Market-Based Mechanisms*

Quitzé Valenzuela-Stookey[†] Francisco Poggi[‡]

September 23, 2024

Abstract

We propose a framework for studying the use of aggregate market outcomes, such as security prices, to inform decision making. A principal designs a mechanism which conditions their action on market outcomes as a proxy for an unobserved payoff-relevant state. The market anticipates the principal's actions, which creates a two-way feedback between the mechanism and the market. Additionally, market participants may attempt to manipulate the market to influence the principal's actions. In a general setting, we characterize the set of joint distributions of the market outcome, principal action, and state that can be implemented as the unique equilibrium using a mechanism that is robust to manipulation. Our characterization admits a tractable representation, and significantly simplifies the principal's design problem. We apply our results to study optimal bailout policies.

*We are grateful to Marciano Siniscalchi, Jeff Ely, Eddie Dekel, Piotr Dworczak, Alessandro Pavan, Ludvig Sinander, Nicolas Schutz, 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, 2021 MaCCI Summer Institute in Competition Policy, and University of Warwick for helpful comments.

[†]Department of Economics, UC Berkeley, Berkeley, CA 94720, U.S.A., quitze@berkeley.edu

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

The question of how a principal can elicit payoff-relevant information from agents who have a vested interest in the principal's actions has been extensively studied in the mechanism-design literature. However, standard techniques for gathering this information may be infeasible when information is widely dispersed among many, potentially anonymous, agents. In such settings, the principal can exploit a fundamental insight of information economics, going back at least to Hayek (1945), that markets aggregate dispersed information.[1] An extensive theoretical and empirical literature studies the way in which market outcomes (which for brevity we refer to here as prices) aggregate the private information of market participants regarding an underlying state of nature (e.g. Grossman and Stiglitz 1980; Roll 1984). In this paper, we study how a principal can leverage this information.

Economic agents in a number of settings use, or are encouraged to use, market prices to inform their decisions. For example, in a cap-and-trade policy for carbon emissions, the price of carbon credits reveals information about firms' costs for reducing emissions. This information may be useful to regulators when determining how many credits to issue in the future (Flachsland et al., 2020; Karp and Traeger, 2021). Similarly, central banks may base monetary policy decisions on the information about inflation expectations inferred from movements in commodity prices or changes in the spread between nominal and indexed government bonds (Bernanke and Woodford, 1997).

In practice, however, a principal engaging in market-based decision-making faces several challenges. One fundamental issue is the potential feedback loop between actions and the market. When the principal's actions affect the market, prices reflect not only the private information of participants, but also their beliefs about the principal's actions. Such feedback effects have long been recognized as a challenge to market-based policy-making, e.g. Dow and Gorton 1997; Ozdenoren and Yuan 2008; Bond et al. 2010. For instance, in the carbon market example, firms' demand for credits today depends on their beliefs about the volume of future credit issuances. Conditioning policy on the market therefore affects the relationship between the price and agents' private information, and thus changes the nature of the information

[1] Indeed, Hayak appears to have been motivated by precisely this consideration: "We cannot expect that this problem [policy-making under incomplete information] will be solved by first communicating all this knowledge to a central board which, after integrating all knowledge, issues its orders. We must solve it by some form of decentralization" (Hayek, 1945).

revealed by prices.[2] Another critical issue is manipulation: agents may seek to influence the principal's action by distorting the market. Even in fairly competitive markets, this is a concern if the principal's action choices are highly sensitive to prices Goldstein and Guembel (2008); Lee (2019).

Our focus is on how a principal can exploit the information revealed by markets in settings with feedback effects, in a way that is robust to manipulation. We adopt an implementation-theoretic approach to this question. That is, we define a class of market-based mechanisms and seek to characterize the set of joint distributions between the state of nature, prices, and the principal's action which can be induced in equilibrium using such a mechanism. We are especially interested in full implementation, i.e. which joint distributions can be induced uniquely (Maskin, 1999). Beyond the usual normative appeal of full implementation, this focus here is driven by the inherent relationship between equilibrium multiplicity and feedback effects: when actions are conditioned on market outcomes there may be multiple self-fulfilling beliefs that market participants can hold about what the principal will do.

To understand the limits of what market-based mechanisms can achieve, we develop a general framework intended to accommodate a wide range of applications. The starting point of the analysis is a market in which a security is traded by agents with private information about a state of nature, which in the baseline model lives in $\mathbb{R}$.[3] The ex-post value of the security depends on the state, as well as an action chosen by the principal. The principal commits up front to a mechanism, or *policy*, which maps the security price to a (possibly mixed) action. Following the literature on markets with feedback effects, we assume that the market is competitive, i.e. agents are price takers, and study rational expectations equilibria in which agents understand the relationship between the price and the principal's action, as specified by the policy.

An important point of departure from the existing literature is that despite modeling the market as competitive, we nonetheless maintain a concern for market manipulation. We do this by imposing that the principal's policy be continuous in a neighborhood of the set of

[2] Similar concerns are expressed by the maxim "when a measure becomes a target, it ceases to be a good measure", sometimes referred to as Goodhart's law, based on an idea originally expressed in the context of monetary policy by Goodhart (1974).

[3] In Section 5.2.3 we extend the results to multi-dimensional state spaces.

equilibrium prices. This restriction captures, in a reduced-form way, the idea that policies which are overly sensitive to the market outcome will be vulnerable to manipulation. It can be viewed as a minimal implication of large-market incentive compatibility, which allows us to exploit the technical convenience of the competitive model while helping to ensure that the predictions are consistent with those of a model with small but not infinitesimal agents.[4]

As detailed above, our goal is to understand the restrictions imposed by feedback effects and manipulation concerns on the set of (uniquely) implementable outcomes. A challenge to identifying general properties of this set which hold in a wide range of environments is that the relationship between the state and the security price can be highly sensitive to details of the market micro-structure, i.e. the specific features of the game played by market participants (see for example Grossman and Stiglitz 1976; Kyle 1985; Kremer 2002). This is true even when there are no market-based interventions, the addition of which only complicates the analysis.[5] Thus, in order to isolate the general properties of market-based mechanisms, we derive a unified framework for studying a wide range of competitive markets. This consists of a set of representation results which allow us to subsume the details of the market micro-structure into a simple reduced-form representation.

Using this framework, we proceed to characterize equilibrium outcome distributions. It is convenient to describe these via the marginal distributions mapping states to principal actions and prices, which we refer to as the *action function* and *price function* respectively. Our main general characterization results show that a necessary condition for unique implementation of a given outcome distribution is that the associated price function be monotone (Theorem 1), and this condition is essentially sufficient subject to mild technical conditions (Theorem 2).[6]

To understand the significance of these results, it is important to emphasize that they depend on both the no-manipulation and unique-implementation constraints. One obvious requirement for implementability is that the action function be *measurable* with respect

[4] Of course, this minimal continuity restriction may not be sufficient to guarantee that equilibrium outcomes in a sequence of finite-agent models converge to those in our competitive limit. This will in general depend on the details of the market micro-structure. To the extent that additional conditions are needed to guarantee convergence, our results provide an upper bound on what market-based mechanisms can achieve.

[5] In essence, this complication is the subject of the Lucas critique (Lucas et al., 1976).

[6] There is no formal connection between this monotonicity condition and those of Myerson (1981) or Maskin (1999) which appear elsewhere in the implementation literature in very different settings.

to the price function, i.e. if the principal is to take distinct actions in two states, then the equilibrium security prices must differ to allow the principal to distinguish between them. Without any concern for manipulation or equilibrium multiplicity, this measurability condition fully characterizes the implementable set. In fact, we show that under either the no-manipulation or unique-implementation constraint, but not both, measurability remains (essentially) the only restriction on outcomes (Proposition 1). In contrast, monotonicity of the price function, a restriction which goes well beyond measurability, is necessitated by a subtle interplay between the two constraints.

Our characterization results clarify the limits of what market-based mechanisms can achieve. Beyond this conceptual insight, the results are also useful for design purposes: rather than optimize over policies mapping prices to principal actions, our characterization results allow us to reformulate the principal's problem as the much simpler one of choosing an action function (mapping the state to the principal's action) subject to the constraint that the induced price function be monotone. In Section 3 we develop a procedure for solving this problem, which is a modest generalization of the ironing technique of Toikka (2011).[7]

As an application, in Section 4 we combine our characterizations of the feasible outcome set with the ironing techniques developed in Section 3 to study market-based bailout policies. We provide a complete characterization of optimal policies and show how these can be understood via a simple a graphical argument derived from our modified ironing procedure.[8]

Finally, we explore additional properties and extensions of the baseline model (Section 6 and Appendix D). We show that the constraints of unique implementation and robustness to manipulation imply a natural notion of robustness to model misspecification. This means the principal's payoff is not highly sensitive to their potentially limited understanding of market fundamentals. We also analyze optimal policy when the requirement of unique implementation is relaxed. In particular, we use our characterization of the implementable set to show

[7]Existing analyses of market-based mechanism design optimize over the space of policy rules (see for example Hauk et al. (2020)). Generally, this approach requires one to impose restrictions on the environment and parameterize the space admissible policies in to solve for equilibrium outcomes in closed form as a function of decision-rule parameters.

[8]In Appendix C we study a version of this problem for a more complicated market micro-structure, the noisy REE model of Grossman and Stiglitz (1980). Further applications to carbon cap-and-trade and monetary policies are presented in Section G of the working paper version (Valenzuela-Stookey and Poggi, 2022).

that if the principal takes a worst-case view of equilibrium multiplicity then the restriction to unique implementation is generally without loss of optimality.

Related literature

This paper belongs to the literatures on mechanism design and implementation theory. In the language of this literature, we are interested in full implementation of a social choice function under a large-market notion of incentive compatibility. The distinguishing feature of the current paper is that the interaction between the principal and the agent is mediated by a market: the principal cannot observe the actions of the agents directly, but instead can only condition their action on the aggregate market outcome (price). In this sense we relate to the literature on mechanism design with limited communication (Segal, 2010; Mookherjee and Tsumagari, 2014). In the current paper the communication technology is a market which coarsens agents' individual actions into a single aggregate message. Given our focus on competitive markets we do not impose a standard incentive compatibility constraint, but consider a more limited non-manipulation condition.

In focusing on manipulation-proof mechanisms, we also contribute to recent work on falsification- and fraud-proof mechanism design (Perez-Richet and Skreta, 2022, 2024). As in these studies, preventing manipulation is an explicit objective of the designer, rather than an implication of the revelation principle, and non-manipulability is determined by the sensitivity of the mechanism to individual reports.[9] The specific no-manipulation criterion we consider is similar in spirit to large-market notions of incentive compatibility, such as Budish (2011) and Azevedo and Budish (2019), in that agents are assumed to be price takers and we require only a limiting notion of robustness to small manipulations. The connection between robustness to manipulation, equilibrium uniqueness (i.e. full implementation), and structural uncertainty, discussed in Appendix D.1, parallels Oury and Tercieux (2012), albeit in a very different setting. Our results on mechanism design under worst-case equilibrium selection relate to the large literature on robust mechanism design (Kajii and Morris, 1997; Carroll, 2015; Dworczak and Pavan, 2022).

The characterization results allow us to solve for optimal mechanisms in outcome space,

[9] The no-manipulation constraint in Perez-Richet and Skreta (2024) is essentially a bounded-differences condition on the mechanism, of which our continuity constraint can be viewed as a limit case.

i.e. the space of action and price functions, subject to the constraint the the price function be monotone. In fact, we can formulate the problem in terms of the price function alone. The problem of maximizing a real-valued functional over a space of monotone functions has been widely studied in the mechanism design literature (Mussa and Rosen, 1978; Myerson, 1981; Nöldeke and Samuelson, 2005; Hellwig, 2008; Toikka, 2011). However the standard ironing approach cannot be applied in the current setting since the set of feasible prices may vary with the state. We extend the results of Toikka (2011) to deal with this complication.[10] This extension makes a modest technical contribution to the growing literature on ironing under additional constraints, e.g. Jullien (2000); Kleiner et al. (2021); Loertscher and Muir (2022); Akbarpour et al. (2024); Kang (2024).

Finally, our work is also related to a large literature in macroeconomics and finance on the two-way feedback between financial markets and the real economy, beginning with Baumol (1965). We are by no means the first to recognize the presence of two-way feedback effects when policy is conditioned on markets. These forces are important in many contexts, such as monetary policy. 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). For a survey of this literature see Bond et al. (2012).[11]

Broadly, our contribution relative to this literature is to bring a design and implementation perspective to policy-making in these settings. We formalize the problem of policy design under commitment in a general setting and provide a full characterization of the feasible outcomes, while accounting for manipulation and equilibrium multiplicity concerns. This implementation question has not previously been studied in the literature. Thus, while we study familiar environments, our analysis is formally quite different from the literature on feedback effects in macroeconomics and finance.

We are also the first in this literature, to our knowledge, to study market-based mechanism design under no-manipulation and unique-implementation constraints simultaneously.

[10]In the standard mechanism-design language, the extension is to allow the set of feasible allocations to depend on the state, whereas feasible allocations are state-independent in Toikka (2011).

[11]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.

Other authors 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. The discontinuous nature of the policy considered in Lee (2019) opens the door to manipulation.

The literature has also documented the fact that feedback effects may induce equilibrium multiplicity (see among others Dow and Gorton (1997), Bernanke and Woodford (1997), Angeletos and Werning (2006), Glasserman and Nouri (2016)). However necessary and sufficient conditions for a policy to induce a unique equilibrium have not been established in a general setting, nor under the additional constraint of robustness to manipulation. To our knowledge, we are the first to study market-based policy under commitment subject to equilibrium uniqueness as a constraint (with or without manipulation-proofness).[12]

On a methodological level, our results characterizing the implementable set in *outcome space* (the set of action and price functions) greatly simplify the principal's problem. The alternative is to maximize directly over the set of policy rules. This approach has been applied fruitfully by Siemroth (2019) and Lee (2019). Hauk et al. (2020) develops variational techniques for optimizing in the space of policy rules. However this approach generally requires strong assumptions on the environment and/or exogenous restrictions to a parametric set of policies, so as to be able to solve for equilibrium outcomes in closed form as a function of decision-rule parameters. Our approach to optimizing over action and price functions is often much more tractable, allowing for flexible policy design and weaker assumptions on the market. This tractability is illustrated in Section 4.

The remainder of the paper is organized as follows. We introduce the general model in Section 1. In Section 2 we first present the main results in a baseline model which includes a few additional simplifying assumptions, but is nonetheless sufficient to communicate the key

[12]A related concern is equilibrium non-existence, which is a particularly salient issue when the principal lacks commitment power, see for example Bond et al. (2010) and Siemroth (2019). These papers focus on the fact that (in the language of the current paper) if the first-best action and price functions violate measurability (i.e. there are distinct states with the same price but different principal actions) then the first-best is not an equilibrium outcome, and an equilibrium may fail to exist. Siemroth (2019) identifies conditions for existence in a model without commitment.

economic message. We then develop an ironing technique for finding optimal mechanisms (Section 3) and apply this technique to study optimal bailout policies (Section 4). In Section 5 we return to the general model and extend the main characterization results. Further extensions and additional properties are discussed in Section 6.

1 General model

The principal must choose an action $a \in \mathcal{A}$, but does not observe the realization of a payoff-relevant state $\theta \in \Theta$. Both $\mathcal{A}$ and Θ are convex subsets of Euclidean spaces, and $\mathcal{A}$ is assumed to be compact.[13] The state is drawn from a distribution ν which is absolutely continuous with respect to the Lebesgue measure.

There is a market where an asset is traded by a unit mass of agents, each having private information about the state. It is notationally convenient to denote the set of asset prices by $\mathcal{P} := \mathbb{R}$. We refer to a function $P : \Theta \to \mathcal{P}$ as a *price function*. While the details of agents' information and the market structure are allowed to vary, the timing of the interaction is as follows throughout the paper:

1. The principal publicly commits to a *policy* $M : \mathcal{P} \to \mathcal{A}$.
2. The asset is traded by agents and a market-clearing price is determined.
3. If the price is p, the principal takes the action $M(p)$.[14]

In order to close the model, it remains to specify what information agents have, how the market operates, and how equilibrium is defined. Our objective is to provide a framework which can be applied to a wide range of markets. To that end, consider first the following general "reduced form" representation of the market: there exists a function $R : \mathcal{A} \times \Theta \to \mathcal{P}$ such that an equilibrium given policy M is defined as a price function P satisfying

$$P(\theta) = R(M \circ P(\theta), \theta) \quad \forall\, \theta \in \Theta. \tag{1}$$

We refer to eq. (1) as the *rational expectations equilibrium* (REE) condition. Crucially, the function R which defines an equilibrium is independent of the principal's policy M.

To understand this definition, consider the extreme case in which all agents in the market

[13] The Euclidean structure is not necessary: it suffices to assume that Θ is a convex subset of a topological vector space endowed with the Borel σ-algebra, and that $\mathcal{A}$ is a subset of a Banach space.

[14] Alternatively, one can interpret the model as representing the steady state of a repeated interaction, in which steps 2 and 3 alternate indefinitely.

observe the state, which we refer to as *symmetric information*.[15] If $R(a, \theta)$ describes the ex-post cash flow generated by the asset when the state is θ and the principal takes action a, then eq. (1) is simply the standard no-arbitrage condition: given the principal's policy and the asset price agents can perfectly predict the principal's action, and so the asset must trade at a price equal to its cash flow given this action.

Models with symmetric information in which equilibrium is defined by the REE condition in eq. (1) are widely studied in the literature on feedback effects (e.g. Bond et al. (2010); Hauk et al. (2020)). However, the symmetric information assumption is clearly restrictive in applications. Moreover, it raises the question of why, if each individual agent perfectly observes the state, the principal must use a market-based policy in the first place.

To study market-based mechanisms under richer market structures we develop a general framework which nonetheless retains the tractability of the symmetric information case. We do this by proving a set of representation results which show that for a broad class of markets there exists a function R such that equilibrium outcomes are characterized by the fixed-point condition in eq. (1). In other words equilibrium outcomes look "as if" there is symmetric information, but for a suitably modified cash-flow function. The function R, which we refer to as the *invariant representation* of the market, is derived from the true cash flows, the payoffs of agents, and the information structure. However, for the purposes of studying market-based mechanisms in such markets, we can simply treat R as we would the cash flow in a model with symmetric information.

The representation results are an important secondary contribution of this paper. However, the details of this approach are orthogonal to our primary focus, which is the design of market-based mechanisms. We therefore begin by studying the design question, taking the invariant representation R and the REE solution concept from eq. (1) as given. In Section 5, we return to the question of how R can be derived from the primitive market microstructure. The domain of applicability for our mechanism-design analysis is the set of markets which admit an invariant representation, and we show that this set includes many commonly studied models, such as the canonical noisy REE model of Grossman and Stiglitz (1980).

[15] As long as the agents observe the same signal, assuming that agents perfectly observe the state is without loss of generality.

Appendix B contains additional details and discussion.

2 Baseline model

We begin by presenting the key insights in a simplified version of the model where the principal's action is the probability of a binary intervention, agents have symmetric information, and the state space is one-dimensional. This model generalizes that in Bond et al. (2010).[16] In Section 5, we return to the setting of the general model.[17]

Consider a principal who must decide whether to take an intervention action $A \in \{0, 1\}$ in a publicly traded company. The principal's payoff from intervention depends on the underlying fundamentals of the firm, represented by a state $\theta \in (\underline{\theta}, \bar{\theta})$.[18] For example, the principal could be a government regulator who is considering a bailout for the company, but prefers to intervene only if the company is sufficiently distressed.

The principal does not observe the state θ, but the state is perfectly observed by agents who trade the firm's security.[19] To exploit the information of market participants, the principal announces a policy $M : \mathbb{R}_+ \to [0, 1]$ where $M(p)$ indicates the probability of intervention as a function of the price of the firm's security. Market participants observe the function M and state θ and trade the security in a competitive market, resulting in a market-clearing price p. The principal then observes the market-clearing price and takes the intervention according to the policy M.

For any probability of intervention $a \in [0, 1]$, we denote $R(a, \theta)$ the expected cash flow generated by the security in state θ, i.e.

$$R(a, \theta) := a \cdot r_1(\theta) + (1 - a) \cdot r_0(\theta) \tag{2}$$

Assume that $r_A(\cdot)$ is continuous and strictly increasing for $A \in \{0, 1\}$. The substantive implication of this assumption is that the state has the same qualitative effect on the security

[16] Bond et al. (2010) focus on market-based policy-making without commitment. They also consider the case where the principal has access to an exogenous signal of the state, which we do not consider here.

[17] While we maintain the assumption that the state space is one-dimensional for most of Section 5, we show in Section 5.2.3 that the key insights extend to settings with multi-dimensional state spaces.

[18] Assuming the state space is open rather than closed has no substantive impact, but simplifies the statement of the results, some of which would otherwise include conditions that apply only at the boundaries.

[19] As noted above, this assumption is not needed for the results presented here, but simplifies the exposition.

value, whether or not the principal intervenes.[20] We also impose the technical conditions that $r_1(\underline{\theta}) \neq r_0(\underline{\theta})$ and $r_1(\bar{\theta}) \neq r_0(\bar{\theta})$; and that the set of intersections of r_0 and r_1 is countable. These assumptions simplify the statement of some results, but are not essential. Equilibrium is defined by the standard no-arbitrage condition.

Definition 1. Fixing a policy M, an *equilibrium* consists of a price function $P : \Theta \to \mathcal{P}$ satisfying the rational expectations condition

$$P(\theta) = R(M \circ P(\theta), \theta) \qquad \forall \theta \in \Theta \tag{3}$$

Given a policy M and an associated equilibrium price function P, let $Q := M \circ P$ be the equilibrium action function. That is, Q is the induced equilibrium map from states to action probabilities. Together with the prior distribution on θ, the pair (Q, P) fully describes the equilibrium joint distribution of states, prices, and actions. It is convenient to formulate the design problem directly over action and price functions, i.e. the equilibrium outcomes, rather than the policy M. To do this, we must characterize the set of pairs (Q, P) that are equilibrium outcomes for some policy M. This is the focus of this section.

Definition 2. An action and price function pair (Q, P) is implementable if there exists a policy M such that P is an equilibrium price given M and $Q = M \circ P$.

Implementability of (Q, P) can be split into two intuitive conditions: a rational expectation condition and a measurability condition.

Observation 1. (Q, P) is implementable if and only if it satisfies

- **RE**: for every state θ, $P(\theta) = R(Q(\theta), \theta)$.
- **Measurability:** $P(\theta) = P(\theta')$ implies $Q(\theta) = Q(\theta')$.

The measurability condition derives from the fact that the principal's action is a function of the price. We say that Q is implementable if there is a P such that (Q, P) is implementable; the RE condition implies that the price function must then be given by $P(\theta) = R(Q(\theta), \theta)$.

Suppose the principal's payoff from action A when the state is θ be $v_A(\theta)$, and denote the expected payoff from intervening with probability a by $v(a, \theta) = av_1(\theta) + (1 - a)v_0(\theta)$.

[20]The assumption that r_A is strictly as opposed to weakly increasing simplifies the analysis, but is not essential here. Moreover in the general model of Section 5.2 we drop the assumption that R is linear, or even monotone, in a; and allow $\theta \mapsto R(a, \theta)$ to be weakly increasing.

