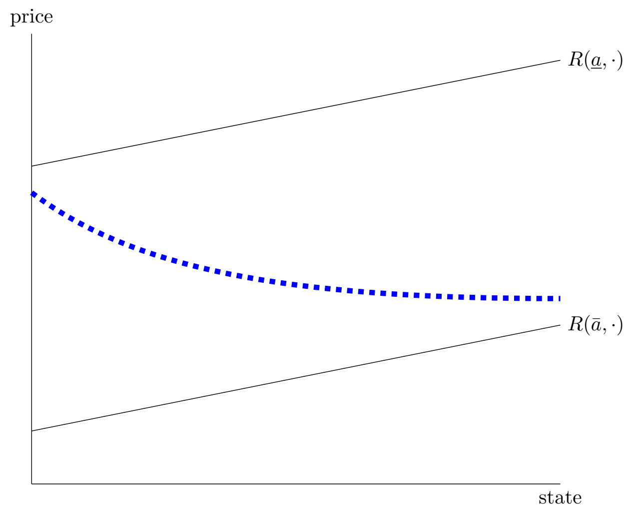

Figure 9: Decreasing price function

The following observation shows that implementing an increasing price function in this setting is easy.

Lemma 7. *If $a \mapsto R(a, \theta)$ is strictly decreasing for all θ then any strictly increasing $M \in \mathcal{M}$ induces an increasing and continuous price function as the unique equilibrium.*

Proof. Proof in Appendix D.1. $\square$

An equilibrium exists for any increasing M by Tarski's fixed point theorem. That the price function is increasing follows from the fact that $a \mapsto R(a, \theta)$ is decreasing and $\theta \mapsto R(a, \theta)$ is increasing. If P is increasing and M is increasing, there will be no equilibrium involving prices above $P(\bar{\theta})$ or below $P(\underline{\theta})$. Moreover, we show that M cannot have a discontinuity on $[P(\underline{\theta}), P(\bar{\theta})]$, which implies that P is continuous.

Decreasing price functions are more interesting in this setting. Non-monotonicity of M will be necessary to robustly implement a decreasing price.

Lemma 8. *Assume $a \mapsto R(a, \theta)$ is strictly decreasing for all θ, and let P be a decreasing price function. If $M \in \mathcal{M}$ uniquely implements P then*

[36] There is empirical evidence that food staples are Giffen goods for extremely poor households (Jensen and Miller, 2008).

i. $M(p)$ *is decreasing and continuous on an open interval containing* $(P(\bar{\theta}), P(\underline{\theta}))$,

ii. M *has discontinuities in* $(P(\underline{\theta}), R(\underline{a}, \underline{\theta})]$ *and* $(R(\bar{a}, \bar{\theta}), P(\bar{\theta})]$.

iii. *There exist* $p'' > p' > P(\underline{\theta})$ *such that* $M(p'') > M(p')$.

iv. *There exist* $p' < p'' < P(\bar{\theta})$ *such that* $M(p'') > M(p')$.

Proof. Proof in Appendix D.2. □

Lemma 8 shows that discontinuous and non-monotone M is necessary to implement a decreasing price. The intuition comes from the fact that the government is attempting to move against the market. Suppose the principal uses a strictly decreasing decision rule. If the lower bound $\underline{a}$ on the action is reached at some price p, then the principal will no longer have the capacity to move against the market for prices below p. Thus for such prices, the market forces generating an increasing price will dominate, and there will be multiple equilibria.

More formally, there are only two ways to guarantee that $\theta_M(p) = \varnothing$, i.e. that there are no equilibria with price p. Either M must specify an action that is too high, meaning $R(M(p), \bar{\theta}) < p$, or too low, so that $R(M(p), \underline{\theta}) > p$. If neither of these hold then there will be some θ such that $R(M(p), \theta) = p$, by continuity of R. The only way to ensure that there are no equilibria with prices in $[R(\underline{a}, \underline{\theta}), R(\underline{a}, \bar{\theta})]$ is to take a high enough action for such prices; it must be that $R(M(p), \bar{\theta}) < p$ for all such prices. At the same time, M must be decreasing on $(P(\bar{\theta}), P(\underline{\theta}))$ in order to implement a decreasing P. This tension is what necessitates discontinuities and non-monotonicities in M.

Lemma 8 is important in applications because it highlights the danger of artificially restricting the class of permissible decision rules. If, for example, one restricts attention to monotone decision rules, it will not be possible to uniquely implement a decreasing price. It is nonetheless common practice in the literature to focus on monotone, or even linear, decision rules (see for example Bernanke and Woodford (1997)). Most papers which make this type of linearity assumption do so in models where the action space is unbounded. The fact that the action space is bounded here is an important driver of the non-monotonicity result in Lemma 8. However in reality there are often bounds on available set of actions.[37] In the grain reserves example, the government cannot sell more grain than it has in reserve. Similarly, central banks in developing countries cannot make unlimited asset purchases without creating significant balance sheet risks (Crowley, 2015). Lemma 8 shows that such restrictions on the feasible actions can interact in surprising ways with conditions on the decision rules used to gain tractability. Our general framework allows us to avoid the need to impose such conditions.

[37] Notice that the conclusion of Lemma 8 does not depend on how "tight" the bounds are; non-monotone decision rules are necessary even if the range of admissible actions is arbitrarily large.

7 Directions for future work

In this paper we provide a full characterization of the set of price and action functions that can be implemented uniquely using a market-based decision rule that satisfies a notion of robustness to manipulation. We hope that the results of this paper, along with the framework that we develop, will be useful in a range of applications, and will facilitate new theoretical insights. Here we briefly describe a few extensions which we plan to study in future work.

In the current paper, we took the market as given. However, it may be possible for the principal to directly shape the market game, beyond their choice of decision rule. For example, the principal could create derivatives based on some primitive assets, and condition policy on the prices for these derivatives. If the market admits a reduced-form representation, this translates into changes to the market-clearing function R. Our results can be applied immediately to such problems.

Another natural extension is to settings in which the principal can condition on multiple prices. For example, a central bank may condition policy decisions on both inflation and the unemployment rate. In preliminary work, we use the tools developed in the current paper to analyse the case of a multidimensional prices. Bond et al. (2010) study a principal who uses a one-dimensional price to inform policy, but has access an *exogenous* signal about the state. The authors show that exogenous information can help the principal circumvent measurability constraints. Additional prices may play a similar role, however the fact that their informativeness is endogenously determined affects their usefulness to the principal.

We are also interested in dynamic market-based design problems. The current paper restricted attention to a one-off decision by the principal. However it will be necessary to model inter-temporal effects if the principal must take actions in multiple periods, and future actions affect the market today. We hope that the results of the current paper will be useful in this setting. Since we have placed very few restrictions on the action space $\mathcal{A}$, it may be possible to use our results to analyse the recursive formulation of the principal's dynamic program. Under this approach, an action a of the principal would consist of the current period decision, as well as the plan for how future decision will react to (present and future) prices.

The current paper focuses on markets in which information is dispersed, and the principal is able to elicit this information only by observing the market outcome. In some markets, however, there may be large identifiable agents with whom the principal can contract directly. For example, in the corporate bailouts example, the regulator may be able to contract directly with the firm to elicit the firm's private information about its degree of distress. This is is a classical mechanism design problem. We are interested in settings in which the principal seeks to contract directly with some agents, while also engaging in market-based design of the type studied in the current paper. The interaction between these two types of design seems a promising ground for future work.

Appendix

A Omitted proofs from Section 3

We will use $\theta_M(p)$ to indicate the states that are consistent with a price p, given a policy function M, i.e., $\theta_M := \{\theta \in \Theta : P_M(\theta) = p\}$. For markets that admit a reduced-form representation,

$$\theta_M(p) := \{\theta \in \Theta : R(M(p), \theta) = p\}$$

Sometimes we will drop the subscript M when the policy function is fixed.

A.1 Preliminary results

Lemma 9. *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 $\theta \in (\theta', \theta'')$. □

Lemma 10. *Fix a continuous M. Assume $\theta \mapsto R(a, \theta)$ is continuous for all a. Then each p such that $\theta(p) = \varnothing$ is of one and only one of the following two types:*

- *Type L:* $\quad R(M(p), \theta') > p \qquad \forall \theta' \in \Theta$.
- *Type H:* $\quad R(M(p), \theta') < p. \qquad \forall \theta' \in \Theta$.

