Listen to the market, hear the best policy decision, but don't always choose it.*

David Reinstein† and Joon Song‡

November 16, 2014

Abstract

Policymakers must often decide whether to pursue a policy that has uncertain benefits. The response of asset markets to proposed policy changes can be a valuable source of information for policy-setting. However, policymakers must take into account that an informed trader may anticipate this and profitably manipulate the market. We show that it is optimal for policymakers to listen to asset markets, but they must commit (e.g., through "political capital") to sometimes pursuing a policy even when the expected welfare effects are negative. Surprisingly, allowing traders to short-sell can make it easier for policymakers to induce truth-telling actions.

1 Introduction

In some cases, a small set of private individuals may be better informed than policymakers about the effect of a potential policy. If these individuals buy and sell assets in publicly observable settings, policymakers might "listen to" these markets to help determine the best policies.

*We thank Holger Breinlich, Oliver Hart, David Hugh-Jones, Santiago Oliveros, Motty Perry, Jeroen van de Ven, Tianxi Wang, and various seminar participants for helpful tips. Mateusz Gatkowski offered helpful research assistance. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†Department of Economics, University of Essex, drein@essex.ac.uk

‡corresponding author, Department of Economics, Sungkyunkwan University, joonsong.econ@gmail.com

We offer an example to motivate this idea. Suppose the Department of Energy (DOE) is considering spending \$100 million on an R&D subsidy to spur innovation and increase profits in the solar energy sector. The subsidy mainly targets a new photovoltaic cell being developed by a firm "Soylent" (SOY). However, only a few engineers working for the large investment firm Golden Sax (GS) have the expertise to judge whether this new cell is technically viable and thus profitable, or technically difficult and thus less likely to be profitable. Suppose the DOE announced that the funding had attained preliminary (but not yet final) institutional approval. If SOY immediately rose substantially, the DOE might attribute this to a GS purchase on the advice of their engineers, implying the project is likely viable, and decide to fund it. If, instead, SOY did not rise much, the DOE might see this as news that the project had a small chance of success, and decide not to fund it. However, this "naive listening" may give GS an incentive to purchase SOY even if its engineers advised against the project. A GS purchase, by boosting SOY greatly, would lead to DOE funding, further increasing SOY's expected returns and yielding a profit for GS. With such behavior by GS the SOY price reaction would be uninformative. However, we will show later that there is always a more sophisticated way of *listening* that will make the price informative.

These concerns are relevant to a variety of contexts, where there is private information about a policy's impact, and where a successful policy will increase the value of an asset.[1] A government may consider a trade agreement to only merit signing if it will bring a large enough profit for domestic industries (perhaps the agreement involves a costly aid package for the partner country). We might consider a measure to protect intellectual property, or a policy subsidizing technical education, in the same vein: firms in the targeted industry may always benefit, but unless these benefits are large they may not justify the costs. There may also be private information about the impacts of policies whose success depends on significantly *reducing* rents in some industry, in order to increase consumer surplus or government revenue. For example, health-care reforms may aim to decrease rents in the insurance, drug, and hospital industries.[2] While *listening to the markets* may

[1]The policymaker's goal need not involve the firm's profit or the asset's value. An exogenous factor, e.g., the progress of technology or the potential supply of some natural resource, may happen to determine both whether a policy will be successful for the policymaker and whether a particular asset is profitable.

[2]For example, the USA's 2010 Affordable Care Act included an individual mandate and subsidies to purchase insurance as well as the establishment of "insurance exchanges" with regulations intended to reduce prices through

be helpful in each of these cases, this might be undermined by possibility that investors will act to conceal their information (i.e., "manipulate" asset prices).

Our setup is broadly as follows. At a precise point in time, a Policymaker (PM) publicly considers whether to implement a specific new policy. We say that the policy is "good" when it will increase the PM's welfare, and it is "bad" otherwise.[3] Next, an *Informed Investor* ($\mathcal{II}$) receives a *signal* of the policy's merit, and acts to maximize his profit. We model the $\mathcal{II}$ as a single investor, or, equivalently, a set of investors who coordinate and can thus drive the market.[4] $\mathcal{II}$ may trade with an *Uninformed Trader* (henceforth "UT").[5] We assume that the $\mathcal{II}$ has no inherent interest in the asset's value, but may buy the asset, short sell it, or do nothing. The asset's price will change in response to the $\mathcal{II}$'s choice; the PM observes this and decides whether or not to execute the policy.

For *listening* to always achieve the best policy, the interests of the PM and the relevant informed investor must be naturally aligned. For some parameters this will be the case, and $\mathcal{II}$ will reveal his private information through his profit-seeking behavior even knowing that the PM would certainly act on this information. However, this alignment condition need not be satisfied in general. For example, although the PM wants the policy to be executed if and only if it is *good*, an investor may *always* want the policy to be executed, or he may *never* want it to be executed.[6]

encouraging competition and transparent pricing. It was widely assumed that that the mandate would have a side effect of boosting insurance company profits, while the exchanges and regulations were meant to counter this. Presumably policymakers hoped in net to reduce (or at least not increase) insurance company profits, but opponents of the bill argued that it would be a giveaway to the insurance industry.

[3]Obviously, the PM's incentives may or may not be aligned with the public interest; this central issue of public choice is out of the scope of this paper; future work may extend our model to explicitly derive and consider a rent-seeking PM's payoff function.

[4]We do not focus on the case where information is diffusely held. In such a case the informed investors might fail to coordinate on the equilibrium that collectively yields them higher profit given the PM's commitment. We return to this discussion in the conclusion.

[5]If your eyes see $\mathcal{II}$ as the Roman numeral "two" you can think of him as a "type $\mathcal{II}$" investor, in comparison to the "type I" uninformed trader "UT". We will occasionally refer to him as just "$\mathcal{II}$" without the article "the" when it sounds better.

[6]It is often more intuitive to consider the incentives of the affected firm or asset-holder, as these will only depend on the direct impact of the policy on the asset. However, the party with inside information need not be the initial

Nonetheless, the PM may still be able to extract information from $\mathcal{II}$ if she can make a binding commitment that with a certain probability she will ignore the asset price and execute (or *not* execute) the policy anyways. This type of commitment changes $\mathcal{II}$'s incentives through two channels. First, it can reduce (or increase) the effect of the $\mathcal{II}$'s action on the probability the policy is executed, hence reducing or increasing the asset's relative expected value after each of $\mathcal{II}$'s actions. Second, if the probabilistic commitment is common knowledge, it will also affect the ex-ante expected returns to the asset, and hence increase or decrease the price the $\mathcal{II}$ must pay the UTs for the asset.

Returning to our first example, a *sophisticated* DOE could build its institutional process in a way that makes it less responsive to the SOY price. Suppose the DOE had publicly committed its resources and reputation to this project in a way that, even if SOY rose little after the preliminary approval, the DOE would, with some probability, feel compelled to fund the project even knowing it is likely a bad policy. For example, long before the preliminary approval the DOE might have hinted that the subsidy would be a central part of its new "Green Initiative"; or DOE might have begun signing contracts with firms involved in this technology. These prior actions may imply that with a certain probability, the DOE will fund the subsidy no matter what the market does. This would have boosted SOY even before the preliminary approval. The higher *initial* share price would decrease the expected profit for GS from buying and inducing the policy after the preliminary approval. This might deter GS from buying when its engineers advise against the project, but yet *not* deter GS from buying when its engineers endorse the project. This informative "separating behavior" is possible only because the funding adds *more* real value when the project is highly viable than when it is technically difficult, thus more profit for GS.

As in the above example, the mechanisms we propose are *indirect*: the PM does not directly pay the $\mathcal{II}$ as a function of his revealed signal and the outcome. Instead, she uses existing asset markets as a tool, and the efficiency of the mechanism depends on the structure and parameters of the environment. We justify this in the conclusion.

We present a simple model with binary policies, signals, and outcomes. We solve for six intuitive

asset-holder; furthermore the initial asset holders may be able sell off their holdings. As we show later, the incentives of an asset-holder and an informed investor may differ, but concerns about manipulation are present for each.

cases which map out the entire parameter space of incentive alignments, offering a taxonomy of these cases, and motivating them with real-world examples. Given ex-ante policy indifference we find that inducing truth-telling behavior and listening to markets is generically beneficial for the PM. We find an interesting asymmetry: where the PM is ex-ante indifferent, and where her commitment is common knowledge, she commits to occasionally (or never) executing after a *bad* signal, but she always executes after a *good* signal. We also find the surprising result that for a positive-measure parameter space, allowing the $\mathcal{II}$ to short-sell – giving $\mathcal{II}$ a greater set of options – may make implementation *easier*, thus increasing the PM's payoff.

