DOI: 10.1111/jmcb.12799

CHRISTOPH SIEMROTH

When Can Decision Makers Learn from Financial Market Prices?

I analyze a general setting where a policymaker needs information that financial market traders have in order to implement optimal policy, and prices can potentially reveal this information. Policy decisions, in turn, affect asset values. I derive a condition for the existence of fully revealing equilibria in competitive financial markets, which identifies all situations where learning from prices for policy purposes works. I discuss the possibility of using market information for banking supervision and central banking, and the general problem of asset design. I also demonstrate that some corporate prediction markets are ill-designed, and show how to fix it.

JEL codes: D53, D82, D84, G10
Keywords: asset design, asymmetric information, financial markets, policy, price informativeness, rational expectations equilibrium, self-defeating prophecy

Economists have long recognized that markets can aggregate and reveal diverse information among market participants via market prices (e.g., Hayek 1945). More recently, scholars and practitioners have called for using market information to improve real decisions and policymaking. There are many areas where market information might help. For example, a central bank (CB) may use asset prices to infer information about inflation expectations or future demand shocks, and adapt policy in response (Bernanke and Woodford 1997). A regulator may learn about the financial health of a bank from bond prices, and use this information for

I am grateful to Klaus Adam, Marco Bassetto, Pierre Boyer, Johannes Bubeck, Antonio Cabrales, Jayant Ganguli, Hans Peter Grüner, Felix Jarman, Xavier Lambin, Justin Leduc, Stefan Niemann, Edward S. Prescott, Andreas Rapp, Philipp Zahn, and seminar/conference participants at Mannheim, University College London, Tilburg, Nuffield, IESE, Rotterdam, Edinburgh, Essex, Surrey, Frankfurt School, Autonoma Barcelona, Carlos Madrid, BI Oslo, ENTER Jamboree 2015, VfS 2015 Münster, SAET 2015 Cambridge, EEA/ESEM 2016 Geneva, GEA 2016 Berlin for very helpful comments. This research was supported by the German Research Foundation (DFG) via SFB 884. Parts of this paper were previously included in a draft titled "The informational content of prices when policy makers react to financial markets."

Christoph Siemroth *is at the Department of Economics, University of Essex* (*E-mail:* christoph.siemroth@essex.ac.uk).

Received November 14, 2017; and accepted in revised form December 2, 2019.

Journal of Money, Credit and Banking, Vol. 0, No. 0 (xxxx 2021)

regulatory purposes (e.g., contingent capital with market trigger, Sundaresan and Wang 2015). Or a company could use internal prediction markets—where asset values depend on the launch date of a new product—to predict whether deadlines can be met, and react if forecasts indicate major delays (e.g., Cowgill and Zitzewitz 2015).

In all of these examples, one or several agents—which I shall call policymakers—react to the information contained in asset prices, and the reaction in turn affects asset values. However, this simultaneous feedback from prices to asset values and from asset values to prices presents problems both in theoretical and practical terms. In practical terms, the policy reaction may punish traders, which diminishes their incentives to reveal information, thus making market prices less informative and reducing their usefulness for policymakers. A theoretical problem is that such policymaker–trader interactions may not have an equilibrium. The main goal of this paper is to identify if and when it is possible for markets to inform policymakers if traders correctly anticipate the policymaker reaction.[1]

To address this question, I develop a general yet simple model of trader and policymaker interaction in competitive financial markets, and derive a necessary and sufficient condition where traders reveal their information by trading, and policymakers use this information for policy purposes in equilibrium. Thus, the condition identifies all situations where we can expect financial markets to work as policymaker tools without compromising the informativeness of prices, and all situations where we cannot. In the setting considered here, full revelation has positive welfare consequences and leads to a Pareto-optimal outcome, as it allows the policymaker to adopt the optimal policy.

Informally, the condition specifies whether the pricing problem is a self-defeating prophecy. The condition is simple and requires invertibility of the expected asset value given the optimal policy reaction to trader information in a sufficient statistic of trader information. If the condition is not fulfilled, then traders anticipate that informative prices would trigger a policymaker reaction that leads to trader losses, thus revelation of trader information is not incentive compatible.[2]

To illustrate the self-defeating prophecy problem, consider the "deadline securities" in corporate prediction markets, which are designed to forecast whether a project will be completed on time (e.g., Cowgill and Zitzewitz 2015). In the applications, I show these provide improper incentives for traders to share information about the project if management reacts to these forecasts. In such a market, employees can buy/sell an asset that pays off 1 if and only if the deadline of a certain project is missed, and 0 otherwise. Buying the asset is equivalent to betting on the deadline being missed. Thus, a high price of the asset indicates that many employees think the

1. Clearly, an unanticipated policymaker intervention does not diminish the incentives to trade on information. This paper instead focuses on the possibility of policymakers *systematically* using the information revealed by market prices.

2. Technically, as this is a competitive model, traders are not willing to clear the market for any fully revealing price function.

deadline will be missed, and the efficient market hypothesis suggests that the prices should be good forecasts of whether the deadline will be missed.

But now suppose management assigns additional resources if the price is high in order to prevent the deadline from being missed. Employees working on the project know whether the deadline is realistic, and can reveal that information by buying/selling the asset. Yet if an employee truthfully reveals that the deadline will be missed—by buying the asset and driving up the price—then management learns the deadline is not realistic and intervenes, thereby making the deadline realistic. Thus, traders who bet on the deadline being missed are punished: they bought the asset at a high price, but due to the reaction its value will be 0. In other words, traders forecast an asset value of 1 (deadline missed), but management reacts to this forecast and falsifies it by assigning more resources to the project. Thus, the pricing problem is a self-defeating prophecy.

Similarly, the necessary and sufficient condition can be used in any other context to check whether existing markets or assets promote information revelation by traders. And perhaps more importantly, if an existing setting features a self-defeating prophecy, the condition can be used to design different kinds of assets. Regarding the project deadline prediction markets, I show that a different asset design can fix the problem and restore the trader incentives for information revelation.

Using the model, I formally investigate the proposal that banking regulation could benefit from price information (e.g., Flannery et al. 2010). I show that banking regulation can potentially benefit from financial market information (e.g., from bank bond prices or credit default swap spreads) and help prevent bank default. However, a general lesson from the analysis is that regulatory action informed by market prices can never perfectly prevent any default in equilibrium. This is because if there is never default due to regulatory intervention, then asset prices cannot work as an early warning system for regulators (which would require prices to change with bank health) and at the same time accurately reflect the zero probability of default (which would require prices not to change in bank health). Thus, market-based intervention has to find a balance between obtaining more information from the market and exploiting that information.

In the context of central banking, I build on the example of Bernanke and Woodford (1997), who showed that an inflation targeting CB cannot use information from asset prices, and show that selection of different assets solves the self-defeating prophecy problem. Going beyond the three applications of banking regulation, corporate prediction markets, and central banking, the final part of the paper discusses asset design more generally. In particular, I obtain first answers to the more general questions of when there exists a possible asset design that allows for information revelation in equilibrium and what general guidelines for asset design one should follow.

These applications illustrate that the results of this paper are useful to identify situations where policymakers can use financial market information, and in designing institutions/assets that allow for better information revelation. Moreover, the problem raised here is a more fundamental challenge for the possibility of informationally efficient markets. Financial market prices may not reflect trader information even if all traders have perfect information, are perfectly rational, and obtain their perfect information for free. Thus, the self-defeating prophecy problem is a nonbehavioral challenge for informational efficiency. Finally, I also derive a necessary and sufficient condition for the existence of partially revealing equilibria, which adds to the existing literature which almost exclusively investigated fully revealing equilibria.

Related literature. Probably the most related paper is Bond et al. (2010). The authors consider the problem where a board of directors has no or only imperfect information about the quality of their agent, the company CEO, whereas traders have perfect information. A low-quality CEO reduces the firm value, hence should be replaced to increase the firm value, whereas medium- and high-quality CEOs should not be replaced, since the intervention is costly. In this setting, there is a difficulty in inferring CEO quality from the company stock price, which is a function of the company value, if traders know that the board might react to it. In the language of this paper, they describe a self-defeating prophecy and situations where it can be resolved. Bond et al. (2010) is specific in several dimensions such as perfect information for traders, binary intervention, or additive separability in the asset payoff function. The current paper generalizes to more flexible trader information structures (allowing traders individually or collectively to be imperfectly informed), asset payoff functions, and arbitrary policymaker preferences. Thus, the model presented here unifies their particular model and others in the literature and derives the necessary and sufficient condition for the possibility of revealing equilibria in this more general setting. It thereby explains the underlying reason why financial markets may fail to reveal information: the policymaker reaction makes the pricing problem a self-defeating prophecy.

Bernanke and Woodford (1997), in an extension of the Woodford (1994) model, presented the first example (to my knowledge) where revelation may fail in financial market/policymaker interaction. They consider a CB that attempts to infer a state variable θ from private forecasts or forecasts implicit in asset prices to reach a constant inflation target. Forecasters directly observe θ, the CB does not. In their static model, there is no rational expectations equilibrium (REE) that fully reveals θ to the CB. Again, their application is an example of a self-defeating prophecy.

Birchler and Facchinetti (2007) address a similar problem in banking supervision, and give a nice description of the "double endogeneity" problem of asset values affecting prices and prices affecting asset values via policy. They model a kind of prediction market that predicts bank failure, and the banking supervisor can react to information contained in these asset prices. As in the models above, full revelation may fail to occur because forward looking traders take into account that the supervisor will react to prices.

The setting considered here is related to the recent literature on contingent capital with market triggers. The idea of contingent capital with market triggers is that information in prices (typically about financial health of a bank) is used for real decisions (convert debt into equity, helping struggling banks raise equity), but this in turn affects asset values (returns to equity are diluted). The argument for market triggers is that they provide more current information than accounting measures, which tend to have a considerable lag. The contingent capital models (Prescott 2012, Sundaresan and Wang 2015) can also suffer from equilibrium nonexistence similar to the self-defeating prophecies problem described here. The main difference is that real decisions in these models are not made by a utility maximizing policymaker, but by a mechanical rule that reacts to market prices.

All of the above papers consider the problem of information revelation to a real decision maker by a noiseless financial market. Siemroth (2019) considers the problem of self-defeating prophecies in markets with noise, where a standard CARA-normal noisy rational expectations model is extended with a decision maker and asset values are endogenous to the real decision. While this model with noise may be conceptually superior to a model without noise, the model here is more general in various dimensions such as information structures and asset payoff functions. Bond and Goldstein (2015) also consider price informativeness in a noisy financial market with policymaker intervention, but their model does not feature self-defeating prophecies.

The current paper contributes to the growing literature of the real effects of financial markets via an informational channel, which mostly consists of studies without self-defeating prophecies. In most of this literature, the "real effect" is the financial market information impact on corporate decisions, as in Dow and Gorton (1997), Subrahmanyam and Titman (1999), Goldstein and Guembel (2008), Foucault and Gehrig (2008), Goldstein et al. (2013), Edmans et al. (2015), and Dow et al. (2017).

1. THE MODEL

1.1 Setup

Assets. Consider a financial market with a single riskless asset with return normalized to 1 ("cash"), and a single risky asset. The optimal policy and the risky asset value depend on a state variable θ, which is the realization of a random variable distributed according to a common prior distribution on support Θ, where Θ contains at least two elements. A policymaker sets a policy $i \in I$. The value of the risky asset is a function $a : \Theta \times I \to \mathbb{R}$, determined by state θ and policy i. Throughout I assume $\Theta \subseteq \mathbb{R}$, whereas policy I may be any set with possibly multidimensional elements.