Proof. If p is of neither type, there exists a pair of states θ', θ'' such that $R(M(p), \theta') - p > 0 > R(M(p), \theta'') - p$. Then by continuity, there is a state $\theta \in (\theta', \theta'')$ such that $R(M(p), \theta') - p = 0$. But then $\theta(p)$ is not empty. □

Corollary 2. *For any continuous M, the set of prices $\{p : \theta_M(p) = \varnothing\}$ is open.*

Lemma 11. *(Generalized intermediate value theorem). Let $F : [0, 1] \to [0, 1]$ 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. Assume that $y_2 > y_1$ (the case with $y_2 = y_1$ is trivial an $y_2 < y_1$ is symmetric). We prove by contrapositive: assume that there exists a $\tilde{y} \in (y_1, y_2)$ such that $\tilde{y} \notin F(p)$ for all $p \in [p_1, p_2]$. Since $F(p)$ is convex, for every p either $\max F(p) < \tilde{y}$ or $\min F(p) > \tilde{y}$. Let $p^* = \sup\{p \in [p_1, p_2) : \max F(p) < \tilde{y}\}$.

Suppose that $\max F(p^*) < \tilde{y}$. Notice that this is only compatible with $p^* < p_2$. Consider the open set $V := (\min F(p^*) - \epsilon, \max F(p^*) + \epsilon)$ with $\epsilon < \tilde{y} - \max F(p^*)$. By upper hemicontinuity, there exists a neighbourhood of p^* such that $F(p) \subset V$ for all p in such neighbourhood. Thus, in a neighbourhood of p^*, $F(p) < \tilde{p}$, what violates the definition of p^*.

Suppose that $\min F(p^*) > \tilde{y}$. Notice that this is only compatible with $p^* > p_1$. Using upper hemicontinuity as before, we get that there is a neighbour of p^* such that $F(p) > \tilde{p}$ for all p in that neighbourhood, what violates the definition of p^*. $\square$

Lemma 12. *$p \mapsto \theta_M(p)$ is compact-valued. If M is continuous at p' then $p \mapsto \theta_M(p)$ is upper hemicontinuous at 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 $\theta_M(p)$. 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 $\theta_M(p') \subseteq V$, for all p'. Thus $p \mapsto \theta_M(p)$ is upper hemicontinuous. $\square$

A.2 Proof of Theorem 1

Proof. Suppose that P is non-monotone. Then there are states $\theta_1 < \theta_2 < \theta_3$ such that either $P(\theta_2) > \max\{P(\theta_1), P(\theta_2)\}$ or $P(\theta_2) < \min\{P(\theta_1), P(\theta_2)\}$. Suppose the latter, and assume that $P(\theta_3) < P(\theta_1)$ (the other cases are symmetric). The result follows from two intermediate claims.

First, it is convenient expand the range of θ_M. Define $\tilde{\theta}_M = \{\theta \in [\underline{\theta}, \bar{\theta}] : R(M(p), \theta) = p\}$, where $R(a, \underline{\theta}) = \lim_{\theta \to \underline{\theta}} R(a, \theta)$ and $R(a, \bar{\theta}) = \lim_{\theta \to \bar{\theta}} R(a, \theta)$ are well defined and continuous given our assumptions on R.

Claim 1. There exists $p'' < P(\theta_2)$ such that at least one of the following holds:

1. For all $\theta \in (\theta_2, \bar{\theta})$ there exists $p \in [p'', P(\theta_2)]$ such that $\theta \in \theta_M(p)$.

2. For all $\theta \in (\underline{\theta}, \theta_2)$ there exists $p \in [p'', P(\theta_2)]$ such that $\theta \in \theta_M(p)$.

Proof of Claim 1: There are two cases to consider. First, suppose $\{p \leq P(\theta_2) : \tilde{\theta}_M(p) = \varnothing\} = \varnothing$. Let $\underline{p} = \min_{a \in \mathcal{A}} R(a, \underline{\theta})$. Then $R(M(p), \underline{\theta}) \leq p$ for all $p \in [\underline{p}, P(\theta_2)]$, otherwise there would be some p in this interval with $\tilde{\theta}_M(p) = \varnothing$. Moreover, by definition $\underline{p} \leq R(M(\underline{p}, \underline{\theta})$. Therefore it must be that $\underline{p} = R(M(\underline{p}, \underline{\theta})$, so $\underline{\theta} \in \tilde{\theta}_M(\underline{p})$. Condition 2 of the claim follows from Lemma 11.

Second, $\{p \leq P(\theta_2) : \tilde{\theta}_M(p) = \varnothing\} \neq \varnothing$, and let $p'' = \sup\{p \leq P(\theta_2) : \tilde{\theta}_M(p) = \varnothing\}$. Then $p'' < P(\theta_2)$ since M is continuous in a neighborhood of $P(\theta_2)$. Additionally, either $\bar{\theta} \in \tilde{\theta}_M(p'')$ or $\bar{\theta} \in \tilde{\theta}_M(p'')$, since $\tilde{\theta}_M(p) = \varnothing$ iff either $R(M(p), \underline{\theta}) > p$ or $R(M(p), \bar{\theta}) < p$ and M is continuous around p''. If $\bar{\theta} \in \tilde{\theta}_M(p'')$ then condition 1 in the claim follows from Lemma 11. If $\underline{\theta} \in \tilde{\theta}_M(p'')$ then condition 2 in the claim follows from Lemma 11.

Claim 2. There exists $p' \in [P(\theta_3), P(\theta_1)]$ such that at least one of the following holds:

1. For all $\theta \in (\theta_1, \theta_3)$ there exists $p \in [p', P(\theta_1)]$ such that $\theta \in \theta_M(p)$.

2. For all $\theta \in (\underline{\theta}, \theta_1)$ there exists $p \in [p', P(\theta_1)]$ such that $\theta \in \theta_M(p)$.

Proof of Claim 2: If $\theta_M(p) \neq \varnothing$ for all $p \in [P(\theta_3), P(\theta_1)]$ then let $p' = P(\theta_3)$. Condition 1 of the claim then follows from Lemma 11. Otherwise, let $p' = \sup\{p \leq P(\theta_1) : \tilde{\theta}_M(p) = \varnothing\}$. The proof is then identical to that of Claim 1.

Combining Claims 1 and 2, the necessity of monotonoicity of P follows from the fact that there is multiplicity on a positive measure set of states for any combination of the conditions in the claims. □

A.3 Proof of Theorem 2

Lemma 13. *Assume R is weakly increasing in θ. For any $M \in \mathcal{M}$ that is robust to multiplicity, Let $p_1 < p_2$ such that there are states $\underline{\theta}$ and $\bar{\theta}$ with $\underline{\theta} < \theta < \bar{\theta}$ for each $\theta \in \theta(p_1) \cup \theta(p_2)$. Then $[p_1, p_2] \in P(\Theta)$.*

Proof. By Theorem 1, the price function P is monotone, so without loss of generality assume that it is increasing, and let $p_1, p_2 \in P(\Theta)$ with $p_2 > p_1$. Assume towards a contradiction that there exists $p \in (p_1, p_2)$ such that $p \notin P(\Theta)$. By Lemma 10 p is either type H or type L. Suppose it is type L, i.e. $R(M(p), \theta) - p > 0$ for all θ. 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 12 (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 11 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. □

Proof. ($\Rightarrow$) Part 1 stems trivially from the market clearing condition necessary for implementation (see Observation 1).

Theorem 1 states that P must be *weakly* monotone. To prove strict monotonicity (part 2) consider $P(\theta) = P(\theta')$. Then, $R(Q(\theta), \theta) = R(Q(\theta'), \theta')$. By measurability, $Q(\theta) = Q(\theta')$, so $R(Q(\theta), \theta) = R(Q(\theta), \theta')$ what, since R is strictly increasing in θ implies that $\theta = \theta'$. Thus, P is strictly monotone.

Now we prove that Q is continuous for any interior state. Since $R(a, \theta)$ is strictly monotone in θ, we have $|\theta_M(p)| \leq 1$ for all p. To see this, consider $\theta, \theta' \in \theta_M(p)$. This means that $R(M(p), \theta) = p = R(M(p), \theta')$ which, by strict monotonicity of R, means that $\theta = \theta'$.

For some interior state θ', let $p^- := \lim_{\theta \searrow \theta'} P(\theta)$ and $p^+ := \lim_{\theta \nearrow \theta'} P(\theta)$. Since M is essentially continuous, M is continuous in an open neighbourhood N of $P(\theta')$. This, together with continuity of R, implies that $\theta_M(p)$ is continuous on N. Thus, there is a neighbourhood of θ' such that $P(\theta) \cap N$ is not empty for all θ in the neighbourhood. Therefore, p^- and p^+ must be equal to $P(\theta)$ or multiplicity would be violated.