Our work is largely a theoretical benchmark; we consider the case where the investor(s) have the maximum potential to manipulate a naive PM and we describe how the optimal mechanism involves randomization, and define the frontier of what the PM can achieve.

Our paper may be seen as a normative proposal. However, these considerations may also be reflected in current practices, and decision-makers are already taking prediction markets, as well as standard asset markets, into account (Arrow et al., 2008). There is anecdotal evidence of some policymaker "market-watching." Although PMs make policy announcements and float trial balloons, there is little evidence that they *explicitly* commit to tying policy to asset prices. However, as we discuss in the conclusion, legislative processes, policy trials, and "committing political capital" may also enable *listening* and informal conditional commitments, in essence a purification of the mixed/behavioral strategy.

Our paper proceeds as follows. Section 2 reviews the literature. We specify our formal model in section 3. In section 4 we give general lemmas and results. We conclude in section 6, highlighting the broader academic and practical policy implications, ways in which the model could be extended, and suggestions for empirical work.

2 Literature Review

Previous work has analyzed the relationship between the values of assets in conventional markets and the policy predictions of information markets, considering the implications of using such analysis to set policy.[7] A more extensive literature has performed "event studies" to measure the effect of policy announcements on conventional asset prices, considering the implications that can be drawn about the policies. A major concern is that if future policy itself reacts to the market's response to the policy announcement, the market response may be hard to interpret; this has been called the "circularity problem" (Bernanke and Mishkin, 1997; Sumner and Jackson, 2008).

To circumvent this circularity, others have proposed the use of "conditional prediction markets"(Hanson, 2013; Abramowicz, 2004). In such a market one asset takes a value, tied to some outcome, if a policy is executed, another asset takes a value if it is *not* executed, and otherwise trades are canceled. For example, Hahn and Tetlock (2003) consider assets whose values are tied to the level of GDP in the event (or non-event, for the second asset) of a carbon emissions cap, and consider what the difference in these asset prices reveal about the likely effect of such a cap.

However, according to Hahn and Tetlock (2003), "[a] general concern is that information markets are susceptible to price manipulation by those with a vested interest in the policy decision". This suggests the that use of prediction markets, which tend to be thin and illiquid, may be limited.[8] When real assets such as firms and physical capital are strongly affected by a policy it may be difficult to make prediction markets large enough to deter manipulation. Furthermore, if the main impact of a policy is only known by a single individual or a coordinated group with the *ability to heavily invest or short sell*, bringing large numbers of uninformed traders into the prediction market will not deter manipulation.

It may be more effective for policymakers to learn from the movement of *real* large-scale asset markets and from real investment decisions when a policy is announced, explicitly recognizing the connection between the probability of executing the policy and the movement in the asset prices. Thus, policymakers could explicitly take into account the incentive for manipulation and use a

[7] *Information markets* are markets which do not represent direct claims on tangible assets (Wolfers and Zitzewitz, 2006). Our analysis may also apply to corporate policy, in cases where the corporate management may know less about outcomes than the outside investors and speculators. Kau et al. (2008) find that, "on average, managers listen to the market: they are more likely to cancel investments when the market reacts unfavorably to the related announcement."

[8] Sumner and Jackson (2008) note that conditional and prediction markets are likely to be thin and hence unreliable, arguing government subsidies to trading are needed to combat this.

mechanism to induce truth-telling behavior.

In a related context, Hanson and Oprea (2009) model a market microstructure with noise traders, potentially informed traders, a competitive market maker, and a thin prediction market. As in our model they have a single rational "manipulator"; however, in contrast to us they assume the manipulator has a specific preference over the market price or over "the beliefs of neutral observers influenced by the price." In other words, their manipulator is essentially a particular type of "noise trader" who has a specific goal unrelated to the asset's true value. In contrast, our $\mathcal{II}$ seeks only to make a profit off of the policy outcome he induces.[9]

2.1 Empirical work and examples

There are several recent cases in which policymakers seem to have listened to the market, or where others have suggested that they should have done so.

During the debate over the US Affordable Care Act, Milani (2010) tracked the stock returns of health insurance companies against a prediction market security whose payoff was tied to the inclusion of the "public option" in the bill. He concluded "the results reveal the market expectation of a negative effect of the public option on the value of health insurance companies ... of around 13%, but it does not support more calamitous scenarios." Friedman (2009) performed event studies on pharmaceutical firms' share prices as they introduced new drugs, comparing the implied profitability of (low versus high Medicare share) drugs before and after the introduction of the Medicare Part D prescription drug benefit. He used this to impute that the bill would lead to $205 billion in additional drug company profits.

Wolfers and Zitzewitz (2006) presented evidence, in the context of the Iraq war in the mid-2000's, that spot and futures market oil prices moved in line with a prediction market for a security that paid off if Saddam Hussein were removed from power by a certain date.[10] They used this to estimate the distribution of investors' beliefs for the impact of the war on the economy, imputing

[9] Our model differs in other important ways. Our $\mathcal{II}$ has unique private information; they assume a large number of traders who can pay a cost to learn about the asset's true value and about the manipulator's preferred price. Another difference: we explicitly model the policy choice.

[10] They interpret this as reflecting the probability of a US attack on Iraq in 2002-2003.

"a substantial probability of an extremely adverse outcome."

As Wolfers and Zitzewitz (ibid) argue, the above sort of evidence could be used to "better understand the consequences of a prospective policy decision ... [and] to inform decision-making in real time." In light of the above evidence, the public option might have been scrapped, the drug benefit repealed or reformed, the Iraq war reconsidered, and the CAFTA agreement reinforced or cancelled (depending on whether PMs thought the gains were large enough).[11]

Poland's 2011-2012 experience seems to be a particularly clear-cut example. In November 2011, Polish PM Donald Tusk announced a new minerals tax. Share prices of KGHM (Poland's sole copper producer) fell by roughly 25% in the two days following the announcement. On January 3, 2012, the Finance Ministry lowered this proposed tax rate after negative reactions from the Economy and Treasury ministries and from the company, which argued the tax would make output at one of its three mines inviable. Economy Minister Waldemar Pawlak explicitly mentioned the share price reaction in his criticism.[12]

3 Model

In this section we present the incentive compatibility constraints and a complete set of results *when an informed investor chooses between buying and not buying the asset.* We will *also* present the setup and the asset market allowing short selling, and we derive a feasible welfare improving (but perhaps not optimal) mechanism in which $\mathcal{II}$ can buy, do nothing, or short sell. As we note, most

[11]The Pentagon also attempted to use markets to predict geopolitical risks such as terrorism. The Defense Advanced Research Projects Agency proposed introducing a policy market in 2003. This project was cancelled, allegedly in light of concerns that bad actors might themselves invest, commit terrorist acts, and profit from this (Hanson, 2005). This would be a distinct form of manipulation from our discussion; such concerns involve *physical actions* taken in order to manipulate *markets*; we are considering *financial* transactions and investments made in order to influence *policy* (which in turn influences asset values). We assume that the PM does not derive welfare from the investments; if these "investments" are acts of terrorism this is obviously the PM's concern.

[12]Pawlak: "If the tax had been more clearly presented during the prime minister's expose, there would be less unrest and less fluctuation in KGHM's share value. Now we have to prepare [the new tax] properly so that it benefits the state but doesn't kill KGHM" (http://www.wbj.pl/article-57461-deputy-pm-criticizes-copper-tax.html, accessed on 1 Jan. 2014)

of our qualitative results hold for both setups.

We first define our notation. The state is denoted by $s \in S$. There is an informed investor "$\mathcal{II}$", who receives a binary signal, $\sigma \in \{\gamma, \beta\}$, in which γ is the *good* signal and β is the *bad* signal. We will refer to the informed investor who has received signal σ by $\mathcal{II}(\sigma)$. The signal is correlated with the state of the world $s \in \{G, B\}$, i.e., the *good* state and the *bad* state (these states reflect the welfare consequences of a policy, as formalized below). $P(s, \sigma)$ represents the joint probability over (s, σ), and $P(s|\sigma)$ is the conditional probability of s given σ. The unconditional probabilities are $P(s) := \sum_{\sigma \in \{\gamma,\beta\}} P(s, \sigma)$ and $P(\sigma) := \sum_{s \in \{G,B\}} P(s, \sigma)$.[13]

3.1 Timing

1. The PM commits to probabilities of execution $q(\hat{\sigma}) \in (0, 1)$ as a function of the signal to be revealed $\hat{\sigma} \in \{\gamma, \beta\}$. These commitments are publicly observed or deduced by all parties.