Traders and trader information. The financial market consists of a mass 1 continuum of risk-neutral traders[3] with a common prior distribution over θ. Every trader $j \in [0, 1]$ receives a signal s_j on the realization of the state variable θ. The vector of signals $s = \{s_j\}_{j\in[0,1]}$ is distributed according to density $f(s|\theta) \neq f(s|\theta')\forall\theta \neq \theta' \in \Theta$. This formulation allows for the standard case of i.i.d. signals, or possibly correlated signals among traders. Since different trader signal profiles s can contain the same information, denote the summary statistic of the signal profile s by $\boldsymbol{s}$, and the

3. The results would be the same if we assumed the same well-behaved risk-averse preferences for all traders with the risky asset being in *zero* net supply. As is well known, a fully revealing equilibrium in this case is a no trade equilibrium with the asset price equal to the expected asset value. With risk-averse preferences and *positive* asset supply, the analysis would be qualitatively similar, although with a risk premium in the price function.

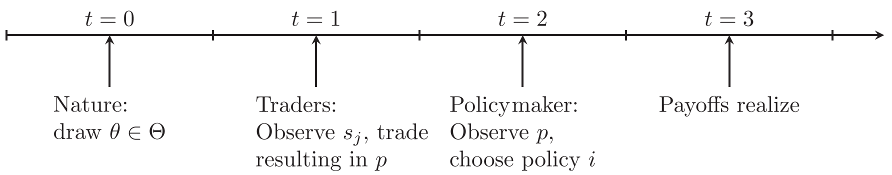

Fig 1. Timeline with Traders and Policymaker.

set of all possible unique realizations of $\boldsymbol{s}$ by $\boldsymbol{S}$, so that $\forall \boldsymbol{s} \neq \boldsymbol{s}' \in \boldsymbol{S} : h(\theta|\boldsymbol{s}) \neq h(\theta|\boldsymbol{s}')$, where h is the conditional probability density function of θ. $\boldsymbol{s}$ is a sufficient statistic for signal profile s if and only if $h(\theta|\boldsymbol{s}) = h(\theta|s)$. The following are three examples of commonly used information structures with corresponding summary statistic that are captured by this setup.

Example 1.

1. All traders receive perfect signals, that is, $s_j = \theta$ for all j, as, for example, in Bernanke and Woodford (1997) or Bond et al. (2010). The summary statistic is $\boldsymbol{s} = \theta$.
2. State $\theta \in \{0, 1\}$ is binomially distributed, $s_j \in \{0, 1\}$, and traders receive imperfect i.i.d. signals, that is, $1 > \Pr(s_j = 1|\theta = 1) = \Pr(s_j = 0|\theta = 0) > 1/2$ for all j. The sufficient statistic is $\boldsymbol{s} = \int s_j dj$.
3. The state space is the entire real line, $\theta \in \mathbb{R}$, traders receive normal i.i.d. signals $s_j \sim \mathcal{N}(\theta, \sigma^2)$ with $\sigma^2 > 0$. The sufficient statistic is $\boldsymbol{s} = \int s_j dj$.

Price function. In the remainder of the paper, I will focus on the revelation of trader information $\boldsymbol{s}$ to the policymaker. Let $p(\boldsymbol{s}) : \boldsymbol{S} \to \mathbb{R}$ be a price function mapping the sufficient statistic of trader information $\boldsymbol{s}$ into an asset price. For example, $a(\theta, i)$ may represent the value of a company, depending on economic fundamentals θ and some actions of the CEO i, while $p(\boldsymbol{s})$ is the price of the publicly traded company stock if traders received information $\boldsymbol{s}$ about θ.

Timing. The timing of decisions is illustrated in Figure 1: first, trading among all j leads to a market price $p(\boldsymbol{s})$, then, observing the price, the policymaker sets i. The results would be the same for simultaneous trading and policymaking, since the policy can condition on the price.

Policymaker information. In this paper, the policymaker knows the prior distribution of state θ, but is otherwise uninformed about the realization of θ. Since the focus is on determining if and when policymakers can learn from financial markets, it is clearly necessary that policymakers do not directly observe θ, otherwise there would be no need to learn from asset prices. But even if the policymaker were imperfectly informed about θ, it would typically not change the problem of inferring information from market prices. The more general case with imperfectly informed policymaker is analyzed in the online appendix; the seemingly small addition of the policymaker signal complicates the analysis without adding much economic insight. As a general rule, whenever the policymaker signal does not rule out any state $\theta \in \Theta$,[4] then full revelation of trader information is possible in the augmented model of the online appendix (with imperfectly informed policymaker) if and only if it is also possible in this model with uninformed policymaker (Corollary 4 in the online appendix). In short, the setting here assumes that traders know something about θ that the policymaker would like to know, which motivates the question if and when traders can reveal that information via asset prices.

Policymaker preferences. The utility function u represents the rational preference ordering of the policymaker over the tuple (θ, i). Thus, the utility maximizing policy i can—and in the interesting cases will—depend on the realization of the state θ. The utility function also includes possible costs of intervention. Consequently, if trader information $\boldsymbol{s}$ were known to the policymaker, she would choose policy

$$i(\boldsymbol{s}) \in \arg\max_{i} \mathbb{E}[u(\theta, i)|\boldsymbol{s}]. \tag{1}$$

To simplify the exposition, I assume that the optimal policy $i(\boldsymbol{s})$ exists and is unique for every $\boldsymbol{s} \in \boldsymbol{S}$. Results can be adapted for multiple solutions and mixed policy strategies in a straightforward manner. For nontriviality, I assume the optimal policy $i(\boldsymbol{s})$ varies for all realizations of $\boldsymbol{s}$, so the policymaker is interested in the trader information $\boldsymbol{s}$.

Welfare interpretation. The policymaker utility may be thought of as representing the welfare of all (unmodeled) nontrader agents in the economy. This would be the case, for example, with the policymaker as a benevolent public agent who aims at maximizing nontrader social welfare. In this interpretation, more information is better for nontrader welfare, so the question addressed here is not merely about the informativeness of prices, but also about social welfare in the real sector.

The crucial object. Define $v(\boldsymbol{s}) : \boldsymbol{S} \to A$, the expected asset value at the optimal policy based on $\boldsymbol{s}$,

$$v(\boldsymbol{s}) := \mathbb{E}[a(\theta, i(\boldsymbol{s}))|\boldsymbol{s}]. \tag{2}$$

The objects $i(\boldsymbol{s})$ and $v(\boldsymbol{s})$ are defined as if $\boldsymbol{s}$ is known to the policymaker, even though it is not. The reason is that once prices are fully revealing (see Definition 2), then $\boldsymbol{s}$ is known to the policymaker. Hence, the policymaker will implement policy $i(\boldsymbol{s})$ leading to expected asset value $v(\boldsymbol{s})$. These are reactions that forward looking traders are going to anticipate if prices are fully revealing.

Given the model, policy i cannot actually be conditioned on $\boldsymbol{s}$ directly, only on $p(\boldsymbol{s})$, so the resulting asset value is $a(\theta, i(p(\boldsymbol{s})))$. The difference is crucial, as $i(\boldsymbol{s})$ (and by extension $v(\boldsymbol{s})$) conditions on the exogenous $\boldsymbol{s}$, whereas $i(p)$ conditions on the endogenous price function. Equilibrium (to be formally defined later) will

4. As is the case, for example, if $\theta \in \mathbb{R}$ and the policymaker signal s_p is normally distributed, putting a positive density on each $\theta \in \mathbb{R}$.

require policy given beliefs to be optimal. Policymaker behavior, in particular u, is common knowledge.

The next definitions introduce two properties of price functions.

DEFINITION 1. *A price function $p(\boldsymbol{s})$ is accurate if and only if*

$$p(\boldsymbol{s}) = \mathbb{E}[a(\theta, i(p))|p(\boldsymbol{s}) = p, s_j] \quad \text{for all } j. \tag{3}$$

In words, an accurate price function requires that asset prices equal expected asset values from the perspective of all traders, where the information set of trader j is both the information contained in prices and his private information s_j. The condition can be interpreted as requiring no systematic mispricing, which becomes clearer in Lemma 1.

In the present setting, $\boldsymbol{s}$ can only be indirectly revealed to the policymaker via price $p = p(\boldsymbol{s})$. If the policymaker knows the price function, which she does in equilibrium, then $\boldsymbol{s}$ can always be inferred from price p if the price function $p(\boldsymbol{s})$ is invertible.

DEFINITION 2. *A price function $p(\boldsymbol{s})$ is fully revealing if and only if $p(\boldsymbol{s})$ is invertible.*

The next lemma establishes that if a price function is fully revealing, then accuracy implies $p(\boldsymbol{s}) = v(\boldsymbol{s})$ and vice versa (all proofs are in Appendix A). Hence, if prices are fully revealing and if the asset should not be mispriced from the perspective of any of the traders, then prices must equal the expected asset value given all trader information $v(\boldsymbol{s})$.

LEMMA 1. *A fully revealing price function $p(\boldsymbol{s})$ is accurate if and only if $p(\boldsymbol{s}) = v(\boldsymbol{s})$.*

1.2 The Possibility of Information Revelation via Prices

This section asks under which conditions a function $p(\boldsymbol{s})$ exists that fulfills two properties: full revelation and accuracy, given that the policymaker makes inferences from this price function. There is no microfoundation for this price function yet, that is, the section does not explain how the price function arises in some specified trading game or equilibrium concept. This foundation is provided in Section 1.4. The analysis is separated in this way to highlight that the impossibility of fully revealing and accurate prices does not depend on this microfoundation nor the equilibrium concept. Instead, under some conditions it is mathematically impossible to find a price function that is both fully revealing and accurate.

The main question is: In which economic environments is it possible for a price function to reveal $\boldsymbol{s}$ (Definition 2) and price accurately (Definition 1) at the same time? Without full revelation, the policymaker has inferior information and may implement suboptimal policies, and without accurate prices, traders might lose money, hence might be better off not trading. Given correct policymaker beliefs about the price function $p(\boldsymbol{s})$, Theorem 1 shows that this is possible if and only if $v(\boldsymbol{s})$, the expected asset value given optimal policy based on trader information $\boldsymbol{s}$, is invertible.

THEOREM 1 (Possibility of full revelation and accurate prices). *Suppose the policymaker knows function $p(\boldsymbol{s})$ and maximizes expected utility. Then a fully revealing and accurate price function exists if and only if $\{a(\theta, i), u(\theta, i), f(s|\theta)\}$ are such that $v(\boldsymbol{s}) = \mathbb{E}_\theta[a(\theta, i(\boldsymbol{s}))|\boldsymbol{s}]$ is invertible.*

Theorem 1 completely characterizes the existence of fully revealing and accurate price functions in terms of policymaker preferences $u(\theta, i)$ (which determine the reaction function $i(\boldsymbol{s})$), asset payoff functions $a(\theta, i)$, and trader information structures $f(s|\theta)$ (defining $\boldsymbol{s} \in \boldsymbol{S}$). If there exists no price function that is both accurate and fully revealing, then such a price function cannot arise in equilibrium no matter the equilibrium concept.