Given that P is continuous for interior states, a discontinuity of Q in a interior state will necessary imply a discontinuity of M for a price in $\bar{P}$, what would violate essential continuity.

Thus, Q must be continuous for all interior states.

P is monotone and bounded (below by $\min_{a\in\mathcal{A}} \underline{R}(a)$ and above by $\max_{a\in\mathcal{A}} \bar{R}(a)$), so $\underline{P} := \lim_{\theta\searrow\underline{\theta}} P(\theta)$ and $\bar{P} := \lim_{\theta\nearrow\bar{\theta}} P(\theta)$ exist. Let M be the policy function that continuously uniquely implements (Q, P). By essential continuity, M is continuous at $\underline{P}$, so $\lim_{p\searrow\underline{P}} M(p) = M(\underline{P})$. But then, since $Q(\theta) = M(P(\theta))$ for all θ, $\lim_{\theta\searrow\underline{\theta}} Q(\theta) = \lim_{\theta\searrow\underline{\theta}} Q(\theta) = Q(M(\underline{P}))$. Same arguments holds for the other extreme state $\bar{\theta}$.

Finally, for the case in which P is strictly decreasing, we need to show that $\underline{Q}$ is not maximal at the bottom, and $\bar{Q}$ is not minimal at the top. Since P is decreasing, for prices right above $\underline{P}$, $\theta_M(p)$ should be empty. $\bar{Q}$ is maximal at the bottom so $R(\cdot, \underline{\theta})$ has a local maximum at $\underline{Q}$. This means that there is a neighborhood around $\underline{Q}$ such that $R(q', \underline{\theta}) < p$ for all q' in the neighbourhood. By essential continuity, for prices slightly above p the action is in such neighbourhood. So for any $\varepsilon > 0$ there exists a $p' \in (p, p+\varepsilon)$ such that $R(M(p'), \underline{\theta}) \leq p$. Since $\theta \mapsto R(a, \theta)$ is strictly increasing and R is continuous, for ε small enough we will also have $R(M(p'), \bar{\theta}) > p'$. But then by continuity of R there exists θ such that $R(M(p'), \theta) = p'$, so $\theta_M(p')$ is not empty. A symmetric argument rules out Q being minimal at the top.

($\Leftarrow$) M can be easily defined on $P(\Theta)$ as follows. Since P is injective, define M on $P(\Theta)$ as $M(p) = Q(P^{-1}(p))$. Notice that M is continuous (by 1 and 3).

The challenge is to define the function M for prices outside $P(\Theta)$. The constructions will differ for increasing and decreasing P.

If P is increasing then define $M(p) = \bar{Q}$ for all prices above $\bar{P}$ and $M(p) = \underline{Q}$ for all prices below $\underline{P}$. We want to check that for all these prices $\theta_M(p) = \emptyset$. For prices above $\bar{P}$, $p \geq \bar{P} = \bar{R}(\bar{Q}) = \bar{R}(M(\bar{P})) > R(M(p), \theta)$ where the last inequality holds for all $\theta \in \Theta$. A symmetric argument proves that $p < R(M(p), \theta)$ for prices below $\underline{P}$. Thus, (Q, P) is CUI.[38]

Now for decreasing P, we need to show that there exists a continuous actions for prices right above $\underline{P}$ so that $\theta_M(p)$ is empty.

Consider a finite partition $\{A_i\}_{i=1}^k$ of $\mathcal{A}$ such that the sets $\{A_i \cap \underline{R}^{-1}(\underline{P})\}_{i=1}^k$ are connected. Moreover, by continuity, we can pick the partition $\{A_i\}_{i=1}^k$ such that the distance between two of the subsets is greater than zero: For A, A' two elements of the partition, if the distance between $A \cap \underline{R}^{-1}(\underline{P})$ and $A' \cap \underline{R}^{-1}(\underline{P})$ is zero, then there is a sequence of actions $\{a_i\}_{i=1}^\infty$ such that $a_i \in A \cap \underline{R}^{-1}(\underline{P})$ and $a = \lim_{i\to\infty} a_i \in A' \cap \underline{R}^{-1}(\underline{P})$. By continuity, the sets are connected. Thus, in a neighbourhood of $\underline{Q}$, $\underline{R}^{-1}(\underline{P})$ is connected. By continuity, this splits the neighbourhood of $\underline{Q}$ in sets for which $\underline{R}(a) > \underline{P}$ and sets for which $\underline{R}(a) < \underline{P}$. Since $\underline{Q}$ is not a local maximum, there exists at least one set for which $\underline{R}(a) > \underline{P}$ that is at a distance 0 of $\underline{Q}$.

Pick a continuous path $\hat{a} : [0,1] \to \mathcal{A}$ such that $\hat{a}(0) = \underline{Q}$ and $\hat{a}(t) \in \mathcal{A}^-$ for all $t > 0$. There exists an increasing function $h : [0,1] \to \mathcal{P}$ such that $h(t) < \underline{R}(\hat{a}(t))$. Thus, we can make $M(\underline{P} + t\epsilon) = h(t)$. Then for all $\tilde{P} \in (\underline{P}, \underline{P} + \epsilon)$, $R(M(\tilde{P}), \theta) > Q(M(\tilde{P}), \underline{\theta}) > \tilde{P}$.

[38] Moreover, any (Q, P) that is CUI and such that P is increasing, can be implemented by an M that is continuous.

As usual, a symmetric argument works to construct M for prices right below $\bar{P}$. Beyond these prices at the neighbourhood of $P(\Theta)$, essential continuity is not binding, so any actions that do not generate equilibria work for the construction. Notice that if for a price all actions generate an equilibria, then that price must be in $P(\Theta)$ by 1.

□

A.4 Closed state space

The following is the alternative to Theorem 2 when Θ is a closed interval. The price function can be discontinuous for the extreme states and violate monotonicity, as long as the function is still injective.

Proposition 9 (Closed Θ)**.** *Assume R is strictly increasing in θ. Then (Q, P) is CUI iff*

1. *$P(\theta) = R(Q(\theta), \theta)$ for all θ,*
2. *P is monotone at the interior of Θ, injective, and such that $P(\underline{\theta}) \neq \lim_{\theta \to \bar{\theta}} P(\theta)$, $P(\bar{\theta}) \neq \lim_{\theta \to \underline{\theta}} P(\theta)$.*
3. *Q is continuous at the interior of Θ and $\bar{Q} := \lim_{\theta \nearrow \bar{\theta}}$ and $\underline{Q} := \lim_{\theta \searrow \underline{\theta}}$ exist. Moreover, if P is decreasing in the interior of Θ, then $\underline{Q}$ is not maximal at the bottom and $\bar{Q}$ is not minimal at the top.*
4. *If Q is discontinuous at $\underline{\theta}$ then $Q(\underline{\theta})$ is not maximal at the bottom. If Q is discontinuous at $\bar{\theta}$ then $Q(\bar{\theta})$ is not minimal at the top.*

Proof. ($\Rightarrow$)

1 to 3 are essentially the same as the previous theorem and can be proven in the same way. The main difference in 2 is that P is not necessarily monotone. However, by continuity of Q in the interior of Θ and R, P will be monotone in the interior of Θ. P has to be injective for the same measurability reason as in the previous proof. Moreover, if $P(\underline{\theta}) = \bar{P} := \lim_{\theta \to \bar{\theta}} P(\theta)$. Then $R(M(\bar{P}), \underline{\theta}) = \bar{P}$. However, $P(\theta) = R(M(P(\theta)), \theta)$ so taking limits, $\bar{P} = R(M(\bar{P}), \bar{\theta}) > R(M(\bar{P}), \underline{\theta})$ by strict monotonicity of R.

Finally, there is an extra condition for the extreme states. If Q is discontinuous at $\underline{\theta}$ then so must be P (otherwise essential continuity would be violated). Moreover, we showed that $P(\underline{\theta}) \neq \bar{P}$. Thus, there exists a neighbourhood N of $P(\underline{\theta})$ such that $N \cap P(\Theta) = P(\underline{\theta})$. If $Q(\underline{\theta})$ is maximal at the bottom, then for prices slightly above $P(\underline{\theta})$ there would be multiple equilibria, for the reasons analyzed in part 3 of the previous result. Thus, $Q(\underline{\theta})$ is not maximal at the bottom, nor $Q(\bar{\theta})$ minimal at the top.

($\Leftarrow$) As before, the goal is to construct an M that continuously uniquely implements (P, Q).