2. $[t = 0]$ The initial asset price A_0 is formed by the expectation of one or more Uninformed Traders (henceforth, UTs).

3. Nature chooses the state of the world, $s \in \{G, B\}$ as well as the signal σ.

4. The $\mathcal{II}$ receives signal σ and chooses $i \in \{b, nb\}$ or $i \in \{b, nb, sh\}$ (depending on the environment we consider), where b, nb, and sh represent buying one unit, doing nothing, or short selling one unit, respectively. The $\mathcal{II}$'s *action* becomes observable to the PM (through its impact on the asset price), sending a signal $\hat{\sigma}$ to the PM.

5. The PM executes the policy with the pre-committed probability $q(\hat{\sigma})$, where $\hat{\sigma} = \sigma$ in a truth telling equilibrium.

6. $[t = 1]$ Payoffs are realized.

Consider stages 1 and 2. As the probability that a good or a bad policy is chosen, and the effects of these on asset prices, is common knowledge, all parties can infer the incentive compatible

[13] Although we restrict the set of states and the set of signals to be two elements, we could easily extend the current setup into many states and many possible signals.

probabilities of execution. For the third stage, our preferred real-world interpretation is that the PM chooses the details of the policy through an unpredictable process. Once these details are announced they constitute a signal of the policy's merit that only $\mathcal{II}$ can interpret.

Note that committing to probabilities $\langle q(\gamma) = 1, q(\beta) > 0\rangle$ is equivalent to a commitment that with probability $q(\beta)$ she will ignore the asset price and execute the policy. Similarly $\langle q(\gamma) < 1, q(\beta) = 0\rangle$ is equivalent to a commitment that with probability $1 - q(\gamma)$ she will ignore the asset price and not execute the policy. (In each case, with the remaining probability she will choose her action in light of the signal.) These are the commitments described in the introduction. We will show below that the optimal commitment always involves $q(\gamma) = 1$.

3.2 Policymaker's objective function

The PM either executes a policy (denoted $p = e$) or does not (denoted $p = ne$). Her payoff with the policy choice p and the state s is $W(p, s)$. Note that the investor's decision is not directly in the PM's objective function. We assume the PM wants to execute the policy if and only if the true signal is γ; i.e., γ represents the "good news" about the policy. Formally,

Assumption 1 *(i)* $\sum_s P(s|\gamma)W(e, s) > \sum_s P(s|\gamma)W(ne, s)$, *(ii)* $\sum_s P(s|\beta)W(e, s) < \sum_s P(s|\beta)W(ne, s)$.

We denote by $q(\sigma)$, the probability of executing the policy for a given signal σ. For a given $\langle q(\gamma), q(\beta)\rangle$, (implicitly assuming the PM has deduced the true signal) her expected payoff is:

$$\Omega(q(\gamma), q(\beta)) := \sum_{s,\sigma} P(s, \sigma)\Big[q(\sigma)W(e, s) + (1 - q(\sigma))W(ne, s)\Big].$$

For later use, we re-write the PM's payoff in terms of gains/losses from execution in either state:

$$\begin{aligned}&\sum_{s,\sigma} P(\sigma, s)W(ne, s) + q(\gamma)P(\gamma)(\mathbb{E}W(e|\gamma) - \mathbb{E}W(ne|\gamma)) + q(\beta)P(\beta)(\mathbb{E}W(e|\beta) - \mathbb{E}W(ne|\beta))\\ =&\sum_{s,\sigma} P(\sigma, s)W(ne, s) + q(\gamma)P(\gamma)\Delta W(\gamma) + q(\beta)P(\beta)\Delta W(\beta), \qquad (1)\end{aligned}$$

where we define $\mathbb{E}W(p|\sigma) := \sum_s P(s|\sigma)W(p, s)$ and $\Delta W(\sigma) := \mathbb{E}W(e|\sigma) - \mathbb{E}W(ne|\sigma)$.

By Assumption 1, $\Delta W(\gamma)$ is positive, and $\Delta W(\beta)$ is negative. Thus, as long as the $\mathcal{II}$'s incentive compatibility constraints (described in sections 3.3.1 and 3.3.2) are satisfied so that the $\mathcal{II}$ reveals the true signal, the PM wants to maximize $q(\gamma)$ and minimize $q(\beta)$.

As a benchmark, we assume that before learning the signal the PM is indifferent.

Assumption 2 *The policymaker is ex-ante indifferent between policies $\langle q(\gamma) = 1, q(\beta) = 1\rangle$ and $\langle q(\gamma) = 0, q(\beta) = 0\rangle$, i.e.,*[14]

$$\sum_s P(s)W(e,s) = \sum_s P(s)W(ne,s) \quad i.e., \quad \Delta W(\gamma)P(\gamma) + \Delta W(\beta)P(\beta) = 0.$$

This indifference assumption implies that changes to $q(\beta)$ and $q(\gamma)$ have equal and opposite direct effects on PM's welfare (holding the $\mathcal{II}$'s behavior constant).

3.3 Asset market and the informed investor's payoff

The asset's fundamental value, $A_1(p,s)$, represents the discounted stream of future earnings from the asset; this will depend on the state and the policy decision. After these become common knowledge (at $t = 1$), the asset's price will equal its fundamental value.[15]

Determination of A_0: We assume A_0 is based on expected outcomes for the correctly anticipated probabilities of execution, i.e., $A_0 = A_0(q(\gamma), q(\beta))$.[16] This endogeneity is both reasonable and relevant. If, *alternatively,* the initial asset price A_0 were exogenous, then holding the asset could be profitable on average given the announced $\langle q(\gamma), q(\beta)\rangle$.[17] E.g., for an arbitrary A_0, the UTs could be systematically fooled and make negative expected profits, violating rational expectations. We can imagine several justifications for an exogenous A_0 (see footnote); however, these do not seem reasonable or relevant to our empirical examples.[18]

[14] Deriving the equivalence of the two conditions:

$$\Delta W(\gamma)P(\gamma) + \Delta W(\beta)P(\beta) = 0 \Leftrightarrow P(\gamma)(\mathbb{E}W(e|\gamma) - \mathbb{E}W(ne|\gamma)) + P(\beta)(\mathbb{E}W(e|\beta) - \mathbb{E}W(ne|\beta)) = 0$$
$$\Leftrightarrow P(\gamma)\mathbb{E}W(e|\gamma) + P(\beta)\mathbb{E}W(e|\beta) = P(\gamma)\mathbb{E}W(ne|\gamma) + P(\beta)\mathbb{E}W(e|\beta) \Leftrightarrow \sum_s P(s)W(e,s) = \sum_s P(s)W(ne,s)$$

[15] We assume that no earnings accrue from the asset until after time $t = 1$; this is without loss of generality.

[16] The results would be equivalent if we alternatively allowed these probabilities to *not* be correctly anticipated but assumed they were publicly announced before A_0 is set.

[17] The results of the model with exogenous A_0 are available by request; we remark on this in section 5.2.

[18] Two possible justifications: (i) If only the $\mathcal{II}$ could profit from holding the asset at $t = 1$ (e.g., through his own production process) and he holds all the bargaining power, then A_0 would be priced at cost, regardless of the q's. However, this would imply that the PM could only identify the signal if she could identify *who* the $\mathcal{II}$ was *in advance*

Price and market microstructure: We assume the asset is publicly traded at its expected value – *unconditional* on the private signal that $t = 0$, but *conditional* on σ at $t = 1$. Several assumptions over the market's microstructure would justify this; we leave these out of our modeling for simplicity. For example, if the UT's are noise traders, who have a liquidity motivation, as in Bagehot (1971), and if the UT's greatly outnumber the $\mathcal{II}$'s, a profit-maximizing market maker will offer a bid-ask spread tightly bounded around the unconditional (uninformed) expected value. However, as a larger volume of (e.g., buy) orders begins to arrive, it becomes less and less likely that this is coming from a noise trader. Thus the bid and ask prices will rise, approaching the expected value in light of the signal σ (see Madhavan, 2000, section 3.2.3 for a survey of compatible models). Note that our qualitative results still hold if a UT (or market maker) shares *some* surplus when trading with the $\mathcal{II}$.