The result is mathematically almost trivial but economically important. It is mathematically straightforward because a function cannot, at the same time, be invertible and not invertible. If the asset value given the policymaker reaction to trader information $v(\boldsymbol{s})$ is not invertible, then a price function $p(\boldsymbol{s})$ cannot reveal the trader information (which requires invertibility) and accurately price the asset (which requires $p(\boldsymbol{s}) = v(\boldsymbol{s})$, i.e., noninvertibility).

What is the interpretation of the necessary and sufficient condition, that is, when does "learning" from asset prices work and when does it not? Suppose for a moment that the combined trader information is θ, that is, $\boldsymbol{s} = \theta$. It is easiest to start with the standard setup, where the asset value is only a function of the (exogenous) fundamentals: $a(\theta, i) = \theta$. Here the state variable exactly equals the asset value. A larger fundamental always corresponds to a larger asset value. Thus, it is easily possible to find a price function that both reveals θ and equals the asset value, namely, $p(\theta) = \theta$. In the present setting, however, asset values may depend on a policy $i(p(\theta))$ which in turn depends on the information contained in the price function $p(\theta)$. As in the introductory example, a higher price might reveal information that leads the policymaker to enact a policy that decreases the asset value, thus punishing traders for trading at high prices. And a lower price might indicate information that leads to a policy with higher asset value, thus again traders lose money due to the asset price/value differential. Consequently, the crucial difference to the standard setup is that asset values are endogenous to a player who reacts to asset prices, which makes it impossible to match asset values and prices if $v(\theta)$ is not invertible.

The best interpretation of the impossibility of fully revealing and accurate prices is that of a self-defeating prophecy. When trading, traders try to forecast the future asset value. Given risk neutrality and a competitive market, prices should equal expected asset values, so market prices can be interpreted as the market forecast of the future asset value. However, if $v(\boldsymbol{s})$ is not invertible, then it is impossible to make a correct forecast in every state if that forecast reveals the state, because the policymaker reacts to the forecasts (prices) after they have been made and falsifies the forecasts: prediction of a high asset value triggers a low asset value, and prediction of a low asset value triggers a higher asset value. Consequently, noninvertibility of $v(\boldsymbol{s})$ is a situation where policymaker preferences and trader preferences (making zero profits

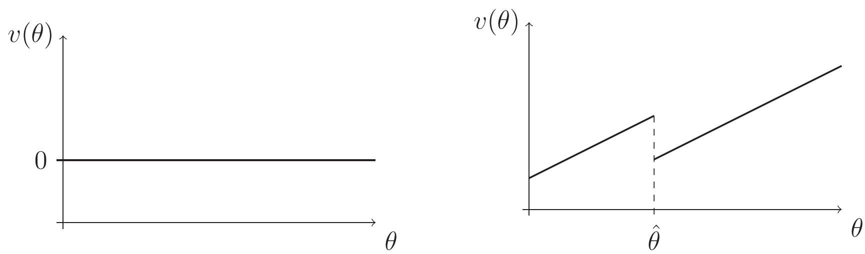

(a) Bernanke and Woodford (1997): $i(\theta) = \arg\min_i \mathrm{Var}(\theta + i)$

(b) Bond et al. (2010), Prescott (2012): intervention below threshold $\hat{\theta}$ increases asset value

Fig 2. Examples for Noninvertible Asset Values at the Optimal Policy $v(\theta)$ under Full Information

in a competitive market) are not aligned: it is impossible that both traders and the policymaker get what they want given this market mechanism.

1.3 Examples from the Literature

In several papers with policymaker/trader interaction from the literature section, traders have perfect information about the state variable θ. It can easily be verified that invertibility of $v(\theta)$ is not fulfilled in these papers (see Figure 2). For example (adapting their notation), in Bernanke and Woodford (1997)'s static model, the CB wants to cancel out all variance due to inflation pressures θ, so that the asset value (assumed to equal the inflation rate) given the optimal policy using trader information is $v(\theta) = \theta + i = c$ for some constant c. Yet if θ were revealed by a price function that changes in θ, then this price function cannot equal the actual asset value c in all states $\theta \in \Theta$. In Bond et al. (2010) (similar in Prescott 2012), the policy variable is binary, and an intervention is value increasing: $a(\theta, i = 1) > a(\theta, i = 0)$. Moreover, the asset value increases in θ. The optimal policy calls for $i = 1$ if and only if $\theta \leq \hat{\theta}$ for some threshold $\hat{\theta}$, hence $v(\theta)$ has a discontinuous downward jump at $\hat{\theta}$, which makes it noninvertible (see Figure 2).

Hence, the underlying problem—noninvertibility of $v(\theta)$—is the same in these papers. Despite the same problem, preferences of the policymaker differ, which shows the problem of full revelation with self-defeating prophecies is not due to specific policymaker goals. In Bernanke and Woodford (1997), a CB wants to minimize the variance of inflation, and in Bond et al. (2010) a board of directors wants to maximize firm value minus intervention cost. In Prescott (2012), the policy is determined by a capital conversion rule.

Despite noninvertibility of $v(\theta)$ in Bond et al. (2010), they show that full revelation may be possible under some strict conditions on the policymaker information. This is the only application in the literature I am aware of where the assumption that the policymaker is informed (in a specific way) makes a difference whether or not prices can be revealing, and the more general necessary and sufficient condition in the case of a partially informed policymaker is derived in the online appendix.

1.4 Existence of Fully Revealing Rational Expectations Equilibria

This section investigates whether full revelation can occur in an REE. That is, under which conditions can the financial market aggregate and reveal private information to the policymaker in a competitive equilibrium?

Recall we have a continuum of risk-neutral traders $j \in [0, 1]$. To obtain finite net demands, assume that each trader is budget constrained with a budget $w > 0$, and assume that each trader owns one unit of the risky asset and cannot short-sell, which sets a lower bound on net demands: $x \geq -1$. A rational expectations equilibrium in this setting is defined as follows.

DEFINITION 3. *An REE with policymaker consists of*

i. *Optimal trader (net) demands for the risky asset given their private information and the information revealed in prices,*

$$x_j(p, s_j) = \arg\max_x \mathbb{E}_\theta[x(a(\theta, i(p)) - p)|p(\boldsymbol{s}) = p, s_j]\forall j \in [0, 1] \\ \text{s.t. } xp \leq w, x \geq -1. \tag{4}$$

ii. *An optimal policy reaction function $i(p)$ given the information revealed in prices,*

$$i(p) = \arg\max_{i \in I} \mathbb{E}_\theta[u(\theta, i)|p(\boldsymbol{s}) = p], \tag{5}$$

iii. *And a price function $p(\boldsymbol{s})$ that clears the market for every $\boldsymbol{s} \in \boldsymbol{S}$, that is,*

$$\int_0^1 x_j(p = p(\boldsymbol{s}), s_j)dj = 0 \forall \boldsymbol{s} \in \boldsymbol{S}. \tag{6}$$

Condition (i) of Definition 3 requires that trader demands maximize the expected utility, based on their private information s_j and the information contained in equilibrium prices $p(\boldsymbol{s})$. This condition is standard except traders recognize that the asset value $a(\theta, i(p))$ is endogenous to the policy reaction to price p, which is the straightforward equilibrium generalization for this setting. Condition (ii) is new since standard models do not feature a nontrader player. The policymaker chooses a policy that maximizes her expected utility given the information contained in equilibrium prices $p(\boldsymbol{s})$. Condition (iii) is a standard market clearing condition.

It turns out that the same condition which determines the existence of a fully revealing and accurate price function is also necessary and sufficient for the existence of a fully revealing REE. This is because Theorem 1 assumed that the policymaker knows $p(\boldsymbol{s})$ and acts optimally given her information, which is now an equilibrium requirement, and accurate prices now follow from utility maximization of traders and market clearing (i.e., traders do not clear the market if asset prices are not accurate).

PROPOSITION 1. *A fully revealing REE exists if and only if* $\{a(\theta, i), u(\theta, i), f(s|\theta)\}$ *are such that* $v(\boldsymbol{s})$ *is invertible.*

Thus, according to the REE concept, financial markets can both aggregate and reveal all trader information $\boldsymbol{s}$, if and only if $v(\boldsymbol{s})$ is invertible.[5] Hence, there are situations where markets cannot be strong-form informationally efficient, that is, prices do not reflect all trader information. If prices fully revealed trader information, then in at least one state there would be mispricing, which introduces incentives to exploit the mispricing, and consequently traders do not support a fully revealing price function in equilibrium. Even in the most extreme case, where all traders perfectly know the state of the world ($s_j = \theta \forall j$) and perfectly know policymaker behavior, prices cannot reflect trader information if invertibility of $v(\boldsymbol{s})$ fails to hold. The problem is not an informational one—traders know everything. Instead, accurate prices and full revelation are mutually exclusive, because the policymaker *de facto* "prefers to falsify trader forecasts." This result is in strong contrast to the standard models without policymaker, where asset values are exogenous and existence of fully revealing REE is generic (see, e.g., Radner 1979 or Allen 1981), that is, markets are (strong-form) informationally efficient.

Note that the fully revealing equilibrium, if it exists, is Pareto-efficient. The price reveals the union of all trader information, so there is symmetric information among traders in equilibrium. The equilibrium is a competitive one, so by the first welfare theorem the allocation among traders is Pareto-efficient (Grossman 1981). Moreover, the policymaker receives all information that is available in the economy and—by the equilibrium requirement—implements the optimal policy, thus maximizing her utility. Consequently, a social planner with access to the combined information in the economy cannot improve on allocation and policy in the fully revealing REE. Thus, full revelation has positive welfare consequences in the setting considered here, and not just for traders, but for the real economy as well, as it benefits from the information of the financial market.

1.5 Partially Revealing Rational Expectations Equilibria

This section discusses equilibrium existence if fully revealing equilibria do not exist, by deriving a necessary and sufficient condition for the existence of partially revealing equilibria when traders are perfectly informed about the state.[6]

5. The online appendix shows that the same condition is necessary and sufficient for full revelation using the Perfect Bayesian equilibrium concept in a continuum economy if traders receive perfectly correlated signals, so the results are relevant beyond the REE concept.

6. Investigating partially revealing equilibria with other signal structures would require us to specify the distribution of information among traders in more detail, which is beyond the scope of this paper. In the main section, focusing on fully revealing equilibria, the distribution of private information among traders did not matter since all information is revealed in equilibrium. This is different when analyzing partially

Discussing partially revealing equilibria will require some additional notation and the concept of set partitions.

Definition 4. *A partition $\mathcal{P}$ of the set of all possible states Θ is a set of subsets, with its subsets ("parts") $\mathcal{I} \in \mathcal{P}$ such that*

- *all parts of the partition together span Θ ($\cup_{\mathcal{I}\in\mathcal{P}}\mathcal{I} = \Theta$),*
- *all parts of the partition are disjoint, so that any state $\theta \in \Theta$ belongs to exactly one part of the partition ($\mathcal{I} \cap \mathcal{I}' = \emptyset \forall \mathcal{I} \neq \mathcal{I}' \in \mathcal{P}$), and*
- *no part of the partition is empty ($\emptyset \notin \mathcal{P}$).*

Hence, a partition divides the set of states Θ into separate and disjoint parts, which are subsets of Θ. A subset $\mathcal{I} \in \mathcal{P}$ may include one or more elements (states). Set partitions are a convenient way to formalize the notion of partial revelation: In a partially revealing equilibrium, the price function reveals a part of the partition $\mathcal{I}$ of Θ, that is, reveals that the realized state is one of the states in $\mathcal{I}$, but not which. Hence, revealing the part of the partition is not equivalent to fully revealing the state if that part includes more than one state. A partially revealing equilibrium therefore must have the same price in all states that are part of the same partition, that is, $p(\theta) = p(\theta')$ for all $\theta \neq \theta' \in \mathcal{I}$.