Since P is injective, M is defined as before on $P(\Theta)$: $M(p) = Q(P^{-1}(p))$. Again, notice that M is continuous for prices associated with interior states (by 1 and 3). As before, the challenge is to define the function M for prices outside $P(\Theta)$.

If P is increasing at the interior of Θ, then define $M(p) = \bar{Q}$ for prices slightly above $\bar{P}$ and $M(p) = \underline{Q}$ for prices slightly below $\underline{P}$. As before, using these prices guarantees that $\theta_M(p) = \emptyset$. For P decreasing in the interior of Θ, we can show that there exists a continuous actions for prices right above $\underline{P}$ so that $\theta_M(p)$ is empty, provided that $Q(\theta)$ is not maximal at the bottom nor minimal at the top.

The extra concern is to define M for prices close to $P(\underline{\theta})$ and close to $P(\bar{\theta})$. When Q is continuous then there is no issue since the prices close to $P(\underline{\theta})$ and $P(\bar{\theta})$ are the same as those close to $\underline{P}$ and $\bar{P}$ that we already accounted for. When Q is discontinuous at $\underline{\theta}$, we can use $\underline{Q}$ for prices slightly below $P(\underline{\theta})$ and find actions for prices slightly above $P(\underline{\theta})$ such that $\theta_a(p)$ is empty, provided that $Q(\underline{\theta})$ is not maximal at the bottom. Symmetric argument for $\bar{\theta}$. For prices outside the neighbourhood of $P(\Theta)$ that was already described, essential continuity does not bind, and we can always find an action such that there is no state consistent with that action and price level.

□

A.5 Proof of Proposition 1

Proof. ($\Rightarrow$) The main difference with Theorem 2 lies in part 4. The first part is the measurability condition that was already necessary for implementation. For the second part, first notice that if $P(\theta) = R(Q(\theta), \theta')$ then $P(\theta') = P(\theta)$: otherwise there are multiple equilibria at the state θ', one with price $P(\theta)$ and one with price $P(\theta')$. Measurability implies that $Q(\theta) = Q(\theta')$.

($\Leftarrow$) We cannot construct the function M that continuously uniquely implements (P, Q) in the same as the one in Theorem 2 since P is not necessarily injective: $P^{-1}(p)$ is not necessarily a singleton anymore. However, the measurability condition in 4 guarantees that $\theta, \theta' \in P^{-1}(p)$ then $Q(\theta) = Q(\theta')$ so $Q(P^{-1}(p))$ is a singleton for all $p \in P(\Theta)$. The construction of M for prices outside of $P(\Theta)$ is the same as in Theorem 2. □

A.6 Proof of Proposition 3

Proof. Conditions *iii* and *iv* are necessary, as discussed in the paragraph preceding Proposition 3. To show necessity of *i* and *ii*, restrict attention to a one-dimensional strictly ordered chain in Θ (e.g. the diagonal). For the restriction of Q to this chain, necessity of *i* and continuity for interior states then follow from the same arguments as in the uni-dimensional case. Under *iii* and *iv*, this implies that *i* holds; if there is a non-monotonicity on some chain then there will be a non-monotonicity on every chain. Similarly Q must be continuous on the interior. □

B Omitted proofs from Section 4

B.1 Proof of Proposition 4

Proof. First, suppose (Q, P) are equilibrium outcomes given M. We want to show that $P(\theta) = P^*(Q(\theta), L_\theta, \theta)$ for almost all θ. Fix a state θ, and let $\mathcal{I}_P(\theta)$ be the public information set to which θ belongs. If $\mathcal{I}_P(\theta) \subseteq L_\theta$ then we are done, so suppose $\mathcal{I}_P(\theta) \setminus L_\theta$ is non-empty. Let $x_i^*(s) = x_i(P(\theta)|Q(\theta), s, \mathcal{I}_P(\theta))$. Under Increasing Differences and Belief Monotonicity, $x_i^*(s)$ is weakly increasing in s. Suppose $x_i^*(s)$ is strictly increasing in s. Then (using the so called "continuum law of large numbers" convention) Belief Monotonicity implies that for any $\theta' \in \mathcal{I}_P(\theta) \setminus L_\theta$, we have $X(P(\theta)|Q(\theta), \mathcal{I}_P(\theta), \theta') > (<) X(P(\theta)|Q(\theta), \mathcal{I}_P(\theta), \theta)$ if $\theta' > (<)\theta$. In either case, $P^*(Q(\theta), \mathcal{I}_P(\theta), \theta') \neq P^*(Q(\theta), \mathcal{I}_P(\theta), \theta)$. But contradicts the assumption that $\theta' \in \mathcal{I}_P(\theta)$. Thus it must be that $s \mapsto x_i^*(s)$ is constant. We now show that his implies the result.

Assume $s \mapsto x_i^*(s)$ is constant, and let $x^* = x_i^*(s)$. Suppose there exists a measurable set $A \subset \mathcal{I}_P(\theta)$, and x' such that

i. $V_i(Q(\theta), \theta', x', P(\theta)) > V_i(Q(\theta), \theta', x^*, P(\theta))$ for all $\theta' \in A$.

ii. $\mu(A|\mathcal{I}_P(\theta)) > 0$.

Then Increasing Differences and Belief Monotonicity imply that $s \mapsto x_i^*(s)$ is not constant, which violates our previous conclusion. Therefore, it must be that no such A, x' exist. If no A, x' satisfy condition i. (the condition $V_i(Q(\theta), \theta', x', P(\theta)) > V_i(Q(\theta), \theta', x^*, P(\theta))$ for all $\theta' \in A$) then there exists x^* such that $V_i(Q(\theta), \theta', x', P(\theta)) \leq V_i(Q(\theta), \theta', x^*, P(\theta))$ for all x and $\theta \in \mathcal{I}_P$. Then it must be that for all $\theta' \in \mathcal{I}_P(\theta)$, we have $P(\theta') = P^*(Q(\theta), L_\theta, \theta) = P^*(Q(\theta'), L_{\theta'}, \theta')$ as desired. The only other possibility is that any A, x' that satisfy condition i., do not satisfy condition ii., so $\mu(A|\mathcal{I}_P(\theta)) = 0$. Let $\{(A_n, x_n')\}_{n\geq 0}$ be the set of all such pairs. These can be divided into two groups: $x_n' > x^*$ and $x_n' < x^*$. Assume that all are of the $x_n' > x^*$ group (a symmetric argument applies to the $x_n' < x^*$ group). Notice that there must exist $\theta^* \in \mathcal{I}_P(\theta)$ such that $V_i(Q(\theta), \theta^*, x^*, P(\theta)) > V_i(Q(\theta), \theta^*, x', P(\theta))$ for all $x' > x^*$ (otherwise x^* could not be optimal under any signal). Moreover, for any n, we have $(\cup_{\theta' \in A_n}(U_{\theta'} \cap \mathcal{I}_P(\theta)), x_n') \in \{(A_n, x_n')\}$ by Increasing Differences, so without loss of generality, assume that $A_n = \cup_{\theta' \in A_n}(U_{\theta'} \cap \mathcal{I}_P(\theta))$ for all n, and assume $A_n \subset A_{n'}$ for $n' > n$. Then we can define a decreasing countable sequence θ_n such that $U_{\theta_n} \subseteq A_n \subseteq U_{\theta_{n+1}}$ for all n and $\cup_n A_n \subseteq \cup_n U_{\theta_n}$. Since $U_{\theta_n} \subseteq A_n = \cup_{\theta' \in A_n}(U_{\theta'} \cap \mathcal{I}_P(\theta))$, Increasing Differences implies that x', U_{θ_n} satisfy condition i., so $\mu(U_{\theta_n}|\mathcal{I}_P(\theta)) = 0$. Then countable additivity of μ implies $\mu(\cup_n U_{\theta_n}|\mathcal{I}_P(\theta)) = 0$, so $\mu(\cup_n A_n|\mathcal{I}_P(\theta)) = 0$.

But then $V_i(Q(\theta), \theta', x', P(\theta)) \leq V_i(Q(\theta), \theta', x^*, P(\theta))$ for all x and all but a conditionally-zero-measure subset of $\mathcal{I}_P(\theta)$. Thus for all $\theta' \in \mathcal{I}_P(\theta) \setminus A$, we have $P(\theta') = P^*(Q(\theta), L_\theta, \theta) = P^*(Q(\theta'), L_{\theta'}, \theta')$ as desired. Thus far, we have reasoned for a fixed public information set $\mathcal{I}_P(\theta)$. However since for any information set. However since $P(\theta) = P^*(Q(\theta), L_\theta, \theta)$ can fail for at most a conditionally-zero-measure subset of any information set, the set of all such θ has zero measure in Θ.