Assume that $v_1(\theta) - v_0(\theta)$ is decreasing and crosses 0 at $\hat{\theta} \in (0, 1)$. That is, the principal prefers to intervene when the state is low.[21] It is easy to see that the first-best action function, which intervenes if and only if $\theta \leq \hat{\theta}$, may not be implementable. Consider, for example, the cases depicted in Figure 1, in which r_1 and r_0 are linear and intersect once at θ^*. The dotted blue lines represent the price function P that is consistent with rational expectations under the first-best action function. Then the first-best is implementable if and only if $\hat{\theta} \geq \theta^*$: for $\hat{\theta} < \theta^*$, the first-best action function induces a price function, P, that violates measurability, since for any price in $p \in (p', p'')$ there are states $\theta < \hat{\theta} < \theta'$ with $P(\theta) = P(\theta') = p$.

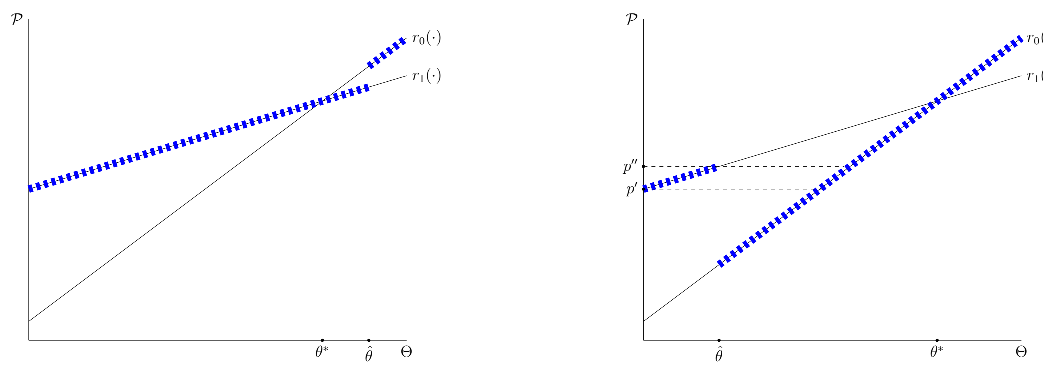

(a) $\hat{\theta} > \theta*$, implementable

(b) $\hat{\theta} < \theta^*$, not implementable

Figure 1: Price functions induced by the first-best action function

A natural question is whether there is an easy way to fix the measurability issue that arises when $\hat{\theta} < \theta^*$. Figure 2 illustrates such an attempt. The action function here coincides with the first-best for all states outside of (θ', θ''). On (θ', θ''), the action function specifies intermediate actions so as to induce the depicted price function, which satisfies measurability. If the interval (θ', θ'') has low probability, this modification provides a good approximation to the first-best expected payoff. Moreover, it is easy to see that it is implementable.

There are, however, two issues with implementing such an action function. First, the policy M that implements this outcome is discontinuous at p' and p''. This means that in some states a small manipulation affecting the market price can induce a discrete change in the action of the principal. Second, it turns out that the first best is not uniquely implemented by this policy: there are other equilibria induced by the same mechanism. We next refine the implementability criterion to address the concerns of vulnerability to market manipulation

[21] As in Bond et al. (2010), we assume that the principal does not care directly about the security price.

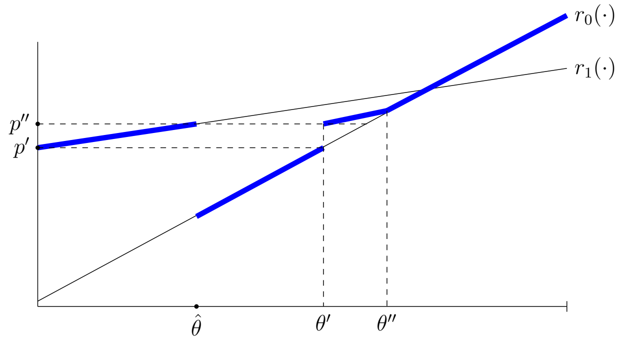

Figure 2: Restoring implementability

and equilibrium multiplicity. It turns out that while it is possible to resolve either the manipulation or the multiplicity issues separately by slightly perturbing the policy in Figure 2, jointly addressing these two concerns requires a significant change.

2.1 Market Manipulation and Equilibrium Multiplicity

Market manipulation is a salient concern in many market-based policy-making environments. Agents in the market may attempt to manipulate the price in order to influence the principal's action by buying/selling an asset, releasing false information, or other means.[22] While we focus on competitive markets in which agents are assumed to behave as price takers, 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 market depends on the sensitivity of the principal's policy. If, for example, the policy is discontinuous then an agent can 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 to restrict the principal to use a continuous policy. However the restricting to everywhere-continuous policies is stronger than is needed to address these concerns. If a discontinuity in the policy of the principal occurs at a price that is far from any which could arise in equilibrium then manipulation via a small price impact is not possible. Moreover, as discussed below, requiring continuity even at such prices artificially restricts the set of implementable

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

outcomes. We therefore impose a weaker condition than continuity.

Definition 3. A policy M is *essentially continuous* if for any equilibrium price function P given M, M is continuous on an open set containing the closure of $P(\Theta)$.[23]

In other words, an essentially continuous policy can have discontinuities only where there are no nearby equilibrium prices. Let $\mathcal{M}$ be the set of essentially continuous policies. We also refer to such policies as "robust to manipulation". As discussed above, this constraint can be thought of as the large-market analog to the no-falsification constraint in the literature on falsification- and fraud-proof mechanism design (Perez-Richet and Skreta, 2022, 2024).

It is important to emphasize that essential continuity (or even continuity) of the policy does not imply continuity of the endogenously-determined equilibrium price function. If one were to assume continuity of the price function, the measurability condition would immediately imply monotonicity of the price. However, there is no principled reason to directly impose continuity of the price function, and such exogenous restrictions on endogenous outcomes are at odds with our mechanism-design approach to the problem. We want to understand what are the restrictions driven fundamentally by the concern for manipulation.

A second concern is that 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. Such multiplicity is pervasive in market-based policy problems (Bernanke and Woodford, 1997). Moreover, given the decentralized nature of the policy the principal may be unable to select which equilibrium will be played, and the existence of multiple equilibria can 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). We are therefore interested primarily in unique implementation.[24]

Definition 4. A policy M is *robust to multiplicity* iff there is a unique equilibrium price function given M.

[23] Given a function $f : X \to Y$ and set $Z \subset X$, we write $f(Z)$ for the set $\{f(z) : z \in Z\} \subset Y$.

[24] In other words, full implementation of the desired action and price functions. In work subsequent to the current paper, Sturm (2023) and Dovis and Kirpalani (2023) also study unique implementation in alternative market-based policy-making contexts.

Unique implementation is desirable in many settings, especially those in which non-fundamental volatility is a concern. Moreover, we show in Section 6 that this restriction is without loss of optimality if the principal takes a worst-case view of equilibrium multiplicity.

In some cases, it will also be useful to consider a slightly weaker notion of robustness to multiplicity: the principal might be willing to have multiple equilibria as long as different equilibria lead to the same outcomes almost surely.

Definition 5. A policy M is *weakly robust to multiplicity* iff for any two equilibrium price function P and P' given M, it holds that $P(\theta) = P'(\theta)$ for almost every state θ.[25]

With the definitions of robustness to manipulation and multiplicity, we are ready to define the refined implementation criterion, which is our primary focus.

Definition 6. A pair of action and price functions (Q, P) is *continuously uniquely implementable* (CUI) if it constitutes an equilibrium under some policy M that is essentially continuous and robust to multiplicity. (Q, P) is *continuously weakly uniquely implementable* (CWUI) if it constitutes an equilibrium under some policy M that is essentially continuous and weakly robust to multiplicity.

It is also useful to introduce a notion of approximate implementation.

Definition 7. An action function Q is said to be *virtually implementable* with policies in set $X \subset \mathcal{A}^{\mathcal{P}}$ if there exist a sequence of action functions Q_n and policies $M_n \in X$ such that Q_n is implemented by M_n for all n, and $Q_n \to Q$ almost everywhere.

We first observe that neither essential continuity nor robustness to multiplicity significantly change the implementable set when considered singly.

Proposition 1. *Assume (Q, P) is implementable and P has countably many discontinuities. Then (Q, P) is also virtually implementable with an essentially continuous policy, as well as with a policy that is weakly robust to multiplicity.*

Proof. Proof in Appendix A.1. □

[25] This also implies that the corresponding action functions are equal almost everywhere.

Jointly however, the constraints imposed by manipulation and multiplicity concerns significantly restrict the set of implementable outcomes to those with monotone price functions.

Proposition 2. *If (Q, P) is CWUI, then P is monotone. If (Q, P) is CUI, then P is monotone and continuous.*

Proof. The proof, which is instructive, is presented in Section 2.2 below. □

Proposition 2 rules out the non-monotone price function depicted in Figure 2. It is possible to use an essentially continuous policy to implement a slight perturbation of this policy, using the construction from the proof of Proposition 1. However, Proposition 2 shows that this construction will necessarily induce multiple equilibria. The proof in Section 2.2 below illustrates why this is the case.

Conversely, monotonicity of the price function is effectively the only restriction imposed by essential continuity and robustness to manipulation.[26] Say that Q is virtually CWUI if Q is virtually implementable with policies that are both essentially continuous and weakly robust to multiplicity. Say that a price function is *essentially monotone* if it is monotone except for a zero-measure set of point discontinuities.[27]

Proposition 3. *Q is virtually CWUI if and only if $R(Q(\theta), \theta)$ is essentially monotone.*

Proof. Proof in Appendix A.4. □

It is worth emphasizing that both increasing and decreasing price functions are feasible, despite the fact that $\theta \mapsto R(a, \theta)$ is increasing for all $a \in [0, 1]$. Strictly decreasing P are implementable when $\max\{r_1(\underline{\theta}), r_0(\underline{\theta})\}$ is larger than the $\min\{r_1(\bar{\theta}), r_0(\bar{\theta})\}$. In such policies the intervention by the principal inverts the relationship between the state and market price. We show in Section 4 that this may be optimal in some cases. It is also interesting to note

[26] The monotonicity condition here bears no direct relation to the types of monotonicity conditions that are common in the implementation literature (Maskin (1999) and Myerson (1981)). This should not be surprising given the distinctive features of our setting, such as the focus on competitive markets.

[27] In other words, it is almost-everywhere equal to a monotone function. Formally, $P(\theta) \in [\min\{\lim_{\theta' \nearrow \theta} P(\theta), \lim_{\theta' \searrow \theta} P(\theta)\}, \max\{\lim_{\theta' \nearrow \theta} P(\theta), \lim_{\theta' \searrow \theta} P(\theta)\}]$ for almost all θ, and $\theta \mapsto \lim_{\theta' \nearrow \theta} P(\theta)$ is weakly monotone. We could alternatively strengthen the definition of virtual implementation to sup-norm convergence, and obtain a version of Proposition 3 where essential monotonicity is replaced simply with monotonicity. However the statement of this alternative result involves additional technical conditions, and does not differ substantively from Proposition 3.

that, as illustrated in the construction in the proof of Proposition 3, in order to implement a decreasing price function it is necessary to use a policy that is discontinuous, albeit outside of the convex hull of the set of equilibrium prices. This is the sense in which restricting attention to continuous policies, as opposed to essentially continuous ones, would be too restrictive: doing so would preclude decreasing price functions.

These results also provide a foundation for the use of relatively simple policies. Consider the non-monotone price function illustrated in Figure 2. Absent manipulation and multiplicity constraints, or even with only one of these conditions (as shown in Proposition 1), a policy of this form may well be optimal. However, implementing such an outcome requires a rather complicated non-monotone policy: for prices below p' the policy specifies no intervention; between p' and p'', it jumps to intervention with probability 1; just above p'', it intervenes with an interior probability; and then eventually returns to no intervention for higher prices. It is difficult to imagine a policymaker announcing such a rule in practice. By contrast, policies that induce monotone price functions tend to be much less complex. Moreover, we show in Section 4 that optimal CWUI policies in this setting take a particularly simple form.

2.2 Proof of Proposition 2

Proposition 2 is a special case of Theorem 1, which holds in a more general setting. However, the proof in the baseline model is more transparent and illustrates the key ideas.

We start by proving the first part of the proposition. In a slight abuse of notation, we use $R(a, \underline{\theta})$ and $R(a, \bar{\theta})$ to denote the limits of $R(a, \cdot)$ when approaching the extreme states. Let (Q, P) be CWUI. For any policy function M, let S_M be the set of pairs (θ, p) that satisfy the rational expectation condition $R(M(p), \theta) = p$ and let P_M be the set of prices consistent with equilibrium, i.e. $P_M := \{p \in \mathbb{R} : (\theta, p) \in S_M \text{ for some } \theta \in \Theta\}$. We begin with a technical lemma, the proof of which is deferred to the appendix.

Lemma 1. *If M is essentially continuous, then the set of equilibrium prices P_M is open.*

Proof. Proof in Appendix A.2. □

We now show that if (Q, P) is CWUI then P must be strictly monotone. The idea behind the proof is illustrated in Figure 3. The solid blue line represents an implementable but non-monotone price function, P. The question is whether we can construct an essentially continuous map from prices to actions which implements this function (a.e) uniquely. Given a candidate policy M, consider the correspondence $\theta_M(p) := \{\theta \in (\underline{\theta}, \bar{\theta}) : R(M(p), \theta) = p\}$. If M implements P, the graph of this correspondence, S_M, must contain that of P. Moreover, if M implements P uniquely then $\theta_M(p)$ should be empty for $p \notin P(\theta)$. However, we show that for prices in the convex hull of $P(\Theta)$ but not in $P(\Theta)$, the graph of θ_M must look something like the dashed line in Figure 3. In particular, as shown in Lemma 1, if $\theta_M(p)$ is non-empty then it is also non-empty in a neighborhood of p. Thus, it is not possible to specify the policy without inducing multiplicity on an interval of states; in Figure 3 there are multiple market-clearing prices for all $\theta \in [\theta', \bar{\theta})$. The formal argument is divided into two parts.

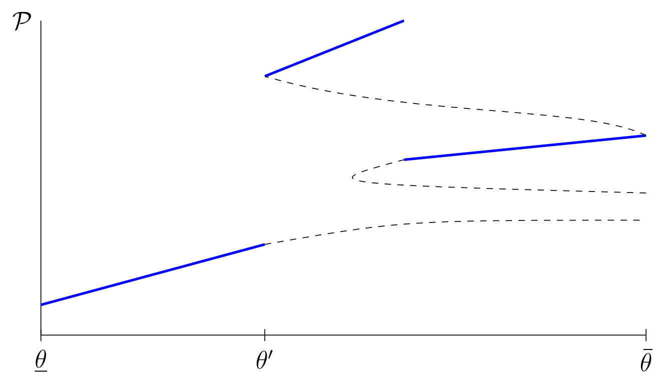

Figure 3: Multiplicity with a non-monotone price function

Lemma 2. *Let M be a essentially continuous policy that is weakly robust to multiplicity. The set of equilibrium prices P_M is an open interval $(\underline{p}, \bar{p})$. Moreover, either*

- $(\underline{\theta}, \underline{p}), (\bar{\theta}, \bar{p}) \in S_M$
- $(\bar{\theta}, \underline{p}), (\underline{\theta}, \bar{p}) \in S_M$

Proof. The proof is split in two parts. First, we show that any price p in the boundary ∂P_M must be an equilibrium price for a boundary state $\underline{\theta}$ or $\bar{\theta}$. Then, we show that for the same limit state $\theta \in \{\underline{\theta}, \bar{\theta}\}$, there cannot be different prices p_1 and p_2 in the boundary of P_M such that (θ, p_1) and (θ, p_2) are in S_M. Since the set P_M is bounded, there must be two boundary prices, and one has to correspond to each of the limit states.

To prove the first part, consider a price $p \in \partial P_M$. Since P_M is open, $p \notin P_M$, which means

$$p \leq R(M(p), \underline{\theta}) \qquad \text{or} \qquad p \geq R(M(p), \bar{\theta}). \tag{4}$$

Since p is in the boundary of P_M, there exists a sequence $\{p_n\}$ of prices in P_M that converges to p. For each of this prices, $R(M(p_n), \underline{\theta}) < p_n < R(M(p_n), \bar{\theta})$. Taking limits, since by essential continuity M is continuous in a neighbourhood of p, we obtain

$$R(M(p), \underline{\theta}) \leq p \leq R(M(p), \bar{\theta}) \tag{5}$$

Equations 4 and 5 jointly imply that either $p = R(M(p), \underline{\theta})$ or $p = R(M(p), \bar{\theta})$.

Now we prove the second part. Let $p_\circ \in \partial P_M$ and lets consider the case in which $p_\circ = R(M(p_\circ), \underline{\theta})$ (the case in which $p_\circ = R(M(p_\circ), \bar{\theta})$ is symmetric). By essential continuity, M is continuous in a neighbourhood B of $p_\circ$. Moreover, there is a sequence $\{p_n\}$ of prices in P_M that converges to $p_\circ$. Let's have n large enough so that $p_n \in B$ and let θ_n be the state such that $R(M(p_n), \theta_n) = p_n$.

Consider an arbitrary state $\theta \in (\underline{\theta}, \theta_n)$ note that the function $f(p) := R(M(p), \theta)$ is continuous in B. Moreover,

$$f(p_n) = R(M(p_n), \theta) - p_n < R(M(p_n), \theta_n) - p_n = 0$$
$$f(p_\circ) = R(M(p_\circ), \theta) - p_\circ > R(M(p_\circ), \underline{\theta}) - p_\circ = 0$$

Where the inequality holds by strict monotonicity of R in the second argument. Thus, applying the Intermediate Value Theorem, there is a price p in B such that $(\theta, p) \in S_M$. Robustness to multiplicity implies that there is a unique price $p_\circ$, since otherwise we would have multiple prices that are consistent with equilibrium in a neighbourhood of $\underline{\theta}$.

Observe that the previous result implies that there is at most two boundary points of P_M. Since P_M is bounded, there are exactly two boundary points, each of which corresponds to equilibrium at one of the extreme states. □

Lemma 3. *Let (Q, P) be CWUI. Then P is strictly monotone.*

Proof. We show that for any $\theta_L < \theta_H$ and $\hat{p}$ in between $P(\theta_L)$ and $P(\theta_H)$ there is a state $\theta \in (\theta_L, \theta_H)$ such that $(\theta, \hat{p}) \in S_M$.

Since P_M is convex, $\hat{p}$ must be an equilibrium price, i.e. there is a $\hat{\theta}$ such that $(\hat{\theta}, \hat{p}) \in S_M$.

Suppose without loss that $\hat{\theta} > \theta_H$. Then, consider the states in $\theta \in (\theta_H, \hat{\theta})$. Define $f(p) := R(M(p), \theta) - p$. This function is continuous in between $P(\theta_L)$ and $P(\theta_H)$ and strictly negative at both $P(\theta_L)$ and $P(\theta_H)$. Moreover, f is strictly positive at $\hat{p}$ therefore by the intermediate value theorem it must have at least two different zeros for each $\theta \in (\theta_H, \hat{\theta})$. This violates weakly unique implementation. $\square$

The proof of the second part of the proposition, namely that CUI implies monotonicity and continuity, is less illuminating and is deferred to Appendix A.3.

3 Solving the design problem

Having characterized the set of implementable outcomes (action and price functions) we can solve the principal's design problem directly in this space. To do so we develop a general ironing technique, which is applicable beyond the setting of Section 2.

Let R be the invariant representation of the market (in need not take the linear form in eq. (2)). Assume that R is differentiable and $\theta \mapsto R(a, \theta)$ is non-decreasing for all $a \in \mathcal{A}$, so the state has the same qualitative effect on the price, regardless of the principal's action. This assumption is satisfied in the illustrative model of Section 2 and maintained in the general model in Section 5. Let $v : \mathcal{A} \times \mathcal{P} \times \Theta \to \mathbb{R}$ be the principal's payoff. For each $p \in \mathcal{P}$ and $\theta \in \Theta$, let $\tilde{v}(p, \theta) = \max\{v(a, p, \theta) : R(a, \theta) = p\}$. Given $\tilde{v}$, we can solve the principal's program directly in terms of the price function P. Proposition 2 and Proposition 3 show that monotonicity of P is necessary and (essentially) sufficient for CWUI implementation in the baseline model, and Theorem 2 shows that this continues to hold in more general settings.

To solve the principal's program we consider separately the case of P non-increasing and P non-decreasing, and then compare the solutions. Since we maintain the assumption that $\theta \mapsto R(a, \theta)$ is non-decreasing for all a, these two cases are not symmetric. We consider first the non-decreasing case, which turns out to be the more difficult.

3.1 The non-decreasing case

Let $\mathcal{X}$ be the set of non-decreasing functions, P, such that $P(\theta) \in R(\mathcal{A}, \theta)$ for all $\theta \in \Theta$. For the case of P non-decreasing, the principal's problem can be written in terms of the price function as

$$\max_{P \in \mathcal{X}} \int \tilde{v}(P(\theta), \theta) d\nu(\theta). \tag{6}$$

Note that in order to maintain compactness of the policy space we do not impose that the price function be strictly increasing. While mechanisms with weakly increasing P may violate the measurability condition, arbitrarily small perturbations can be used to restore measurability.

Our approach of first characterizing the set of feasible outcomes and then maximizing directly over this set is analogous to the auction design procedure of Myerson (1981).[28] Indeed, we also arrive at a control problem with a monotonicity constraint, albeit via a very different path. If it were the case that $R(\mathcal{A}, \theta)$ was independent of θ, i.e. the set of possible prices did not depend on the state, then the program in (6) would be standard. Such problems are widespread in the mechanism design literature, and solutions using various techniques are well understood (Mussa and Rosen, 1978; Myerson, 1981; Nöldeke and Samuelson, 2005; Hellwig, 2008; Toikka, 2011).

The analysis here is complicated by the fact that in general the set of possible prices may depend on the state. For example, if $\mathcal{A} = [\underline{a}, \bar{a}]$ and $a \mapsto R(a, \theta)$ is increasing then $R(\mathcal{A}, \theta)$ is independent of θ if and only if $R(\bar{a}, \cdot)$ and $R(\underline{a}, \cdot)$ are constant. This assumption is overly restrictive in many settings, including the bailouts application studied below. Thus we must extend the standard ironing procedure to problems with state-dependent action sets (where prices here play the role of actions in the usual formulation). This extension parallels closely the results of Toikka (2011), but requires a few modifications.

Remark 1. The standard ironing approach (e.g. Toikka (2011)) derives an "ironed objective" function which can the be maximized state-by-state to obtain a solution. A natural conjecture for how to incorporate state-dependent constraints on prices is to derive the ironed objective function precisely as in Toikka (2011), and then simply impose the constraints on prices in the maximization step. Unfortunately, this simple fix does not work; it is necessary to modify the construction of the ironed objective itself, as we do here. The application studied

[28] In Myerson (1981), the set of feasible (IR and IC) outcomes are the direct revelation mechanisms for which the allocation rule is monotone and transfers satisfy the envelope condition (a characterization sometimes referred to as the Spence-Mirrlees lemma).

in Section 4 illustrates this point (see Remark 2).