Adapting the earlier notation, let $i(\mathcal{I})$ denote the optimal policy choice if the policymakers knows the realized state is part of $\mathcal{I}$ (but not which state in that part of the partition). We can now state the necessary and sufficient condition for the existence of a partially revealing equilibrium (a special case of which is a nonrevealing equilibrium).

Proposition 2. *Suppose traders are perfectly informed ($s_j = \theta \forall j$). A partially revealing equilibrium exists if and only if there exists a partition $\mathcal{P}$ of Θ such that*

1. $\exists \mathcal{I} \in \mathcal{P} : |\mathcal{I}| > 1$, *and*
2. $v(\mathcal{I}) = \mathbb{E}[a(\theta, i(\mathcal{I}))|\theta \in \mathcal{I}]$ *is invertible for all* $\mathcal{I} \in \mathcal{P}$, *and*
3. $a(\theta, i(\mathcal{I})) = a(\theta', i(\mathcal{I}))$ *for all* $\theta \neq \theta' \in \mathcal{I}$, *for all* $\mathcal{I} \in \mathcal{P}$.

Hence, a partially revealing equilibrium exists if and only if the state space can be partitioned such that three conditions are fulfilled: First, at least one part of the partition includes more than one state, otherwise it would be a fully revealing equilibrium.

Second, revealing parts of the partition leads to a policymaker reaction such that the expected asset values for each part of the partition is different. This condition (and the reason why it is necessary for equilibrium) is very similar to the case of fully revealing equilibria, and indeed this is not surprising when partially revealing equilibria are viewed as "fully revealing parts of a partition of Θ" (rather than fully revealing the state). That is, the equilibrium fully reveals groups of states rather than individual states. The condition is necessary because without it, traders would not clear the

revealing equilibria that may retain asymmetry of information among traders, so that the distribution of information matters for market clearing.

market for such a price function, as prices differ between parts of the partition but asset values do not. An interpretation of this is the self-defeating prophecy problem as before: Partially revealing information leads to a policy reaction that makes traders not want to reveal that information.

Third, within each part of the partition, the asset values from the perspective of the traders need to be constant, that is, cannot change with the state. This is a requirement that did not appear in the discussion of fully revealing equilibria, because it refers to the nonrevealing component of the equilibrium. In a partially revealing equilibrium, there exists at least one part of the partition with multiple states such that the price is the same for all states in that part. Hence, the policymaker cannot distinguish the states in that part of the partition, and must use the same policy in all of these states. But traders can distinguish these states, so for this to be an equilibrium, the asset value—evaluated at the partially informed policy choice $i(\mathcal{I})$—needs to be the same for all states $\theta \in \mathcal{I}$ in the part. If it was not, then traders would not clear the market for such a price function, because asset values diverge from prices in at least one state in that part. Thus, this condition ensures that traders are willing to sustain a price function that withholds information about which state in the part realized. To illustrate this, consider the following example.

EXAMPLE 2. Bond and Goldstein (2015) consider a setting (simplified here) where the asset value is $a(\theta, i) = c + i$ with c being independent of θ, that is, the asset value does not directly depend on the state θ, but may nevertheless be affected through the policy action i. In their interpretation, the asset is the stock of a firm whose value depends on a government intervention (such as a cash injection) which may vary by state.

If the policymaker is not informed in this setting, then if the price reveals no information, the policymaker has to take the same action $i(\emptyset)$ in every state. Thus, the asset value would be $a(\theta, i) = c + i(\emptyset)$, which is independent of the state and thus fulfills condition 3 of Proposition 2. Hence, there is a nonrevealing equilibrium with $p(\theta) = c + i(\emptyset) \forall \theta \in \Theta$. And, if policymaker preferences are such that $i(\theta)$ is invertible, then a fully revealing equilibrium exists in addition. □

Consequently, a partially revealing equilibrium may exist alongside a fully revealing one, since the conditions in Propositions 1 and 2 are not mutually exclusive. If they coexist, then a partially revealing equilibrium can be interpreted as the market withholding some information that the policymaker needs, which is Pareto-inferior since traders make zero expected profits in any competitive equilibrium, but the policymaker is better off in a fully revealing one. It may also be that only one type of equilibrium exists or neither.[7] For specific applications, it is a matter of checking the conditions in Propositions 1 and 2, which will be illustrated in the following section. A very easy-to-check sufficient condition for nonexistence of any partially revealing

7. And multiple different partially revealing equilibria may exist, whereas the fully revealing equilibrium is unique if the policy reaction $i(s)$ is unique.

equilibrium, and one for the existence of a nonrevealing equilibrium, follows from Proposition 2:

COROLLARY 1. *Suppose traders are perfectly informed* ($s_j = \theta \forall j$).

i. *If the asset is such that* $a(t, i) \neq a(t', i) \forall t \neq t' \in \Theta, \forall i \in I$, *then there exists no partially revealing REE.*
ii. *If the asset is such that* $a(t, i) = a(t', i) \forall t \neq t' \in \Theta, \forall i \in I$, *then there exists a nonrevealing REE.*

2. APPLICATIONS

2.1 Design of Assets in Corporate Prediction Markets

Corporate prediction markets are designed to elicit information dispersed among employees about business-relevant future outcomes such as whether project deadlines can be met, what next quarter's demand for a product will be, or whether a competitor will enter a particular market segment. Prediction markets typically trade simple assets whose value depends on these outcomes. The information revealed by these markets is most useful if it is used to improve corporate decisions. However, this "policymaker" reaction to market prices is exactly what can create self-defeating prophecies, because it also affects asset values.

One example where these prediction markets affect company policy is mentioned in Cowgill and Zitzewitz (2015). Ford decided against introducing several new products after prediction market forecasts revealed that these would not be popular among consumers. Indeed, improving decisions was the main reason for using these markets: "Ford Motor Company [turned to prediction markets] to improve their ability to make decisions that would be in line with customer interests" (HPC Wire 2011).

The following example investigates the project deadline assets that are used in corporate prediction markets to determine whether a project will finish on time and whether intervention may be necessary. The empirical paper of Cowgill and Zitzewitz (2015) describes these markets with institutional details. The analysis here shows for the first time (to the best of my knowledge) that these markets are designed in a way that undermine the incentives of traders to reveal their information. An alternative asset design can fix this problem.

EXAMPLE 3. A project in a company may miss the deadline ($\theta = 1$) or it may meet the deadline ($\theta = 0$) with the currently available resources. This is the "state of the project." Insiders—for example, the employees who work on the project—know the state ($s_j = \theta$), while the manager (policymaker) does not. The deadline asset in the corporate prediction market pays 1 if and only if the deadline is missed, and 0 otherwise. The manager can react to information about whether the deadline will be met: $i = 0$ means the project does not receive additional resources (more manpower, funds, etc.), and $i = 1$ means the project receives additional resources that definitely ensure completion on time. The manager does not want to commit additional resources unless it is necessary, since it is costly. Consequently, the asset value is $a(\theta, i) = \mathbf{1}\{i = 0 \wedge \theta = 1\}$, that is, the project misses the deadline if and only if the manager does not commit additional resources ($i = 0$) and the project misses the deadline without additional resources ($\theta = 1$).

To determine whether full revelation is possible in equilibrium, calculate $v(\theta)$, that is, the asset value if θ (the trader information) were revealed to the manager. Clearly, $v(\theta = 1) = 0$, since the manager prevents missing the deadline by assigning more resources, $i(\theta = 1) = 1$. Moreover, $v(\theta = 0) = 0$ with $i(\theta = 0) = 0$. Thus, the deadline is never missed, the $v(\theta)$-function is not invertible in θ, and a self-defeating prophecy prevents revelation of trader information: Traders anticipate that revelation of $\theta = 1$ (a missed deadline) triggers a policy reaction that prevents the deadline being missed, hence they would lose money by buying the asset.

How can corporations solve this problem? Since they are free to design other assets in their own markets, a simple adjustment fixes the problem. Consider another asset with value $a(\theta, i) = i$, that is, the asset pays \$1 if and only if the company commits additional resources. The $v(\theta)$-function is $v(\theta = 1) = i(\theta = 1) = 1$ and $v(\theta = 0) = i(\theta = 0) = 0$, that is, it is invertible. Thus, instead of designing an asset that predicts the outcome (deadline missed yes/no), which the company might seek to change depending on state and information, another asset simply predicts the intervention decision. In equilibrium, traders have incentives to forecast the optimal policy for the policymaker, and this forecast is self-*fulfilling*, since the policymaker wants to follow the "recommendation."

While this alternative asset is preferable in terms of trader incentives, this solution is not perfect, since there exists another equilibrium, which is not fully revealing (see conditions in Proposition 2). Suppose the prior beliefs of the manager are such that she commits additional resources unless she receives information that it is unnecessary. The price function $p(\theta = 1) = p(\theta = 0) = 1$ is uninformative, since the price is identical in both states, and accurate, since the manager will commit additional resources without additional information; it is an equilibrium. Yet clearly the manager does not implement the optimal policy of committing additional resources if and only if it is necessary. Similarly, if the prior beliefs imply not to commit additional resources, then $p(\theta = 1) = p(\theta = 0) = 0$ is an uninformative and suboptimal equilibrium.

The best theoretical solution, the asset $a(\theta, i) = \theta$, is not practically implementable in the case of project deadline prediction markets, since θ (the information to be revealed) is the outcome *in the absence of intervention*, that is, a counterfactual that is not *ex post* contractible if additional resources are committed.

One caveat is that the model uses a competitive equilibrium concept, where traders act as price takers. In very small corporate prediction markets, this may not be realistic: A single trader can affect prices, which possibly introduces incentives for price manipulation. A lot of these corporate markets at Google and Ford are quite large (Cowgill and Zitzewitz 2015)—at Google, all employees were eligible to participate—so the price taker assumption is plausible in many cases. Moreover, there is the potential problem that employees might seek to manipulate outcomes for financial gain, which is beyond the scope of this paper, but discussed in Ottaviani and Sørensen (2007). Overall, the lesson here is that even if employees do not manipulate these markets, then they might not work as management tools if they are poorly designed due to self-defeating prophecies.

2.2 Market Information for Banking Supervision

There is considerable asymmetric information between banks and the authorities that are tasked to supervise and regulate these banks. With costly audits, supervisors can reduce this information asymmetry, but still empirical evidence shows that financial markets prices—especially of bank equity and bank debt (bonds)—contain information that bank supervisors do not have (e.g., Berger, Davies, and Flannery 2000 or Gunther et al. 2001). Thus, these papers argue that financial markets as a whole could potentially provide information that regulators would like to have, just as is assumed in this paper. Indeed, this is also recognized and practiced by regulators: "[U]sing market information in supervision, the Federal Reserve and other regulatory agencies already monitor subordinated debt yields and issuance patterns in evaluating the condition of large banking organizations" (Greenspan 2001). Moreover, market prices such as those of credit default swaps appear to incorporate crucial information earlier than agency credit ratings (e.g., Flannery et al. 2010). But this paper shows that whether traders share their information in equilibrium depends on what regulators do with the information inferred from financial market prices.