We further assume that the signal and the state are *specific* to the policy, i.e., if the policy is *not* executed, then the asset's value does not depend on the state:

Assumption 3

$$A_1(ne) := A_1(p = ne, s = G) = A_1(p = ne, s = B).$$

Thus, the expectation of $A_1(ne)$ is invariant to the signal (and state), i.e., $\mathbb{E}A_1(ne|\sigma) = A_1(ne)$ for any $\sigma \in \{\gamma, \beta\}$. We define $\mathbb{E}A_1(e|\sigma) := \sum_s P(s|\sigma)A_1(p = e, s)$. Using this, we derive the initial asset price:

$$\begin{aligned} A_0(q(\gamma), q(\beta)) =& P(\gamma)q(\gamma)\sum_s P(s|\gamma)A_1(e,s) + P(\beta)q(\beta)\sum_s P(s|\beta)A_1(e,s) \\ &+ [P(\gamma)(1-q(\gamma)) + P(\beta)(1-q(\beta))]A_1(ne) \\ =& P(\gamma)q(\gamma)[\mathbb{E}A_1(e|\gamma) - A_1(ne)] + P(\beta)q(\beta)[\mathbb{E}A_1(e|\beta) - A_1(ne)] + A_1(ne) \\ =& q(\gamma)P(\gamma)\Delta A_1(\gamma) + q(\beta)P(\beta)\Delta A_1(\beta) + A_1(ne), \end{aligned} \tag{2}$$

and identify his precise choice, a difficult proposition. (ii) If the policy was considered a zero-probability event, and the policy as well as the q functions were announced to the $\mathcal{II}$'s before being publicly announced, or the $\mathcal{II}$'s could react to this information before the UTs, then A_0 might also be unaffected by the q's; it also seems unlikely that the PM could orchestrate this. Furthermore, neither of these scenarios seem to reflect the empirical cases we describe.

where $\Delta A_1(\sigma) := \mathbb{E}A_1(e|\sigma) - A_1(ne)$ is the expected benefit of execution (relative to non-execution) to an asset-holder given signal σ. Note that price A_0 does not depend on the realized signal, as the signal is not common knowledge at $t = 0$.

$\mathcal{II}$'s payoffs: $\mathcal{II}$'s payoff is his net return from buying, doing nothing, or short selling the asset. We assume that he can buy the first unit at price A_0 before his action is detected, and the price rises to A_1.[19] If he chooses to buy at time $t = 0$, setting $i = b$, he pays A_0 to the UT, and earns the asset's fundamental value A_1 at $t = 1$. Thus, for either policy choice $p \in \{e, ne\}$ and for either signal $s \in \{G, B\}$, the $\mathcal{II}$'s payoff from buying is:

$$V(p, s, i = b) = A_1(p, s) - A_0.$$

If he short sells at time $t = 0$ (i.e., setting $i = sh$), he gets paid A_0 by the UT, and then must buy the asset at $t = 1$. Thus, holding the policy constant, the payoff from short selling is the negative of the payoff from buying, i.e.,

$$V(p, s, i = sh) = A_0 - A_1(p, s) = -V(p, s, i = b).$$

If the $\mathcal{II}$ does nothing at time $t = 0$ (i.e., setting $i = nb$), he neither pays nor receives anything at $t = 0$ and neither owns nor owes the asset at time $t = 1$; thus

$$V(p, s, i = nb) = 0.$$

We next define $\mathcal{II}$'s expected payoff from buying when he receives signal σ and he leads the PM to believe the signal was $\hat{\sigma}$ (henceforth, we will say "his action reports $\hat{\sigma}$")

$$\begin{aligned}
\mathbb{E}V(i = b|\sigma, \hat{\sigma}) &= \sum_s P(s|\sigma)[q(\hat{\sigma})V(e, s, b) + (1 - q(\hat{\sigma}))V(ne, s, b)] \\
&= q(\hat{\sigma})[\mathbb{E}V(e|\sigma) - \mathbb{E}V(ne)] + \mathbb{E}V(ne) = q(\hat{\sigma})[\mathbb{E}A_1(e|\sigma) - \mathbb{E}A_1(ne)] + \mathbb{E}A_1(ne) - A_0(q(\gamma), q(\beta)) \\
&= q(\hat{\sigma})\Delta A_1(\sigma) + A_1(ne) - A_0(q(\gamma), q(\beta)). \qquad (3)
\end{aligned}$$

In other words, when the $\mathcal{II}$'s signal is σ and his action reports $\hat{\sigma}$, his expected payoff from buying is $q(\hat{\sigma})\Delta A_1(\sigma)$, the expected increase in asset value given the true and reported signals, plus the baseline value under non-execution $A_1(ne)$, less the cost of buying the share $A_0(q(\gamma), q(\beta))$.

[19] Of course, the *size* of this first unit will not affect our results.

When the $\mathcal{II}$'s signal is σ and the PM believes it was $\hat{\sigma}$, the $\mathcal{II}$'s profit from short selling is

$$\mathbb{E}V(i = sh|\sigma, \hat{\sigma}) = -q(\hat{\sigma})\Delta A_1(\sigma) - A_1(ne) + A_0(q(\gamma), q(\beta)) = -\mathbb{E}V(i = b|\sigma, \hat{\sigma}). \tag{4}$$

Expanding $A(ne) - A_0(q(\gamma), q(\beta))$, equations (3) and (4) become:

$$\mathbb{E}V(i = b|\sigma, \hat{\sigma}) = q(\hat{\sigma})\Delta A_1(\sigma) - [q(\gamma)P(\gamma)\Delta A_1(\gamma) + q(\beta)P(\beta)\Delta A_1(\beta)], \tag{5}$$

$$\mathbb{E}V(i = sh|\sigma, \hat{\sigma}) = -q(\hat{\sigma})\Delta A_1(\sigma) + [q(\gamma)P(\gamma)\Delta A_1(\gamma) + q(\beta)P(\beta)\Delta A_1(\beta)]. \tag{6}$$

The first term on the right hand side of (5) and (6) is the expected gain or loss in asset value from probabilistic execution given the true and reported signals. The bracketed term is the *ex-ante* expectation of this gain or loss, i.e., $A_0 - A_1(ne)$.

3.3.1 Incentive compatibility constraints: shorting not allowed

For $\mathcal{II}$'s behavior to be truth-telling, he must prefer to take an action $i \in \{b, nb\}$ whenever he sees a good signal and take the other action whenever he sees a bad signal. Formally, the PM may interpret $i = b$ and $i = nb$ as (i) the investor having received signal γ and β respectively, or as (ii) the investor having received signal β and γ respectively. Thus, there are two possible sets of incentive compatibility constraints.

The 1st set: Buy if and only if the signal is γ. Here, the PM interprets $i = b$ as $\mathcal{II}$ having received γ. The two incentive compatibility constraints are:

$$q(\gamma)\Delta A_1(\gamma) - [q(\gamma)P(\gamma)\Delta A_1(\gamma) + q(\beta)P(\beta)\Delta A_1(\beta)] \geq 0$$

$$\Leftrightarrow q(\gamma)\Delta A_1(\gamma) \geq q(\beta)\Delta A_1(\beta) \tag{7}$$

$$0 \geq q(\gamma)\Delta A_1(\beta) - [q(\gamma)P(\gamma)\Delta A_1(\gamma) + q(\beta)P(\beta)\Delta A_1(\beta)]$$

$$\Leftrightarrow P(\gamma)q(\gamma)\Delta A_1(\gamma) \geq (q(\gamma) - P(\beta)q(\beta))\Delta A_1(\beta) \tag{8}$$

Constraint (7) above ensures that the $\mathcal{II}$ prefers to buy after getting a good signal, and constraint (8) ensures he prefers not to buy after getting a bad signal. I.e., the first constraint deters $\mathcal{II}(\beta)$ from mimicking $\mathcal{II}(\gamma)$, and the second deters $\mathcal{II}(\gamma)$ from mimicking $\mathcal{II}(\beta)$.