For each $\theta \in \Theta$ write $R(\mathcal{A}, \theta) = [\underline{p}(\theta), \bar{p}(\theta)]$. Under the maintained assumption that $\theta \mapsto R(a, \theta)$ is non-decreasing and differentiable, both $\underline{p}(\theta)$ and $\bar{p}(\theta)$ are as well. Assume moreover that either $\underline{p}$ is constant or it is strictly increasing, and the same for $\bar{p}$. This assumption simplifies the analysis, but is not essential. Finally, assume that $\tilde{v}$ is continuously differentiable.

We take ν to be the uniform distribution on $[0, 1]$; this is without loss of generality as we can always make a change of variables to redefine the problem in terms of the quantiles of ν, which are uniformly distributed. Since this transformation is standard we omit the details (see for example Myerson 1981).

For $P \in \mathcal{X}$, define the generalized inverse of $P^{-1} : [\underline{p}(0), \bar{p}(1)] \to \Theta$ by

$$P^{-1}(x) := \inf\{\theta \in \Theta : P(\theta) \geq x\},$$

where we adopt the convention $\inf \varnothing = 1$. Define $\rho(x) := \sup\{\theta \in \Theta : \underline{p}(\theta) \leq x\}$, where we adopt the convention $\sup\{\varnothing\} = 0$ (note the distinction between $\rho(x)$ and $\underline{p}^{-1}(x)$). Under the maintained assumptions ρ and $\bar{p}^{-1}$ are continuous. Define

$$H(x, \theta) := \int_{\bar{p}^{-1}(x)}^{\theta} \tilde{v}_1(x, r) dr$$

and for all $x \in [\underline{p}(0), \bar{p}(1)]$

$$G(x, \cdot) := conv H(x, \cdot).$$

Importantly, since the domain of $H(x, \cdot)$ is $[\bar{p}^{-1}(x), \rho(x)]$, the convex hull of $H(x, \cdot)$ is

$$\begin{aligned} conv H(x, \theta) := \min_{\lambda, \theta_1, \theta_2} & \lambda H(x, \theta_1) + (1 - \lambda) H(x, \theta_2) \\ s.t. \quad & \theta_1, \theta_2 \in [\bar{p}^{-1}(x), \rho(x)] \\ & \lambda \in [0, 1], \\ & \theta = \lambda\theta_1 + (1 - \lambda)\theta_2. \end{aligned}$$

Let $g(x, \theta) := G_2(x, \theta)$ and define $\bar{v}(x, \theta) = \tilde{v}(\underline{p}(\theta), \theta) + \int_{\underline{p}(\theta)}^{x} g(s, \theta) ds$. Let Γ be the set of measurable functions P from Θ to $\mathbb{R}$ such that $P(\theta) \in [\underline{p}(\theta), \bar{p}(\theta)]$ for all θ. Let $\Phi(\theta) := \arg\max\{\bar{v}(p, \theta) : p \in [\underline{p}(\theta), \bar{p}(\theta)]\}$. Define $P^*(\theta) := \max \Phi(\theta)$ and $P_*(\theta) := \min \Phi(\theta)$.

Proposition 4. *(Values) If $x \mapsto \tilde{v}(x,\theta)$ is weakly concave then*

$$\sup_{P\in\mathcal{X}}\left\{\int_0^1 \tilde{v}(P(\theta),\theta)d\theta\right\} = \sup_{P\in\Gamma}\left\{\int_0^1 \bar{v}(P(\theta),\theta)d\theta\right\}.$$

Proposition 5. *(Maximizers) Let $x \mapsto \tilde{v}(x,\theta)$ be weakly concave. Then P^* and P_* are monotone and obtain the maximum in eq.* (6). *Furthermore, if $P \in \mathcal{X}$ solves (6) then $P_*(\theta) \leq P(\theta) \leq P^*(\theta)$ and $P(\theta) \in \Phi(\theta)$ a.e.*

The proofs of Proposition 4 and Proposition 5 are contained in Appendix A.5. Say that $\tilde{v}$ is *separable* if $\tilde{v}(p,\theta) = c(p)h(\theta)+b(p)+l(\theta)$, for continuously differentiable functions c, h, b, l and c strictly increasing.[29] For separable problems we can simplify the construction above. Define $D : [0,1] \to \mathbb{R}$ by

$$D(\theta) := \int_0^\theta h(r)dr$$

and for any x and any $\theta \in [\bar{p}^{-1}(x), \rho(x)]$ define

$$\begin{aligned} T(x,\theta) := \min_{\lambda,\theta_1,\theta_2} \ & \lambda D(\theta_1) + (1-\lambda)D(\theta_2) \\ s.t. \quad & \theta_1,\theta_2 \in [\bar{p}^{-1}(x), \rho(x)] \\ & \lambda \in [0,1], \\ & \theta = \lambda\theta_1 + (1-\lambda)\theta_2. \end{aligned}$$

In other words, $T(x,\theta)$ is the convex hull of the restriction of D to the interval $[\bar{p}^{-1}(x), \rho(x)]$. Let $t(x,\theta) = T_2(x,\theta)$.

Proposition 6. *For a separable problem we have*

$$\bar{v}(x,\theta) = c(\underline{p}(\theta))h(\theta) + b(x) + l(\theta) + \int_{\underline{p}(\theta)}^x c'(s)t(s,\theta)ds.$$

Proof. Proof in Appendix A.6 □

Note that in contrast the the standard case of state-independent action sets (where again, prices here play the role of actions in the standard formulation) separability of the objective does not obviate the need to iron action-by-action. In the separable problems of Myerson (1981) and Mussa and Rosen (1978) for example, one can identify $\bar{v}$ by convexifying

[29] Differentiability can be disposed of without great difficulty, at the cost of additional notation.

a function that does not depend on the action, over a fixed domain. In other words, if $\bar{p}^{-1}(x)$ and $\rho(x)$ are independent of x, then we can see from the construction above that $t(x, \theta)$ is also independent of x. This convenient property is lost when there are state-dependent action constraints. Instead, we must take Toikka's action-contingent ironing approach. Nonetheless, Proposition 6 simplifies the task of identifying $\bar{v}$. In particular, x only affects $t(x, \theta)$ via the changes in the domain of the function being ironed. This observation is useful for establishing properties of the solution, as illustrated by the graphical argument in Section 4.

3.2 The non-increasing case

The case where we require P to be non-increasing is more straightforward. For non-increasing P to be possible, it must be that $\underline{p}(1) < \bar{p}(0)$. Moreover, any non-increasing price function must take values in the interval $[\underline{p}(1), \bar{p}(0)]$. Thus we are back to the case in which the action set is independent of the state, and we can apply the usual ironing procedure. Let $\check{v}(x, \theta) = \tilde{v}(\bar{p}(0) - x, \theta)$. Then the results for the non-decreasing case apply with $\check{v}$ replacing $\tilde{v}$, where Γ is the set of measurable functions taking values in $[0, \bar{p}(0) - \underline{p}(1)]$ and $\mathcal{X}$ is the non-decreasing subset of Γ (in this case ironing technique collapses to that of Toikka 2011).

4 Application: optimal bailout policy

To illustrate how the characterization of CWUI outcomes can be combined with the results of Section 3 to solve the principal's program, we return to the setting introduced in Section 2. We interpret the principal as a government agency considering extending a bailout to a publicly-traded company, as in Bond and Goldstein (2015). We are particularly interested in qualitative features of the optimal policy, such as whether the principal is more responsive to the market when prices are high or low.

Let the action $a \in [0, 1]$ be the probability of intervention (alternatively, we can interpret a as the size of the bailout). Recall that the invariant representation is given by $R(a, \theta) = ar_1(\theta) + (1 - a)r_0(\theta)$, where r_1 and r_0 are strictly increasing. Assume that $r_1 > r_0$, i.e intervention always increases the company's gross cash flow.[30] Thus $\bar{p}(\theta) = r_1(\theta)$ and $\underline{p}(\theta) = r_0(\theta)$. Recall that the principal's expected payoff is $v(a, \theta) = av_1(\theta) + (1 - a)v_0(\theta)$, where

[30] The case where r_1 and r_0 cross can also be handled, but is slightly more involved.

$v_1(\theta) - v_0(\theta)$ is decreasing and crosses 0 at $\hat{\theta} \in (0,1)$. Then we can write

$$\tilde{v}(p,\theta) = \frac{p - r_0(\theta)}{r_1(\theta) - r_0(\theta)} v_1(\theta) + \left(1 - \frac{p - r_0(\theta)}{r_1(\theta) - r_0(\theta)}\right) v_0(\theta)$$
$$= \underbrace{v_0(\theta) - (v_1(\theta) - v_0(\theta)) \frac{r_0(\theta)}{r_1(\theta) - r_0(\theta)}}_{l(\theta)} + \underbrace{\frac{v_1(\theta) - v_0(\theta)}{r_1(\theta) - r_0(\theta)}}_{h(\theta)} \cdot p.$$

For simplicity we maintain here there assumption that θ is uniformly distributed. As noted above, the non-uniform case can be easily handled by making a change of variables to redefine the problem in quantile space. We study the solutions when P is restricted to be increasing and decreasing separately, and then compare them.

Price functions with flat segments cannot be induced exactly in equilibrium as they would violate measurability. Since we are ultimately interested in describing the policies that can be used to (approximately) implement the optimal policy, it is useful to first clarify how this approximation is done. If we have a price function that is non-decreasing (non-increasing) and flat at x over the interval $[\theta', \theta'']$ then for any $\varepsilon \neq 0$ we can let $M(x) = \{a : R(a, \theta') = x\}$ and $M(x+\varepsilon) = \{a : R(a, \theta'') = x + \varepsilon\}$, and take M to be a continuous and increasing (decreasing) interpolation in between x and $x + \varepsilon$ (the inverse of R in a is well defined under our maintained assumptions). Thus, we interpret a price function that is flat at level x as a policy which is highly sensitive to the price around x. Similarly, a price function which is discontinuous at θ is approximated by one which gradually adjusts the action over the range of the discontinuity, so as to induce a price function which approximates the jump at θ.

4.1 The non-decreasing case

We now turn to characterizing the optimal mechanisms. Note that the objective is separable, with $c(p) = p$, so we can apply Proposition 6. In particular

$$\bar{v}(x,\theta) = r_0(\theta)h(\theta) + l(\theta) + \int_{r_0(\theta)}^{x} t(s,\theta)ds,$$

where $t(x,\theta) = T_2(x,\theta)$, and

$$
\begin{aligned}
T(x,\theta) := \min_{\lambda,\theta_1,\theta_2} \quad & \lambda D(\theta_1) + (1-\lambda) D(\theta_2) \\
s.t. \quad & \theta_1, \theta_2 \in [r_1^{-1}(x), r_0^{-1}(x)] \\
& \lambda \in [0,1], \\
& \theta = \lambda\theta_1 + (1-\lambda)\theta_2.
\end{aligned}
$$

and $D(\theta) = \int_0^\theta h(r)dr$. The shape of the function D is the key to understanding the structure of optimal mechanisms. Surprisingly, there is always an optimal non-decreasing price function which takes one of three simple forms, illustrated in Figure 4.

1. **Responsive at the bottom.** $P(\theta) = r_0(\theta)$ for $\theta \geq \theta'$ and $P(\theta) = r_0(\theta')$ for $\theta < \theta'$, where $\theta' > \hat{\theta}$.
2. **Responsive at the top.** $P(\theta) = r_1(\theta)$ for $\theta \leq \theta'$ and $P(\theta) = r_1(\theta')$ for $\theta > \theta'$, where $\theta' < \hat{\theta}$.
3. **Responsive in the middle.** $P(\theta) = r_1(\theta)$ for $\theta \leq \theta'$, $P(\theta) = r_1(\theta')$ for $\theta \in (\theta', \theta'')$, and $P(\theta) = r_0(\theta)$ for $\theta > \theta''$, where $\theta' < \hat{\theta} < \theta''$ and $r_1(\theta') = r_0(\theta'')$.

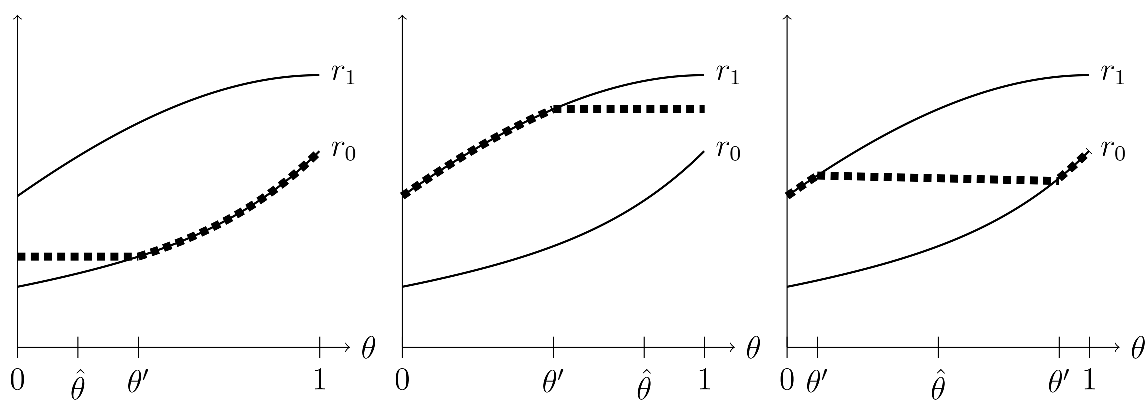

(a) Responsive at bottom (b) Responsive at top (c) Responsive in middle

Figure 4: Possible non-decreasing price functions for the no-gap case

The names of these classes of mechanisms are derived from the policies used to (approximately) induce them: A policy that is responsive at the bottom is such that the probability of intervention is decreasing and highly sensitive to the price for low prices, and is constant at no intervention for higher prices. Conversely, a policy that is responsive at the top is highly sensitive to the price for high prices and constant at full intervention for low prices.

Policies that are responsive in the middle are constant for both low and high prices, and highly sensitive over a narrow range of intermediate prices.

We say that there is a *gap* if $r_1(0) > r_0(1)$ (this is also the only case in which mechanisms with decreasing price functions are possible). Define $r_1^{-1}(x) := \inf\{\theta : r_1(\theta) \geq x\}$ and $r_0^{-1}(x) = \rho(x) := \sup\{\theta : r_0(\theta) \leq x\}$. Define

$$\phi(\theta) := \arg\max \left\{ \int_{r_0(\theta)}^{x} t(s,\theta)ds \ : \ x \in [r_0(\theta), r_1(\theta)] \right\}.$$

Proposition 7. *There is an optimal increasing price function, P^*, which is*

1. *Responsive at the bottom if $D\big(r_0^{-1}(r_1(0))\big) \leq 0$.*
   - *In particular, P^* is flat at the x^* such that $D(r_0^{-1}(x^*)) = 0$.*
2. *Responsive at the top if $D\big(r_0^{-1}(r_1(0))\big) \geq 0$ and $D(1) \geq D\big(r_1^{-1}(r_0(1))\big)$.*
   - *In particular, P^* is flat at the x^* such that $D(r_1^{-1}(x^*)) = D(1)$.*
3. *Responsive in the middle if $D\big(r_0^{-1}(r_1(0))\big) \geq 0$ and $D(1) \leq D\big(r_1^{-1}(r_0(1))\big)$*
   - *In particular, P^* is flat at the x^* such that $D(r_1^{-1}(x^*)) = D(r_0^{-1}(x^*))$.*

Moreover, the optimal increasing price function is unique unless there is a gap and $D(0) = D(1)$. If these two conditions hold then all and only increasing selections from $[r_0(1), r_1(0)]$ are optimal. (Note that when there is a gap $0 = r_1^{-1}(r_0(1))$ and $1 = r_0^{-1}(r_1(0))$.)

Proof. Proof in Appendix A.7. □

The conclusions of Proposition 7 can essentially be read directly from the figures depicted in Figure 5. As the left panel shows, these figures cover the case where $D\left(r_0^{-1}(r_1(0))\right) \geq 0$ (moreover we can see from the figures that $D(1) > D\left(r_1^{-1}(r_0(1))\right)$, although this is not explicitly illustrated in either panel). Since $T(r_1(0), \cdot)$ is strictly increasing, $t(r_1(0), \theta)$ is strictly positive for all $\theta \leq r_0^{-1}(r_1(0))$. Moreover, it is not hard to see from the figure that $t(x, \cdot)$ will also be strictly positive for $x < r_1(0)$. Thus to solve

$$\max \left\{ \int_{r_0(0)}^{x} t(s,0)ds \ : \ x \in [r_0(0), r_1(0)] \right\}$$

we want to choose x as high as possible, i.e. $P(0) = r_1(0)$.

The right panel of Figure 5 shows $T(x, \cdot)$ for $x > r_1(0)$. As x increases, the interval $[r_1^{-1}(x), r_0^{-1}(x)]$ shifts to the right. Eventually we reach an x^* such that $T(x, \cdot)$ is flat. If this occurs for $x^* < r_0(1)$ then the optimal policy will be responsive in the middle. From the figures we can see that $x^* < r_0(1)$ precisely when $D(1) < D\left(r_1^{-1}(r_0(1))\right)$. Otherwise the optimal mechanism responsive at the top. The proof of Proposition 7 simply formalizes this graphical argument. The graphical argument also facilitates comparative statics, although we do not pursue this direction further here.

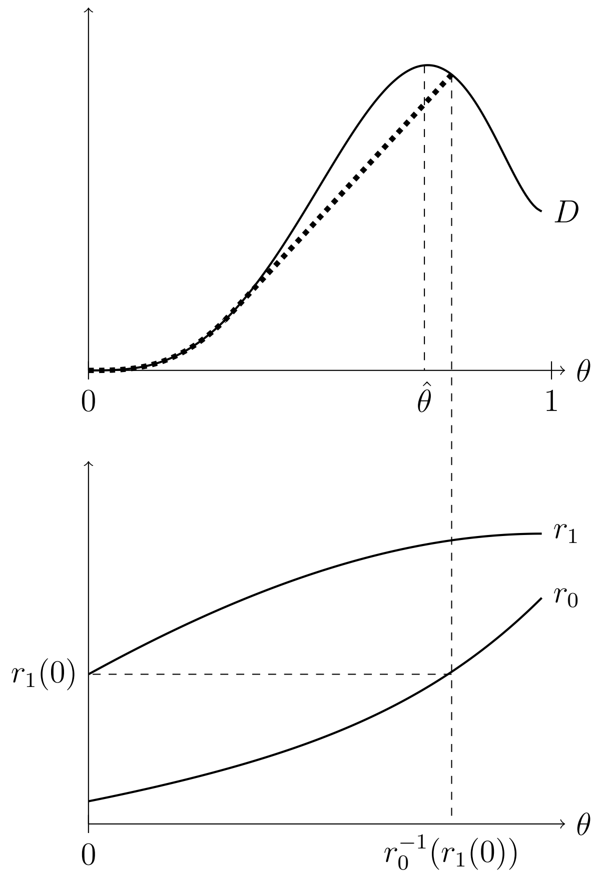

Top: D (solid) and $T(r_1(0),\cdot)$ (dashed)
Bottom: depicting $r_0^{-1}(r_1(0))$

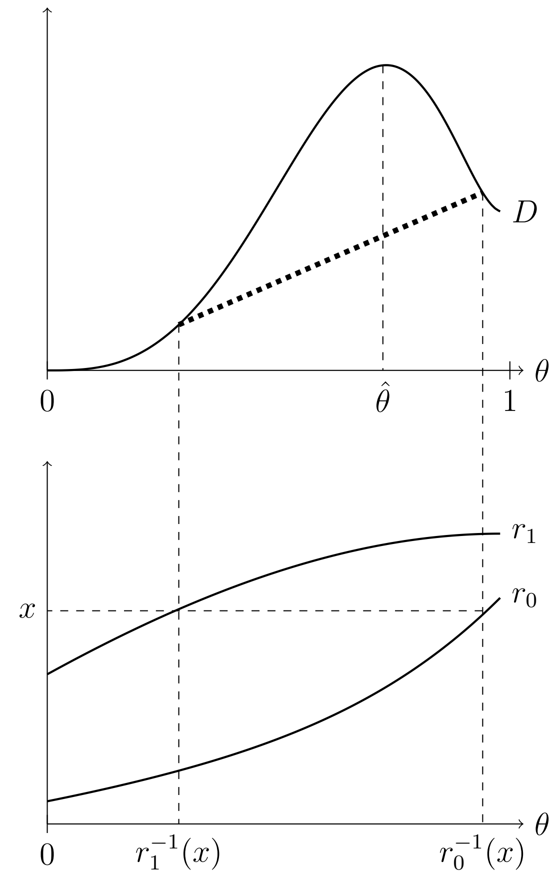

Top: D (solid) and $T(x,\cdot)$ (dashed)
Bottom: depicting $r_1^{-1}(x)$ and $r_0^{-1}(x)$)

Figure 5: Constructing the function T

Remark 2. This example illustrates why it is necessary to modify the ironing procedure to account for state-dependent action constraints. Simply ironing as in Toikka (2011) and then solving the state-by-state maximization problem subject to the action bounds does not work. This approach would mean taking the convex hull of D on $[0,1]$ in order to determine T and t. We would conclude that that $t(x,\theta) \geq (\leq)\ 0$ for all x, θ if $D(1) \geq (\leq)\ D(0)$. Then maximization of $\bar{v}(\cdot,\theta)$ over $[r_0(\theta), r_1(\theta)]$ would yield $P(\theta) = r_1(\theta)$ for all θ if $D(1) >\ D(0)$ and $P(\theta) = r_0(\theta)$ for all θ if $D(1) <\ D(0)$. As Proposition 7 shows, neither is ever optimal.

4.2 The non-increasing case

Strictly decreasing price functions are only possible if there is a gap, i.e. $r_1(0) > r_0(1)$. Assuming this holds, define

$$\begin{aligned}\check{v}(x,\theta) &:= l(\theta) + h(\theta)(r_1(0) - x) \\ &= \underbrace{l(\theta) + h(\theta)r_1(0)}_{\check{l}} \underbrace{-h(\theta)}_{\check{h}} x.\end{aligned}$$

Then replacing $\tilde{v}$ with $\check{v}$, and using the constant action bounds $[0, r_1(0) - r_0(1)]$ we obtain

$$\bar{v}(x,\theta) = \check{l}(\theta) + \int_0^x \check{t}(s,\theta)ds$$

where $\check{t}$ is the derivative of the convexification of $\check{D}(\theta) := \int_0^\theta \check{h}(r)dr = -\int_0^\theta h(r)dr$ over the interval $[0, r_1(0) - r_0(1)]$.

Proposition 8. *The unique optimal decreasing mechanism sets $P(\theta) = r_1(0)$ for $\theta < \hat{\theta}$ and $P(\theta) = r_0(1)$ for $\theta > \hat{\theta}$.*

Proof. Since $-h$ crosses 0 once from below, $\check{D}$ is quasi-convex with a unique minima at $\hat{\theta}$. Then $\check{t}(s,\theta) < (>) 0$ for all $s \in [0, r_1(0) - r_0(1)]$ and all $\theta < (>)\ \hat{\theta}$. □

In terms of the policy, the (approximately) optimal policy does not intervene at $r_0(1)$ and then features a rapidly increasing probability of intervention for prices just above $r_0(1)$, but only up to a level a such that $ar_1(\hat{\theta}) + (1-a)r_0(\hat{\theta})$ is approximately equal to $r_0(1)$. The probability of intervention then gradually increases as a function of the price, so as to induce a P that is nearly vertical around $\hat{\theta}$. Then for prices just below $r_1(0)$ the probability of intervention again rapidly increases in the price, reaching 1 at $r_1(0)$. In other words, policy is highly responsive to the price at the top and bottom, and gradually adjusts in the middle. Note that this policy inverts the meaning of high and low prices, relative to the market without intervention: high prices are now associated with lower states.