The main interest of banking supervisors is to maintain the stability of the banking sector and to prevent failures of systemically important banks that might require a costly bailout. Thus, potentially market information could help indicate problems with particular banks, which ideally helps the regulator to intervene before it is too late. Consider the following simple example of market information in banking supervision.

EXAMPLE 4. First consider the case of no regulator. Suppose a bank issued bonds, whose value is $a(\theta, i) = \min\{\theta, F\}$, where $\theta \geq 0$ is bank health and $F > 0$ is the face value of the bond. If $\theta < F$, then the bank cannot repay all of its obligations and the bond value falls. Suppose traders observe bank health $s_j = \theta$. While there is no equilibrium that reveals θ, there is one where the market price reveals whether there will be (partial) default. The asset value based on trader information is $v(\theta) = \min\{\theta, F\}$, so prices $p(\theta) = v(\theta)$ reveal whether $\theta > F$ or $\theta < F$, that is, whether the bank is unable to repay in full, and if so it reveals bank health θ.

Now suppose we introduce an uninformed regulator who can intervene with action $i \geq 0$. Action i could be any of the possible interventions such as forced recapitalization or mandating dividend restrictions to maintain equity cushions, etc. The intervention improves the bank's health, so the asset value is $a(\theta, i) = \min\{\theta + i, F\}$ (i.e., the regulator improves health to F if $\theta < F$ or does not intervene if $\theta \geq F$). Hence, the interpretation of θ is now "bank health absent intervention." Suppose the regulator wants to prevent losses to debt holders (e.g., to prevent contagion), which could be described by the reaction function $i_1(\theta) = \max\{F - \theta, 0\}$. Given such a reaction function, forward looking traders—anticipating the regulatory response if their information was revealed—value the bond as

$$v_1(\theta) = \min\{\theta + i_1(\theta), F\} = F, \tag{7}$$

that is, at face value independent of bank health, so there exists no revealing equilibrium (Proposition 1). Thus, without further features in the model, traders will not share their information knowing there is a regulator that perfectly reacts to market information about bank health to prevent default. The assumption that the regulator can simply "fix" any bank issues she knows about θ is a strong one, yet the more general point stands: Any intervention that is informed by market prices and thereby successfully prevents any default cannot work in equilibrium, because the final asset value would not vary with the state θ, yet market prices must vary in bank health θ in order to reveal it (this is also recognized by Birchler and Facchinetti 2007). This point applies also to other assets whose value depends on bank health, such as credit default swaps. Thus, interestingly, imperfect interventions have the upside that they might keep the information-revelation incentives of traders alive.

To see this, suppose the regulator cannot perfectly restore bank health to full solvency, but rather imperfectly reacts with $i_2(\theta) = \max\{F - \theta - \varepsilon(\theta), 0\}$, where $\varepsilon(\theta) > 0$ means there is underreaction relative to perfect intervention. This imperfection could be due to costs associated with interventions. Also assume $-1 < \varepsilon'(\theta) < 0$ for technical reasons. Then we have an asset value (anticipating the policy reaction if θ was revealed) of

$$v_2(\theta) = \begin{cases} F & \text{if } \theta \geq F \text{ (no intervention needed)}, \\ \theta & \text{if } F > \theta \geq F - \varepsilon(\theta) \text{ (needed intervention does not occur)}, \\ F - \varepsilon(\theta) & \text{if } F - \varepsilon(\theta) > \theta \text{ (imperfect intervention occurs)}. \end{cases} \tag{8}$$

This function is plotted in Figure 3. Hence, since $F - \varepsilon(\theta)$ is invertible in θ (assured by $\varepsilon'(\theta) < 0$ and $\varepsilon'(\theta) > -1$, so that $F - \varepsilon(\theta) = \theta$ has a unique solution), the price function $p(\theta) = v(\theta)$ reveals the realization of θ in the subintervals of θ where intervention occurs, namely, in $\theta < F - \varepsilon(\theta)$ (see Figure 3). This is a partially revealing equilibrium. The fact that θ is not revealed for $\theta \geq F$ is not critical since the regulator would not want to change her action in this subinterval ($i(\theta) = 0$ for $\theta \geq F$), thus she does not need to know the exact θ-realization. In summary, imperfect reactions from the policymaker can make prices more informative, which might ultimately be better for the policymaker, while perfect interventions may destroy the basis upon which they operate.

In summary, while the empirical evidence suggests that the market as a whole possesses useful information for banking authorities, this information may not be revealed to the same degree if the authorities rely more on this information in the future, as it can damage trader incentives. A general point here is that market information

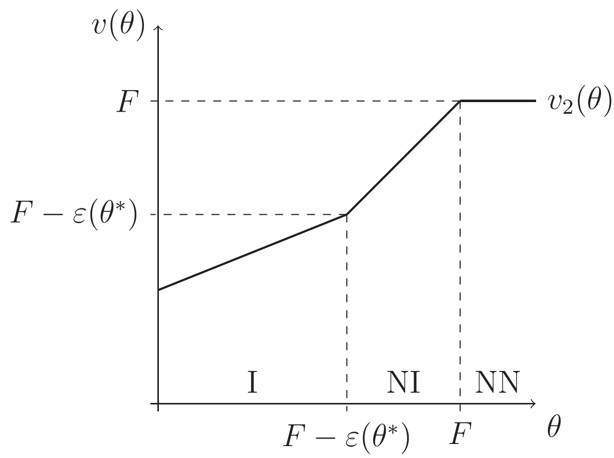

Fig 3. The Asset Value $v_2(\theta)$ Given Imperfect Intervention $i_2(\theta)$ by the Regulator is Invertible in Bank Health θ for $\theta < F$.

NOTES: θ^* is defined as $F - \varepsilon(\theta^*) = \theta^*$. I is the region of θ where (imperfect) Intervention occurs, NI is the region where No Intervention occurs but would be required, and NN is the region where intervention is Not Necessary.

cannot facilitate interventions that achieve full price stability, since the information needed to intervene appropriately can only be transmitted by changing prices.

2.3 Information from Asset Prices for Monetary Policy

It is well known that CBs monitor asset prices, which can potentially reveal information about inflation expectations or future inflation shocks (e.g., Bernanke and Woodford 1997). Moreover, a large literature on Taylor rules finds that CBs react to asset prices, housing prices, or oil prices (e.g., Rigobon and Sack 2003; L'œillet and Licheron 2012; Finocchiaro and Heideken 2013). Bernanke and Woodford (1997) showed how a self-defeating prophecy (my term) can arise if the CB is inflation targeting and reacts to trader information incorporated in asset prices. The next example briefly makes their argument in simple terms and (new) shows how a different asset can fix the problem.

EXAMPLE 5. Suppose next period's inflation rate is given by $\pi = \theta - i$, where $\theta \in \mathbb{R}$ is a fundamental of the macroeconomy and i is the CB's policy instrument (interest rate). The CB is inflation targeting, so it has a target of $t \in \mathbb{R}$, which can be represented with the utility function $u(\theta, i) = -(\pi - t)^2$.

Suppose traders observe the state of the economy θ, that is, $s_j = \theta$. Moreover, suppose that there is an asset in the financial market whose (real) value depends on the future inflation rate, like a government bond: $a_1(\theta, i) = f(\pi)$ with $f(.)$ being strictly decreasing. Bernanke and Woodford (1997) showed there is no REE in this case. Using the methods from this paper, the optimal policy function is $i(\theta) = \theta - t$.

The expected asset value given this policy is $v_1(\theta) = f(\pi = t)$, which is not invertible in θ (indeed, it is constant). Thus, there is no informative equilibrium, because if the CB were to learn θ, then the future inflation rate would equal the target rate, and any asset whose value depends on the future inflation rate would have a constant value (i.e., not change in θ). But if the asset value does not change in θ, then a price function that does change in θ (as required by full revelation) cannot be an equilibrium.

As in the case of the deadline corporate prediction market, using a different asset with a different underlying can fix the problem of the self-defeating prophecy. Consider the asset $a_2(\theta, i) = i$ instead. The expected asset value given optimal policy based on trader information is $v_2(\theta) = i(\theta) = \theta - t$, which is invertible in θ. Thus, by Proposition 1, a fully revealing equilibrium exists. In other words, rather than using assets whose value depends on the future inflation rate, an inflation targeting CB should look at interest rate futures. This is because the prices of these securities can at the same time reveal information to the CB and be consistent with the policy reaction to this information. However, by Corollary 1, there also exists a nonrevealing equilibrium for a_2, so the desirable equilibrium is not the only one.

Note that the model so far considered the existence of one asset whose value might depend on (θ, i). Generalizing, if assets a_1 and a_2 exist in parallel, then in equilibrium the price of a_2 would fully reveal θ, and then a_1 could be priced at $v_1(\theta)$ even though it is not invertible, since the trader information θ is already revealed by the other asset. A similar point with an additional asset is made in Bond et al. (2010). Thus, since interest rate futures do exist in most currencies, the problem raised by Bernanke and Woodford (1997) may not be that severe. In general, a self-defeating prophecy only arises if *none* of the existing assets have an invertible v-function. Consequently, the self-defeating prophecy problem could be viewed as a consequence of incomplete markets.

2.4 How to Design Assets that Promote Information Revelation

The previous subsections discussed specific applications whose lessons may not always carry over to other applications. This section gives more general results on asset design.

A natural question to ask is whether there always exists an asset $a(.)$ for every policymaker utility function $u(.)$ so that there is a fully revealing equilibrium (i.e., so that $v(\boldsymbol{s})$ is invertible). In the case where the combined trader information is $\boldsymbol{s} = \theta$, and assuming away problems of contractibility for now, the answer is trivially "yes." The asset value function $a(\theta, i) = \theta$ is always invertible in $\boldsymbol{s} = \theta$ for any policymaker utility function $u(\theta, i)$. Hence, since the problem of self-defeating prophecies results from the policy reaction to the information revealed by asset prices, we can fix the problem by making asset values independent of the policy (i.e., revert back to the standard case of exogenous asset values).

From a practical perspective, this simple answer may not be satisfactory, as θ may not be *ex post* verifiable or contractible. Hence, a financial asset—essentially a contract promising payoffs depending on certain contingencies—could not directly condition on θ. For example, θ may be the CEO quality or effort that is never directly observed but affects the firm (and stock) value. Thus, the question whether there always exists an asset that supports information revelation cannot be answered at this level of generality, since the set of feasible assets depends on the specific application.

Still, the proposition in this section provides some general guidelines that can serve as a recommendation for asset design. Moving beyond the trivial case where an asset can directly condition on θ, assume now that an outcome $o : \Theta \times I \to \mathbb{R}$ is contractible. This outcome is exogenous and cannot be affected by the asset designer. Still, we can design an asset $A : \mathbb{R} \to \mathbb{R}$ which maps the observable outcome $o(\theta, i)$ into an asset value.

Denote the set of all invertible functions $A : \mathbb{R} \to \mathbb{R}$ by $\mathcal{A}$ and the set of noninvertible functions by $\mathcal{A}'$. Moreover, denote the set of strictly increasing functions by $\overline{\mathcal{A}} \subset \mathcal{A}$ and the set of strictly decreasing functions by $\underline{\mathcal{A}} \subset \mathcal{A}$.