The 2nd set: Buy if and only if the signal is β

The second set of incentive compatibility constraints are similarly stated:

$$q(\beta)\Delta A_1(\beta) \geq q(\gamma)\Delta A_1(\gamma) \tag{9}$$

$$P(\beta)q(\beta)\Delta A_1(\beta) \geq (q(\beta) - P(\gamma)q(\gamma))\Delta A_1(\gamma) \tag{10}$$

Remark 1: We do not allow the PM to set direct payments after any combination of revealed signal and realized state. However, $A_0(q(\gamma), q(\beta))$ plays a role similar to a direct payment in the standard mechanism model with transferable utility. As shown in the $\mathcal{II}$'s payoff (e.g., $q(\gamma)\Delta(A) + A_1(ne) - A_0(q(\gamma), q(\beta))$), the choice of randomized policy execution $\langle q(\gamma), q(\beta)\rangle$ changes the amount the $\mathcal{II}$ has to pay/receive (i.e., $A_0(q(\gamma), q(\beta))$). More specifically, if $\mathcal{II}$ chooses to buy, his net payoff is $q(\gamma)\Delta(A) + A_1(ne) - A_0(q(\gamma), q(\beta))$, and he has to pay $A_0(q(\gamma), q(\beta))$. However, unlike the direct payment in the standard mechanism design literature, the monetary allocation $A_0(q(\gamma), q(\beta))$ depends on the allocation of a non-money commodity (i.e., the randomized allocation $\langle q(\gamma), q(\beta)\rangle$).

Remarks 2: Our paper can also be compared with the literature on costly signaling models following Spence (1973): $\mathcal{II}$ sends a costly signal by buying/not buying (or short-selling) the asset, and upon observing $\mathcal{II}$'s action, the PM chooses execution/non-execution. There are two important differences. Most obviously, the PM commits to the probability of execution, $(q(\gamma), q(\beta))$. Furthermore, the strategy $q(\sigma)$ after observing signal σ also influences $A_0(q(\gamma), q(\beta))$, i.e, the equilibrium play determines the payoffs (including out-of-equilibrium payoffs).

3.3.2 Incentive compatibility constraints: shorting allowed

In an alternative environment where short selling is allowed (which we return to in section 5), the PM may prefer to implement a mechanism where she interprets $i = b$ and $i = sh$ (i) as the investor having received signal γ and β respectively, or (ii) β and γ respectively. We refer to this as a "buy/short" implementation. We impose that $\mathcal{II}$ cannot choose "do nothing". However, we verify in online appendix 1.4.[20] that under the implementation we derive, $\mathcal{II}$ will not want to "do nothing". Again, there are two sets of IC constraints.

[20] https://davidreinstein.files.wordpress.com/2010/07/online-appendix.pdf

The 1st set: Buy if and only if the signal is γ, otherwise short sell.

$$q(\gamma)\Delta A_1(\gamma) - \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \geq -q(\beta)\Delta A_1(\gamma) + \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \tag{11}$$

$$-q(\beta)\Delta A_1(\beta) + \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \geq q(\gamma)\Delta A_1(\beta) - \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \tag{12}$$

where $\mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) := q(\gamma)P(\gamma)\Delta A_1(\gamma) + q(\beta)P(\beta)\Delta A_1(\beta)$. Constraint (11) ensures the $\mathcal{II}$ prefers buying to short selling after getting a good signal, and (12) ensures he prefers short selling to buying after a bad signal.

The 2nd set: Buy if and only if the signal is β, otherwise short sell.

$$-q(\gamma)\Delta A_1(\gamma) + \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \geq q(\beta)\Delta A_1(\gamma) - \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \tag{13}$$

$$q(\beta)\Delta A_1(\beta) - \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \geq -q(\gamma)\Delta A_1(\beta) + \mathbb{E}(\Delta A_1|q(\gamma), q(\beta)) \tag{14}$$

4 Analysis

4.1 General results

With the aforementioned commitment $\langle q(\gamma), q(\beta)\rangle$, the PM's problem is:

$$\max_{0\leq q(\cdot)\leq 1} \sum_{s,\sigma} P(\sigma, s)\Big[q(\sigma)W(e, s) + (1 - q(\sigma))W(ne, s)\Big]$$

$$s.t. \ \text{set 1 [(7) and (8)] or set 2 [(9) and (10)].} \tag{15}$$

Definition 1 (First best policy) $\langle q(\gamma) = 1, q(\beta) = 0\rangle$ *is the* first best policy.

Definition 2 (Blind policy) $\langle q(\gamma), q(\beta)\rangle$ *with* $q(\gamma) = q(\beta)$ *is a* blind policy.

Definition 3 (Incentive-constrained optimal policy) *An* incentive-constrained optimal policy *solves the PM's maximization problem (15).*

Considering PM's payoff in equation (1), the first best policy trivially maximizes the policymaker's welfare since $\Delta W(\gamma) > 0$ and $\Delta W(\beta) < 0$. However, it may not be feasible since it may not be incentive compatible.

A PM who employs a blind policy *does not listen to markets* as the signal does not alter her probability of executing the policy, i.e., $q(\gamma) = q(\beta)$. However, this policy is always incentive compatible (i.e., it leads $\mathcal{II}$ to take a distinct action after each signal).

Proposition 4 *Any blind policy is incentive compatible (i.e., it satisfies one of the sets of incentive compatibility constraints). This holds whether or not short selling is allowed.*

Proof. If $\Delta A_1(\gamma) \geq \Delta A_1(\beta)$, or $\Delta A_1(\gamma) \leq \Delta A_1(\beta)$, it is trivial that the first set, or the second set, respectively, of incentive compatibility constraints are satisfied with $q(\gamma) = q(\beta)$. ∎

Since a blind policy is incentive compatible, it represents a lower-bound on the welfare that the PM can achieve; moreover, she will always be worse off if she executes the policy more frequently conditional on a bad signal than she does conditional on a good signal.

Lemma 5 *The incentive-constrained optimal* $\langle q(\gamma), q(\beta) \rangle$ *satisfies* $q(\gamma) \geq q(\beta)$*, whether or not short selling is allowed.*

Proof. Considering expression (1) in light of the PM's indifference (Assumption 2), $q(\gamma) < q(\beta)$ implies that the sum of the last two terms in expression (1) is negative; while $q(\gamma) = q(\beta)$ (i.e., *not listening to markets*) implies this sum is zero. Thus we must have $q(\gamma) \geq q(\beta)$. ∎

We show that $q(\gamma) = 1$ for any implementation that does strictly better then a blind policy. [21] Moreover, we show later that generically, the optimal policy *is* superior to a blind policy.

Proposition 6 *Any policy that yields higher welfare for the PM than a blind policy satisfies* $1 = q(\gamma) \geq q(\beta)$*, whether or not short selling is allowed.*

Proof. Given that $\mathcal{II}$'s payoff is a linear function of $q(\gamma)$ and $q(\beta)$ without a constant term, incentive compatibility constraints are written as $Cq(\gamma) \gtreqless q(\beta)$, with or without short sale. Suppose **a** constraint being considered is of the form $Cq(\gamma) \geq q(\beta)$. If $C \geq 0$ then the first best $\langle q(\gamma) =$

[21] This result is consistent with asymmetry of the problem; the PM "naturally" wants $q(\beta) = 0$ and $q(\gamma) = 1$, and zero and one do not have symmetric properties. Note that this proposition depends critically on the assumption of *endogenous* A_0; as we demonstrate in section 5.2, if A_0 were exogenous, this proposition would not hold.

$1, q(\beta) = 0\rangle$ trivially satisfies this, implying that this constraint is non-binding. On the other hand, if $C < 0$, $q(\gamma) = q(\beta) = 0$ is the only possible solution, which is a blind policy.

Thus, any binding IC constraint for an implementation yielding PM welfare greater than a blind policy must be expressed as $Cq(\gamma) \leq q(\beta)$.

Since $q(\gamma) \geq q(\beta)$ from Lemma 5, $C \leq 1$ must be the case (if not, we have a contradiction: $Cq(\beta) \leq Cq(\gamma) \leq q(\beta)$ with $C > 1$). If $C = 1$ we must have $q(\gamma) = q(\beta) = 1$, again a blind policy, due to Lemma 5. If neither constraint is binding the PM will set the first-best policy $1 = q(\gamma) > q(\beta) = 0$. If a constraint binds, then it is satisfied with $Cq(\gamma) = q(\beta) < 1$ where $C < 1$, then PM could do better by increasing $q(\gamma)$ by ϵ and increasing $q(\beta)$ by a smaller amount $C\epsilon$, which is a contradiction to the supposition that $\langle q(\gamma), q(\beta)\rangle$ is incentive-constrained optimal. ∎

Given that $q(\gamma) = 1$, this commitment is equivalent to a commitment that with probability $q(\beta)$ she will ignore the asset price and execute the policy (and with the remaining probability she will choose her action in light of the signal). Given that $q(\gamma) = 1$, we focus only on $q(\beta)$ for the characterization of incentive-constrained efficiency. With two constraints for each set and a single choice variable $q(\beta)$, it is trivial that at most one constraint binds generically.