4.2.1 Possibility of both increasing and decreasing optimal price functions

Without engaging in a detailed comparison to determine when the optimal policy features an increasing or decreasing price function, we can nonetheless see that either case may be optimal for some parameter values. This follows from the fact that when there is a gap and $D(0) = D(1)$, there are optimal increasing-price policies which have a constant, and thus non-increasing, price (Proposition 7). In such cases the optimal price function is therefore non-decreasing. Moreover, by Proposition 8 the uniquely optimal price function is not constant. Thus decreasing-price policies remain optimal when $D(0)$ is close to $D(1)$.

5 Results for the general model

We show here that the characterizations outlined in Section 2 hold more generally, and explore additional extensions.

5.1 The market

One important restriction in Section 2 was the assumption that agents had symmetric information about the state. This assumption precludes many applications of interest. In order to have a unified framework for studying different markets without the symmetry assumption, we begin by deriving a convenient way to succinctly summarize equilibrium outcomes in a market. Informally, a market *admits an invariant representation* if there exists a function $R : \mathcal{A} \times \theta \to \mathcal{P}$ such that P is an equilibrium price function if and only if it satisfies the REE condition in eq. (1). In other words, the state and the principal's action pin down the price in any equilibrium. The important restriction inherent in this representation is that the function R is independent of the principal's policy. While this appears to a be a restrictive condition, it turns out to hold in a range of competitive markets, including the canonical noisy rational expectations equilibrium model of Grossman and Stiglitz (1980) which has been widely used to study market-based policy. To preview, the following are examples of markets covered by our analysis.

1. *Private values.* A continuum of price-taking agents have private values for an asset. Let $x_i(s_i, p, a)$ be the demand of agent i with value s_i when the price is p and the agent anticipates the principal will take action a. Assume $p \mapsto x_i(s_i, p, a)$ is strictly decreasing for all s_i, a. Let $\theta = \{s_i\}_{i \in I}$ be the profile of types. An equilibrium is defined by a price function $P : \Theta \mapsto \mathcal{P}$ satisfying the market-clearing condition

$$0 = \int_{i \in I} x_i(s_i, P(\theta), M \circ P(\theta)) di.$$

For example, the agents may be firms trading carbon credits (as in Section G of Valenzuela-Stookey and Poggi (2022)). The firm's type is its abatement cost (of reducing emissions). The principal's action is a policy that affects the abatement cost.

2. *Labor market.* There are a continuum of firms and workers. The principal is a government agency which commits to a map M from the unemployment rate p to a level of unemployment benefits a. Each worker's search effort is determined by the level of benefits, and their probability of being unemployed depends on their search effort and labor demand conditions θ. Given benefit level a, worker i's probability of being unemployed can therefore be summarized by a function $u_i(a, \theta)$. An equilibrium is defined by a function $P : \Theta \mapsto \mathcal{P}$ such that $\int_i u_i(M \circ P(\theta), \theta) = P(\theta)$.[31]
3. *Common values, imperfect information.* Consider a standard rational expectations equilibrium (REE) model of an asset market. The asset's ex-post cash flow given state $\theta \in \mathbb{R}$ and principal action a is $\pi(\theta, a)$. Each agent observes a private signal $s_i = \theta + \varepsilon_i$. Agents then submit a demand schedule $X_i(p, s_i)$ to a market maker, and have utility $u_i : \mathbb{R} \to \mathbb{R}$. Given a policy M, P is called a REE price function iff there exist demand functions $X_i(p, s_i)$ s.t.
   (a) $X_i(p, s_i) = \arg\max_x \mathbb{E}\left[u_i\big(x \cdot (\pi(M(p), \theta) - p)\big) \mid s_i, P(\theta) = p\right]$, and
   (b) $\int X_i\left(P_M(\theta), s_i\right) di = 0 \quad \forall \quad \theta \in \Theta$.
4. *Noisy REE (Grossman and Stiglitz, 1980).* The model is as in the previous example, except that there is aggregate uncertainty in the form of a supply shock z. This is a workhorse model of asymmetric information in asset markets. The asset value is $\pi(a, \omega)$ and signals are $s_i = \omega + \varepsilon_i$. The state includes the payoff-relevant state and supply shock: $\theta = (z, \omega)$. Market clearing means
$$\int X_i\left(P(\theta), s_i\right) di = z \quad \forall \quad \theta \in \Theta.$$

We now define formally the conditions characterizing the set of markets to which our analysis applies. In anticipation of the fact that we will only be interested in a subset of all possible policies, e.g. those which satisfy essential continuity, let $\mathcal{W} \subseteq \mathcal{A}^{\mathcal{P}}$ be a set of *admissible* policies.[32] We maintain that the assumption that all constant policies are

[31] We can interpret this fixed-point condition as a steady-state condition in a dynamic model.
[32] $\mathcal{A}^{\mathcal{P}}$ is the set of functions from $\mathcal{P}$ to $\mathcal{A}$.

admissible: letting M_a be the constant policy for action $a \in \mathcal{A}$, we have $M_a \in \mathcal{W}$ for all $a \in \mathcal{A}$.

In general, after the principal chooses the policy $M \in \mathcal{W}$, a set of market participants play some continuation game. The actions of the market participants in this continuation game, together with the state, determine a price in $\mathcal{P}$. Let this relationship be represented by $\tilde{P}(\theta, \psi)$ where ψ is a profile of actions taken by the market participants.

The game is paired with a solution concept, which for each M defines a set of equilibria in the continuation game. We call the game/solution-concept pair a *market* if for any policy M and equilibrium strategy profile ψ^*_M, the map $\theta \mapsto \tilde{P}(\theta, \psi^*_M)$ is a deterministic function.[33] In other words, given a policy M, for any equilibrium in the market there exists a *price function* $P \in \mathcal{P}^\theta := \{P : \Theta \to \mathcal{P}\}$ describing the price which realizes in each state. We maintain the assumption that there exists at least one equilibrium for every constant policy, i.e. every fixed principal action.

There may be multiple equilibria given M, with distinct price functions. We say that P is an *equilibrium price function* given M if it is the price function in some equilibrium given policy M. The fact that the set of equilibrium price functions depends on M is the source of the feedback effect.

Definition 8. A market *admits an invariant representation* given $(\mathcal{W}, \hat{\mathcal{P}})$, where $\hat{\mathcal{P}} \subset \mathcal{P}^\Theta$, if there exists a function $R : \mathcal{A} \times \Theta \to \mathcal{P}$ such that

1. If P is an eq. price function given $M \in \mathcal{W}$, then $P(\theta) = R(M \circ P(\theta), \theta)$ $\forall \theta$.
2. If $P \in \hat{\mathcal{P}}$ and $P(\theta) = R(M \circ P(\theta), \theta)$ for all θ then P is an eq. price function given M.

The key general observation for our analysis is that, for the purposes of choosing M, a market with an invariant representation can be fully summarized by R: that is, $P \in \hat{\mathcal{P}}$ is an equilibrium price function given M iff $P(\theta) = R(M \circ P(\theta), \theta)$.[34]

[33] The assumption that this function is deterministic just means that θ captures all the relevant uncertainty in the market. This definition of a market does not imply that the agents' strategies must be measurable with respect to the state. The state may contain dimensions that are not directly payoff relevant for the principal. For example, in a noisy REE model of an asset market, as in Grossman and Stiglitz (1980), the state will include the supply shock and the "payoff relevant state", but not the agents' private signals. In other settings, the state can represent the entire profile of agents' private signals, as in Jordan (1982).

[34] Moreover, since R does not depend on the policy 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 policy was used. Thus

It may be that a given market only admits an invariant representation given a $\hat{\mathcal{P}}$ that is a strict subset of $\mathcal{P}^{\Theta}$, and $\mathcal{W}$ that is a strict subset of $\mathcal{A}^{\mathcal{P}}$. In this case, the extent to which the property is useful depends on whether $\hat{\mathcal{P}}$ and $\mathcal{W}$ contain the action and price functions which are relevant for the principal's design problem (see Appendix B.1 for further discussion). The smaller the sets $(\mathcal{W}, \hat{\mathcal{P}})$, the easier it is to satisfy the conditions for an invariant representation. Indeed, if R is an invariant representation given $(\mathcal{W}'', \mathcal{P}'')$ and $\mathcal{W}' \subset \mathcal{W}''$, $\mathcal{P}' \subset \mathcal{P}''$, then R is also an invariant representation given $(\mathcal{W}', \mathcal{P}')$. On the other hand, the smaller is $(\mathcal{W}, \mathcal{P})$, the less useful the invariant representation is.

Determining whether or not a given market admits an invariant representation with respect to the desired $(\mathcal{W}, \hat{\mathcal{P}})$ is the first step for applying our analysis. Based on Definition 8, it is not immediately obvious which markets admit an invariant representation. In part, this is because of the existential qualifier in the definition. In Appendix B.1, we provide an axiomatic characterization of the markets that admit an invariant representation. This makes it relatively easy to check whether a given market satisfies this property.

In most applications it is enough, for practical purposes, that there exists a function R which describes the equilibrium for almost all states. The market admits an *a.e. invariant representation* given $(\mathcal{W}, \hat{\mathcal{P}})$ if Condition 1 in Definition 8 holds for almost all θ (with Condition 2 unchanged). To facilitate applications, we show the following.

Proposition 9. *In the above examples*[35]

- *Markets 1-2 admit an invariant representation in* $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$.
- *Market 3 admits an a.e invariant representation in* $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$. *The a.e. qualifier can be dropped if* $\theta \mapsto \pi(a, \theta)$ *is strictly increasing for all* a.
- *Market 4 admits an a.e invariant representation, where* $\mathcal{W}$ *is the set of policies that induce a unique equilibrium, and* $\hat{\mathcal{P}}$ *is the set of price functions with "non-intersecting level sets".*[36] *(Under the assumptions of CARA utility,* π *affine in* ω*, additive nor-*

a principal contemplating the introduction of a market-based policy can use historical aggregate data to estimate the function R and design the policy, without being subject to the Lucas critique that a change in the policy regime will change the relationship between the fundamentals (state and principal action) and the price (Lucas et al., 1976).

[35] Many of the parametric assumptions imposed on these markets can be relaxed. See Appendix B.3.

[36] These sets $\mathcal{W}$ and $\mathcal{P}$ are discussed in detail in Appendix C. We are interested in unique implementation, so the restriction to such policies is without loss. Moreover, we show that the "non-intersecting level sets" property is a necessary condition for unique implementation, and so the restriction to this set of price

mal signal structure and normally distributed supply shocks, which are standard in the literature.)

Proof. In Appendix B.5. □

Again, the important takeaway from Proposition 9 is that in all the markets in question, we know that P is an equilibrium price function given $M \in \mathcal{W}$ if $P \in \hat{\mathcal{P}}$ and $P(\theta) = R(M \circ P(\theta), \theta)$; and that in any equilibrium this condition holds almost everywhere (so for design purposes, we can focus on equilibria such that it holds for all θ).

Appendix B contains further discussion of the invariant representations. It is worth noting that not all markets admit an invariant representation. The key property is that the market be competitive, in the sense formalized in Appendix B.1.

5.2 General characterization results

Aside from assuming that agents in the market had symmetric information, as discussed above, the model of Section 2 embedded a few additional simplifications relative to the general model laid out in Section 1. First, we imposed that the principal's action was one-dimensional. We also assumed that $R(a, \theta)$ was linear in a and strictly increasing in θ. Finally, we restricted attention to a one-dimensional state space. In this section we show that the key results continue to hold without these restrictions. Specifically, we show that a monotone price function is necessary when the invariant representation R is weakly increasing in the state (Theorem 1), and essentially sufficient under additional mild conditions (Theorem 2).

For simplicity, we assume in this section that the market admits and invariant representation in$(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$. This assumption can be easily relaxed, and we do so in Section 5.2.3. To simplify the statement of the results, we also assume that the state space is an open bounded interval $\Theta = (\underline{\theta}, \bar{\theta})$.[37] We then extend the results to multidimensional Θ. We also assume that the invariant representation R is continuous in (a, θ) and non-decreasing in θ.[38] Finally, we add two technical conditions on the invariant representation at the extreme states. Let $R(a, \underline{\theta}) := \inf_{\theta \in \Theta} R(a, \theta)$ and $R(a, \bar{\theta}) := \sup_{\theta \in \Theta} R(a, \theta)$. First, we assume that R converges

functions is also without loss.

[37]The results for closed θ are the same, except that it is necessary to modify the boundary conditions in Theorem 2. We omit this result in the interest of brevity.

[38]In many micro-foundations both monotonicity and continuity of R can be derived from natural conditions on primitives, as discussed in Appendix B.

uniformly to the extremes. 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 $R(a,\underline{\theta})$ and $R(a,\bar{\theta})$. Second, we assume that for all $p \in \mathcal{P}$ and $\theta \in \{\underline{\theta},\bar{\theta}\}$, the set of actions for which $R(a,\theta) = p$ is the union of finitely many connected subsets of $\mathcal{A}$.[39] These technical assumptions are satisfied in all applications we consider.

Theorem 1 (Necessity)**.** *If (Q,P) is CWUI, then P is monotone.*

Proof. In Appendix A.8. □

A monotone price function, together with the market-clearing condition $P(\theta) = R(M \circ P(\theta),\theta)$, is nearly, but not exactly, sufficient for CUI. We require an additional technical condition.

Definition 9. The action function Q satisfies boundary condition 1 (BC1) if there are $\bar{Q},\underline{Q} \in \mathcal{A}$ such that $\bar{Q} = \lim_{\theta\to\bar{\theta}} Q(\theta)$ and $\underline{Q} = \lim_{\theta\to\underline{\theta}} Q(\theta)$. Moreover, Q satisfies boundary condition 2 (BC2) if it satisfies BC1, $R(\underline{Q},\underline{\theta}) \neq \inf \mathcal{P}$ implies that $R(\cdot,\underline{\theta})$ doesn't have a local maximum at $\underline{Q}$, and $R(\bar{Q},\bar{\theta}) \neq \sup \mathcal{P}$ implies $R(\cdot,\bar{\theta})$ doesn't have a local minimum at $\bar{Q}$.

Theorem 2. *Assume R is strictly increasing in θ.*[40] *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 satisfies BC1. Moreover, if P is decreasing, Q satisfies BC2.*

Proof. In Appendix A.10. □

5.2.1 Characterization results: technical discussion

The first point in Theorem 2 is simply the rational expectations condition that was necessary for implementation even without additional constraints. It is worth noting that continuity of Q is not implied by the continuity of M, but is instead a consequence of requiring unique implementation. In Section 5.2.2 we show that by slightly relaxing to weak robustness to

[39] Given continuity, this assumption means that the market-clearing price at the extremes does not oscillate too frequently (as a function of the action).

[40] In Proposition 19 we extend the result to weakly increasing R.

multiplicity we get a characterization that allows for discontinuous Q, so we do not view condition 3 as a critical characteristic of implementable pairs. The monotonicity of P is the essential point.

Notice that for any (Q, P) that is CUI, the continuity of Q implies continuity of P on Θ, 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.

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 is a way to define a continuous M for prices slightly above $\sup P(\Theta)$ and slightly below $\inf P(\Theta)$ such that these prices can never occur in equilibrium.

5.2.2 Extension: characterizing CWUI outcomes

To characterize the set of CWUI outcomes, i.e. those (Q, P) that are implemented by a policy in $\mathcal{W}^-$, we make the assumption that the market is *fully bridgeable*:

Definition 10 (Full bridgeability)**.** For any state $\theta \in \Theta$ and actions a, a' such that $R(a, \theta) \neq R(a', \theta)$, there exists a continuous function $\gamma : [0, 1] \to \mathcal{A}$ such that

- $\gamma(0) = a$, $\gamma(1) = a'$.
- $x \mapsto R(\gamma(x), \theta)$ is strictly monotone.

If $\mathcal{A} = [0, 1]$ then full bridgeability is satisfied iff $a \mapsto R(a, \theta)$ is monotone for every θ. For more general action spaces, weaker notions of monotonicity suffice.[41]

[41] 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 $\pi(z, \theta)$ (for example, Z may be a set of conditions that the principal can attach to a bailout, and π the cash-flow of the company). Fixing the state θ, any action $a \in \Delta(Z)$ induces a distribution over the set of values $\pi(Z, \theta)$. 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 full bridgeability. This result, along with more general sufficient conditions for bridgeability, is discussed in Appendix E, where we also relax the assumption.

Proposition 10. *Assume R is strictly increasing in θ and the market is fully bridgeable. 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 is also discontinuous at θ^*.*
4. *Q is BC1. Moreover, if P is decreasing, Q is BC2.*

Proof. In Appendix A.11 □

The main substantive difference between CUI and CWUI outcomes is that the action function need not be continuous.

5.2.3 Extension: multidimensional state space

Suppose that Θ is an open subset of $\mathbb{R}^N$, endowed with the usual product partial order. When working with a one-dimensional state space, we were able to prove the existence of an invariant representation given $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$ by imposing monotonicity of some primitive objects, such as asset dividends, in θ (see Appendix B.3). The difficulty with moving to a multi-dimensional state space is that we cannot in general identify a complete order on Θ for which such monotonicity conditions hold. It is therefore useful, for example for the application of Appendix C, to relax the assumption that the market admits an invariant representation given $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$.

We first introduce an additional condition: say the market has *level sets represented by* a (possibly empty valued) set function $L : \mathcal{A} \times \mathcal{P} \to 2^{\Theta}$ if $M(p) = a$ implies that there is an equilibrium with $P(\theta) = p$ for all $\theta \in L(a, p)$. We can think of the states in $L(a, p)$ as "payoff equivalent in equilibrium", given action a.

Notice that if the market admits an invariant representation given $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$ then there is an L that represents the level sets; define $L(a, p) = \{\theta : R(a, \theta) = p\}$. The point of identifying L is precisely to relax the assumption that there is an invariant representation given $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$.

Suppose M is a policy for which there exists a unique equilibrium. Let (Q_M, P_M) be the action and price functions in this equilibrium. Then it must be that if $P_M(\theta) = p$ then $P_M(\theta') = p$ for all $\theta' \in L(Q_M(\theta), p)$; if not then by definition of L there is equilibrium multiplicity in the states $L(Q_M(\theta), p) \setminus \{\theta' : P_M(\theta') = p\}$. Thus we can define an invariant representation R for the market by $R(a, \theta) = \{p : \theta \in L(a, p)$. This is a well-defined function iff there is a unique p such that $\theta \in L(a, p)\}$; in which case we say that the market has *unique level sets represented by* L. In Appendix B.4 we show that it is satisfied in the noisy REE model.

Given a market with unique level sets represented by L, say that a price function P has *complete level sets* if for all $p \in P(\Theta)$ there exists $a \in \mathcal{A}$ such that $\{\theta \in \Theta : P(\theta) = p\} = L(a, p)$. Let $\tilde{P}$ be the set of price functions with complete level sets. Let $\mathcal{W}^U$ be the set of policies for which there exists a unique market equilibrium. Then the preceding discussion is summarized by the following lemma.

Lemma 4. *If the market has unique level sets represented by* L*, then it admits an invariant representation given* $(\mathcal{W}^U, \tilde{\mathcal{P}})$*, defined by* $R(a, \theta) := \{p : \theta \in L(a, p)\}$.

Notice that in deriving the representation in Lemma 4, we are using the restrictions implied by the unique implementation requirement. Such a market need not admit an invariant representation given $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$.

Define $\bar{T}(a, \theta) = \{\theta' : R(a, \theta') = R(a, \theta)\}$. The following result is analogous to Theorem 2 in the uni-dimensional case.

Proposition 11. *Assume the market level sets uniquely represented by* L*, and so admits an invariant representation in* $(\mathcal{W}^U, \tilde{\mathcal{P}})$*. Assume moreover that the representation* R *is strictly increasing.*[42] *If* (Q, P) *is CUI, then*

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* $\theta' \in \bar{T}(Q(\theta), \theta)$.

v. $Q(\theta) \neq Q(\theta') \Rightarrow \bar{T}(Q(\theta), \theta) \cap \bar{T}(Q(\theta'), \theta') = \varnothing$.

Proof. In Appendix A.12. □

The conditions of Proposition 11 are also sufficient, except we require analogs to BC1 and BC2 for the multi-dimensional space. We omit the details for brevity. Since any market

[42] By this we mean strictly increasing in the usual product partial order on $\mathbb{R}^N$.

that admits an invariant representation in $(\mathcal{A}^{\mathcal{P}}, \mathcal{P}^{\Theta})$ has unique level sets represented by some L, Proposition 11 applies to such markets.

6 Properties and additional results

In Appendix D we discuss properties of CUI policies and study optimal policy when the unique implementation requirement is relaxed. Briefly, these results can be summarized as follows.

Structural uncertainty. In many settings there may be uncertainty regarding the relationship between the principal action, price, and the state. For example, there may be noise traders in the market who induce variability into the price.[43] Additionally, the principal may have limited data with which to estimate the representation R. It is therefore desirable to use a policy such that outcomes are suitably continuous with respect to small perturbations of R. We show that in fact any CUI action and price functions are implementable in a way that is *robust to structural uncertainty*, i.e. such that outcomes are suitably continuous in R (Theorem 3). (CWUI outcomes can be implemented so as to satisfy a weaker notion of robustness).

Beyond uniqueness. In some cases, the principal may be willing to tolerate the existence of multiple equilibria. The restriction to unique implementation would be especially unappealing if the principal could choose some policy which induces multiple equilibria, but such that all of these equilibria dominate, from the principal's perspective, the best equilibrium that can be implemented uniquely. We show that this is never the case. The key insight is that even if a policy induces multiple equilibria, at least one of these will be weakly uniquely implementable (Theorem 4). Thus for a principal who evaluates a set of possible equilibria according to the worst case, the restriction to (weak) unique implementation is without loss of optimality (Corollary 3). In other words, unique (i.e. full) implementation is without loss of optimality in a "robust mechanism design" sense (e.g. Carroll (2015)).[44]

[43] As shown in appendix B.4, an alternative way to deal with such noise is to fold it into the state.

[44] In Valenzuela-Stookey and Poggi (2022) (Proposition 9), we also study the case where the principal takes a less extreme approach to multiplicity, as in Dworczak and Pavan (2022).

A Omitted proofs

A.1 Proof of Proposition 1

Proof. For the first part, assume that Q has no discontinuities on the set $\{\theta \in \Theta : r_0(\theta) = r_1(\theta)\}$: if this does not hold then we can approximate Q arbitrarily well with a policy that has no such discontinuities (see Lemma 7 for an explicit construction). Under this assumption on Q, if we can find a convergent sequence of price functions P_n implemented by continuous policies M_n, such that the set $\{\theta \in \Theta : |P_n(\theta) - P(\theta)| < 1/n\}$ has measure less than $1/n$, then we can do the same for Q by defining $Q_n = M_n \circ P_n$. We now construct such a sequence.