For the converse direction, we want to show that if $R(a,\theta) = P^*(a, L_\theta, \theta)$ and Q, P, M satisfy commitment and market clearing then (Q, P) are equilibrium outcomes given M. Thus we need to check that $X(P(\theta)|M(P(\theta)), \mathcal{I}_P(\theta), \theta) = 0$ for all θ. Fix a public information set $\mathcal{I}_P(\theta)$. The first part of the above proof for the other direction continues to hold: it must be that $s \mapsto x_i^*(s) \equiv x_i(P(\theta)|Q(\theta), s, \mathcal{I}_P(\theta))$ is constant, otherwise P could not be constant on $\mathcal{I}_P(\theta)$. But then second part of the above proof tells us that for all but a conditionally-zero-measure subset of $\mathcal{I}_P(\theta)$, we have $V_i(Q(\theta), \theta', x^*, P(\theta)) > V_i(Q(\theta), \theta', x', P(\theta))$ for all $x' \neq x^*$. Let θ'' be a state such that this inequality holds. Then $x_i(P(\theta)|Q(\theta), s, \mathcal{I}_P(\theta)) = x_i(P(\theta)|Q(\theta), s, L_{\theta''})$ for all s, so markets clear in state θ if and only if $P = P^*(a, L_\theta, \theta)$. □

B.1.1 Proof of Lemma 2

Proof. First note that $s_i \mapsto x_i(p|a, \mathcal{I}, s_i)$ is Lipschitz continuous since Ω is bounded and $s_i = \omega + \varepsilon_i$ for a normally distributed ε_i. Increasing ω by δ has the same effect on aggregate demand as increasing s_i by δ for all i. Then $\omega \mapsto X(p|a, \mathcal{I}, \omega)$ is Lipschitz continuous since σ_i and τ_i are bounded in the population. □

B.1.2 Proof of Lemma 3

Proof. First, note if (ω'', z'') and (ω', z') are elements of $\mathcal{I}$, with $\omega'' > \omega'$ then it must be that $z'' > z'$. This follows from the fact that aggregate demand is strictly increasing in ω and strictly decreasing in p.

The function $\omega \mapsto X(p|a, \mathcal{I}, \omega)$ is Lipschitz continuous by Lemma 2. So for any $\kappa > 0$ there exists $\delta > 0$ such that for any $(\omega'', z''), (\omega', z') \in \mathcal{I}$, we have $|\omega'' - \omega'| < \delta$ implies $|z'' - z'| < \kappa$. In other words, there is uniform bound on the "slope" of $\mathcal{I}$ in $\Omega \times \mathcal{Z}$ space. Since the prior distribution on $\Omega \times \mathcal{Z}$ is absolutely continuous, this implies the desired result. □

B.1.3 Proof of Lemma 4

Proof. Define the random variable $\tilde{V}^a := \pi(a,\theta) = \beta_0^a + \beta_1^a \theta$. Then define $\tilde{S}_i^a := \beta_1^a s_i + \beta_0^a = \tilde{V}^a + \beta_1^a \varepsilon_i$. Thus conditional on knowing the principal's action, investor i's observation of s_i is equivalent to observing a signal $\tilde{S}_i^a$ which is equal to the true dividend $\tilde{V}^a$ plus normal random noise, where the variance of the noise term depends on a; it is given by $\sigma_{ai}^2 = (\beta_1^a)^2 \sigma_i^2$. The results then follows from the proof of Proposition 2.2 in Breon-Drish (2015) (Online Appendix). The proposition in Breon-Drish (2015) pertains to the information sets revealed by equilibrium price functions which are continuous and satisfy a differentiability assumption. However for the relevant direction of the proof, these conditions are only needed to guarantee that the distribution of $\tilde{V}^a$ conditional on $\mathcal{I}$ has a density, which is implied here by Lemma 3. □

B.1.4 Proof of Proposition 5

Proof. Given Lemma 4, we just need to identify what the coefficients on the linear statistic are.

Fix M, and let $L_M : \Omega \times \mathcal{Z} \times \mathcal{A} \to \mathbb{R}$ be the equilibrium statistic in a generalized linear equilibrium in which the price reveals exactly a hyperplane. Define the random variable $\tilde{V}^a := \pi(a,\omega) = \beta_0^a + \beta_1^a\omega$. Then define $\tilde{S}_i^a := \beta_1^a s_i + \beta_0^a = \tilde{V}^a + \beta_1^a\varepsilon_i$. Thus conditional on knowing the principal's action, investor i's observation of s_i is equivalent to observing a signal $\tilde{S}_i^a$ which is equal to the true dividend $\tilde{V}^a$ plus normal random noise, where the variance of the noise term depends on a; it is given by $\sigma_{ai}^2 = (\beta_1^a)^2\sigma_i^2$. Let $\tilde{L}^a$ be the random variable $L_M(\omega, z, a)$.

We first fix the principal's action at a, and generalize Breon-Drish (2015) Proposition 2.1 to allow for supply shocks with a truncated normal distribution. We will therefore suppress dependence of $\tilde{S}_i^a, \tilde{V}^a, \tilde{L}^a$ on the action a for the time being. Abusing notation, write the statistic L in terns of v, rather than ω; that is, $L(v, z|a) = \alpha v - z$, suppressing the dependence on M.[39] For fixed a, the truncation is the only difference between the current setting and that of Breon-Drish (2015) Proposition 2.1. By the same steps as the proof for Proposition 2.1 in Breon-Drish (2015) Online Appendix, we can show that the conditional distribution of $\tilde{V}^a$ conditional on $\tilde{S}_i^a = s_i$ and $\tilde{L}^a = \ell$ is given by

$$dF_{\tilde{V}|\tilde{S},\tilde{L}}(v|s_i,\ell) = \frac{\mathbb{1}[\ell - \alpha v \in (-b,b)] \exp\left\{\left(\frac{1}{\sigma_{ai}^2}s_i + \frac{\alpha}{\sigma_Z^2}\ell\right)v - \frac{1}{2}\left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2}\right)v^2\right\} dF_{\tilde{V}}(v)}{\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp\left\{\left(\frac{1}{\sigma_{ai}^2}s_i + \frac{\alpha}{\sigma_Z^2}\ell\right)x - \frac{1}{2}\left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2}\right)x^2\right\} dF_{\tilde{V}}(x)}, \quad (4)$$

where $\mathbb{1}[\cdot]$ is the indicator function. This is not in the *exponential family* of distributions, as defined in Breon-Drish (2015) Assumption 10. Nonetheless, it will have similar properties. We can write the conditional distribution in (4) as

$$\mathbb{1}[\ell - \alpha v \in (-b,b)] \exp\left\{\hat{L}(s_i,\ell)v - g\left(\hat{L}(s_i,\ell);\alpha,\ell\right)\right\} dH(v;\alpha),$$

where

$$\begin{aligned}
\hat{L}(s,\ell) &= \left(\frac{1}{\sigma_{ai}^2}s_i + \frac{\alpha}{\sigma_Z^2}\ell\right) \\
g_i(\hat{L};\alpha,\ell) &= \log\left(\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp\left\{\left(\frac{1}{\sigma_{ai}^2}s_i + \frac{\alpha}{\sigma_Z^2}\ell\right)x - \frac{1}{2}\left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2}\right)x^2\right\} dF_{\tilde{V}}(x)\right) \\
dH_i(v;\alpha) &= \exp\left\{-\frac{1}{2}\left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2}\right)v^2\right\} dF_{\tilde{V}}(v)
\end{aligned}$$

[39] This abuse of notation is done to match the notation of Breon-Drish (2015). Note that in that paper "a" is used in place of α to denote the slope of the equilibrium statistic. The reader examining Breon-Drish (2015) should not confuse this with the notation for the principal action used in the current paper.