Corollary 7 *For a given set, at most one of the two incentive compatibility constraints will bind (whether or not short selling is allowed).*

Inequalities (7) and (8) are summarized as:

$$q(\gamma)\Delta A_1(\gamma) - \mathbb{E}(\Delta A_1 | q(\gamma), q(\beta)) \geq 0 \geq q(\gamma)\Delta A_1(\beta) - \mathbb{E}(\Delta A_1 | q(\gamma), q(\beta))$$

The set of $\langle q(\gamma), q(\beta)\rangle$ satisfying the above inequalities is non-empty only if $\Delta A_1(\gamma) \geq \Delta A_1(\beta)$.

Summarizing (9) and (10) similarly yields:

$$q(\gamma)\Delta A_1(\gamma) - \mathbb{E}(\Delta A_1 | q(\gamma), q(\beta)) \leq 0 \leq q(\gamma)\Delta A_1(\beta) - \mathbb{E}(\Delta A_1 | q(\gamma), q(\beta)).$$

The set of $\langle q(\gamma), q(\beta)\rangle$ satisfying the above two inequalities is non-empty only if $\Delta A_1(\gamma) \leq \Delta A_1(\beta)$.

Note that $\Delta A_1(\gamma) \geq \Delta A_1(\beta)$ and $\Delta A_1(\beta) \geq \Delta A_1(\gamma)$ are exclusive (excepting the knife-edge case) and exhaustive. Thus, only one set can be implemented generically, so the relative size of

$\Delta A_1(\gamma)$ and $\Delta A_1(\beta)$ fully determines which set of incentive compatibility constraints are used. The same holds when short-selling is allowed and the PM implements a mechanism getting the investor to choose $i \in \{b, sh\}$ after a good signal and choose the other element of $\{b, sh\}$ after a bad signal.[22] (We will show later that not doing anything, i.e., $i = nb$, is never preferred in this implemented mechanism.) From this, we derive the following proposition.

Proposition 8 *The PM induces Set 1 ("buy only if good") if $\Delta A_1(\gamma) > \Delta A_1(\beta)$ and Set 2 ("buy only if bad") if $\Delta A_1(\gamma) > \Delta A_1(\beta)$, both where short selling is prohibited and in the buy/short implementation.*

Finally, we show that the PM always does better by listening to the market.

Proposition 9 *The incentive-constrained efficient policy is generically superior to blind policies, whether or not short selling is allowed.*

Proof. To prove this we show that generically, for incentive-constrained optimality, $q(\beta) < 1$. From Proposition 6, this implies generically $1 = q(\gamma) > q(\beta)$, which is superior to a blind policy.

Generically, $\Delta A_1(\gamma) \neq \Delta A_1(\beta)$. (i) Assume $\Delta A_1(\gamma) > \Delta A_1(\beta)$. Suppose that $q(\gamma) = q(\beta) = 1$. Then Set 1 becomes: $P(\beta)\Big[\Delta A_1(\gamma) - A_1(\beta)\Big] \geq 0$, $P(\gamma)\Big[\Delta A_1(\beta) - A_1(\gamma)\Big] \leq 0$, whether short selling is allowed or not, as shown in the proof for Proposition 4. We can see that both of these two constraints hold with strict inequality, implying $q(\beta) < 1$ would also satisfy these constraints, and thus $q(\beta) = 1$ must be suboptimal. (ii) Assuming $\Delta A_1(\gamma) < \Delta A_1(\beta)$, we can apply the same logic to the second set of incentive compatibility constraints to derive a contradiction. ∎

[22] Inequalities (11) and (12) and Inequalities (13) and (14) are written as follows with short selling allowed, for the mechanism considered:

$$[q(\beta) + q(\gamma)]\Delta A_1(\beta) - 2\mathbb{E}(\Delta A_1) \geq 0 \geq [q(\beta) + q(\gamma)]\Delta A_1(\gamma) - 2\mathbb{E}(\Delta A_1) \quad (16)$$

$$[q(\beta) + q(\gamma)]\Delta A_1(\beta) - 2\mathbb{E}(\Delta A_1) \leq 0 \leq [q(\beta) + q(\gamma)]\Delta A_1(\gamma) - 2\mathbb{E}(\Delta A_1) \quad (17)$$

The first set (the second set) is non-empty only if $\Delta A_1(\beta) \geq \Delta A_1(\gamma)$ ($\Delta A_1(\beta) \leq \Delta A_1(\gamma)$). Note that $\Delta A_1(\gamma) \geq \Delta A_1(\beta)$ and $\Delta A_1(\beta) \geq \Delta A_1(\gamma)$ are exclusive (excepting the knife-edge case) and exhaustive. By the same reason in the environment without short selling, the relative size of $\Delta A_1(\gamma)$ and $\Delta A_1(\beta)$ fully determines which set of incentive compatibility constraints are used.

The proof works by showing that generically, for incentive-constrained efficiency, $q(\beta) < 1$. This is then combined with Proposition 6 to show that generically $1 = q(\gamma) > q(\beta)$, which is superior to a blind policy. In other words, listening is generically better then not listening. Thus – given our assumption that the signal is informative of the true state – ex-ante indifference implies that the PM is willing to make commitment $\langle q(\gamma), q(\beta)\rangle$ in order to learn the signal.[23]

4.2 Full characterization: Six cases

Figure 1 describes the binding constraints in each region. (Note that $\emptyset$ means that no constraint is binding, i.e., the PM achieves the first best).

[Figure 1 about here.]

We divide the parameter sets into six cases in terms of the policy's effect on the *asset's value* after each signal: (i) $\Delta A_1(\gamma) > \Delta A_1(\beta) > 0$, (ii) $\Delta A_1(\beta) > \Delta A_1(\gamma) > 0$, (iii) $\Delta A_1(\beta) > 0 > \Delta A_1(\gamma)$, (iv) $0 > \Delta A_1(\beta) > \Delta A_1(\gamma)$, (v) $0 > \Delta A_1(\gamma) > \Delta A_1(\beta)$, and (vi) $\Delta A_1(\gamma) > 0 > \Delta A_1(\beta)$. This classification, along with a taxonomy we will explain, is displayed in figure 2.

[Figure 2 about here.]

We can interpret these six cases in terms of the alignment of incentives of the PM and the *asset owner*. To give insight, we solve for case 1 in detail; the other five cases are solved in footnotes. Results are summarized in figure 3 in section 5.1.

Further simplification of the two sets of incentive compatibility constraints

With $q(\gamma) = 1$, we can rewrite Set 1 as

$$\Delta A_1(\gamma) \geq \Delta A_1(\beta) q(\beta) \geq \Big[\frac{\Delta A_1(\beta) - \Delta A_1(\gamma)}{P(\gamma)} + \Delta A_1(\gamma)\Big], \tag{18}$$

and Set 2 as

$$\Delta A_1(\beta) q(\beta) \geq \Delta A_1(\gamma) \geq \Big[\frac{\Delta A_1(\gamma) - \Delta A_1(\beta)}{P(\beta)} + \Delta A_1(\beta)\Big] q(\beta). \tag{19}$$

[23] In some cases, no commitment will be necessary, i.e., $\langle q(\gamma) = 1, q(\beta) = 0\rangle$ may be optimal.

Case (i), "A Treat": $\Delta A_1(\gamma) > \Delta A_1(\beta) > 0$, *i.e., the asset's value increases when the policy is executed, more so with the good signal.*

This policy represents a *treat* for the asset-owner; execution will always increase the asset's value, but from the PM's perspective, it is only worth "buying this treat" where the signal suggests if it will be *very* beneficial to the asset holder. The "Soylent" R&D example from the introduction fits this case. Alternately, the current case may reflect a trade policy that involves costly concessions for the government but will be worth executing if it boosts a particular export sector by a sufficient amount (e.g., see Breinlich, 2011).[24] Alternatively, it might reflect a public research and development funding plan that will certainly stimulate some new inventions and raise profits somewhat, but will require dramatic results to justify its large costs. A macroeconomic stimulus or a bank or international bailout may have similar properties.[25]

We consider Set 1, inequalities (18).

$$\Delta A_1(\gamma) \geq \Delta A_1(\beta) q(\beta) \geq \Big[\frac{\Delta A_1(\beta) - \Delta A_1(\gamma)}{P(\gamma)} + \Delta A_1(\gamma)\Big],$$

As the PM wants to minimize $q(\beta)$, the only constraint that might bind is the second one. Note that the bracket term can be positive or negative depending on the sign of $\Delta A_1(\beta) - P(\beta)\Delta A_1(\gamma)$, as $\frac{\Delta A_1(\beta)-\Delta A_1(\gamma)}{P(\gamma)} + \Delta A_1(\gamma) = \frac{1}{P(\gamma)}(\Delta A_1(\beta) - (1 - P(\gamma))\Delta A_1(\gamma)) = \frac{1}{P(\gamma)}(\Delta A_1(\beta) - P(\beta)\Delta A_1(\gamma))$.