Note that since $\theta \mapsto R(a, \theta)$ is strictly increasing for all a, if P is implementable, in particular if it satisfies measurability, then P^{-1} is a well-defined function on $P(\Theta)$ (and P^{-1} is discontinuous only if the M that induces it is discontinuous). For any $\varepsilon > 0$, let $\rho_\varepsilon : \mathcal{P} \mapsto \Theta$ be a continuous function which approximates P^{-1} on $P(\Theta)$; formally, we require

$$\rho_\varepsilon^{-1}(\theta) \cap (P(\theta) - \varepsilon, P(\theta) + \varepsilon) \cap R(\mathcal{A}, \theta) \neq \varnothing$$

for all $\theta \in \Theta$.[45] Such a function exists since P has countably many discontinuities by assumption, and so by measurability P^{-1} has countably many discontinuities as well.

Let $\hat{M}_\varepsilon(p) = \{a \in [0,1] : R(a, \rho_\varepsilon(p)) = p\}$. $\hat{M}_\varepsilon$ is single-valued and continuous on any open interval such that $r_0(\rho_\varepsilon(p)) \neq r_1(\rho_\varepsilon(p))$, by continuity of R and strict monotonicity of $a \mapsto R(a, \theta)$ at such θ. Moreover if $r_0(\rho_\varepsilon(p)) \neq r_1(\rho_\varepsilon(p))$ then $M(p) = [0, 1]$. Thus there exists a continuous selection from $\hat{M}_\varepsilon$, which we denote by M_ε. Now by definition of ρ_ε, for each $\theta \in \Theta$ we can define $P_\varepsilon(\theta) \in (P(\theta) - \varepsilon, P(\theta) + \varepsilon) \cap R(\mathcal{A}, \theta)$ such that $P_\varepsilon(\theta) \in \rho_\varepsilon^{-1}(\theta)$. By definition of M_ε, we have $R(M_\varepsilon \circ P_\varepsilon(\theta), \theta) = P_\varepsilon(\theta)$. Thus M_ε implements P_ε as desired.

For the second part, note that for any price $p \in P(\Theta)$ there is a unique $a \in [0, 1]$ and state θ such that $R(a, \theta) = p$. Thus, constructing M for those prices is straightforward. We need to determine what actions to assign to prices outside of $P(\Theta)$. For prices outside of $r_1(\Theta) \cap r_0(\Theta)$, we can find an action such that $R(a, \theta) \neq p$ for all $\theta \in \Theta$. For prices in $r_1(\Theta) \cap r_0(\Theta)$, this is not possible. Thus, we have to construct a set of states that has measure zero and are going to target with the policy function. Let A be a countable set of measure

[45] Where $\rho_\varepsilon^{-1}(\theta) := \{p \in \mathcal{P} : \rho_\varepsilon(p) = \theta\}$.

zero such that $\{p : R(a, \theta) = p \quad \text{for some } a \in [0,1] \text{ and } \theta \in A\}$. We can always construct this set with a step function $g : r_1(\Theta) \cap r_0(\Theta) \to \Theta$. For any $p \in r_1(\Theta) \cap r_0(\Theta)$ such that $p \notin P(\Theta)$, let $M(p) \in A$. Then, the set of states for which there is multiple equilibrium prices is a subset of A, and thus has measure zero. The M constructed implements (Q, P) and is weakly robust to multiplicity. □

A.2 Proof of Lemma 1

Proof. Let $\hat{p} \in P_M$. We want to show that there is a neighbourhood around $\hat{p}$ that is in P_M. By definition, there is a state $\hat{\theta}$ such that $\hat{p} = R(M(\hat{p}), \hat{\theta})$. Due to the strict monotonicity of R in its second argument, we have:

$$R(M(\hat{p}), \hat{\theta} - \epsilon) < \hat{p} < R(M(\hat{p}), \hat{\theta} + \epsilon) \tag{7}$$

for some small $\epsilon > 0$. Since M is essentially continuous, there is a neighbourhood around $\hat{p}$ within which M is continuous. By the continuity of M within this neighbourhood and the continuity of R in its first argument, the functions $R(M(p), \hat{\theta} - \epsilon)$ and $R(M(p), \hat{\theta} + \epsilon)$ are continuous in p within a neighbourhood around $\hat{p}$. As a result, there is a $\delta > 0$ such that the inequalities from Eq. 7 are maintained for any $p \in (\hat{p} - \delta, \hat{p} + \delta)$. Then by the Intermediate Value Theorem, there exists a $\theta \in (\hat{\theta} - \epsilon, \hat{\theta} + \epsilon)$ such that $p = R(M(p), \theta)$. □

A.3 Proof for second part of Proposition 2

Lemma 5. *If (Q, P) is CUI, then P is monotone and continuous.*

Proof. Let fix some state $\hat{\theta}$ and let S be the set of limit points of P at $\hat{\theta}$, i.e.

$$S = \{p \in \mathbb{R} : \text{ there exists a sequence } \{\theta_n\} \text{ s.t. } \theta_n \to \hat{\theta} \text{ and } P(\theta_n) \to p\}$$

Since the function P is bounded, the set S must be non-empty. Let $p \in S$ be an arbitrary limit point and $\{\theta_n\}$ be a sequence that converges to $\hat{\theta}$ and such that $P(\theta_n)$ converges to p.

By essential continuity, M is continuous at p. Moreover, Since $P(\theta_n) - R(M(P(\theta_n)), \theta_n) = 0$ for every n in the sequence, by taking limits we get that $p - R(\hat{\theta}, M(p)) = 0$. By robustness to multiplicity, we know that there is a unique price satisfying this equation, and therefore there is a unique limit point of P at $\hat{\theta}$. Therefore, the function P is continuous.

Moreover, P is a injection: let $\theta, \theta' \in \Theta$. If $P(\theta) = P(\theta')$ then $R(M(P(\theta)), \theta) = P(\theta) = P(\theta') = R(M(P(\theta')), \theta') = R(M(P(\theta)), \theta')$. By strict monotonicity of R in it's first argument, this implies that $\theta = \theta'$. Since P is a continuous injection, it must be strictly monotone. $\square$

A.4 Proof of Proposition 3

Proof. Let $P(\theta) := R(Q(\theta), \theta)$. First, assume that (P, Q) is virtually CWUI. That means that there exists a sequence $\{(P_n, Q_n)\} \to (Q, P)$ such that for all n, (P_n, Q_n) is CWUI. By Proposition 2, P_n is monotone. Since the set of monotone functions is closed, the limit $P^* := \lim_{n\to\infty} P_n$ is monotone. Since $P_n \to P$ almost everywhere, P must be essentially monotone.

For the other direction, we need to show that if P is essentially monotone then (Q, P) virtually CWUI. Without loss of generality assume that P is weakly monotone, since we can otherwise replace $P(\theta)$ with $\lim_{\theta' \nearrow \theta} P(\theta')$. We first show that if P is strictly monotone and satisfies some additional properties then (Q, P) is CWUI. We then show that if P is weakly monotone then (Q, P) can be approximated arbitrarily well by some (Q', P') which satisfy these additional properties, with P' strictly monotone. Together, this proves the result. Let $T := \{\theta \in \Theta : r_1(\theta) = r_0(\theta)\}$.

Lemma 6. *Let $\bar{r} := \min R(\mathcal{A}, \bar{\theta})$ and $\underline{r} := \max R(\mathcal{A}, \underline{\theta})$, and let $P(\theta) := R(Q(\theta), \theta)$.*

- *If P is strictly increasing and Q has no discontinuities on T, then (Q, P) is CWUI.*
- *If P is strictly decreasing, $P(\bar{\theta}) > \bar{r}$, and $P(\underline{\theta}) < \underline{r}$, then (Q, P) is CWUI.*

Proof of Lemma 6. First, notice that if $T \neq \varnothing$, then no strictly decreasing P is implementable, since both r_1 and r_0 are strictly increasing.

Let $D \subset \Theta$ be the set of states at which P is discontinuous. Since P is monotone, D is countable and has zero measure. We construct a correspondence $\hat{P} : \Theta \rightrightarrows \mathcal{P}$ which is equal to P on $\theta \setminus D$. For $\theta \in D$, if P is increasing let $\hat{P}(\theta) = [\lim_{\theta' \nearrow \theta} P(\theta), \lim_{\theta' \searrow \theta} P(\theta)]$, and if P is decreasing let $\hat{P}(\theta) = [\lim_{\theta' \searrow \theta} P(\theta), \lim_{\theta' \nearrow \theta} P(\theta)]$. Then $\hat{P}^{-1}(p) := \{\theta \in \Theta : \hat{P}(\theta) = p\}$ is a monotone function from $co(P(\Theta))$ to Θ (where $co(P(\Theta))$ is the convex hull of the range of P). Flat segments of $\hat{P}^{-1}$ correspond to discontinuities in P.

For each $p \in co(P(\Theta))$ define $\hat{M}(p) := \{a \in \mathcal{A} : R(a, \hat{P}^{-1}(p)) = p\}$. Note that $a \mapsto R(a, \theta)$

is strictly monotone for any θ such that $r_0(\theta) \neq r_1(\theta)$, so $\hat{M}(p)$ is single valued for any p such that $\hat{P}^{-1}(p) \notin T$. By continuity of R, we also have that $\hat{M}$ is continuous in an open neighborhood of any such p. Define $M(p) = \hat{M}(p)$ for all p such that $\hat{P}^{-1}(p) \notin T$. Note that if in addition $p \in P(\Theta)$ then $M(p) = Q(P^{-1}(\theta))$.

If $\hat{P}^{-1}(p) \in T$ then $\hat{M}(p) = [0, 1]$. Given our assumption that Q is continuous on an open neighborhood of any $\theta \in T$, such a p must be in $P(\Theta)$. Then let $M(p) = Q(P^{-1}(p))$. By assumption Q is continuous on an open neighborhood of any $\theta \in T$, so M is continuous on an open neighborhood of any such p.

In summary, we have defined a function M on $co(P(\Theta))$ that is continuous and satisfies $Q(\theta) = M(P(\theta))$ for all θ. Moreover, the only states for which there are multiple prices satisfying the RE condition are those in D, a zero-measure set.

It only remains to define M for $p \notin co(P(\Theta))$ so as to ensure that Q is implemented uniquely. If P is increasing then let $M(p) = Q(\bar{\theta})$ for $p > P(\bar{\theta})$ and $M(p) = Q(\underline{\theta})$ for $p < P(\underline{\theta})$. Since $\theta \mapsto R(a, \theta)$ is strictly increasing, this ensures that $p = R(M(p), \theta)$ cannot be satisfied for any $\theta \in \Theta$ and $p \notin co(P(\theta))$, as desired.

If P is decreasing then by assumption $P(\bar{\theta}) > \bar{r}$ and $P(\underline{\theta}) < \underline{r}$. We construct M for $p < P(\bar{\theta})$; the construction for $p > P(\underline{\theta})$ is symmetric. Let $\underline{r} = \min R(\mathcal{A}, \bar{\theta})$. For any $p \in [\bar{r}, P(\bar{\theta})]$ there is a unique $a \in \mathcal{A}$ such that $R(a, \bar{\theta}) = p$, which we denote by $\hat{M}(p)$. If $a \mapsto R(a, \bar{\theta})$ is increasing (resp. decreasing) then $a < \hat{M}(p)$ (resp. $a > M(p)$) implies that there is no $\theta \in \Theta$ such that $R(a, \theta) = p$. In either case, for some $\varepsilon > 0$ we can extend M continuously for $p \in [\bar{r} + \varepsilon, P(\bar{\theta})]$ such that there is no $\theta \in \Theta$ with $R(M(p), \theta) = p$. For $p < \bar{r} + \varepsilon$ we let $M(p)$ be the unique $a \in \mathcal{A}$ such that $R(a, \underline{\theta}) = P(\bar{\theta})$. This ensures that there is no $\theta \in \Theta$ such that $R(M(p), \theta) = p$. □

Lemma 7. *If $P(\theta) = R(Q(\theta), \theta)$ is weakly monotone then for any $\varepsilon > 0$ there exists Q' such that $|Q - Q'|_\infty < \varepsilon$, $R(Q'(\theta), \theta)$ is strictly monotone, and moreover*

- *If P is weakly increasing, Q' has no discontinuities on T.*
- *If P is weakly decreasing $P(\bar{\theta}) > \bar{r}$ and $P(\underline{\theta}) < \underline{r}$.*

Proof of Lemma 7. Suppose P is increasing and Q has a discontinuity at $\theta \in T$. For any $\varepsilon > 0$ small enough so that $[\theta - 2\varepsilon, \theta) \cap T = \varnothing$, let $Q'' = \lim_{\theta' \searrow \theta} Q(\theta)$ on $[\theta - \varepsilon, \theta]$, and let

$Q'' = Q$ otherwise. Then $P'' = R(Q''(\theta), \theta)$ is monotone on $[\underline{\theta}, \theta - \varepsilon]$ and on $[\theta - \varepsilon, \bar{\theta}]$, but may be discontinuous at $\theta - \varepsilon$, and thus non-monotone. However we can approximate P'' with a P' that is continuous and strictly increasing, and which coincides with P outside of $(\theta - 2\varepsilon, \bar{\theta})$. Since $(\theta - 2\varepsilon, \bar{\theta}) \cap T = \varnothing$, for each θ in this interval there is a unique $Q'(\theta)$ such that $R(Q'(\theta), \theta) = P'(\theta)$. Moreover Q' is continuous, as desired.

For the decreasing case, we simply construct P' to approximate P near $\underline{\theta}$ and $\bar{\theta}$, such that $P'(\bar{\theta}) > \bar{r}$ and $P'(\underline{\theta}) < \underline{r}$. Then we construct Q' as above, recalling the maintained assumption that $\underline{\theta}, \bar{\theta} \notin T$. □

□

A.5 Proof of Proposition 4 and Proposition 5

We begin with some preliminary observations.

Lemma 8.

i. The function g is continuous in the price p.

ii. The function $\bar{v}$ is continuous, continuously differentiable in the price, and has increasing differences in (p, θ).

iii. If $\tilde{v}$ is weakly concave in p then so is $\bar{v}$.

Proof. We show that H is continuous. Given this, the proof is identical to that of Lemmas 4.9, 4.10, and 4.11 in Toikka (2011), and is omitted.

Fix any θ and $p \in [\underline{p}(\theta), \bar{p}(\theta)]$. For any sequence (p_n, θ_n) such that $p_n \in [\underline{p}(\theta_n), \bar{p}(\theta_n)]$ converging to (p, θ),

$$\lim_{n\to\infty} H(p_n, \theta_n) = \lim_{n\to\infty} \int_0^1 \tilde{v}_1(p_n, r) \mathbb{1}_{\{\bar{p}^{-1}(p_n) \le r \le \theta_n\}} dr.$$

by definition of H. Note that $|\tilde{v}_1(p_n, r)\mathbb{1}_{\{\bar{p}^{-1}(p_n) \le r \le \theta\}}|$ is bounded, by continuity of $\tilde{v}_1$ and compactness of the domain. Moreover, since $\bar{p}^{-1}$ is continuous under the maintained assumptions, $\lim_{n\to\infty} \tilde{v}_1(p, r)\mathbb{1}_{\{\bar{p}^{-1}(p_n) \le r \le \theta\}} = \tilde{v}_1(p, r)\mathbb{1}_{\{\bar{p}^{-1}(p) \le r \le \theta\}}$ for almost all r. Then $\lim_{n\to\infty} H(p_n, \theta_n) = H(p, \theta)$ by Lebesgue's dominated convergence theorem. □

Given the properties of $\bar{v}$ established in Lemma 8, Berge's Maximum theorem implies that Φ is upper hemi-continuous, non-empty, and compact valued. Moreover since $\bar{v}$ has increasing differences and $\theta \mapsto [\underline{p}(\theta), \bar{p}(\theta)]$ is increasing in the strong set order, the selections P^* and P_* defined above exist and are non-decreasing.

For $P \in \mathcal{X}$ define

$$\Delta(P) = \int_0^1 (\tilde{v}(P(\theta), \theta) - \bar{v}(P(\theta), \theta)\, d\theta.$$

Lemma 9. $\sup_{P \in \mathcal{X}} \Delta(P) \leq 0$.

Proof. For any P in $\mathcal{X}$

$$\begin{aligned}
\Delta(P) &= \int_0^1 \tilde{v}(P(\theta), \theta) - \bar{v}(P(\theta), \theta) d\theta \\
&= \int_0^1 \left[\tilde{v}(\underline{p}(\theta), \theta) + \int_{\underline{p}(\theta)}^{P(\theta)} \tilde{v}_1(s, \theta) ds - \left(\tilde{v}(\underline{p}(\theta), \theta) + \int_{\underline{p}(\theta)}^{P(\theta)} g(s, \theta) ds \right) \right] d\theta \\
&= \int_0^1 \int_{\underline{p}(\theta)}^{P(\theta)} \big(\tilde{v}_1(s, \theta) - g(s, \theta) \big) ds d\theta \\
&= \int_{\underline{p}(0)}^{\bar{p}(1)} \int_{P^{-1}(s)}^{\rho(s)} \big(\tilde{v}_1(s, \theta) - g(s, \theta) \big) d\theta ds \\
&= \int_{\underline{p}(0)}^{\bar{p}(1)} \left(\int_{\bar{p}^{-1}(s)}^{\rho(s)} \big(\tilde{v}_1(s, \theta) - g(s, \theta) \big) d\theta - \int_{\bar{p}^{-1}(s)}^{P^{-1}(s)} \big(\tilde{v}_1(s, \theta) - g(s, \theta) \big) d\theta \right) ds \\
&= \int_{\underline{p}(0)}^{\bar{p}(1)} \Big(\big(H(s, \rho(s)) - G(s, \rho(s)) \big) - \big(H(s, P^{-1}(s)) - G(s, P^{-1}(s)) \big) \Big) ds \\
&= \int_{\underline{p}(0)}^{\bar{p}(1)} \big(G(s, P^{-1}(s)) - H(s, P^{-1}(s)) \big) ds \\
&\leq 0
\end{aligned}$$

where the equality in the third-to-last line follows from the definition of H and G and the fact that $H(s, \bar{p}^{-1}(s)) = G(s, \bar{p}^{-1}(s))$ (by definition of the convex hull); the equality in the second-to-last line follows from $H(s, \rho(s)) = G(s, \rho(s))$; and the final inequality follows by definition of the convex hull. $\square$

Lemma 10. *If $\tilde{v}(p, \theta)$ is weakly concave in p then $\Delta(P^*) = \Delta(P_*) = 0$.*

Proof. For any P in $\mathcal{X}$ define

$$A_P := \{x \in [\underline{p}(0), \bar{p}(1)] : G(x, P^{-1}(x)) < H(x, P^{-1}(x))\}.$$

As shown above, $\Delta(P) = \int_{\underline{p}(0)}^{\bar{p}(1)} \big(G(s, P^{-1}(s)) - H(s, P^{-1}(s)) \big) ds$, so it suffices to show that A_{P^*} and A_{P_*} have countably many elements.

First, let P be any monotone selection from Φ, let $x \in A_P$, and let $\theta_x := P^{-1}(x)$. Then $\theta_x \in (\bar{p}^{-1}(x), \rho(x))$, since by definition of the convex hull $G(x, \bar{p}^{-1}(x)) = H(x, \bar{p}^{-1}(x))$ and $G(x, \rho(x)) = H(x, \rho(x))$ for all x.

Claim. Let P be a monotone selection from Φ and let $x \in A_P$. If $\underline{p}(P^{-1}(x)) < x < \bar{p}(P^{-1}(x))$ then there exists an open neighborhood U of $P^{-1}(x)$ such that $x \in \Phi(\theta)$ for all $\theta \in U$.

Proof of Claim. Since P is monotone $\lim_{\theta \nearrow \theta_x} P(\theta) \leq P(\theta_x) \leq \lim_{\theta \searrow \theta_x} P(\theta)$. Then upper-hemicontinuity of Φ implies $\lim_{\theta \nearrow \theta_x} P(\theta), \lim_{\theta \searrow \theta_x} P(\theta) \in \Phi(\theta_x)$. Since $\Phi(\theta_x)$ is convex by Lemma 8, we have $x \in \Phi(\theta_x)$.

By definition of A_P we have $G(x, \theta_x) < H(x, \theta_x)$, and since both functions are continuous there is an open neighborhood U' of θ_x such that $G(x, \theta) < H(x, \theta)$ for $\theta \in U'$. Then $g(x, \cdot)$ is constant on U'.

If $x \in (\underline{p}(\theta_x), \bar{p}(\theta_x))$ then $g(x, \theta_x) = 0$ since $\bar{v}$ is concave and continuously differentiable in p by Lemma 8. Then $g(x, \theta) = 0$ for all $\theta \in U'$. Moreover, since $\underline{p}$ and $\bar{p}$ are continuous, there exists an open neighborhood U of θ_x, $U \subset U'$, such that $x \in (\underline{p}(\theta), \bar{p}(\theta))$ for all $\theta \in U$. Thus $x \in \Phi(\theta)$ for all $\theta \in U$. $\square$

We can now complete the proof of Lemma 10. Let $x \in A_{P_*}$. Suppose $x = \bar{p}(P_*^{-1}(x))$. Since $P_*^{-1}(x) > \bar{p}^{-1}(x)$, this implies that x is a point of discontinuity of $\bar{p}^{-1}$. Since $\bar{p}^{-1}$ is monotone there can be at most countably many such points. By a symmetric argument, there can be at most countably many $x \in A_{P_*}$ such that $x = \underline{p}(P_*^{-1}(x))$. Finally, if $\underline{p}(P_*^{-1}(x)) < x < \bar{p}(P_*^{-1}(x))$ then by the above claim there exists an open neighborhood U of $P_*^{-1}(x)$ such that $x \in \Phi(\theta)$ for all $\theta \in U$. Then $P_*(\theta) = \min \Phi(\theta) \leq x$ for all $\theta \in U$. Moreover by definition of the inverse, $P_*(\theta) \geq x$ for all $\theta > \theta_x$. Thus $P_*(\theta) = x$ for all $\theta \in U$ such that $\theta \geq P_*^{-1}(x)$, so x is a point of discontinuity of P_*^{-1}. Since P_*^{-1} is monotone there can be at most countable many such points.

Consider now $x \in A_{P^*}$. Let $\theta_x := (P^*)^{-1}(x)$. The cases for $x = \bar{p}(\theta_x)$ and $x = \underline{p}(\theta_x)$ are as above. If $\underline{p}(\theta_x) < x < \bar{p}(\theta_x)$ then by the above claim there exists an open neighborhood U of θ_x such that $x \in \Phi(\theta)$ for all $\theta \in U$. Then $P^*(\theta) = \max \Phi(\theta) \geq x$ for all $\theta \in U$. But the definition of the inverse implies $P^*(\theta) < x$ for all $\theta < \theta_x$, which is a contradiction. $\square$

Proof of Proposition 4.

$$\begin{aligned}\sup_{P\in\Gamma}\int_0^1 \bar{v}(P(\theta),\theta)d\theta &\geq \sup_{P\in\mathcal{X}}\int_0^1 \bar{v}(P(\theta),\theta)d\theta\\ &\geq \sup_{P\in\mathcal{X}}\int_0^1 \bar{v}(P(\theta),\theta)d\theta + \Delta(P)\\ &= \sup_{P\in\mathcal{X}}\int_0^1 \tilde{v}(P(\theta),\theta)d\theta\end{aligned}$$

where the third line follows from Lemma 9 and the second from the definition of Δ. Conversely

$$\begin{aligned}\sup_{P\in\mathcal{X}}\int_0^1 \tilde{v}(P(\theta),\theta)d\theta &\geq \int_0^1 \tilde{v}(P^*(\theta),\theta)d\theta\\ &= \int_0^1 \bar{v}(P^*(\theta),\theta)d\theta\\ &= \sup_{P\in\Gamma}\int_0^1 \tilde{v}(P(\theta),\theta)d\theta\end{aligned}$$

where the second line follows from Lemma 10 and the last line from the definition of P^*. $\square$

Proof of Proposition 5.