This has the following important implication (essentially the same as Lemma A6 in Breon-Drish (2015)). Since the conditional distribution must integrate to 1, i.e.

$$\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp\left\{\hat{L}(s_i,\ell)v - g\left(\hat{L}(s_i,\ell);\alpha,\ell\right)\right\} dH(v;\alpha) = 1$$

we have that

$$\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp\left\{\hat{L}(s_i,\ell)v\right\} dH(v;\alpha) = \exp\left\{g\left(\hat{L}(s_i,\ell);\alpha,\ell\right)\right\}.$$

As a result, for any $t \in \mathbb{R}$ we have

$$\mathbb{E}\left[\exp\{t\tilde{V}\}|s,\ell\right] = \exp\left\{g\left(t+\hat{L}(s_i,\ell);\alpha,\ell\right) - g\left(\hat{L}(s_i,\ell);\alpha,\ell\right)\right\}.$$

The remainder of the proof for the fixed-action case proceeds as in Breon-Drish (2015) Proposition 2.1. In particular, this shows that in any generalized linear equilibrium with fixed action a,

$$\alpha = \int_i \frac{\tau_i}{\sigma_{ai}^2} di.$$

Since $v = \beta_0^a + \beta_1^a \omega$ and $\sigma_{ai}^2 = (\beta_1^a)^2 \sigma_i^2$ we have

$$L^*(\omega, z|a) = \beta_0^a \int_i \frac{\tau_i}{\sigma_{ai}^2} di + \left(\frac{1}{\beta_1^a}\int_i \frac{\tau_i}{\sigma_{ai}^2} di\right) \cdot \omega - z$$

Since information revelation is characterized by the level sets of L^*, we can ignore the first term.

We now show that the result holds under feedback as well. Given M, the investor knows which action the principal will take conditional on the price. In a generalized linear equilibrium, the investor's demand is therefore determined by maximizing utility given that the price is p, the action is $M(p)$, the observed signal is $\tilde{S}_i^a$, and the extended state is in $\{(\omega, z) : L_M(\omega, a|a) = \ell\}$ for the value of ℓ corresponding to price level p. The remaining question is which $L_M(\cdot|a)$ could constitute equilibrium statistics given action a and decision rule M. The first part of the proof shows that if the principal's action is fixed at a then there is a unique equilibrium statistic $L^*(\omega, z|a)$. Since all investors know the principal's action once they observe the price, this L^* must be the equilibrium statistic, regardless of M. □

B.1.5 Proof of Proposition 6

Proof. Proposition 5 says that the equilibrium price must reveal at least the linear statistic. We want to show that the price can reveal no more than this. For $p \in \tilde{P}(\Omega, \mathcal{Z})$ let $l^*(p)$ be the linear statistic revealed by p. Suppose that $\mathcal{I}(p) := \{(\omega, z) : \tilde{P}(\omega, z) = p\} \neq l^*(p)$, so that the price reveals more than the linear statistic. We show that in this case there will be multiplicity. This follows from the fact that the set of states $\{(\omega, z) : X(p|M(p), \mathcal{I}(p), \omega) = z\}$ is the entire linear statistic $l^*(p)$. This follows from the proof of Lemma 3 and Proposition 2.2 in Breon-Drish (2015) (Online appendix), which shows that individual demands will be linear in signals for any price. □

C Proofs for Section 5

C.1 Proofs for Section 5.1

Lemma 14. *Given a function $F : X \times [0, 1] \to X$ on a compact subset X of a Euclidean space, define the function*

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

Assume $t \mapsto F(x, t)$ is continuous. If $G(t)$ is single valued and $x \mapsto F(x, t)$ is continuous on an open neighborhood of $G(t)$ 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 14, we need to show that there exists an open neighborhood U of t such

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

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 $F : X \times \Theta \times (0, 1) \to X$ on a compact subset X of a Euclidean space, define the function*

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

Let S be any compact subset of Θ such that $G(t, \theta)$ is single valued for all $\theta \in S$. Then $t \rightrightarrows G(t, \theta)$ is upper and lower hemicontinuous at t, uniformly over S.

Proof. Since $G(t,\theta)$ is single valued 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 $\delta - \rho$.[40] By Lemma 14, 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 14 that for each $\theta \in S$ there exists $\varepsilon_\theta, \kappa_\theta > 0$ such that

$$|t' - t| < \varepsilon_\theta \Longrightarrow |F(x,\theta,t') - x| > \kappa_\theta \ \ \forall\, x \in X \setminus V(\theta). \tag{5}$$

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 (5).

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 (5) we know that $z(\theta,\varepsilon_\theta) > \kappa_\theta$ for all $\theta \in S$. Then for

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

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.

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. $\square$

C.1.1 Proof of Theorem 3

Proof. There are two cases to consider: $P(\underline{\theta}) \leq P(\bar{\theta})$ or $P(\underline{\theta}) > P(\bar{\theta})$.

If $P(\underline{\theta}) \leq P(\bar{\theta})$ P is weakly increasing by Theorem 1. Then as noted in Section 3, (Q, P) can be implemented by a decision rule M that is continuous. 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. Moreover, $G(t, \theta) = \tilde{Q}(\theta, t)$ will be single valued on all but a zero-measure set of states when M is weakly robust to multiplicity, and single valued everywhere when M is robust to multiplicity. Therefore for any $\varepsilon > 0$ we can find a compact set S such that $G(t, \theta)$ is single valued for all $\theta \in S$. When M is robust to multiplicity let $S = \Theta$. Then Proposition 10 applies, which gives the result.

If $P(\underline{\theta}) > P(\bar{\theta})$ then Theorem 1 implies that P is weakly decreasing. As shown in the proof of Theorem 1, there exists a closed set $C \supset [P(\bar{\theta}), P(\underline{\theta})]$ such that M is continuous on C, but may have discontinuities outside of C. We are free to define M outside of C, so long as there are no $p \notin C$ such that $R(M(p), \theta) = p$. Let $M(p) = Q(\bar{\theta})$ if $p \notin C$ and $p > P(\underline{\theta})$, and let $M(p) = Q(\underline{\theta})$ if $p \notin C$ and $p < P(\bar{\theta})$. Since $P(\underline{\theta}) > P(\bar{\theta})$ by assumption, and $\theta \mapsto R(a, \theta)$ is weakly increasing for all a, there exists $\varepsilon > 0$ such that *(i)* $p - R(M(p), \theta) > \varepsilon$ for all θ and all $p \notin C$, $p > P(\underline{\theta})$, and *(ii)* $R(M(p), \theta) - p < \varepsilon$ for all θ and all $p \notin C$, $p < P(\bar{\theta})$. Therefore, conditions *(i)* and *ii* will continue to hold for some $\varepsilon' > 0$ and any R' that is sufficiently close to R in the sup-norm. This implies that it is sufficient to establish upper and lower hemicontinuity of $R \rightrightarrows \tilde{Q}_R$ for the restriction of M to C. Since M is continuous on C the argument applied to above the $P(\underline{\theta}) \leq P(\bar{\theta})$ P case holds here as well. $\square$

C.1.2 Proof of Lemma 6

Proof. Suppose M is discontinuous at p', and let $\theta' \in \theta_M(p'|R)$. First, suppose that $p \mapsto R(M(p), \theta')$ is continuous at p'. Since M is discontinuous, there exists an open neighborhood U or $M(p')$ such that for any $\varepsilon > 0$ there exists $p'' \in N_\varepsilon(p')$ with $M(p) \notin U$. Since $p \mapsto R(M(p), \theta')$ is continuous at p', for any $\delta > 0$ we can choose ε small to guarantee $|R(M(p''), \theta') - R(M(p'), \theta')| < \delta$. But then let $\hat{R}$ be a continuous function in a δ-neighborhood of R such that $\hat{R}(M(p''), \theta') = p'$, so $M(p'') \in \tilde{Q}_{\hat{R}}(\theta'|M)$. Therefore we cannot have upper hemicontinuity of $R \mapsto \tilde{Q}_R(\theta'|M)$ at R.

Now, suppose $p \mapsto R(M(p), \theta')$ is discontinuous at p'. Assume M is left-continuous at p' (symmetric argument for right-continuous, and similar for removable discontinuity). Then there exists $\varepsilon > 0$ such that either $R(M(p), \theta') < p$ for all $p \in [p' - \varepsilon, p')$ or $R(M(p), \theta') > p$ for all $p \in [p' - \varepsilon, p')$. Assume without loss of generality that the former holds. Then let $\hat{R}$ be a continuous function such that $\hat{R}(M(p), \theta') > R(M(p), \theta')$ for all $p \in [p' - \varepsilon, p')$. For $\hat{R}$ close to R there will be a neighborhood U or p' such that $\hat{R}(M(p), \theta') \neq p$ for all $p \in U$. This is because M is discontinuous at p'. Then $R \mapsto \tilde{Q}_R(\theta'|M)$ cannot be lower hemicontinuous at R. $\square$

C.2 Section 5.2

C.2.1 Proof of Proposition 7

Proof. Claim 0. For any $\theta' \in (\underline{\theta}, \bar{\theta})$ and p' such that $\theta' \in \theta_M(p')$, there exist p'' such that $\theta_M(p'') \cap \{\underline{\theta}, \bar{\theta}\} \neq \varnothing$, $\theta_M(p) \neq \varnothing$ for all $p \in (\min\{p', p''\}, \max\{p', p''\})$ and M is continuous on $(\min\{p', p''\}, \max\{p', p''\})$ (when this interval is non-empty).