PROPOSITION 3 (Asset design and full revelation). *Consider the case of an observable outcome $o : \Theta \times I \to \mathbb{R}$ where the asset designer can choose any $a(\theta, i) = A(o(\theta, i))$.*

i. Suppose the combined trader information is perfect ($\mathbf{s} = \theta$). If a noninvertible function $A' \in \mathcal{A}'$ allows full revelation and accurate prices, then so does any invertible function $A \in \mathcal{A}$, but the converse does not hold.

ii. Suppose $o(\theta, i(\mathbf{s}))$ conditional on $\mathbf{s}$ can be ranked along $\mathbf{s}$ in a first-order stochastic dominance sense, that is, $o(\theta, i(\mathbf{s}))|\mathbf{s}$ first-order stochastic dominates $o(\theta, i(\mathbf{s}'))|\mathbf{s}'$ or vice versa for any $\mathbf{s} \neq \mathbf{s}' \in \mathbf{S}$. Then any strictly increasing $\overline{A} \in \overline{\mathcal{A}}$ or strictly decreasing $\underline{A} \in \underline{\mathcal{A}}$ allows for fully revealing and accurate prices, but any nonmonotone $A \notin (\overline{\mathcal{A}} \cup \underline{\mathcal{A}})$ (and in particular any $A' \in \mathcal{A}'$) might not.

iii. Suppose $\mathbf{s}$ and θ have discrete distributions. Then there exists an asset payoff function $A : \mathbb{R} \to \mathbb{R}$ with a fully revealing REE if and only if the probability mass function of $o(\theta, i(\mathbf{s}))$ conditional on $\mathbf{s}$ is distinct for all $\mathbf{s} \in \mathbf{S}$.

Part (i) shows that invertible asset payoff functions A are never worse, but may be better suited to promote fully revealing prices than noninvertible ones (in the sense that a fully revealing equilibrium exists) if $f(s|\theta)$ is such that the combined trader information reveals θ, that is, if $\mathbf{s} = \theta$. Thus, in this case, we may say that invertible asset payoff functions are "weakly better" in terms of information revelation. This is because a noninvertible asset payoff function may "bunch" or "pool" several states $\theta \in \Theta$ in a single asset value, thereby making it impossible to infer the state from the asset value, precluding invertibility of $v(\mathbf{s})$.

EXAMPLE 6. To illustrate part (i), consider prediction markets (e.g., Wolfers and Zitzewitz 2004, Page and Siemroth 2017) which typically feature Arrow-securities of the form

$$A(o(\theta, i)) = \begin{cases} 1 & \text{if } o(\theta, i) \geq T, \\ 0 & \text{if } o(\theta, i) < T \end{cases} \tag{9}$$

for a threshold $T \in \mathbb{R}$, because the prices of these assets have a natural interpretation as market probability forecasts. For example, such an asset could be used to forecast future sales $o(\theta, i) = \theta \cdot i$ of a product in a specific time period. But if the range of possible demand is $\Theta = \mathbb{R}_+$, with $i \geq 0$ for example being marketing activity, then the state space dimension is infinitely larger than the asset value dimension. Consequently, no fully revealing equilibrium exists, and such a prediction market can at best reveal whether or not predicted sales are above the target, not what the underlying demand for the product is exactly.

Part (i) states that invertible asset payoff functions are weakly better. For example, an asset that pays off linearly with sales would work better, such as $A_1(o(\theta, i)) = o(\theta, i)$. Linear assets have in fact been used in Ford motor company sales forecasts (see Cowgill and Zitzewitz 2015). Or if the asset value should be bounded, then $A_2(o(\theta, i)) = 1 - 1/\exp(o(\theta, i))$ is a possibility. Importantly, both of these asset value functions are strictly increasing, and as long as the marketing activity $i(\theta)$ is such that $o(\theta, i(\theta))$ is strictly increasing, then these assets can fully reveal demand θ while (9) cannot for any $i(\theta)$.

Intuition might suggest that invertible asset payoff functions $A \in \mathcal{A}$ should always be at least as good as noninvertible asset payoff functions $A' \in \mathcal{A}'$, since the asset value given optimal policy based on trader information $\boldsymbol{s}$ should be invertible (Proposition 3) for a fully revealing equilibrium to exist. However, this intuition is not correct in general, as can be shown by way of counterexamples if the combined trader information is $\boldsymbol{s} \neq \theta$.[8]

Still, part (ii) shows that strictly increasing or strictly decreasing asset payoff functions A are weakly better than other functions (in the sense that a fully revealing equilibrium exists) if the distribution of outcome $o(\theta, i(\boldsymbol{s}))$ conditional on $\boldsymbol{s}$ has a first-order stochastic dominance ranking, that is, if $o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}$ first-order stochastic dominates $o(\theta, i(\boldsymbol{s}'))|\boldsymbol{s}'$ or vice versa for any $\boldsymbol{s} \neq \boldsymbol{s}' \in \boldsymbol{S}$. In this case, any strictly monotone asset payoff function A preserves the ordering of conditional means depending on $\boldsymbol{s} \in \boldsymbol{S}$. This result is more general than the previous one when it comes to trader information structures, because it holds true even when the combined trader information is not equal to the state, but still induces a first-order stochastic dominance ordering on outcomes.

EXAMPLE 7. To illustrate part (ii), let us revisit the Bond et al. (2010) application but change the information structure. Let the observable outcome be the firm's balance sheet total $o(\theta, i) = \theta + T(\theta) \cdot i$ where θ are the financial conditions of the firm, $i \in \{0, 1\}$ is a costly but value-increasing intervention, and $T(\theta) = \beta\theta$ is a linear function representing the effect of the intervention depending on the state of the firm. Bond

8. Consider this simple artificial example. Suppose θ is continuously distributed so that $\boldsymbol{s} \in (0, 1)$ and $o(\theta, i(\boldsymbol{s}))|\boldsymbol{s} \sim \mathcal{N}(0, \boldsymbol{s})$. With the invertible asset payoff function $A(x) = x$, the conditional expected asset value is $\mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s})))|\boldsymbol{s}] = \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}] = 0 \forall \boldsymbol{s} \in \boldsymbol{S}$. But with the noninvertible function $A'(x) = x^2$, the conditional expected asset value is $\mathbb{E}_\theta[A'(o(\theta, i(\boldsymbol{s})))|\boldsymbol{s}] = \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}))^2|\boldsymbol{s}] = \boldsymbol{s}$ (since $\mathbb{E}[x^2] = \mathrm{Var}(x) + \mathbb{E}[x]^2 = \mathrm{Var}(x)$), which fully reveals the trader information $\boldsymbol{s}$.

et al. (2010) assume $0 > T'(\theta) > -1$ (positive effect of intervention weakens as θ increases), hence $0 > \beta > -1$.

Suppose traders collectively possess the information $\boldsymbol{s} = \theta + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ and $\sigma_\varepsilon^2 > 0$, so that the combined trader information is a noisy signal of conditions θ. The common prior distribution is $\theta \sim \mathcal{N}(\mu, \sigma^2)$ with $\sigma^2 > 0$.

Suppose further that the decision maker, who does not have information about θ, wishes to maximize firm value subject to the private intervention cost, $u(\theta, i) = \theta + (T(\theta) - C)i$, with $C > 0$. If trader information $\boldsymbol{s}$ was revealed, the optimal policy $i(\boldsymbol{s})$ would be to intervene whenever

$$\begin{aligned} &\mathbb{E}[\theta + (T(\theta) - C)|\boldsymbol{s}] \geq \mathbb{E}[\theta|\boldsymbol{s}] \iff \mathbb{E}[\theta|\boldsymbol{s}]\beta \geq C \\ \iff & \frac{\mu/\sigma^2 + \boldsymbol{s}/\sigma_\varepsilon^2}{1/\sigma^2 + 1/\sigma_\varepsilon^2}\beta \geq C \iff \boldsymbol{s} \leq \frac{\sigma_\varepsilon^2 C(1/\sigma^2 + 1/\sigma_\varepsilon^2)}{\beta} - \mu\sigma_\varepsilon^2/\sigma^2 =: \hat{\boldsymbol{s}}, \end{aligned} \tag{10}$$

where the last inequality changes because $\beta < 0$. The outcome $o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}$ is normally distributed and trader information $\boldsymbol{s}$ shifts the mean. Consequently, the firm stock value $A(o(\theta, i)) = o(\theta, i)$ if the trader information $\boldsymbol{s}$ were known is

$$v(\boldsymbol{s}) = \begin{cases} \mathbb{E}[\theta|\boldsymbol{s}] & \text{if } \boldsymbol{s} > \hat{\boldsymbol{s}}, \\ \mathbb{E}[\theta|\boldsymbol{s}](1 + \beta) & \text{if } \boldsymbol{s} \leq \hat{\boldsymbol{s}}, \end{cases} \tag{11}$$

with the explicit expression for $\mathbb{E}[\theta|\boldsymbol{s}]$ in (10). If the parameters $\mu, \sigma^2, \sigma_\varepsilon^2$ are such that $\mathbb{E}[\theta|\hat{\boldsymbol{s}}] < 0$, then $v(\boldsymbol{s})$ is not invertible, so there is no full revelation in equilibrium. Graphically $v(\boldsymbol{s})$ looks similar to $v(\theta)$ in Figure 2(b). If, on the other hand, $\mathbb{E}[\theta|\hat{\boldsymbol{s}}] \geq 0$, then $o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}$ has a first-order stochastic dominance ordering: larger realizations of $\boldsymbol{s}$ increase the mean of the posterior distribution over the firm balance. Part (ii) then implies that *all* strictly monotone asset payoff functions, for example, assets which pay strictly more as the firm's balance total increases, have a fully revealing equilibrium.

If θ and $\boldsymbol{s}$ have discrete distributions, then part (iii) of Proposition 3 states that an asset payoff function $A(.)$ that yields a fully revealing equilibrium always exists if and only the conditional probability mass functions $f(o(\theta, i(\boldsymbol{s}))|\boldsymbol{s})$ are distinct for all $\boldsymbol{s} \in \boldsymbol{S}$. Informally, the result says that if different trader information induces different distributions over observable outcomes (assuming the trader information was known to the policymaker), then an asset payoff function with fully revealing equilibrium is guaranteed to exist, otherwise no such asset exists.

If the outcomes given optimal policy $i(\boldsymbol{s})$ do not have different distributions for different $\boldsymbol{s}$, then the conditional expected values of the outcomes cannot differ, and neither can any expected asset value that conditions on these outcomes. For example, if outcomes can take values $\{1, 2\}$, and $\Pr(o(\theta, i(\boldsymbol{s})) = 1|\boldsymbol{s} = \boldsymbol{s}_k) = \Pr(o(\theta, i(\boldsymbol{s})) = 1|\boldsymbol{s} = \boldsymbol{s}_l)$ for $\boldsymbol{s}_k \neq \boldsymbol{s}_l$, then any transformation A that maps from outcomes into asset values will yield the same expected asset value whether $\boldsymbol{s} = \boldsymbol{s}_k$ or $\boldsymbol{s} = \boldsymbol{s}_l$, since both the outcome values and the probability weights are the same. Thus, the $v(\boldsymbol{s})$ invertibility condition from Proposition 3 cannot be fulfilled by *any* asset value function. The Bernanke and Woodford (1997) setting with $\boldsymbol{s} = \theta$ (see Section 2.3) also violates this condition, since the outcome (inflation) given $\boldsymbol{s}$ is constant across $\boldsymbol{s}$, because the inflation targeting CB would always hit the inflation target, so that the probability mass function of outcomes is the same for every s. Hence, *no* asset that pays off depending on the inflation rate could possibly reveal $\boldsymbol{s}$.