That P^* obtains the maximum is immediate from the proof of Proposition 4, and similarly for P_*. If P obtains the maximum in eq. (6) then

$$\int_0^1 \bar{v}(P(\theta),\theta)d\theta \geq \int_0^1 \tilde{v}(P(\theta),\theta)d\theta = \sup_{y\in\Gamma}\int_0^1 \tilde{v}(y(\theta),\theta)d\theta$$

where the inequality follows from Lemma 9. Thus $P(\theta)\in\Phi(\theta)$, which implies $P_*(\theta)\leq P(\theta)\leq P^*(\theta)$, for almost all θ. $\square$

A.6 Proof of Proposition 6

Proof. Using the definitions above $H(x,\theta) = c'(x)\int_{\bar{p}^{-1}(x)}^{\theta} h(r)dr + (\theta-\bar{p}^{-1}(x))b'(x)$. From the definition of the convex hull we have

$$\begin{aligned}convH(x,\theta) = (\theta-\bar{p}^{-1}(x))b'(x) + c'(x)\min\Big\{\lambda\int_{\bar{p}^{-1}(x)}^{\theta_1} h(r)dr + (1-\lambda)\int_{\bar{p}^{-1}(x)}^{\theta_2} h(r)dr\\ s.t.\quad &\theta_1,\theta_2\in[\bar{p}^{-1}(x),\rho(x)]\\ &\lambda\in[0,1],\\ &\theta = \lambda\theta_1+(1-\lambda)\theta_2\Big\}.\end{aligned}$$

Then $convH(x,\theta) = c'(x)T(x,\theta)+(\theta-\bar{p}^{-1}(x))b'(x)$, by the definition of D and T. The result follows. □

A.7 Proof of Proposition 7

Proof. First, observe that D is strictly increasing below $\hat{\theta}$ and strictly decreasing above $\hat{\theta}$. This function is depicted as the solid line in Figure 5. For any price x, the function $T(x,\cdot)$ is defined as the convex hull of the restriction of D to $[r_1^{-1}(x), r_0^{-1}(x)]$.

Claim. If $t(x'',\theta) \geq 0$ for some θ then $t(x',\theta) > 0$ for all $x' < x''$ and all θ in the domain of $t(x',\cdot)$, i.e. $[r_1^{-1}(x'), r_0^{-1}(x')]$. Similarly, if $t(x',\theta) \leq 0$ for some θ then $t(x'',\theta) < 0$ for all $x' < x''$ and all θ in the domain of $t(x'',\cdot)$.

Proof of Claim. I prove the first part of the claim, for $t(x'',\theta) \geq 0$, the other case is symmetric.

First, note that if $r_0^{-1}(x) \leq \hat{\theta}$ then D is strictly increasing on $[r_1^{-1}(x), r_0^{-1}(x)]$, so $t(x,\theta)$ is strictly positive for all θ for which it is defined. Thus the claim is immediate if $r_0^{-1}(x'') \leq \hat{\theta}$

If $r_0^{-1}(x'') > \hat{\theta}$ then there exists θ such that $t(x'',\theta) \geq (>)\ 0$ if and only if $D(r_0^{-1}(x'')) \geq (>)\ D(r_1^{-1}(x''))$, in which case $t(x'',\theta) \geq (>)\ 0$ for all θ in its domain. This follows from the fact that D is quasi-concave. Note that $D(r_0^{-1}(x'')) \geq D(r_1^{-1}(x''))$ can hold only if $r_0^{-1}(x'') < \hat{\theta}$. Then for $x' < x''$ we have $D(r_0^{-1}(x')) > D(r_0^{-1}(x''))$ and $D(r_1^{-1}(x')) < D(r_1^{-1}(x''))$, since r_0^{-1}, r_1^{-1} are strictly increasing and D is strictly increasing (decreasing) below (above) $\hat{\theta}$. Thus $D(r_0^{-1}(x')) > D(r_1^{-1}(x'))$, so $t(x',\theta) \geq (>)0$ for all θ in its domain. □

Suppose $D\big(r_0^{-1}(r_1(0))\big) < 0$. Then the convex hull of the restriction of D to $[0, r_0^{-1}(r_1(0))]$ is downward slopping, so $t(r_1(0),\theta) \leq 0$ for all $\theta \in [0, r_0^{-1}(r_1(0))]$. Since $D(\hat{\theta}) > 0$ it must be that $r_0^{-1}(r_1(0)) > \hat{\theta}$. Since $D(0) = 0$ and $D(\hat{\theta}) > 0$, the convex hull of the restriction of D to $[0,\hat{\theta}]$ is strictly increasing. Since $x \mapsto r_0^{-1}(x)$ is strictly increasing, there exists some $x^* \in (r_0(\hat{\theta}), r_1(0))$ such that $D(r_0^{-1}(x^*)) = 0$, by the intermediate value theorem. Then $t(x^*,\theta) = 0$ for all $\theta \in [0, r_0^{-1}(x^*)]$. Moreover $x > x^*$ implies $t(x,\theta) < 0$ for all θ, and $x < x^*$ implies $t(x,\theta) > 0$ for all θ. The unique solution to

$$\Phi(\theta) = \max\left\{\int_{r_0(\theta)}^{x} t(s,\theta)ds \ :\ x \in [r_0(\theta), r_1(\theta)]\right\}$$

is therefore to set $x = x^*$ for $\theta \in [0, r_0^{-1}(x^*)]$, and $x = r_0(\theta)$ for $\theta > r_0^{-1}(x^*)$.

Suppose $D\big(r_0^{-1}(r_1(0))\big) \geq 0$. Then $T(r_0(1), \cdot)$ is strictly increasing, so $t(r_0(1), \theta) > 0$ for all θ in its domain. Then by the above claim, for all θ and all $x < r_1(0)$ we have $t(x, \theta) > 0$. Thus any element of $\Phi(\theta)$ must be greater than $r_1(0)$ for all θ.

Now suppose $D(1) > D\big(r_1^{-1}(r_0(1))\big)$. Then be the same argument, any element of $\phi(\theta)$ must be greater than $r_0(1)$, so an optimal P cannot be responsive in the middle. To show that it must be responsive at the top, note that there exists $r_0(1) < x < r_1(1)$ such that $r_1^{-1}(x) \geq \hat{\theta}$. Then $D(r_1^{-1}(x)) > D(r_0^{-1}(x)) = D(1)$. Then there exists $r_0(1) < x^* < x$ such that $D(r_1^{-1}(x^*)) = D(r_0^{-1}(x^*)) = D(1)$ by the intermediate value theorem (and this x^* is unique given the shape of D and strict monotonicity of r_1^{-1}, r_0^{-1}). Then $\phi(\theta) = x^*$ for all $\theta \in [r_1^{-1}(x^*), 1]$, and $\phi(\theta) = 1$ for all $\theta < r_1^{-1}(x^*)$.

If instead $D(1) \leq D\big(r_1^{-1}(r_0(1))\big)$ then it follows by a similar argument that there exists $r_1(0) < x^* < r_0(1)$ such that x^* is the unique solution in $\phi(\theta)$ for all $\theta \in [r_1^{-1}(x^*), r_0^{-1}(x^*)]$. Moreover by the above claim $t(x, \theta) < 0$ for all $x > x^*$, so the unique solution for $\theta > r_0^{-1}(x^*)$ is to set $P(\theta) = r_0(\theta)$.

Finally, if there is a gap and $D(0) = D(1)$ then $t(x, \theta) = 0$ for all $x \in [r_0(1), r_1(0)]$ and all θ; and if $x > r_1(0)$ ($x < r_0(1)$) then $t(x, \theta) < 0$ ($t(x, \theta) > 0$) for all θ. $\square$

A.8 Proof of Theorem 1

We begin with some preliminary results. Use $\theta_M(p)$ to indicate the states that are consistent with a price p, given a policy function M, i.e., the states for which there exists a P implementing $\theta_M := \{\theta \in \Theta : M \text{ implements } P \text{ and } P(\theta) = p\}$.

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

Lemma 11. *If R is weakly increasing in θ then $\theta_M(p)$ is convex valued.*

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

Lemma 12. *For any M, each p such that $\theta_M(p) = \varnothing$ is of one and only one of the following two types:*

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

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. $\square$

Lemma 13. *(Generalized intermediate value theorem). Let $F : [0,1] \to [0,1]$ be a non-empty, 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 and $y_2 < y_1$ 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 neighborhood of p^* such that $F(p) \subset V$ for all p in such neighborhood. Thus, in a neighbourhood of p^*, $F(p) < \tilde{y}$, which 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 neighbor of p^* such that $F(p) > \tilde{y}$ for all p in that neighborhood, what violates the definition of p^*. $\square$

Lemma 14. *For any M, $p \mapsto \theta_M(p)$ is compact-valued. If M is continuous at p' and $\theta_M(p')$ is not empty for every p in a neighborhood of p' then $\theta_M(\cdot)$ is upper hemicontinuous at p'.*

Proof. Compact valuedness 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 $\bar{\Theta}$ be the closure of Θ, and let $\bar{\theta}_M(p) := \{\theta \in \bar{\Theta} : R(M(p), \theta) = p\}$ (where for $\theta \in \bar{\Theta} \setminus \Theta$ we define $R(a, \theta)$ as the limit as $\theta' \to \theta$ for $\theta' \in \Theta$). Let V be an open set containing $\bar{\theta}_M(p)$. Since $\bar{\theta}_M(p)$ is compact Then $\Theta \setminus V$ is compact, so there exists $\kappa > 0$ such that $|R(M(p), \theta) - p| > \kappa$ for all $\theta \in \bar{\Theta} \setminus V$. Then by continuity of R, M there exists an open neighborhood U of p such that $|R(M(p'), \theta) - p'| \geq \kappa$ for all $p' \in U, \theta \in \bar{\Theta} \setminus V$. Thus $\theta_M(p') \subseteq V$. Thus $p \mapsto \bar{\theta}_M(p)$ is upper hemicontinuous at p'.

Now since $\theta_M(p')$ is non empty, if $\bar{\theta} \in \bar{\theta}_M(p')$ then there is some $\theta_1 < \bar{\theta}$ such that $[\theta_1, \bar{\theta}) \subset \theta_M(p')$. Similarly if $\underline{\theta} \in \bar{\theta}_M(p')$. Then any open set containing $\theta_M(p')$ also contains $\bar{\theta}_M(p')$. Thus θ_M is upper hemicontinuous at p' as well. □

A.9 Proof of Theorem 1

Proof. Let (Q, P) be CWUI and M a policy that implements (Q, P). Let $P_M(\theta) := \{p \in \mathbb{R} : R(M(p), \theta) = p\}$. For any $\theta \in \Theta$, Let $a_\theta := \sup\{p \leq P(\theta) : \theta_M(p) = \emptyset\}$ and $b_\theta := \inf\{p \geq P(\theta) : \theta_M(p) = \emptyset\}$.

Lemma 15. *For any θ^*, if $a_{\theta^*} = b_{\theta^*} = P(\theta^*)$ then either $R(M(P(\theta^*)), \theta) = P(\theta^*)$ for all $\theta \leq \theta^*$, or $R(M(P(\theta^*)), \theta) = P(\theta^*)$ for all $\theta \geq \theta^*$ (or both).*

Proof. Since M is essentially continuous, it is continuous in a neighborhood of $P(\theta^*)$. Let $(\underline{p}, \bar{p})$ be such a neighborhood.

Claim 1. Either (1.1) For every $p \in (P(\theta^*), \bar{p})$ there exists $p \in (P(\theta^*), p)$ such that $R(M(p'), \theta) > p'$ for all $\theta \in \Theta$, and/or (1.2) For every $p \in (P(\theta^*), \bar{p})$ there exists $p' \in (P(\theta^*), p)$ such that $R(M(p'), \theta) < p'$ for all $\theta \in \Theta$. Similarly, either (1.3) For every $p \in (\underline{p}, P(\theta^*))$ there exists $p' \in (p, P(\theta^*))$ such that $R(M(p'), \theta) > p'$ for all $\theta \in \Theta$, and/or (1.4) For every $p \in (\underline{p}, P(\theta^*))$ there exists $p' \in (p, P(\theta^*))$ such that $R(M(p'), \theta) < p'$ for all $\theta \in \Theta$.

Claims (1.1)-(1.4) follow from Lemma 12 and the assumption that $a_{\theta^*} := \sup\{p \leq P(\theta) : \theta_M(p) = \emptyset\} = P(\theta^*) = b_{\theta^*} := \inf\{p \geq P(\theta) : \theta_M(p) = \emptyset\}$.

Claim 2. $R(M(P(\theta^*)), \theta) = P(\theta^*)$ for all $\theta \leq \theta^*$ if either (1.1) or (1.3) hold. $R(M(P(\theta^*)), \theta) = P(\theta^*)$ for all $\theta \geq \theta^*$ if either (1.2) or (1.4) hold.

Suppose (1.1) holds. Then we can find a sequence $\{p_n\}$ with $p_n < P(\theta^*)$ and $p_n \to P(\theta^*)$ such that $R(M(p_n), \theta) > p_n$ for all θ. By continuity, $R(M(P(\theta^*)), \theta) \geq P(\theta^*)$, and since $\theta \mapsto R(M(P(\theta^*)), \theta)$ is non-decreasing, we have $R(M(P(\theta^*)), \theta) = P(\theta^*)$ for all $\theta \leq \theta^*$.

By a symmetric argument, if (1.3) holds, then $R(M(P(\theta^*)), \theta) \geq P(\theta^*)$, and since $\theta \mapsto R(M(P(\theta^*)), \theta)$ is non-decreasing and $R(M(P(\theta^*)), \theta^*) = P(\theta^*)$, we have $R(M(P(\theta^*)), \theta) = P(\theta^*)$ for all $\theta \leq \theta^*$. A symmetric argument applies to (1.2) and (1.4). Combined, Claims 1 and 2 complete the proof. □

Lemma 16. *For any θ^*, if $a_{\theta^*} < b_{\theta^*}$ then $P_M(\theta) \cap (a_{\theta^*}, b_{\theta^*}) \neq \emptyset$ for all $\theta \in \Theta$.*

Proof. Note that, by definition of a_θ and b_{θ^*}, $(a_{\theta^*}, b_{\theta^*}) \subseteq P_M(\Theta)$. By definition of essential continuity, M is continuous in $(a_{\theta^*}, b_{\theta^*})$. First, we prove that either (i) $R(M(a_{\theta^*}), \underline{\theta}) = a_{\theta^*}$ and $R(M(b_{\theta^*}), \bar{\theta}) = b_{\theta^*}$ or (ii) $R(M(a_{\theta^*}), \bar{\theta}) = a_{\theta^*}$ and $R(M(b_{\theta^*}), \underline{\theta}) = b_{\theta^*}$.

Consider a sequence of prices p_n such that $p_n \in (a_{\theta^*}, b_{\theta^*})$ and p_n converges to a_{θ^*}. For every n , since $p_n \in (a_{\theta^*}, b_{\theta^*}) \subset P_M(\Theta)$, there exists $\theta_n \in \Theta$ such that $R(M(p_n), \theta_n) = p_n$. Thus, for each n, $R(M(p_n), \underline{\theta}) - p_n \leq R(M(p_n), \theta_n) - p_n = 0$ where the first inequality holds by monotonicity of R.

Since M is continuous on a neighbourhood of a_{θ^*}, taking limits side-by-side, we get that $R(M(a_{\theta^*}), \underline{\theta}) \leq a_{\theta^*}$. Likewise, we can prove that $a_{\theta^*} \leq R(M(\bar{a}_{\theta^*}), \theta)$.

It remains to show that a_{θ^*} is not in the interior. Suppose by contradiction that it is. Then $0 \in (R(M(a_{\theta^*}), \underline{\theta}) - a_{\theta^*}, R(M(a_{\theta^*}), \bar{\theta}) - a_{\theta^*}))$. For an $\epsilon > 0$ small enough, for every $p \in (a_{\theta^*} - \epsilon, a_{\theta^*})$, $0 \in (R(M(p), \underline{\theta}) - p, R(M(p), \bar{\theta}) - p))$. Thus, applying the intermediate value theorem, $(a_{\theta^*} - \epsilon, a_{\theta^*}] \subset P_M(\Theta)$, what violates the definition of a_{θ^*}. The same argument can be applied to prove that $b_{\theta^*} \in \{R(M(b_{\theta^*}), \underline{\theta}), R(M(b_{\theta^*}), \bar{\theta})\}$.

Next we show that it cannot be the case that $R(M(a_{\theta^*}), \underline{\theta}) = a_{\theta^*}$ and $R(M(b_{\theta^*}), \underline{\theta}) = b_{\theta^*}$. If that is the case, since continuity of M in $(a_{\theta^*}, b_{\theta^*})$ implies that $\theta_M(p)$ is upper hemicontinuous in $(a_{\theta^*}, b_{\theta^*})$, $P_M(\tilde{\theta})$ is not a singleton for a small $\delta > 0$, $\tilde{\theta} \in [\underline{\theta}, \underline{\theta} + \delta)$. A symmetric argument rules out the case in which $R(M(a_{\theta^*}), \bar{\theta}) = a_{\theta^*}$ and $R(M(b_{\theta^*}), \bar{\theta}) = b_{\theta^*}$.

We finish the proof by showing that $P_M(\theta') \cap (a_{\theta^*}, b_{\theta^*}) \neq \emptyset$ for all $\theta' \in \Theta$. Suppose case (i) holds, i.e. that $\underline{\theta} \in \theta_M(a_{\theta^*})$ and $\bar{\theta} \in \theta_M(b_{\theta^*})$. By the intermediate value theorem in Lemma 13, for every $\theta' \in (\underline{\theta}, \bar{\theta})$ there exists a $p \in (a_{\theta^*}, b_{\theta^*})$ such that $\theta \in \theta_M(p)$. Thus, $p \in P_M(\theta') \cap (a_{\theta^*}, b_{\theta^*})$. If case (ii) holds, a symmetric argument proves the claim. $\square$

Suppose, towards a contradiction, that (Q, P) is implemented by $M \in \mathcal{W}^*$ and P is not monotone. Assume in particular that there is $\theta_1 < \theta_2 < \theta_3$ such that $P(\theta_1) > P(\theta_3) > P(\theta_2)$. (The other cases of non-monotonicity are symmetric.)

Claim 1. Either (1.1) $a_{\theta_1} = a_{\theta_2} = a_{\theta_3}$ and $b_{\theta_1} = b_{\theta_2} = b_{\theta_3}$, or (1.2) $a_{\theta_1} = b_{\theta_1}$, $a_{\theta_2} = b_{\theta_2}$, and $a_{\theta_3} = b_{\theta_3}$.

To shop Claim 1, either there exists $i \in \{1, 2, 3\}$ such that $a_{\theta_i} < b_{\theta_i}$, or not. Suppose such an i does exist and there exists $j \neq i$ such that $a_{\theta_i} \neq a_{\theta_j}$ (which also implies $b_{\theta_i} \neq b_{\theta_j}$). Then by Lemma 16, for any $\theta \in \Theta$ we have $P_M(\theta) \cap (a_{\theta_i}, b_{\theta_i}) \neq \varnothing$. If $a_{\theta_j} < b_{\theta_j}$ then by Lemma 16 for any $\theta \in \Theta$ we have $P_M(\theta) \cap (a_{\theta_j}, b_{\theta_j}) \neq \varnothing$, so there is multiplicity in all states. If $a_{\theta_j} = b_{\theta_j}$ then by Lemma 15 either there is multiplicity above θ_j or below θ_j.

Claim 2. If (1.1) holds, then we have a contradiction.

M is continuous in $(P(\theta_1), P(\theta_2))$. Thus we can apply Lemma 13: for every $\theta \in (\theta_1, \theta_3)$ there exists a price $p_\theta \in (P(\theta_3), P(\theta_1))$ such that $p_\theta \in P_M(\theta)$. Likewise, for every $\theta \in (\theta_2, \theta_3)$ there exists a price in $p'_\theta \in (P(\theta_2), P(\theta_3))$ such that $p'_\theta \in P_M(\theta)$. Thus, for all states in $\theta \in (\theta_2, \theta_3)$, $p'_\theta \neq p_\theta$ and therefore the set $P_M(\theta)$ has more than one element, what implies a violation of robustness to multiplicity.

Claim 3. If (1.2) holds, then we have a contradiction.

In this case, Lemma 15 implies that for each $i \in \{1, 2, 3\}$, either $P(\theta_i) \in P_M(\theta)$ for all $\theta \leq \theta_i$, or for all $\theta \geq \theta_i$. Any combination of these conditions for $i \in \{1, 2, 3\}$ implies multiplicity on a positive measure set. $\square$

A.10 Proof of Theorem 2

Lemma 17. *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 12 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 14 (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 13 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 for p_1. $\square$

Proof. ($\Rightarrow$) Part 1 follows trivially from the rational expectations condition.

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 $R(Q(\theta), \theta) = R(Q(\theta), \theta')$ which, 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 neighborhood N of $P(\theta')$. This, together with continuity of R, implies that $\theta_M(p)$ is continuous on N. Thus, there is a neighborhood of θ' such that $P(\theta) \cap N$ is not empty for all θ in the neighborhood. 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 an interior state will necessarily imply a discontinuity of M for a price in $\bar{P}$, which would violate essential continuity. Thus, Q must be continuous for all interior states.

P is monotone and bounded (below by $\min_{a \in \mathcal{A}} R(a, \underline{\theta})$ and above by $\max_{a \in \mathcal{A}} R(a, \bar{\theta})$), 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, $\lim_{\theta \searrow \underline{\theta}} Q(\theta) = \lim_{\theta \searrow \underline{\theta}} Q(\theta) = Q(M(\underline{P}))$. The same arguments hold 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 neighborhood. By essential continuity, for prices slightly above p the action is in such neighborhood. 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 differs 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 $\bar{P}$, $p \geq \bar{P} = R(\bar{Q}, \bar{\theta}) = R(M(\bar{P}), \bar{\theta}) > R(M(p), \theta)$ 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.[46]

Now for decreasing P, we need to show that there exists a continuous policy for prices right above $\underline{P}$ so that $\theta_M(p)$ is empty. Let $\underline{R}(a) = R(a, \underline{\theta})$ and 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 $\tilde{P} \in (\underline{P}, \underline{P} + \epsilon)$, $R(M(\tilde{P}), \theta) > Q(M(\tilde{P}), \underline{\theta}) > \tilde{P}$.

Use a symmetric construction for M below $\bar{P}$. Beyond these prices, at the neighborhood of $P(\Theta)$, essential continuity is not binding, so any actions that do not generate equilibria work for the construction. By 1, if for a price all actions generate an equilibrium, then that price must be in $P(\Theta)$. $\square$

A.11 Proof of Proposition 10

Proof. Necessity for condition 1 is immediate, and condition 2 follows from Theorem 1. The argument for condition 4 is the same as in Theorem 2.