Let $\theta' \in (\underline{\theta}, \bar{\theta})$ be arbitrary, and let p' be such that $\theta' \in \theta_M(p')$. If $\{p \leq p' : \theta_M(p) = \varnothing\}$ is empty then $p'' = \arg\min_{a \in \mathcal{A}} R(a, \underline{\theta})$ satisfies the conditions of the claim. Similarly, if $\{p \geq p' : \theta_M(p) = \varnothing\}$ is empty then $p'' = \arg\max_{a \in \mathcal{A}} R(a, \bar{\theta})$ satisfies the conditions of the claim. Assume that $\{p \leq p' : \theta_M(p) = \varnothing\} \neq \varnothing$ and $\{p \geq p' : \theta_M(p) = \varnothing\} \neq \varnothing$. Let $\underline{p} = \sup\{p \leq p' : \theta_M(p) = \varnothing\}$ and $\bar{p} = \inf\{p \geq p' : \theta_M(p) = \varnothing\}$. Since $M \in \mathcal{M}$, we have $\underline{p} < p' < \bar{p}$. Since M must be continuous on $(\underline{p}, \bar{p})$, we have $\theta_M(\underline{p}) \cap \{\underline{\theta}, \bar{\theta}\} \neq \varnothing$ and $\theta_M(\bar{p}) \cap \{\underline{\theta}, \bar{\theta}\} \neq \varnothing$. This proves Claim 0.

Claim 1. Let $\theta' \in (\underline{\theta}, \bar{\theta})$ and p' be such that $\theta' \in \theta_M(p')$. Let p'' be such that $\theta_M(p) \neq \varnothing$ for all $p \in (\min\{p', p''\}, \max\{p', p''\})$ and M is continuous on $(\min\{p', p''\}, \max\{p', p''\})$ (when this interval is non-empty). Then if $\underline{\theta} \in \theta_M(p'')$ and $p'' \leq p'$ ($p'' \geq p'$) there exists an equilibrium with a price function that is increasing (decreasing) on $[\underline{\theta}, \theta']$. Similarly, if $\bar{\theta} \in \theta_M(p'')$ and $p'' \geq p'$ ($p'' \leq p'$) there exists an equilibrium with a price function that is increasing (decreasing) on $[\theta', \bar{\theta}]$.

We will show the claim for $\bar{\theta} \in \theta_M(p'')$ and $p'' \geq p'$; all others cases are symmetric. For any θ, the set $\theta_M^{-1}(\theta)$ is compact: if $R(M(p), \theta) \neq p$ then this holds for all $\tilde{p}$ in a neighborhood p, since $M \in \mathcal{M}$ is continuous around equilibrium prices. If $p' = p''$ then we are done: convexity of $\theta_M(p)$ (Lemma 9) implies that there is a constant, and thus monotone, equilibrium price function on $[\theta', \bar{\theta}]$. Assume instead that $p'' > p'$. If there exists $\theta^* \in (\theta', \bar{\theta})$ such that $p^* > p''$ for any $p^* \in \theta_M^{-1}(\theta'')$ then there exists $\tilde{\theta} \in (\theta', \bar{\theta})$ such that $p'' \in \theta_M^{-1}(\tilde{\theta})$, by continuity of M on (p', p'') and Lemma 11. Then convexity of $\theta_M(p'')$ implies that we can construct a flat price function above $\tilde{\theta}$. Therefore assume no such θ^* exists. By a symmetric argument, we can assume that $\theta_M^{-1}(p) \cap [p', p''] \neq \varnothing$ for all $\theta \in [\theta', \bar{\theta}]$.

We want to construct an increasing equilibrium price function on $[\theta', \bar{\theta}]$. Consider an arbitrary price function $\tilde{P}$ such that $\tilde{P}(\theta) \in \theta_M^{-1}(\theta) \cap [p', p'']$ for all $\theta \in [\theta', \bar{\theta}]$, $\tilde{P}(\underline{\theta}) = p'$, and $\tilde{P}(\bar{\theta}) = p''$. We will show that any violations of monotonicity can be ironed without leading to further violations.

Claim 1.2. Suppose $\tilde{P}(\theta_2) < \tilde{P}(\theta_1) < \tilde{P}(\theta_3)$ for $\bar{\theta} > \theta_3 > \theta_2 > \theta_1 >$. Then there exists $p \in \theta_M^{-1}(\theta_2) \cap [\tilde{P}(\theta_1), \tilde{P}(\theta_3)]$.

Claim 1.2 follows immediately from Lemma 11. This in turn shows that Claim 1 holds for $\bar{\theta} \in \theta_M(p'')$ and $p'' \geq p'$, which is what we wished to show.

Claim 0 and Claim 1 together imply the existence of a monotone price function. If R is strictly increasing in θ then measurability of the action with respect to the price implies that P must be strictly monotone. □

C.2.2 Proof of Lemma 1

Proof. Let $\{\theta_n\}$ be an increasing sequence converging to θ'. Suppose P is increasing (the argument is symmetric if P is decreasing). Then $\{P(\theta_n)\}$ is an increasing and bounded sequence, and so converges. Denote this limit by $\bar{p}$. Since M is essentially continuous, it is continuous in a neighborhood of $\bar{p}$. Hence $\lim_{n\to\infty} Q(\theta_n) = \lim_{n\to\infty} M(P(\theta_n)) = M(\bar{p})$. □

C.2.3 Proof of Proposition 2

Proof. Given Theorem 2, we need only show that Q can have a discontinuity at θ^* iff P has a bridgeable discontinuity at θ^*. Clearly Q can be discontinuous at θ^* iff P is continuous at θ^* (otherwise M would need to be discontinuous at $P(\theta^*)$). As shown in the proof of Theorem 1, P can be discontinuous at θ^* only if $\theta_M(p) = \theta^*$ on $(\min\{\lim_{\theta\nearrow\theta^*} P(\theta), \lim_{\theta\searrow\theta^*} P(\theta)\}, \max\{\lim_{\theta\nearrow\theta^*} P(\theta), \lim_{\theta\searrow\theta^*} P(\theta)\})$. This is possible iff there exists γ satisfying the definition of bridgeability (in which case we take $M = \gamma$ on this interval). □

C.2.4 Proof of Theorem 4

Proof. Let (Q, P) be an equilibrium induced by M, such that P is strictly monotone, which exists by Proposition 7. Since $M \in \mathcal{M}$ induces (Q, P), P can have no degenerate discontinuities. Let $\hat{M} = M$ on $P(\Theta)$ and $\mathcal{P} \setminus [\inf P(\Theta), \sup P(\Theta)]$. We show how to define $\hat{M}$ for the remaining prices such that it is essentially continuous and weakly uniquely implements (Q, P).

Suppose P has a non-degenerate discontinuity at θ^*, and let $\underline{p} = \lim_{\theta\nearrow\theta^*} P(\theta)$ and $\bar{p} = \lim_{\theta\searrow\theta'} P(\theta)$. If the discontinuity at θ^* is bridgeable then we can define $\hat{M}$ on $[\min\{\underline{p}, \bar{p}\}, \max\{\underline{p}, \bar{p}\}]$ such that *(i)* $\hat{M}(\underline{p}) = \lim_{\theta\nearrow\theta^*} Q(\theta)$, *(ii)* $\hat{M}(\bar{p}) = \lim_{\theta\searrow\theta^*} Q(\theta)$, and *(iii)* $p = R(\hat{M}(p), \theta^*)$ for all $p \in [\min\{\underline{p}, \bar{p}\}, \max\{\underline{p}, \bar{p}\}]$. Since the environment is fully bridgeable, this can be done for all discontinuities. Thus $\hat{M}$ so defined is continuous on $[\inf P(\Theta), \sup P(\Theta)]$ and coincides with M on $\hat{M} = M$ on $P(\Theta)$ and $\mathcal{P} \setminus [\inf P(\Theta), \sup P(\Theta)]$. Since M was essentially continuous, so is $\hat{M}$. Moreover, there are multiple market-clearing prices only in states at which P had a discontinuity. Since P is monotone, this set has measure zero. □

C.3 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.[41] 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 11. *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 $t : [0,1] \to [\bar{\lambda}, 1] \times [0,1]$ 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.

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 11 makes it easy to identify when a discontinuity will be bridgeable. For example,

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

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 11 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 11 (when $\pi(z', \theta) \neq \pi(z'', \theta)$ for all $z' \neq z''$ it is in fact a lattice).