If instead $\Pr(o(\theta, i(\boldsymbol{s})) = 1|\boldsymbol{s} = \boldsymbol{s}_k) > \Pr(o(\theta, i(\boldsymbol{s})) = 1|\boldsymbol{s} = \boldsymbol{s}_l)$, then $A(x) = x$ works as asset payoff function, since then

$$\mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s}_l)))] = \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}_l))] > \mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s}_k)))] = \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}_k))] \quad (12)$$

so that a price function equal to the expected asset values is fully revealing. Hence, part (iii) is an existence result: if we can find a contractible outcome $o(\theta, i)$ which does not suffer from a self-defeating prophecy problem, then we can *always* find an asset which pays off depending on the outcome and has a fully revealing equilibrium.

3. CONCLUSION

This paper develops a general model that captures many situations where a policymaker maker tries to infer information from financial market prices in order to improve real decisions, and these decisions might affect asset values. The main result establishes a necessary and sufficient condition for the existence of a fully revealing and accurate (nonbiased) price function, which is also necessary and sufficient for the existence of a fully revealing REE where the policymaker obtains all trader information. Thus, the main result identifies all situations where forward looking traders reveal their information via trading and policymakers can use this information, and all situations where information revelation is not incentive-compatible. Moreover, if a fully revealing equilibrium exists, then it is Pareto-optimal, so this setting is one where "informational efficiency" corresponds to Pareto-efficiency, and a functioning financial market improves outcomes in the real sector.

The model unifies several applications from the literature and identifies the common cause of noninformative prices: self-defeating prophecies. In this analogy, risk-neutral traders in a competitive market try to forecast the future asset value and buy the asset up to this price. If the necessary and sufficient condition is not fulfilled, then the economic environment is such that it is impossible to correctly forecast the asset value, since these forecasts (prices) trigger policy reactions that affect asset values and falsify the forecasts.

This problem arises because asset values are endogenous to the policy reaction, which does not occur in standard financial market models where asset values are exogenous. The problem of self-defeating prophecies shows that financial markets may not be informationally efficient even if all traders are perfectly rational, perfectly informed, and markets are perfectly competitive and without noise.

The paper demonstrates how the results are useful for the design of assets that promote information revelation. In particular, it shows how "deadline assets" in corporate prediction markets suffer from the self-defeating prophecy problem and how a different design of assets fixes the problem. The analysis also shows that market information can be useful for banking regulation, but at the same time market-based interventions can never perfectly stabilize the banking sector.

APPENDIX A: PROOFS

Proof of Lemma 1. Necessity. To show: Accuracy implies $p(\boldsymbol{s}) = v(\boldsymbol{s})$. First, abusing notation, $\mathbb{E}[a(\theta, i(p))|p(\boldsymbol{s}) = p, s_j] = \mathbb{E}[a(\theta, i(\boldsymbol{s}))|p(\boldsymbol{s}) = p, s_j]$ since prices are fully revealing by assumption. Second,

$$\mathbb{E}[a(\theta, i(\boldsymbol{s}))|p(\boldsymbol{s}) = p, s_j] = \mathbb{E}[\mathbb{E}[a(\theta, i(\boldsymbol{s}))|\boldsymbol{s}]|p(\boldsymbol{s}) = p, s_j] = \mathbb{E}[v(\boldsymbol{s})|p(\boldsymbol{s}) = p, s_j] \tag{A1}$$

by the law of iterated expectations. Next, plugging this term into the definition of accurate prices, $p(\boldsymbol{s}) = \mathbb{E}[v(\boldsymbol{s})|p(\boldsymbol{s}) = p, s_j]$ for all j, and taking the conditional expectation on both sides, $\mathbb{E}[p(\boldsymbol{s})|\boldsymbol{s}] = \mathbb{E}[\mathbb{E}[v(\boldsymbol{s})|p(\boldsymbol{s}) = p, s_j]|\boldsymbol{s}]$. Again using iterated expectations yields $p(\boldsymbol{s}) = v(\boldsymbol{s})$.

Sufficiency. To show: $p(\boldsymbol{s}) = v(\boldsymbol{s})$ implies accurate prices. This immediately follows:

$$\mathbb{E}[a(\theta, i(p))|p(\boldsymbol{s}) = p, s_j] = \mathbb{E}[a(\theta, i(p))|v(\boldsymbol{s}) = p, s_j] = v(\boldsymbol{s}) = p(\boldsymbol{s}) \quad forall j. \tag{14}$$

□

Proof of Theorem 1. Necessity. To show: Full revelation and accurate prices imply invertibility of $v(\boldsymbol{s})$. By Lemma 1, full revelation and accurate prices imply $p(\boldsymbol{s}) = v(\boldsymbol{s})$. Since full revelation implies invertibility of $p(\boldsymbol{s})$, the equality implies that $v(\boldsymbol{s})$ is invertible.

Sufficiency. To show: Invertibility of $v(\boldsymbol{s})$ implies fully revealing and accurate prices. By construction: The price function $p(\boldsymbol{s}) = v(\boldsymbol{s})$ is fully revealing since $v(\boldsymbol{s})$ is invertible, and accurate (Lemma 1). □

Proof of Proposition 1. Necessity. To show: Existence of a fully revealing REE implies invertibility of $v(\boldsymbol{s})$. If $\boldsymbol{s}$ is revealed, then the policy reaction is $i(\boldsymbol{s})$ and the expected asset value is $v(\boldsymbol{s})$. Since $\boldsymbol{s}$ is also revealed to all traders, which is a sufficient statistic of all trader information, the net demands of all traders are

$$x_j(p, s_j) = \begin{cases} w/p & \text{if } p(\boldsymbol{s}) < v(\boldsymbol{s}) \\ \in [-1, w/p] & \text{if } p(\boldsymbol{s}) = v(\boldsymbol{s}) \\ -1 & \text{if } p(\boldsymbol{s}) > v(\boldsymbol{s}). \end{cases} \tag{A2}$$

Consequently, since we have an REE, we must also have $p(\boldsymbol{s}) = v(\boldsymbol{s})$, which is the only price function that clears the market for all $\boldsymbol{s} \in \boldsymbol{S}$. And since the equilibrium is fully revealing (i.e., $p(\boldsymbol{s})$ is invertible), $v(\boldsymbol{s})$ must be invertible.

Sufficiency. To show: Invertibility of $v(\boldsymbol{s})$ implies existence of a fully revealing REE. By construction: Set $p(\boldsymbol{s}) = v(\boldsymbol{s})$, which fully reveals $\boldsymbol{s}$, leads to policy reaction $i(\boldsymbol{s})$ with expected asset value $v(\boldsymbol{s})$ and therefore clears the market. □

PROOF OF PROPOSITION 2. Sufficiency. If there exists such a partition $\mathcal{P}$, then $p(\theta) = \mathbb{E}_\theta[a(\theta, i(\mathcal{I}))|\theta \in \mathcal{I}]$ is a partially revealing equilibrium which reveals that $\theta \in \mathcal{I}$. Traders clear the market for this price function, since prices equal expected asset values from the perspective of traders for every $\theta \in \Theta$ (guaranteed by conditions 2 and 3). Moreover, since there exists $|\mathcal{I}| > 1$ (condition 1), this price function is partially revealing.

Necessity. If condition 1 fails, then by definition the price function is fully and not partially revealing. If condition 2 fails, then there exists no partially revealing price function which induces a policy reaction $i(\mathcal{I})$ such that traders clear the market for that price function, because partial revelation requires invertibility in $\mathcal{I}$ whereas market clearing requires noninvertibility. If condition 3 fails, then there exists no partition with part $\mathcal{I} \in \mathcal{P}$ with more than one element such that $a(\theta, i(\mathcal{I})) = a(\theta', i(\mathcal{I}))$ for $\theta \neq \theta' \in \mathcal{I}$. Hence, in any partially revealing price function the expected asset value differs by state within some part $\mathcal{I}$ whereas the price remains constant, hence traders do not clear the market. □

PROOF OF PROPOSITION 3.

(i) We consider the case $\boldsymbol{s} = \theta$. Define $v(\boldsymbol{s} = \theta, A) := \mathbb{E}_\theta[A(o(\theta, i(\theta)))|\theta] = A(o(\theta, i(\theta)))$. Thus, since $\boldsymbol{s} = \theta$ removes all uncertainty in the realization of θ, we can drop the expectations. First, given some pair $\theta \neq \theta'$, we will show that if $v(\theta, A') \neq v(\theta', A')$, then we must also have $v(\theta, A) = v(\theta', A)$ for any $A \in \mathcal{A}$. Formally,

$$v(\theta, A') \neq v(\theta', A') \Rightarrow v(\theta, A) \neq v(\theta', A) \forall A \in \mathcal{A}. \tag{A3}$$

To show (16) holds, note that for any function $A' \in \mathcal{A}'$,

$$A'(o(\theta, i(\theta))) \neq A'(o(\theta', i(\theta'))) \Rightarrow o(\theta, i(\theta)) \neq o(\theta', i(\theta')). \tag{A4}$$

Moreover, because $\mathcal{A}$ is the set of invertible functions,

$$o(\theta, i(\theta)) \neq o(\theta', i(\theta')) \iff A(o(\theta, i(\theta))) \neq A(o(\theta', i(\theta'))) \forall A \in \mathcal{A}. \tag{A5}$$

Combining (17) and (18) results in (16).

Now if (16) holds for all possible $\theta \neq \theta' \in \Theta$—which can only happen if the noninvertibilities are outside the support of $o(\theta, i(\theta))$—then by Theorem 1, the asset A' has a fully revealing and accurate price function. Thus, (17) and (18) imply that if (16) holds for all possible $\theta \neq \theta' \in \Theta$, then any invertible asset function $A \in \mathcal{A}$ must also have a fully revealing and accurate price function. We still need to show that $A \in \mathcal{A}$ can allow for full revelation and accurate prices but $A' \in \mathcal{A}'$ might not. This is immediately obvious from

$$A(o(\theta, i(\theta))) \neq A(o(\theta', i(\theta'))) \nRightarrow A'(o(\theta, i(\theta))) \neq A'(o(\theta', i(\theta'))). \tag{A6}$$

(ii) By first-order stochastic dominance in $\boldsymbol{s}$, $\mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}] > \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}'))|\boldsymbol{s}']$ or inequality reversed for any $\boldsymbol{s} \neq \boldsymbol{s}' \in \boldsymbol{S}$. Clearly, any strictly increasing function $\overline{A} \in \overline{\mathcal{A}}$ (and similarly any strictly decreasing function $\underline{A} \in \underline{\mathcal{A}}$) of the outcomes preserves this inequality:

$$\mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}] > \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}'))|\boldsymbol{s}'] \Rightarrow \mathbb{E}_\theta[\overline{A}(o(\theta, i(\boldsymbol{s})))|\boldsymbol{s}] > \mathbb{E}_\theta[\overline{A}(o(\theta, i(\boldsymbol{s}')))|\boldsymbol{s}']. \tag{A7}$$

Thus, rewriting $\mathbb{E}_\theta[\overline{A}(o(\theta, i(\boldsymbol{s})))|\boldsymbol{s}] = v(\boldsymbol{s})$, first-order stochastic dominance of $o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}$ and strictly increasing/decreasing asset payoff functions guarantee an invertible $v(\boldsymbol{s})$ function, which guarantees fully revealing and accurate prices (Theorem 1).