If $\frac{1}{P(\gamma)}(\Delta A_1(\beta) - P(\beta)\Delta A_1(\gamma)) < 0$, the constraint does not bind, and the optimal solution is $< q(\gamma) = 1, q(\beta) = 0 >$. On the other hand if $\frac{1}{P(\gamma)}(\Delta A_1(\beta) - P(\beta)\Delta A_1(\gamma)) > 0$, the constraint is binding; thus, the optimal solution is $\langle q(\gamma) = 1, q(\beta) = \frac{1}{P(\gamma)}(1 - P(\beta)\frac{\Delta A_1(\gamma)}{\Delta A_1(\beta)})\rangle$.

[24] Breinlich (2011) examined stock market reactions to the 1989 Canada-United States Free Trade Agreement (CAFTA). He found that "increases in the likelihood of ratification led to stock market gains of exporting firms relative to non-exporters", and used this to impute increased "expected per-period profits of exporters by around 6-7% relative to non-exporters."

[25] Extending the model, suppose the $\mathcal{II}$ is a large bond speculator, the PM is the European Union, and the policy is a package guaranteeing bonds against default, requiring austerity measures, and giving loans and aid to Greece. The default risk and the effectiveness of the policy are both uncertain. The bond holders (and Greek leaders) might prefer the EU to provide the maximal aid, but it may not be worth the cost to the EU. The EU could "announce consideration" of a policy, implying a certain conditional probability of execution, and see how the markets react. The *direction* of the likely effect is known (bonds will increase in value and yields will decline) but the magnitude of the effect will determine whether to execute the policy.

Solution: In summary, the optimal policy is

$$\langle q(\gamma) = 1, q(\beta) = \max(0, \frac{1}{P(\gamma)}(1 - P(\beta)\frac{\Delta A_1(\gamma)}{\Delta A_1(\beta)}))\rangle,$$

and the possibly binding incentive compatibility constraint is for type β not to mimic type γ. Also, the optimal policy interprets the $\mathcal{II}$'s buying as signal γ and his not buying as signal β.

Case (ii), "Tiger": $\Delta A_1(\beta) > \Delta A_1(\gamma) > 0$, i.e., *the asset's value increases when the policy is executed, less so with the good signal.*

Consider a benefit program such as the USA's Medicare part D that is expected to be somewhat profitable for the drug industry, but will also yield other public benefits. However, depending on the true market structure and true prospects for innovation, the drug industry might be able to use this to reap excess profits at the expense of consumers (see Friedman, 2009).

Solution: induce buy after β, "do nothing" after γ and $\langle q(\gamma) = 1, q(\beta) = \frac{\Delta A_1(\gamma)}{\Delta A_1(\beta)}\rangle$, and only the IC constraint for β not to mimic γ binds.[26]

Contrasting (i) and (ii), we see that the behavior the PM tries to induce after each signal depends on the relative asset gains under the good or bad policy.

Case (iii), "Judo": $\Delta A_1(\beta) > 0 > \Delta A_1(\gamma)$, i.e., *the asset's value increases when policy is executed under the bad signal, and decreases when it is executed under the good signal.*

Here the interaction between the policymaker and the industry is largely zero-sum. The policy may be a tax reform or regulation intended to be harsh and punitive. However, it may backfire, perhaps if the firm finds loopholes, and may actually increase profits (hence the term "Judo").

Solution: induce buy after γ, "do nothing" after β and $\langle q(\gamma) = 1, q(\beta) = \frac{P(\beta)}{1-P(\gamma)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}}\rangle$, and only the incentive compatibility constraint for type γ not to mimic β binds.[27]

[26] Since the policymaker wants to minimize $q(\beta)$, only the first inequality binds in Set 1: $\Delta A_1(\beta)q(\beta) \geq \Delta A_1(\gamma) \geq [\frac{\Delta A_1(\gamma)-\Delta A_1(\beta)}{P(\beta)} + \Delta A_1(\beta)]q(\beta)$, i.e., the constraint for β not to mimic γ. Thus, the solution is derived.

[27] In Set 2: $\Delta A_1(\beta)q(\beta) \geq \Delta A_1(\gamma) \geq [\frac{\Delta A_1(\gamma)-\Delta A_1(\beta)}{P(\beta)} + \Delta A_1(\beta)]q(\beta)$, the only constraint that might bind is the second one, which is re-written as $1 \leq [\frac{1-\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}}{P(\beta)} + \frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}]q(\beta)$. Note that the bracket term is simplified as $\frac{1}{P(\beta)}[1 - \frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)} + P(\beta)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}] = \frac{1-P(\gamma)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}}{P(\beta)}$, which is larger than 1 since $P(\beta) < 1$ and $1 - P(\gamma)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)} > 1$. Thus, we derive the solution.

Case (iv), "Weapon": $0 > \Delta A_1(\beta) > \Delta A_1(\gamma)$, i.e., *the asset's value decreases when the policy is executed, more so with the good signal.*

This describes a policy intended/expected to severely reduce profits, perhaps to reduced these profits are seen as monopoly or monopsony rents, and thus reducing them may increase consumer surplus. There may be some cost to administering this policy. It may require a severe regulatory burden so it will only be worth doing if it has a major "trust busting" effect. Advocates of this policy may argue that it will reduce "excess profiteering" by monopolists and oligopolists. Opponents may argue it will have little effect on rents, as the oligopolists will find ways to evade it, yet it will lead to large bureaucratic costs and negative unintended consequences for consumers. This policy is a "weapon" worth using only if it is fierce enough.

Solution: induce buy after β, "do nothing" after γ and $\langle q(\gamma) = 1, q(\beta) = \frac{P(\beta)}{1-P(\gamma)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}}\rangle$, and only the IC constraint for type γ not to mimic β binds.[28]

Case (v), "Chemotherapy": $0 > \Delta A_1(\gamma) > \Delta A_1(\beta)$, i.e., *The asset's value decreases when the policy is executed, more so with the bad signal.*

This may reflect a tax increase or increased regulatory burden on industry, but one that is not intended to be excessively burdensome; for example, the Polish mining tax (section 2.1). This also may reflect a stricter price cap for a regulated industry such as a utility; the government wants to limit profits but not to bankrupt the firm(s). Macroeconomic policy raising interest rates to respond to inflation might be similarly characterized. Like *chemotherapy*, this policy is expected to do some damage, but it is only successful if it does not harm the patient (or asset) too much.

Solution: induce buy after γ, "do nothing" after β and $\langle q(\gamma) = 1, q(\beta) = \frac{-\Delta A(\gamma)}{-\Delta A_1(\beta)}\rangle$, and only the incentive compatibility constraint for γ not to mimic β binds.[29]

[28] Since the policymaker wants to minimize $q(\beta)$ in Set 2 ($[-\Delta A_1(\beta)]q(\beta) \leq [-\Delta A_1(\gamma)] \leq [\frac{-\Delta A_1(\gamma)+\Delta A_1(\beta)}{P(\beta)} - \Delta A_1(\beta)]q(\beta)$), the second inequality binds, which is re-written as: $1 \leq [\frac{1-\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}}{P(\beta)} + \frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}]q(\beta)$.

Note that the bracket term is simplified as $\frac{1}{P(\beta)}[1 - \frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)} + P(\beta)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}] = \frac{1-P(\gamma)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}}{P(\beta)}$, which is larger than 1 since $(1 - P(\gamma)\frac{\Delta A_1(\beta)}{\Delta A_1(\gamma)}) > 1 - P(\gamma) = P(\beta)$. Thus, we derive the solution.

[29] We consider Set 1: $[-\Delta A_1(\gamma)] \leq [-\Delta A_1(\beta)]q(\beta) \leq [\frac{-\Delta A_1(\beta)+\Delta A_1(\gamma)}{P(\gamma)} - \Delta A_1(\gamma)]$ where the last bracket term is positive. Since the policymaker wants to minimize $q(\beta)$, the first inequality binds. Thus, we derive the solution.