Corollary 3. *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.*

D Proofs for Section 6

D.1 Proof of Lemma 7

Proof. An equilibrium exists for any increasing M by Tarski's fixed point theorem. That the price function will be increasing follows from the fact that $a \mapsto R(a, \theta)$ is decreasing and $\theta \mapsto R(a, \theta)$ is increasing. If P is increasing and M is increasing, there will be no equilibrium involving prices above $P(\bar{\theta})$ or below $P(\underline{\theta})$.

We show that M can have no discontinuities on $[P(\underline{\theta}), P(\bar{\theta})]$, which implies that P is continuous. Suppose, towards a contradiction that there is a non-empty set D of discontinuities in this region, and let $p' = \inf D$. By definition of $\mathcal{M}$, $p' \in (P(\underline{\theta}), P(\bar{\theta}))$. Let $a' = \lim_{p \nearrow p'} M(p)$. For any $p \in (P(\underline{\theta}, p')$ and any $a \in (M(P(\underline{\theta}), a')$ there exists $\theta \in (\underline{\theta}, \bar{\theta})$ such that $R(a, \theta) = p$. This follows from the fact that $a \mapsto R(a, \theta)$ is decreasing. Then for any $p \in (P(\underline{\theta}), p')$ there exists θ such that $R(M(p), \theta) = p$, since M is increasing and continuous on $(P(\underline{\theta}, p')$. This contradicts the definition of p'. $\square$

D.2 Proof of Lemma 8

Proof. Condition i is immediate. For ii, first note that for $p \in (R(\underline{a}, \underline{\theta}), R(\underline{a}, \bar{\theta}))$ it must be that $M(p) > \underline{a}$; if not then $R(M(p), \theta) = p$ for some $\theta \in (\underline{\theta}, \bar{\theta})$. Suppose there is no discontinuity on $(P(\underline{\theta}), R(\underline{a}, \underline{\theta})]$. Then M must be decreasing over this domain to prevent multiplicity, and $\lim_{p \searrow R(\underline{a}, \underline{\theta}} M(p) = \underline{a}$. But for $p \in (R(\underline{a}, \underline{\theta}), R(\underline{a}, \bar{\theta}))$ it must be that $M(p) > \underline{a}$, so there must be a discontinuity. A symmetric argument applies to $(R(\bar{a}, \bar{\theta}), P(\bar{a})]$

Conditions *iii* and *iv* follow from a similar argument. Define $\bar{p}$ by $\bar{p} = \sup\{p : M \text{ is decreasing on } (P(\bar{\theta}), \bar{p})\}$. The argument above implies that $\bar{p} \leq R(\underline{a}, \underline{\theta})$. This implies *iii*. A symmetric argument implies *iv*. $\square$

E No Commitment

We have assumed throughout that the principal is able to commit to a decision rule. In this section we briefly analyze the situation in which the principal cannot commit.

We assume that all market participants understand the principal's preferences, and can thus predict what the principal will do as a function of the principal's information set. In any REE, the price function will reveal some information to the principal, as a function of which the principal will take their preferred action. Thus any equilibrium prince function P will induce a map $m(\cdot; P) : \mathcal{P} \to \mathcal{A}$, where $m(p; P)$ is the principal's optimal action given the information revealed by $P(\theta) = p$ (or some mixture over these in the case of indifference). A rational expectations equilibrium without commitment consists of a price function P and decision rule m such that

i. $P(\theta) = R(m(P(\theta)), \theta)$ for all θ. *(rational expectations)*

ii. For all p, $m(p)$ is an optimal action for the principal conditional on $\{\theta : P(\theta) = p\}$. *(principal optimality)*

The principal optimality condition replaces the commitment condition in the definition of REE used under commitment.

Let $Q^* : \Theta \to \mathcal{A}$ be the principal's first-best action function. That is, Q^* specifies the principal's optimal action in each state. Assume for simplicity that $\theta \mapsto R(a, \theta)$ is strictly increasing for all a. Then given any m, there is at most a single state θ such that $p = R(m(p), \theta)$. Thus any REE price function must be fully revealing. Given this observation, we have the following equivalent definition of a REE without commitment

Lemma 15. *Assume $\theta \mapsto R(a, \theta)$ is strictly increasing for all a. Then (P, m) constitute an REE without commitment if and only if (P, Q^*) are implementable under commitment (as defined in Section 2.1).*

This observation has the following immediate corollary.

Corollary 4. *If Q^* is not implementable then there does not exist an REE without commitment.*

Moreover, we can use the characterization results under commitment to understand equilibrium behavior without commitment. For example, Theorem 1 has the following implication.

Corollary 5. *If $\theta \mapsto R(Q^*(\theta), \theta)$ is non-monotone then either there will be discontinuities at some equilibrium prices or there will be multiple equilibria.*

In other words, the equilibrium will either be vulnerable to manipulation (and not be robust to structural uncertainty), or it will suffer from non-fundamental volatility. Note that there can only be multiple equilibria if there are states for which the principal has multiple optimal actions.

References

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

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.

R. Boleslavsky, D. L. Kelly, and C. R. Taylor. Selloffs, bailouts, and feedback: Can asset markets inform policy? *Journal of Economic Theory*, 169:294–343, 2017.

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.

B. Breon-Drish. On existence and uniqueness of equilibrium in a class of noisy rational expectations models. *The Review of Economic Studies*, 82(3):868–921, 2015.

Y.-K. Che, C. Choe, and K. Rhee. Bailout stigma. *Available at SSRN 3208238*, 2018.

C. Cohen, S. A. Abbas, M. Anthony, T. Best, P. Breuer, H. Miao, A. Myrvoda, E. Togo, et al. The role of state-contingent debt instruments in sovereign debt restructurings. Technical report, International Monetary Fund, 2020.

T. Cordella, L. A. Ricci, and M. Ruiz-Arranz. Debt overhang or debt irrelevance? *IMF Staff Papers*, 57(1):1–24, 2010.

J. Crowley. *Central and commercial bank balance sheet risk before, during, and after the global financial crisis.* International Monetary Fund, 2015.

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.

C. Flachsland, M. Pahle, D. Burtraw, O. Edenhofer, M. Elkerbout, C. Fischer, O. Tietjen, and L. Zetterberg. How to avoid history repeating itself: the case for an eu emissions trading system (eu ets) price floor revisited. *Climate Policy*, 20(1):133–142, 2020.

P. Glasserman and B. Nouri. Market-triggered changes in capital structure: Equilibrium price dynamics. *Econometrica*, 84(6):2113–2153, 2016.

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

A. Greenspan. Harnessing market discipline: the region (september). *Federal Reserve Bank of Minneapolis, MN*, pages 10–11, 2001.

B. Grochulski and R. Wong. Contingent debt and performance pricing in an optimal capital structure model with financial distress and reorganization. *FRB Richmond Working Paper*, 2018.

S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408, 1980.

E. Hauk, A. Lanteri, and A. Marcet. Optimal policy with general signal extraction. *Economic Research Initiatives at Duke (ERID) Working Paper*, 2020.

F. A. Hayek. The use of knowledge in society. *The American economic review*, 35(4):519–530, 1945.

M. Hellwig. On the aggregation of information in competitive markets. *Journal of Economic Theory*, 22(3):477–498, 1980.

R. T. Jensen and N. H. Miller. Giffen behavior and subsistence consumption. *American economic review*, 98(4):1553–77, 2008.

J. S. Jordan. The generic existence of rational expectations equilibrium in the higher dimensional case. *Journal of Economic Theory*, 1982.

L. Karp and C. P. Traeger. Smart cap. *CEPR Discussion Paper No. DP15941*, 2021.

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

R. E. Lucas et al. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46, 1976.

R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

E. Ozdenoren and K. Yuan. Feedback effects and asset prices. *The journal of finance*, 63(4): 1939–1975, 2008.

D. Pálvölgyi and G. Venter. Multiple equilibria in noisy rational expectations economies. *Available at SSRN 2524105*, 2015.

C. Siemroth. The informational content of prices when policy makers react to financial markets. *Journal of Economic Theory*, 179:240–274, 2019.

F. Teschner, D. Rothschild, and H. Gimpel. Manipulation in conditional decision markets. *Group Decision and Negotiation*, 26(5):953–971, 2017.

J. B. Warner, R. L. Watts, and K. H. Wruck. Stock prices and top management changes. *Journal of financial Economics*, 20:461–492, 1988.

M. L. Weitzman. Prices vs. quantities. *The review of economic studies*, 41(4):477–491, 1974.

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.