It remains to show that Q can have a discontinuity at θ^* iff P has a discontinuity at θ^*. To see this, notice that Q can be discontinuous at θ^* iff P is continuous at θ^* (otherwise M would

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

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 $(\min\{\lim_{\theta \nearrow \theta^*} P(\theta), \lim_{\theta \searrow \theta^*} P(\theta)\}, \max\{\lim_{\theta \nearrow \theta^*} P(\theta), \lim_{\theta \searrow \theta^*} P(\theta)\})$ (otherwise there is multiplicity for all states in a neighborhood of θ^*). Such a P can be implemented by an essentially continuous M iff there exists γ satisfying the definition of bridgeability (in which case we take $M = \gamma$ on this interval). Once M is defined over these intervals corresponding to the discontinuities in P, the argument for sufficiency of conditions (1)-(4) is the same as in Theorem 2. □

A.12 Proof of Proposition 11

Proof. Note that $\bar{T}(a, \theta) = L(a, R(a, \theta))$. Then, as outlined in the discussion preceding Proposition 11, conditions *iv* and *v* are necessary given the definition of L, since otherwise there would be multiplicity. *i* is obviously necessary. To show necessity of *ii* 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 monotonicity and continuity for interior states follow from the same arguments as in the uni-dimensional case. The key step is that under *iii* and *iv*, this implies that *ii* 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. □

References

M. Akbarpour, P. Dworczak, and S. D. Kominers. Redistributive allocation mechanisms. *Journal of Political Economy*, 132(6):1831–1875, 2024.

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

S. Athey. Monotone comparative statics under uncertainty. *The Quarterly Journal of Economics*, 117(1):187–223, 2002.

E. M. Azevedo and E. Budish. Strategy-proofness in the large. *The Review of Economic Studies*, 86(1):81–116, 2019.

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.

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.

E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

G. Carroll. Robustness and linear contracts. *AER*, 105(2):536–63, 2015.

A. Dovis and R. Kirpalani. Unique implementation with market-based interventions. *unpublished*, 2023.

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. *Econometrica*, 90(5):2017–2051, 2022.

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.

C. Goodhart. Problems of monetary management: the uk experience. *Papers in Monetary Economics*, 1974.

S. J. Grossman and J. E. Stiglitz. Information and competitive price systems. *The American Economic Review*, 66(2):246–253, 1976.

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. *AER*, 35(4):519–530, 1945.

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

M. F. Hellwig. A maximum principle for control problems with monotonicity constraints. *MPI Collective Goods Preprint*, 2008.

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

B. Jullien. Participation constraints in adverse selection models. *Journal of Economic Theory*, 93(1):1–47, 2000.

A. Kajii and S. Morris. The robustness of equilibria to incomplete information. *Econometrica: Journal of the Econometric Society*, pages 1283–1309, 1997.

Z. Y. Kang. Optimal indirect regulation of externalities. *Available at SSRN 3586050*, 2024.

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

A. Kleiner, B. Moldovanu, and P. Strack. Extreme points and majorization: Economic applications. *Econometrica*, 89(4):1557–1593, 2021.

I. Kremer. Information aggregation in common value auctions. *Econometrica*, 70(4):1675–1682, 2002.

A. S. Kyle. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335, 1985.

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

S. Loertscher and E. V. Muir. Monopoly pricing, optimal randomization, and resale. *Journal of Political Economy*, 130(3):566–635, 2022.

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

E. Maskin. Nash equilibrium and welfare optimality. *ReStud*, 66(1):23–38, 1999.

D. Mookherjee and M. Tsumagari. Mechanism design with communication constraints. *Journal of Political Economy*, 122(5):1094–1129, 2014.

M. Mussa and S. Rosen. Monopoly and product quality. *Journal of Economic theory*, 18(2): 301–317, 1978.

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

G. Nöldeke and L. Samuelson. Optimal bunching without optimal control. *Available at SSRN 785804*, 2005.

M. Oury and O. Tercieux. Continuous implementation. *Econometrica*, 80(4), 2012.

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.

E. Perez-Richet and V. Skreta. Test design under falsification. *Econometrica*, 90(3):1109–1142, 2022.

E. Perez-Richet and V. Skreta. *Fraud-proof non-market allocation mechanisms*. Centre for Economic Policy Research, 2024.

R. Roll. Orange juice and weather. *The American Economic Review*, 74(5):861–880, 1984.

I. R. Segal. Nash implementation with little communication. *Theoretical Economics*, 5(1): 51–71, 2010.

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

J. Sturm. How to fix a coordination failure: A "super-pigouvian" approach. *unpublished*, 2023.

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

J. Toikka. Ironing without control. *Journal of Economic Theory*, 146(6):2510–2526, 2011.

Q. Valenzuela-Stookey and F. Poggi. Market-based mechanisms. *Working Paper*, 2022. URL `https://www.franciscopoggi.com/files/mbm.pdf`.

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

Online Appendix

B Representing the market

In this section we further discuss the invariant representation, and prove existence for an important class of markets.

B.1 Axiomatic characterization for invariant representations

We first provide a general characterization of the markets that admit an invariant representation in terms of primitive properties.

Definition 11. Let $\hat{\mathcal{P}} \subset \mathcal{P}^\Theta$. The market is *competitively identified in* $(\mathcal{W}, \hat{\mathcal{P}})$ if

1. For any $M, M' \in \mathcal{W}$, any equilibrium price functions P given M and P' given M', and any state $\theta \in \Theta$, $M(P(\theta)) = M'(P'(\theta)) \quad \Rightarrow \quad P(\theta) = P'(\theta)$.
2. Let $\{M_a\}_{a \in K}$ be a family of constant policies and $\{P_a\}_{a \in K}$ a family of respective equilibrium price functions, where $K \subset \mathcal{A}$. Let $\{I_a\}_{a \in K}$ be a partition of Θ such that $P_a(I_a) \cap P_{a'}(I_{a'}) = \varnothing$ for all $a, a' \in K$ s.t. $a \neq a'$.[47] Let $P : \Theta \to \mathcal{P}$ be defined by $P := P_a$ on I_a, for each $a \in K$. Then $P \in \hat{\mathcal{P}}$ implies that P is an equilibrium price function for some $M \in \mathcal{W}$.

Both conditions in Definition 11 represent a sense in which equilibrium outcomes are separable across states, which is true in general only if agents are price takers. The first condition says that the equilibrium principal action in state θ uniquely identifies the equilibrium price in state θ, across all $M \in \mathcal{W}$ and all associated equilibria. The second part of Definition 11 says that we can generate new equilibria by stitching together equilibrium price functions, provided the resulting function is in some predetermined set $\hat{\mathcal{P}}$.

Proposition 12. *A market is competitively identified in* $(\mathcal{W}, \hat{\mathcal{P}})$ *iff it admits an invariant representation given* $(\mathcal{W}, \hat{\mathcal{P}})$.

Proof. The "if" is immediate. For the "only if", note that there is at most one equilibrium for any constant policy function M_a: take two equilibrium price functions P and P' given M_a,

[47]For a function $f : X \mapsto Y$ we use the notation $f(I) := \{f(x) : x \in I\}$ for $I \subset X$.

since $M_a(P(\theta)) = M_a(P'(\theta)) = a$, by Part 1 of the definition of a competitively identified, $P(\theta) = P'(\theta)$. Let $R(a, \theta)$ be the equilibrium price function given M_a.

1. Let P be an equilibrium price function given M. Fix θ, Let $a = M(P(\theta))$. Then consider M_a and respective equilibrium price function $R(a, \theta)$. By part 1, since $M(P(\theta)) = a = M_a(R(a, \theta))$ it must be that $P(\theta) = R(a, \theta)$.
2. Let $I_a = \{\theta : M \circ P(\theta) = a\}$. Then $P(I_a) \cap P(I_{a'}) = \varnothing$ for all $a, a' \in M \circ P(\Theta)$. (If $p \in P(I_a) \cap P(I_{a'})$, the definition of I_a says that $M(p) = a' \neq a = M(p)$). The families $\{M_a\}_{a \in M \circ P(\Theta)}$ and $\{P_a\}_{a \in M \circ P(\Theta)}$, where $P_a = R(a, \cdot)$, satisfy the conditions in part 2 of the definition of competitively identified.

□

Proposition 12 and Part 1 in Definition 8 tells us that $P(\theta) = R(M \circ P(\theta), \theta)$ for all θ is a necessary condition for P to be an equilibrium price function. This is immediately implied by part 1 in Definition 11. Part 2 of Definition 8, which is implied by condition 2 in Definition 11, says the converse. In some settings, before identifying an invariant representation, the requirements of robustness (to manipulation and multiplicity) imply restrictions on the set of possible equilibrium price functions. Thus for the purposes of identifying the invariant representation, we can restrict attention to a set of price functions $\hat{\mathcal{P}} \subsetneq \mathcal{P}$. Condition 2 in the definition of competitively identified may not hold for all $P \in \mathcal{P}$, but it is enough to know that it holds on $\hat{\mathcal{P}}$. The utility of this approach is illustrated in Appendix B.4.

B.1.1 Importance of price taking

No all markets admit invariant representations; as formalized in Appendix B.1, the key property is that agents be price takers. Markets in which agents are not price takers generally fail to admit an invariant representation.

Example 1. Consider the standard auction environment. The market participants are the bidders and the seller. An equilibrium consists of a mechanism for the seller and strategies for the bidders, specifying their bids as a function of their type. This environment is a market, where we let the price be that paid by the winning bidder. The principal could be a regulator who commits to taking some action as a function of the price. This action may affect the value of the object.

If $\mathcal{W}$ contains non-constant policies, it is easy to construct examples showing that the market does not admit an a.e. invariant representation. This is because bidders internalize the effect that their actions have on the price (winning bid) and thus respond to global properties of the principal's policy, which implies a violation of Definition 11, part 1.

B.2 Equilibrium inferences and global effects

One reason for the market to fail to admit an invariant representation is if for some policy M, there exist multiple equilibria with the same action function, but different price functions. However, it is relatively straightforward to extend our analysis to this type of market, using a representation via a market-clearing correspondence $R : \mathcal{A} \times \Theta \mapsto 2^{\mathcal{P}}$. The more interesting and challenging scenario, in terms of representing the market, is when there are *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). Nonetheless, we obtain representations.

To illustrate the challenges involved in modeling markets in which global effects are present, and to understand the reasons such effects might arise, consider the REE asset market of Section 5.1, Example 3. Recall that as stated in Proposition 9, this markets admits an invariant representation.

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 of Section 5.1 Example 1, other investors' signals are informative about a payoff-relevant state, and thus investors draw inferences from the price. Formally, this can be seen in the demand optimality condition

$$X_i(p, s_i) = \arg\max_x \mathbb{E}\left[u_i\big(x \cdot (\pi(M(p), \theta) - p)\big) \mid s_i, P_M(\theta) = p\right] \tag{8}$$

The difficulty with analyzing market-based policy in this environment can be seen by examining (8). The principal's policy affects investors in two ways. The first is a direct *forward guidance* effect: the policy determines what action investors anticipate, conditional on the price, and thus affects the anticipated dividend $\pi(M(p), \theta)$. However there's also an 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 policy 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 policy and the equilibrium action and price functions will matter in determining the price in a given state. Such global dependence makes it more difficult to analyze the principal's problem in outcome space (the space of action and price 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 an implementable action and price function. The price function, depicted in Figure 6a, 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 given 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 $X_i(P_2(\theta^*), s_i) = \arg\max_x \mathbb{E}\left[u_i\big(x \cdot (\pi(a^*, \theta) - p)\big) \mid s_i, \theta \in [\theta_1, \theta_2]\right]$

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 be 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 6b,

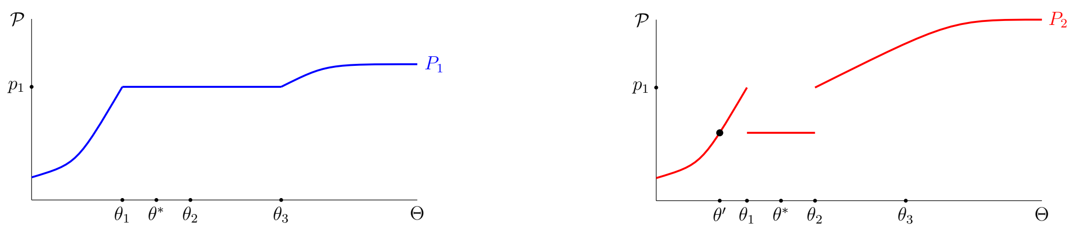

(a) Q_1 equilibrium

(b) Q_2 not implementable

Figure 6: Information effects and global dependence

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.

B.3 Invariant representation in REE

Despite the presence of informational effects, an invariant representation can be derived under quite mild conditions.

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.[48] 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 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 of investor types aggregate demand is deterministic, conditional on the state and the principal action a. Thus we can write market

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

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 policy 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.

In many settings, there is a monotone relationship between investors' private signals and their actions. It turns out that this is sufficient to guarantee the existence of an a.e. invariant 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 $U_\theta = \{\theta' \in \Theta : \theta' \geq \theta\}$.

Single Crossing. $V_i(a, \theta, x, p)$ satisfies single crossing in x, θ. Moreover, if $\theta' \in L_\theta$ then $V_i(a, \theta', x, p) - V_i(a, \theta', x', p) = V_i(a, \theta, x, p) - V_i(a, \theta, x', p)$.

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

The second piece of the Single Crossing assumption says that i has the same preferences over quantities in states $\theta, \theta' \in L_\theta$, conditional on a, p.

To see how these two assumptions imply that the market admits an invariant representation, consider the example illustrated in Figure 6. 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 Single Crossing and Athey (2002)). 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 to prove that the market admits an invariant representation 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 13. *Assume there is a complete order on* Θ *such that Single Crossing and Belief Monotonicity are satisfied. Then the market admits an a.e invariant representation given* $(\mathcal{P}^{\mathcal{A}}, \Theta^{\mathcal{P}})$*. In particular,* $R(a, \theta) = P^*(a, L_\theta, \theta)$*.*

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 Single Crossing 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 this 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 Single Crossing 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. 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 Single Crossing, 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 $U_{\theta_n} \subseteq A_n = \cup_{\theta' \in A_n}(U_{\theta'} \cap \mathcal{I}_P(\theta))$, Single Crossing 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 $R(a, \theta) = P^*(a, L_\theta, \theta)$, $M \circ P = Q$, and $R(Q(\theta), \theta) = P(\theta)$ 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 the 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 $P = P^*(a, L_\theta, \theta)$. □

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

B.4 Noisy REE in asset markets

This section illustrates a market in which restrictions on P derived without reference to an invariant representation are then used to identify the invariant representation. We do this in the canonical noisy REE model of Grossman and Stiglitz (1980) and Hellwig (1980) by extending 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 ω 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 $-\exp\left\{-\frac{1}{\tau_i} x(\pi(a, \omega) - p)\right\}$, 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 is that $(\omega, z) \in \mathcal{I}$, and i's private signal is s_i. Aggregate demand is $X(p|a, \mathcal{I}, \omega)$.

We generalize the standard noisy REE definition. Define a *public information function* as $\lambda : \mathcal{P} \to 2^{(\Omega, \mathcal{Z})}$. Given a pair (P, λ) of price and public information functions, say that markets clear if $X(P(\omega, z)|M(P(\omega, z)), \lambda(P(\omega, z)), \omega) = z$ for all (ω, z). An equilibrium given M consists of a pair (P, λ) such that

1. Markets clear given (P, λ).
2. Information is consistent: there exists a price function P' such that markets clear given (P, λ), and such that $\lambda(p) = \{(\omega, z) : P'(\omega, z) = p\}$

The standard definition of REE would require that $P' = P$, meaning public information is consistent with the equilibrium price function. If this holds then we say that P characterizes a *fully consistent* equilibrium. We will show in the end that under the constraint of equilibrium uniqueness, it is without loss of generality to restrict attention to fully consistent equilibria. However it is precisely because of the concern with equilibrium multiplicity that we introduce the more general notion. To see why, fix the policy M and suppose that (P, λ) is a fully consistent equilibrium. Recall that we interpret this market as one in which investors submit limit orders to a market maker. Suppose there exists another price function $\hat{P} \neq P$ such that markets clear given $(\hat{P}, \lambda)$. Then one may well be concerned that for a realized state (ω, z) the market maker sets the price $\hat{P}(\omega, z)$, rather than $P(\omega, z)$. 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. We want a notion of equilibrium uniqueness which rules out this type of multiplicity, hence our more general equilibrium definition.

On the other hand, the search for truly unique implementation is hopeless in the noisy REE model studied here, even restricting to fully consistent equilibria, since there are multiple (meaningfully different) fully consistent equilibria even when there is no policy feedback, that is, fixing the principal's action (Pálvölgyi and Venter, 2015). This is because even for a fixed principal action, there may be multiple fully consistent price functions (which of course correspond to different public information functions). We therefore need to consider a weaker uniqueness notion.

What we really want to rule out is multiplicity arising from the endogeneity of the principal's action. We therefore in this context say that M is robust to multiplicity if for any equilibrium (P, λ) given, there is no other price function $\hat{P}$ such that $(\hat{P}, \lambda)$ is also an equilibrium. In other words, we fix the public information function, and require uniqueness of equilibrium price functions. We show (Proposition 15) that any P which is uniquely implementable according to this notion must characterize a fully consistent equilibrium.

We turn now to establishing existence of a suitable invariant representation. The complication in this setting is that there is no easy way to narrow down the space of possible public information sets that can be revealed by the price. We therefore analyse directly the problem of characterizing what equilibria can be induced with a policy $M \in \mathcal{M}$ that is robust to multiplicity. To do this, we first need some preliminary results.

Lemma 18. *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. 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. $\square$

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.

Lemma 19. *For any $p \in \mathcal{P}$ and $a \in \mathcal{A}$, let $\mathcal{I}$ be a set satisfying $\mathcal{I} \subseteq \{(\omega, z) : X(p|a, \mathcal{I}, \omega) = z\}$. The distribution of ω conditional on $\mathcal{I}$ is absolutely continuous.*

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 18. 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 a 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. $\square$

Lemma 20. *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. 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 follow 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 19. $\square$

In words, Lemma 20 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. The following proposition identifies exactly which hyperplanes the public information sets can lie in.

Proposition 14. *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*

$$L^*(\omega, z|a) = \left(\frac{1}{\beta_1^a}\int_i \frac{\tau_i}{\sigma_i^2} di\right) \cdot \omega - z \tag{9}$$

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 almost all $(\omega'', z''), (\omega', z') \in \mathcal{I}$

Proof. Given Lemma 20, 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.[49] 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

[49] 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.

of $\tilde{V}^a$ conditional on $\tilde{S}_i^a = s_i$ and $\tilde{L}^a = \ell$ is given by

$$dF_{\tilde{V}|\tilde{S},\tilde{L}}(v|s_i,\ell) = \frac{\mathbb{1}[\ell - \alpha v \in (-b,b)] \exp\left\{\left(\frac{1}{\sigma_{ai}^2}s_i + \frac{\alpha}{\sigma_Z^2}\ell\right)v - \frac{1}{2}\left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2}\right)v^2\right\} dF_{\tilde{V}}(v)}{\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp\left\{\left(\frac{1}{\sigma_{ai}^2}s_i + \frac{\alpha}{\sigma_Z^2}\ell\right)x - \frac{1}{2}\left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2}\right)x^2\right\} dF_{\tilde{V}}(x)}, \tag{10}$$

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 (10) 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 $\hat{L}(s,\ell) = \left(\frac{1}{\sigma_{ai}^2}s_i + \frac{\alpha}{\sigma_Z^2}\ell\right)$ and

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

$$dH_i(v;\alpha) = \exp\left\{-\frac{1}{2}\left(\frac{1}{\sigma_{ai}^2} + \frac{\alpha^2}{\sigma_Z^2}\right)v^2\right\} dF_{\tilde{V}}(v)$$

This has the following important implication (essentially the same as Lemma A6 in Breon-Drish (2015)). The conditional distribution integrates to 1, i.e. $\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp\left\{\hat{L}(s_i,\ell)v - g\left(\hat{L}(s_i,\ell);\alpha,\ell\right)\right\} dH(v;\alpha) = 1$, so $\int_{\frac{\ell-b}{\alpha}}^{\frac{\ell+b}{\alpha}} \exp\left\{\hat{L}(s_i,\ell)v\right\} dH(v;\alpha) = \exp\left\{g\left(\hat{L}(s_i,\ell);\alpha,\ell\right)\right\}$. 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, we have $\alpha = \int_i \frac{\tau_i}{\sigma_{ai}^2} di$. 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 policy 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. □

We now wish to use these properties, in particular Lemma 20, to identify features of equilibrium. Proposition 14 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. 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 15. *Maintain the assumptions of Proposition 14. If $M \in \mathcal{M}$ is robust to multiplicity then (up to zero-measure violations) 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 (9). Moreover, $\tilde{P}$ characterizes a fully consistent equilibrium.*

Proof. Proposition 14 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 19 and Proposition 2.2 in Breon-Drish (2015) (Online appendix), which shows that individual demands will be linear in signals for any price (so aggregate demand is linear in the state). This implies also that $\tilde{P}$ is a fully consistent equilibrium. □

Let $\tilde{\mathcal{P}}$ be the set of price functions P such that every level set of P given by $\{(\omega, z) : L^*(\omega, z|M(p)) = \ell\}$ for some ℓ. We refer to $\tilde{\mathcal{P}}$ as the price functions with *non-intersecting level sets.*

The idea behind Proposition 15 is illustrated in Figure 7. Figure 7a 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 an invariant representation. The representation is saved, however, by the uniqueness requirement. In the situation depicted in Figure 7a, 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.

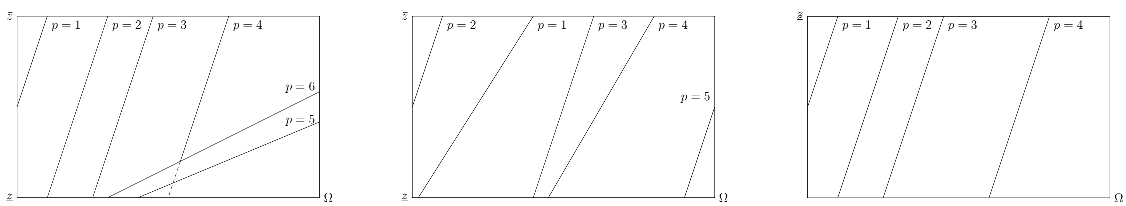

(a) Multiplicity (b) No multiplicity (c) Fixed action $a \in \mathcal{A}$

Figure 7: Linear statistics

From (9) 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 7c. 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. The proof of Proposition 14 also yields an expression for $R(a, \omega, z)$, 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 $\Omega \times \mathcal{Z}$).

Define $L(a, p) := \{(\omega, z) : X(p|a, L^*(\omega, z|a), \omega) = z\}$. Note that by definition, $L(a, p) = \{(\omega, a) : L^*(\omega, z|a) = \ell\}$ for some ℓ.