Clearly, for any nonmonotone functions $A \notin (\overline{\mathcal{A}} \cup \underline{\mathcal{A}})$, the inequality is not necessarily preserved and hence these do not guarantee existence of fully revealing and accurate prices. Formally, for $A \notin (\overline{\mathcal{A}} \cup \underline{\mathcal{A}})$,

$$\mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}))|\boldsymbol{s}] > \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}'))|\boldsymbol{s}'] \nRightarrow \mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s})))|\boldsymbol{s}] > \mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s}')))|\boldsymbol{s}']. \tag{A8}$$

(iii) Given that $\theta \in \{\theta_1, \ldots, \theta_L\}$ and $\boldsymbol{s} \in \{\boldsymbol{s}_1, \ldots, \boldsymbol{s}_K\}$ have discrete distributions, $o(\theta, i(\boldsymbol{s}))$ also has a discrete distribution. Denote the number of different values of $o(\theta, i(\boldsymbol{s}))$ in the support by $1 < M \leq KL$. For the remainder of this proof, I abbreviate $o(\theta, i(\boldsymbol{s}))$ as o (the generic element) or o_m (a specific element) whenever possible.

First, rewrite the probability mass function of the outcomes given $\boldsymbol{s}$ in the form

$$f(o|\boldsymbol{s} = \boldsymbol{s}_k) = \sum_{m=1}^{M} \alpha_{km} b_m(o), \tag{A9}$$

with $\alpha_{km} = f(o_m|\boldsymbol{s}_k)$ and $b_m = 1$ if and only if $o = o_m$ and $b_m = 0$ otherwise. Second, write the asset value function as

$$A(o) = \sum_{n=1}^{M} \beta_n b_n(o) \tag{A10}$$

with the β_n's to be determined later and $b_n(o) = 1$ if and only if $o = o_n$ as before. Since $n = 1, \ldots, M$, we specify one coefficient for every realization of o, thus specifying the vector $\beta = (\beta_1, \ldots, \beta_M)$ is equivalent to specifying the function $A(o)$ for all possible values of o.

Using these definitions to rewrite the conditional expected asset value,

$$\begin{aligned}\mathbb{E}_\theta[A(o)|\boldsymbol{s}_k] &= \sum_{m=1}^{M} A(o) f(o|\boldsymbol{s}_k) = \sum_{m=1}^{M}\left[\left(\sum_{n=1}^{M}\beta_n b_n(o)\right)\left(\sum_{q=1}^{M}\alpha_{kq} b_q(o)\right)\right]\\ &= \sum_{m=1}^{M}\beta_m \alpha_{km}, \end{aligned}\tag{A11}$$

since $b_n(o) = b_q(o) = 1$ if and only if $m = n$ and $m = q$. Using the definition of α_{km}, $\sum_{m=1}^{M}\beta_m\alpha_{km} = \sum_{m=1}^{M}\beta_m f(o_m|\boldsymbol{s}_k)$, which is the conditional expectation with β_m being the value of $A(o_m)$.

Now we can express these objects in matrix form. Define the $K \times M$ matrix X as containing the α_{km} coefficients. Thus, the matrix product of X and the vector β, $X\beta$, equals the vector of conditional expectations V,

$$X\beta = \begin{bmatrix}\alpha_{11} & \alpha_{12} & \alpha_{13} & \dots & \alpha_{1M}\\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \alpha_{K1} & \alpha_{K2} & \alpha_{K3} & \dots & \alpha_{KM}\end{bmatrix}\begin{bmatrix}\beta_1\\ \vdots \\ \beta_M\end{bmatrix} = \begin{bmatrix}\mathbb{E}_\theta[A(o)|\boldsymbol{s}_1]\\ \vdots \\ \mathbb{E}_\theta[A(o)|\boldsymbol{s}_K]\end{bmatrix} =: V. \tag{A12}$$

Recall that $\alpha_{km} = f(o_m|\boldsymbol{s}_k)$. By assumption, these probability mass functions are distinct for all $\boldsymbol{s} \in \boldsymbol{S}$, thus all rows of matrix X are distinct. Hence, mathematically, the problem translates into the following question: Under which conditions can we find a vector β (representing function $A(o)$), such that $V = X\beta$—where X has distinct rows—has distinct elements (so that $v(\boldsymbol{s})$ is invertible)?

The answer is that such a β (and thus $A(.)$) always exists. To see this, take two rows of X, r_i and r_j with $i \neq j$. The conditional expectations given $\boldsymbol{s}_i$ and $\boldsymbol{s}_j$ are the same if and only if

$$r_i\beta = r_j\beta \iff (r_i - r_j)\beta = 0, \tag{A13}$$

that is, if and only if the difference between the rows is orthogonal to β. Since all rows of X are distinct, $(r_i - r_j)$ is not the zero vector, and thus any $\beta \in \mathbb{R}^M$ except a negligible set (on a hyperplane of $\mathbb{R}^M$) will not be orthogonal to $(r_i - r_j)$ and yield different conditional expectations. And since there are finitely many pairwise comparisons of r_i and r_j, the subset of $\mathbb{R}^M$ from which we can select β is $\mathbb{R}^M$ without up to K hyperplanes of $\mathbb{R}^M$, which is a nonempty set. However, from this, it is also obvious that there exists no β such that V has distinct elements if, for any $i \neq j$, $r_i = r_j$. $\square$

LITERATURE CITED

Allen, Beth. (1981) "Generic Existence of Completely Revealing Equilibria for Economies with Uncertainty when Prices Convey Information." *Econometrica*, 49, 1173–99.

Berger, Allen N., Sally M. Davies, and Mark J. Flannery. (2000) "Comparing Market and Supervisory Assessments of Bank Performance: Who Knows What When?." *Journal of Money, Credit and Banking*, 32, 641–67.

Bernanke, Ben S., and Michael Woodford. (1997) "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit, and Banking*, 29, 653–84.

Birchler, Urs W., and Matteo Facchinetti. (2007) "Self-Destroying Prophecies? The Endogeneity Pitfall in Using Market Signals as Triggers for Prompt Corrective Action." Working Paper. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.493.1086&rep=rep1&type=pdf.

Bond, Philip, and Itay Goldstein. (2015) "Government Intervention and Information Aggregation by Prices." *Journal of Finance*, 70, 2777–812.

Bond, Philip, Itay Goldstein, and Edward Simpson Prescott. (2010) "Market-Based Corrective Actions." *Review of Financial Studies*, 23, 781–820.

Cowgill, Bo, and Eric Zitzewitz. (2015) "Corporate Prediction Markets: Evidence from Google, Ford, and Firm X." *Review of Economic Studies*, 82, 1309–41.

Dow, James, and Gary Gorton. (1997) "Stock Market Efficiency and Economic Efficiency: Is There a Connection?." *Journal of Finance*, 52, 1087–129.

Dow, James, Itay Goldstein, and Alexander Guembel. (2017) "Incentives for Information Production in Markets Where Prices Affect Real Investment." *Journal of the European Economic Association*, 15, 877–909.

Edmans, Alex, Itay Goldstein, and Wei Jiang. (2015) "Feedback Effects, Asymmetric Trading, and the Limits to Arbitrage." *American Economic Review*, 105, 3766–97.

Finocchiaro, Daria, and Virginia Queijo Heideken. (2013) "Do Central Banks React to House Prices?." *Journal of Money, Credit and Banking*, 45, 1659–83.

Flannery, Mark J., Joel F. Houston, and Frank Partnoy. (2010) "Credit Default Swap Spreads as Viable Substitutes for Credit Ratings." *University of Pennsylvania Law Review*, 158, 2085–123.

Foucault, Thierry, and Thomas Gehrig. (2008) "Stock Price Informativeness, Cross-Listings, and Investment Decisions." *Journal of Financial Economics*, 88, 146–68.

Goldstein, Itay, and Alexander Guembel. (2008) "Manipulation and the Allocational Role of Prices." *Review of Economic Studies*, 75, 133–64.

Goldstein, Itay, Emre Ozdenoren, and Kathy Yuan. (2013) "Trading Frenzies and Their Impact on Real Investment." *Journal of Financial Economics*, 109, 566–82.

Greenspan, Alan. (2001) "Harnessing Market Discipline." *Region*, 6–7.

Grossman, Sanford J. (1981) "An Introduction to the Theory of Rational Expectations under Asymmetric Information." *Review of Economic Studies*, 48, 541–59.

Gunther, Jeffery W., Mark E. Levonian, and Robert R. Moore. (2001) "Can the Stock Market Tell Bank Supervisors Anything They Don't Already Know?" *Federal Reserve Bank of Dallas: Economic and Financial Review*, 2–9.

Hayek, Friedrich A. (1945) "The Use of Knowledge in Society." *American Economic Review*, 35, 519–30.

HPC Wire. (2011) "Ford Motor Company Turns to Cloud-Based Prediction Market Software." https://www.hpcwire.com/2011/02/22/ford_motor_company_turns_to_cloud-based_prediction_market_software/

L'œillet, Guillaume, and Julien Licheron. (2012) "How Does the European Central Bank React to Oil Prices?." *Economics Letters*, 116, 445–7.

Ottaviani, Marco, and Peter Norman Sørensen. (2007) "Outcome Manipulation in Corporate Prediction Markets." *Journal of the European Economic Association*, 5, 554–63.

Page, Lionel, and Christoph Siemroth. (2017) "An Experimental Analysis of Information Acquisition in Prediction Markets." *Games and Economic Behavior*, 101, 354–78.

Prescott, Edward S. (2012) "Contingent Capital: The Trigger Problem." *Federal Reserve Bank of Richmond Economic Quarterly*, 98, 33–50.

Radner, Roy. (1979) "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices." *Econometrica*, 47, 655–78.

Rigobon, Roberto, and Brian Sack. (2003) "Measuring the Reaction of Monetary Policy to the Stock Market." *Quarterly Journal of Economics*, 118, 639–69.

Siemroth, Christoph. (2019) "The Informational Content of Prices When Policy Makers React to Financial Markets." *Journal of Economic Theory*, 179, 240–74.

Subrahmanyam, Avanidhar, and Sheridan Titman. (1999) "The Going-Public Decision and The Development of Financial Markets." *Journal of Finance*, 54, 1045–82.

Sundaresan, Suresh, and Zhenyu Wang. (2015) "On the Design of Contingent Capital with a Market Trigger." *Journal of Finance*, 70, 881–920.

Wolfers, Justin, and Eric Zitzewitz. (2004) "Prediction Markets." *Journal of Economic Perspectives*, 18, 107–26.

Woodford, Michael. (1994) "Nonstandard Indicators for Monetary Policy: Can Their Usefulness be Judged from Forecasting Regressions?." In *Monetary Policy*, edited by N. Gregory Mankiw, pp. 95–115. Chicago: University of Chicago Press.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Figure 1: Timeline with traders and policy maker.