Corollary 1. Assume CARA utility, π affine in ω, additive normal signal structure and truncated-normally distributed supply shocks. Then the market has unique level sets represented by L. As a result, the market admits an a.e. invariant representation given $(\mathcal{W}^*, \tilde{\mathcal{P}})$. Moreover, $(\omega, z) \mapsto R(a, \omega, z)$ is strictly increasing for all a.

There are several differences between the environment of Proposition 14 and that of Breon-Drish (2015) Proposition 2.1. Most importantly, the current setting features a feedback effect, whereas asset returns follow a fixed distribution in Breon-Drish (2015).[50]

B.5 Proof of Proposition 9

1. Private values: Take two action functions M, M' and respective equilibrium price functions P, P'. Suppose that at θ, $M(P(\theta)) = M'(P'(\theta))$. The optimal demand schedule for agent i is $x^*(p|s_i) := x_i(s_i, p, M(p))$. If $P(\theta) > (<) P'(\theta)$ then for all i, $x_i(P(\theta)) < (>) x_i(P'(\theta))$. But then markets cannot clear in state θ for both P and P'.
2. Labor market: This is immediate, let $R(a,\theta) = \int_i u_i(a,\theta)$.
3. Common values, imperfect information: This follows from Proposition 13.
4. Noisy REE: See Corollary 1.

C Bailouts: imperfect information aggregation

Consider now the bailout problem in a framework with aggregate uncertainty. The market is modeled as in Section 5, example 4. The state $\theta = (\omega, z)$ consists of the payoff-relevant state $\omega \in \Omega$, interpreted as the strength of the company, and the shock to aggregate supply $z \in \mathcal{Z}$.[51] Agents' signals are equal to ω plus a normally distributed noise term with variance σ_i^2. Agent i has CARA utility with risk-aversion coefficient τ_i. The supply shock has a truncated-normal distribution on $(\underline{z}, \bar{z})$, where $\underline{z} = -\bar{z}$ (we allow $\bar{z} = \infty$). Additionally, we make the simplifying assumption that ω is uniformly distributed. The cash flow (the ex-post value of the asset) is assumed to be linear, given by $\pi(a,\omega) = \beta_0(a) + \beta_1(a)\omega$. Assume that $\beta_1(a)$ is strictly positive, strictly decreasing, weakly convex, and twice differentiable.

We normalize the social benefit of intervention to 1, so the principal's ex-post payoff is $u(a,\theta) := (1-a)\omega - ac$. The first-best action function is to make the maximal intervention if $\omega \leq -c$, and make zero intervention otherwise.

[50] Additionally, 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. We allow here for the supply shock to follow a truncated normal distribution, where Breon-Drish (2015) considers only the un-truncated distribution.

[51] We denote the strength of the company as ω, rather than θ as in Section 2, to maintain consistency with the rest of the paper: θ must contain all variables which determine the price.

The invariant representation for this market is discussed in detail in Appendix B.4. In brief, we show that any CWUI price function P must belong to the set $\tilde{\mathcal{P}}$ with the "no intersecting level sets" property. In particular, for any $p \in P(\Theta)$ there exists $\ell \in \mathbb{R}$ and $a \in \mathcal{A}$ such that $\{(\omega, z) : P(\omega, z) = p\} = \{(\omega, z) : L^*(\omega, z|a) = \ell\}$, where

$$L^*(\omega, z|a) := \frac{\kappa}{\beta_1(a)}\omega - z = \ell \tag{11}$$

and $\kappa := \int_i \frac{\tau_i}{\sigma_i^2} di$.[52] Given that $P \in \tilde{\mathcal{P}}$ is a necessary condition for CWUI, it suffices to show that the market admits an invariant representation in $(\mathcal{W}^*, \tilde{\mathcal{P}})$.

Lemma 21. *The above market admits an a.e. invariant representation given $(\mathcal{W}^*, \tilde{\mathcal{P}})$. Moreover, $(\omega, z) \mapsto R(a, \omega, z)$ is strictly increasing for all a.*

Proof. In Appendix B.4. □

C.1 Reformulating the problem

To characterize optimal policy in this setting, we use Proposition 11 to reduce the problem of choosing $Q : \Omega \times \mathcal{Z} \to [0, 1]$ to the simpler one-dimensional problem of choosing the marginal action function $\omega \mapsto Q(\omega, \bar{z})$.[53] For this purpose, it is convenient to define action functions on $\mathbb{R} \times \mathcal{Z}$, rather than just $\Omega \times \mathcal{Z}$.[54]

By Proposition 11, if Q is a CUI action function and $Q(\omega, \bar{z}) = a$, then the state $(\omega, \bar{z})$ belongs to a level set of the equilibrium price function which is defined by $\{(\omega', x') : R(a, \omega', z') = R(a, \omega, \bar{z})\}$. Since the principal's action must be measurable with respect to the price, this implies that $Q(\omega', z') = Q(\omega, z)$ for all (ω', z') in this level set. Let $w(a, z, x) : \mathcal{A} \times \mathcal{Z} \times \mathbb{R} \to \mathbb{R}$ be the unique $\omega \in \mathbb{R}$ such that $R(a, w(a, z, x), z) = R(a, x, \bar{z})$.[55] From eq. (11), $w(a, z, x) := x - \frac{1}{\kappa}\beta_1(a)\,(\bar{z} - z)$.

[52] Note that the higher is β_1, i.e. the more responsive is the cash flow to ω, the less informative is the price about ω. This is because when β_1 is high, investors' private signals of ω are less informative about the cash flow π. As a result, investors put less weight on their private signals relative to the public information contained in the price, and so the price is less informative. Since β_1 is decreasing in the level of intervention, this implies that higher levels of intervention make the price more informative about the payoff-relevant state. Bond and Goldstein (2015) also study how market-based interventions affect the efficiency of information aggregation by prices.

[53] For simplicity we focus on CUI; the analysis for CWUI is nearly identical. In fact, the optimal CUI and CWUI policies coincide.

[54] We still consider $\Omega \times \mathcal{Z}$ to be the state space, for the purposes of defining CUI.

[55] It is convenient to define w for x that are not in Ω.

Thus, any CUI $Q : \mathbb{R} \times \mathcal{Z} \to [0,1]$ is uniquely identified among the set of CUI action functions by its marginal $\omega \mapsto Q(\omega, \bar{z})$ (to visualize this approach, see Figure 8). Conversely, Proposition 11 implies that $\alpha : \mathbb{R} \to [0,1]$ is the marginal of a CUI Q if and only if it satisfies the following properties: *i)* α is continuous, *ii)* $x \mapsto R(\alpha(x), x, \underline{z})$ is monotone, and *iii)* $x \mapsto w(\alpha(x), \underline{z}, x)$ is strictly increasing. Condition *(iii)* says that the linear statistics for the price function do not intersect. Thus we optimize over $\alpha : \mathbb{R} \to [0,1]$ that satisfy conditions (i)-(iii).

C.2 Characterization of optimal policy

We say that the principal is *hawkish* if they are more predisposed to intervene than the market would like. In the model with supply shocks, this would mean that $a \mapsto \pi(a, \omega)$ is decreasing for all $\omega \geq -c$, where $\omega = -c$ is the highest state at which the principal wants to intervene. This means that the principal wants to intervene whenever doing so increases the company's cash flow, but possibly also in states where it does not. To accounts for the imperfection in the aggregation of information by the price, we say here that the principal is hawkish if $a \mapsto \pi(a, \omega)$ is decreasing for all $\omega \geq -\frac{1}{\kappa}\bar{z}\beta_1(1) - c$. The additional term $-\frac{1}{\kappa}\bar{z}\beta_1(1)$ reflects the fact that information aggregation is necessarily imperfect, and so the principal may choose less than full intervention for some ω below $-c$.

The characterization of optimal policy is most easily stated under the following assumption, which says that the range of fundamental uncertainty, i.e. $[\underline{\omega}, \bar{\omega}]$, is large relative to the range of supply shocks $[\underline{z}, \bar{z}]$ and the marginal cost of intervention c. This implies that $a = 0$ is optimal for low ω and $a = 1$ for high ω.[56]

Assumption 1. $\bar{\omega} \geq \frac{1}{\kappa}\bar{z}\left(\beta_1(0) - \beta_1'(0)\right) - c$ and $\underline{\omega} \leq \frac{1}{\kappa}\bar{z}\beta_1(1) - c$.[57]

Without aggregate uncertainty, i.e. if information was perfectly aggregated by the price, then first-best action function would be CUI if the principal is hawkish. This is not the case with aggregate uncertainty.

Proposition 16. *When the principal is hawkish, the optimal CUI action function Q^* is characterized by $\alpha^*(\omega) := Q^*(\omega, \bar{z})$ given by*

[56] This assumption is made purely for the purpose of exposition; it simplifies the proof of Proposition 16. Dropping this assumption does not change the qualitative results.

[57] Since $\beta_1(a)$ is decreasing, $\beta_1(0) - \beta_1'(0) \geq \beta_1(1)$, so the conditions are not redundant.

- $\alpha^*(\omega) = 1$ *for* $\omega \leq \omega^* := \frac{1}{\kappa}\bar{z}\beta_1(1) - c$
- $\alpha^*(\omega) = 0$ *for* $\omega \geq \omega^{**} := \frac{1}{\kappa}\bar{z}\left(\beta_1(0) - \beta_1'(0)\right) - c$
- α^* *is continuous and strictly decreasing on* (ω^*, ω^{**})*, defined as the unique solution to* $\frac{1}{\kappa}\bar{z}\big(\beta_1(a) - (1-a)\beta_1'(a)\big) - c = \omega$.

Thus Q^* *features maximal intervention for low-*ω*/high-*z *states and no intervention for high-*ω*/low-*z *states. Moreover, intermediate actions are taken with positive probability if and only if* $a \mapsto \beta_1(a)$ *is non-constant.*

Proof. Let $\mathcal{C}^+$ be the set of continuous and decreasing functions from $\mathbb{R}$ to $[0,1]$. Then the principal's program can be written as follows.

$$\max_{\alpha\in\mathcal{C}^+} \int_{-\infty}^{\infty}\int_{\underline{z}}^{\bar{z}} u\big(\alpha(x), w(z,\alpha(x),x)\big) f_\omega(w(z,\alpha(x),x)) f_z(z)\ dz\ dx \tag{12}$$

$$s.t.\quad x \mapsto R(\alpha(x), x, \bar{z}) \text{ increasing} \quad \text{and} \quad x \mapsto w(\alpha(x), \underline{z}, x) \text{ increasing.}$$

We solve this program ignoring the constraints that $x \mapsto R(\alpha(x), x, \bar{z}$ and $x \mapsto w(\alpha(x), \underline{z}, x)$ are increasing, and then verify that they are satisfied.

Let $\omega^{**} := \frac{1}{\kappa}\bar{z}\left(\beta_1(0) - \beta_1'(0)\right) - c$, and $\omega^* := \frac{1}{\kappa}\bar{z}\beta_1(1) - c$. First, consider $x \in (\omega^*, \omega^{**})$. Using the assumption that ω is uniformly distributed, the derivative of the objective at x with respect to a is proportional to

$$\begin{aligned}
&\int_{\underline{z}}^{\bar{z}} \left[u_1\big(a, w(a,z,x)\big) + u_2(a, w(a,z,x)) w_1(a,z,x)\right] f_z(z) dz \\
&\quad = \int_{\underline{z}}^{\bar{z}} \left[-\left(x + c - \frac{1}{\kappa}\beta_1(a)(\bar{z}-z)\right) - (1-a)\frac{1}{\kappa}\beta_1'(a)\left(\bar{z}-z\right)\right] f_z(z) dz \\
&\quad = \frac{1}{\kappa}\bar{z}\left(\beta_1(a) - (1-a)\beta_1'(a)\right) - (x+c)
\end{aligned} \tag{13}$$

Under the maintained assumption that β_1 is convex, $a \mapsto \beta_1(a) - (1-a)\beta_1'(a)$ strictly positive and strictly decreasing; and the expression in eq. (13) is strictly decreasing in a. Thus setting eq. (13) equal to 0 defines a decreasing and continuous (in fact, differentiable) $\alpha_* : \Omega \mapsto [0,1]$.

For $x \in \mathbb{R}$ such that either $x > \bar{\omega}$ or $w(0, \underline{z}, x) < \underline{\omega}$ the derivative with respect to a does not take the form in eq. (13). Consider $x > \bar{\omega}$. For any $a \in [0,1]$ and a policy α such that $\alpha(x) = a$, the posterior over $\mathbb{R}$ induced by the public information $\{(x', z) : L(x', z|a) = L(x, \bar{z}|a)\}$ first-order stochastically dominates that induced by $\{(x', z) : L(x', z|0) = L(\omega^{**}, \bar{z}|0)\}$. Since no intervention is optimal at ω^{**}, no intervention remains optimal at $x > \omega^{**}$. Full intervention is optimal for $x < \omega^*$.

We now argue that α_* also satisfies the constraints in the principal's program, and so constitutes an optimal policy. First, we show that $x \mapsto w(\alpha_*(x), \underline{z}, x)$ is increasing.

$$\frac{d}{dx} w(\alpha_*(x), \underline{z}, x) = 1 - \frac{1}{\kappa}\beta_1'(\alpha_*(x))(\bar{z} - \underline{z})\alpha_*'(x)$$

Implicit differentiation of the condition in eq. (13) over the range of x for which α_* is non-constant yields $\alpha_*'(x) = \left(\bar{z}\frac{1}{\kappa}\left(2\beta_1'(\alpha_*(x)) - (1-\alpha_*(x))\beta_1''(\alpha_*(x))\right)\right)^{-1}$. Substituting this into the previous expression and using the fact that $\underline{z} = -\bar{z}$ yields

$$\frac{d}{dx} w(\alpha_*(x), \underline{z}, x) = 1 - \frac{2\beta_1'(\alpha_*(x))}{(2\beta_1'(\alpha_*(x)) - (1-\alpha_*(x))\beta_1''(\alpha_*(x)))}$$

Since $\beta_1' < 0$ and $\beta_1'' \geq 0$, we have $\frac{d}{dx} w(\alpha_*(x), \underline{z}, x) \geq 0$ as desired.

We now show that $\omega \mapsto R(\alpha_*(\omega), \omega, z)$ is increasing. From eq. (13) we have $\alpha_*(\omega) = 1$ for all $\omega \leq \frac{1}{\kappa}\bar{z}\beta_1(z) - c$. Under the assumption that the principal is hawkish, $a \mapsto \pi(a, \omega)$ is decreasing for all $\omega > -w(1, \underline{z}, \frac{1}{\kappa}\bar{z}\beta_1(z) - c) = -\frac{1}{\kappa}\bar{z}\beta_1(z) - c$. Since α_* is decreasing above $\frac{1}{\kappa}\bar{z}\beta_1(z) - c$, this implies that $\omega \mapsto R(\alpha_*(\omega), \omega, z)$ is increasing. $\square$

The level sets of the equilibrium price function are depicted in Figure 8. Northwest of the blue dotted line full intervention occurs. Southeast of the solid red line there is no intervention. In between, Q^* is a continuous function that takes intermediate values. The equilibrium price is increasing in the southeast direction.

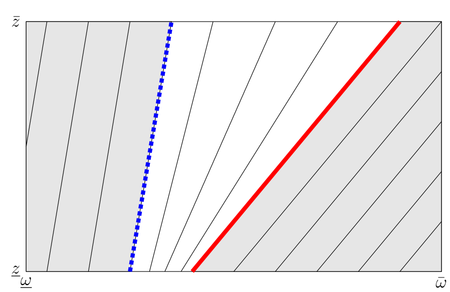

Figure 8: Level sets of optimal action function.

The first-best cannot be implemented when ω is not perfectly revealed. Aside from providing an analytical characterization of the optimal policy, the interesting feature of Proposition 16 is that intermediate actions play a significant role. This would not the case for a hawkish principal when the price is fully revealing, since in that case the first best would be

CUI. Intermediate actions are optimal here—even though the principal's ex-post payoff is linear—because the principal recognizes the impact their action has on the degree to which information is aggregated by the price. This effect operates via changes in the slope of the linear statistic in eq. (11). For higher states, the principal balances the desire to reduce the level of intervention against the reduction in the informativeness of the price. Only when $a \mapsto \beta_1(a)$ is constant does the principal take only extreme actions in equilibrium. The following can be observed from Proposition 16 and eq. (11).

Corollary 2. Under the optimal policy

i. The set of states for which intermediate actions are taken is increasing (in the set inclusion order) in $\beta_1(0) - \beta_1'(0)$, and decreasing in $\beta_1(1)$.
ii. The probability of taking an intermediate action decreases if private signals become more precise, i.e. σ_i decreases for almost all i.

C.2.1 Related work

To understand the value of solving the problem in outcome space, as we do here, it is instructive to contrast our approach with that of Siemroth (2019), who studies market-based policy-making in the same noisy REE setting. Siemroth (2019) optimizes directly over the policy maker's policy.[58] They do not impose robustness to manipulation or multiplicity. However, as it turns out, in order to make the problem of maximizing directly over policies tractable they must impose a number of restrictions which, using the insights of Theorem 2, can be seen to imply that these robustness properties must in fact be satisfied. Most importantly, they exogenously restrict attention to equilibria in which the induced price function is continuous, which immediately necessitates that it be monotone (otherwise there would be a violation of measurability). Moreover, Siemroth (2019) assumes that the ex-post value of the asset is additively separable in the state and the principal's action. As they note, this implies that "price informativeness is constant and unaffected by policy maker preferences". Our approach allows us to relax the additive separability assumption, which in turn reveals the trade-off faced by the principal between responding optimally to their information and

[58]Siemroth (2019) considers a policy maker without commitment power, which yields novel practical insights but is not essential for understanding the comparison with the current paper.

preserving the informativeness of the price (Proposition 16).

D Properties and extensions

D.1 Structural uncertainty

Assume throughout this section that Θ is closed and bounded. Endow the space of invariant representations $R : \mathcal{A} \times \Theta \to \mathbb{R}$ with the sup-norm. For a given policy M and invariant representation R, let $\tilde{Q}_R(\theta|M) := \{a \in \mathcal{A} : M(R(a, \theta)) = a\}$. In words, $\tilde{Q}_R(\theta|M)$ is the set of actions that are consistent with rational expectations in state θ.

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 a 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)* $\tilde{Q}_{\hat{R}}(\theta|M) \cap V(\theta) \neq \varnothing$.

Definition 12. A policy M is **robust to structural uncertainty at** R if $R \rightrightarrows \tilde{Q}_R$ is uniformly continuous at R

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 13. A policy 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 policy should induce almost the same joint distribution of states and actions for small perturbations to the invariant representation. 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 policy that is (weakly) robust to structural uncertainty.

Theorem 3. *If (Q, P) are CUI then they are implementable given invariant representation R with an essentially continuous policy that is robust to multiplicity and structural uncertainty* *at R. If (Q,P) are CWUI then they are implementable with an essentially continuous policy that is weakly robust to multiplicity and weakly robust to structural uncertainty at R.*

To prove Theorem 3, we make use of the following intermediate results.

Lemma 22. *Given a function $F : X \times [0,1] \to X$ on a compact subset X of an 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 $|F(x,t) - x| > \kappa \;\; \forall \;\; x \in X \setminus V$. 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 than zero, so the desired κ exists.

To complete the proof of Lemma 22, we need to show that there exists an open neighborhood U of t such $|F(x,t') - x| > \kappa \;\; \forall \;\; x \in X \setminus V, \;\; t' \in U$. By continuity of $t' \mapsto F(x,t') - x$, for each x there exists a ε_x such that $|t' - t| < \varepsilon_x$ implies $|F(x,t') - x| > \kappa$. For each x, define $\ell(x,\varepsilon) = \min\{|F(x,t') - x| : |t' - t| \le \varepsilon/2\}$, which exists by continuity of F and compactness of $|t' - t| \le \varepsilon/2$.

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 17. *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 a 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: 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 $\delta - \rho$.[59] By Lemma 22, 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 22 that for each $\theta \in S$ there exists $\varepsilon_\theta, \kappa_\theta > 0$ such that

$$|t' - t| < \varepsilon_\theta \Longrightarrow |F(x,\theta,t') - x| > \kappa_\theta \ \ \forall\, x \in X \setminus V(\theta). \tag{14}$$

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 (14). 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

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

1, $\theta \mapsto z(\theta, \varepsilon)$ is continuous at any $\theta \in S$. By (14) we know that $z(\theta, \varepsilon_\theta) > \kappa_\theta$ for all $\theta \in S$. Then for any $\theta \in S$ there exists an open neighborhood $B(\theta)$ of θ such that $\theta' \in B(\theta)$ implies $z(\theta', \varepsilon_\theta) > \kappa_\theta$. This proves Claim 2.

To complete the proof of Proposition 17, 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

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

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. □

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 5.2, (Q, P) can be implemented by a policy 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 17 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 is 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. □

The important implication of Theorem 3 is that small changes in R lead to small changes in the principal's expected payoff. Even though under the perturbed invariant representation 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 is not robust to structural uncertainty, at least when the discontinuity is not essential, i.e. when the left and right limits of M exist.[60] This further motivates the restriction to essentially continuous policies. 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 23. *Assume that M is robust 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. 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. □

[60] 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.

Lemma 23 shows that essential continuity is, to an extent, necessary for robustness to structural uncertainty.

D.2 Beyond uniqueness

The key insight is that even if a policy induces multiple equilibria, at least one of these will be weakly uniquely implementable. This is established via the following intermediate result.

Proposition 18. *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).*

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 11) 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 13. 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 13. 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. $\square$

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 5.2. 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. Let (Q, P) be an equilibrium induced by M, such that P is strictly monotone, which exists by Proposition 18. 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. □

Say that the principal is pessimistic if they evaluate a set of possible equilibria according to the worst case.

Corollary 3. Assume R is strictly increasing in θ and the environment is fully bridgeable. If the principal is pessimistic then the restriction to (Q, P) that are CWUI is without loss of optimality.

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 action and price functions can generally be very well approximated by CUI action and price 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 $(\hat{Q}, \hat{P})$.

D.3 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 a condition to account for actions for which the induced price is constant over an interval of states.

Consider (Q, P) implementable, which must therefore satisfy rational expectations and measurability. 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 19. *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 BC1. Moreover, if* P *is decreasing, then* Q *is BC2.*
4. *For all* θ, θ', $P(\theta) = P(\theta')$ *or* $P(\theta) = R(Q(\theta), \theta')$ *implies* $Q(\theta') = Q(\theta)$.

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 (Q, P) in the same way 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 outside of $P(\Theta)$ is the same as in Theorem 2. □

E Bridgeability and CWUI

To characterize the set of outcomes (Q, P) that are CWUI, first, note that 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, essential continuity of the implementing M gives the following result.

Lemma 24. *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. 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 $\lim_{n\to\infty} Q(\theta_n) = \lim_{n\to\infty} M(P(\theta_n)) = M(\bar{p})$. $\square$

Suppose P has a discontinuity at θ^*, and let $\underline{p} = \lim_{\theta \nearrow \theta^*} P(\theta)$ and $\bar{p} = \lim_{\theta \searrow \theta'} P(\theta)$.

Definition 14. Say that a 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 20. *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. 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). $\square$

Of the conditions in Proposition 20, the bridgeability condition is in theory the most difficult to verify. Fortunately, we show that most relevant environments are fully bridgeable and so this condition can be ignored. If the environment if fully bridgeable then we can replace condition 3 in Proposition 20 with the following: 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.