Case (vi), "Paternalist": $\Delta A_1(\gamma) > 0 > \Delta A_1(\beta)$, i.e., *The asset's value increases when policy is executed under a good signal, and decreases when it is executed after a bad signal.*

This policy may be directly designed to benefit the industry, such as a change in regulations meant to deter destructive competition, or allow coordination on an industry-standard. This may apply to patent reform, or to a complicated change in trade agreements or in immigration law.[30] Another policy that is contested in this way is putting limits on CEO pay. This may also reflect policy with other goals, e.g., an educational reform, but which is only seen as worth doing if it happens to help (and not harm) some key industries. This might be called a *paternalist* policy because at best it helps an industry achieve higher profits than they could achieve alone, but at worst it represents a misguided government overreach that hurts the private sector.

Solution: induce buy after β, "do nothing" after γ and $\langle q(\gamma) = 1, q(\beta) = 0\rangle$; neither constraint binds.[31]

5 Comparison to alternative assumptions

5.1 Comparison: main model versus buy/short implementation

In many real-world settings investors may be able to short sell an asset without large transactions costs. With short selling, there are potentially six ways of implementing truth-telling behavior: through $\mathcal{II}$ choosing any of "buy, do nothing, or short sell" after a good signal, and choosing any *other* action after a bad signal. Each implementation requires that four IC constraints be satisfied, as after each signal, one action must be preferred over the other two. For tractability we focus on the "buy/short" implementations; IC constraints for these were given in section 3.3.2.

[30] E.g., Breinlich (2011) examined stock market reactions to the 1989 Canada-United States Free Trade Agreement (CAFTA). He found that "increases in the likelihood of ratification led to stock market gains of exporting firms relative to non-exporters", and used this to impute increased "expected per-period profits of exporters by around 6-7% relative to non-exporters."

[31] Consider Set 1: $\Delta A_1(\gamma) \geq \Delta A_1(\beta) q(\beta) \geq [\frac{\Delta A_1(\beta) - \Delta A_1(\gamma)}{P(\gamma)} + \Delta A_1(\gamma)]$. To minimize $q(\beta)$, the only constraint that might bind is the second constraint. Dividing both sides by $\Delta A_1(\beta)$, we derive $q(\beta) \leq [\frac{1 - \frac{\Delta A_1(\gamma)}{\Delta A_1(\beta)}}{P(\gamma)} + \frac{\Delta A_1(\gamma)}{\Delta A_1(\beta)}]$ since $\Delta A_1(\beta)$ is negative. Then $q(\beta) = 0$ satisfies the constraint. Thus, we derive the solution.

In this implementation, the binding constraint, and thus the feasible implementation, depends on whether $P(\gamma) > P(\beta)$. Again, an "asymmetric asset gain" may be traded off against the "relative information advantage" of inducing the less-likely PM choice. All results for the buy/short implementation are derived in the online appendix 1.2.[32]

We summarize the optimal policies for each of the six cases, separately for $P(\gamma) > P(\beta)$ and $P(\gamma) < P(\beta)$. Since $q(\gamma) = 1$ for all of our derived implementations, $q(\gamma)$ is left out of the graphs.

[Figure 3 about here.]

[Figure 4 about here.]

Intuition might suggest that restricting the tools available to the informed investor, e.g., by forbidding short sale, would make it harder for him to "manipulate" the policy outcome and conversely, easier to get him to reveal his private information. However, allowing $\mathcal{II}$ to short sell may *lower* the incentive compatible $q(\beta)$ and thus improve the PM's welfare. There are two potentially countervailing effects. Consider case (v), where the implementation involves "buy if $\sigma = \gamma$." A potentially profitable short-selling opportunity means $\mathcal{II}$ may not prefer to buy after a good signal. On the other hand, allowing short-selling may make it *easier* to dissuade him from buying after a bad signal. Although allowing short sale provides a further tool for $\mathcal{II}$'s deviation, it also can provide him more benefit when he is truthful. Short sale makes both the right-hand side and the left hand side of each incentive compatibility constraint larger; hence, it may become easier (or harder) to enforce truth-telling.

It is easy to see that allowing short selling improves welfare in the following cases. In Figures 4 and 3, the buy/short implementation sets $q(\beta) = 0$ if $\Delta A_1(\gamma) < 0 < \Delta A_1(\beta)$ and $P(\gamma) < \frac{1}{2} < P(\beta)$, while $q(\beta) > 0$ in the main model. The same holds for $P(\gamma) > \frac{1}{2} > P(\beta)$, $\Delta A_1(\beta) < \Delta A_1(\gamma) < 0$, $1-2P(\gamma)\frac{\Delta A_1(\gamma)}{\Delta A_1(\beta)} < 0$. Similarly, for $P(\gamma) > \frac{1}{2} > P(\beta)$, $0 < \Delta A_1(\gamma) < \Delta A_1(\beta)$, $1-2P(\gamma)\frac{\Delta A_1(\gamma)}{\Delta A_1(\beta)} < 0$. Figures 5 and 6 depict the optimal policy for specific values of $P(\gamma)$ and $P(\beta)$, allowing a visual comparison of $q(\beta)$ in the two environments (with or without short-selling).

[Figure 5 about here.]

[32] https://davidreinstein.files.wordpress.com/2010/07/online-appendix.pdf

[Figure 6 about here.]

We have demonstrated that for some parameters, allowing short selling helps the PM.[33]

5.2 Comparison: main model versus model with exogenous A_0

Proposition 6 (stating that the PM may set $q(\beta) > 0$ but will always set $q(\gamma) = 1$) depended on the endogeneity of A_0 in $(q(\gamma), q(\beta))$. We now show that $q(\beta) = 0$ and $q(\gamma) < 1$ can be optimal if A_0 is exogenous. We offer a single example (the full case-by-case analysis is available by request). For brevity, we assume short-selling is not allowed. We consider an example resembling case (i) of our main model, letting the $\mathcal{II}$'s payoffs satisfy:

$$\mathbb{E}\, V(e|\gamma) > \mathbb{E}\, V(e|\beta) > 0 > \mathbb{E}\, V(ne|\sigma), \tag{20}$$

where $\mathbb{E}\, V(p|\sigma) := \sum P(s|\sigma)A_1(p, s, i = b) - A_0$. I.e., for this case an $\mathcal{II}$ who buys gets a profit when a good policy is executed, a lower profit when a bad policy is executed, and sustains a loss when the policy is not executed. Inequalities (20) can be derived from the assumption of $\Delta A_1(\gamma) > \Delta A_1(\beta) > 0$ and $A_0 > 0$.

Condition (20) implies that $\mathcal{II}(\gamma)$ prefers *buying* to *doing nothing*, and $\mathcal{II}(\beta)$ prefers *doing nothing* to *buying*. This requires the following incentive compatibility constraints:

$$\sum_s P(s|\gamma)[q(\gamma)(V(e, s, b)) + (1 - q(\gamma))V(ne, s, b)] \geq 0, \tag{21}$$

$$0 \geq \sum_s P(s|\beta)[q(\gamma)V(e, s, b) + (1 - q(\gamma))V(ne, s, b)]. \tag{22}$$

These constraints are simplified into:

$$q(\gamma)[\mathbb{E}\, V(e|\gamma) - \mathbb{E}\, V(ne|\gamma)] \geq -\mathbb{E}\, V(ne|\gamma) \text{ and } -\mathbb{E}\, V(ne|\beta) \geq q(\gamma)[\mathbb{E}\, V(e|\beta) - \mathbb{E}\, V(ne|\beta)]. \tag{23}$$

Given that $q(\beta)$ has no effect on these IC constraints (as it does not affect A_0), the PM sets $q(\beta) = 0$. Note that the same intuition does not go through in the main model with *endogenous*

[33] As we have only solved for the buy/short implementation, we cannot determine whether for some parameter values, the PM becomes worse off by allowing short selling.

A_0; setting $q(\beta) > 0$ in our preferred model will affect the asset's initial price and thus the $\mathcal{II}$'s relative incentives to buy or *not buy*.

On the other hand $q(\gamma)$ cannot be either 0 or 1: if $q(\gamma) = 1$, then the investor with signal $\sigma = \beta$ always succeeds in deceiving the policymaker, and if $q(\gamma) = 0$, then the policymaker never executes the policy, which implies an inefficiency since the policymaker ignores valuable information. Thus $q(\gamma)$ must be a number between 0 and 1. This demonstrates that with a fixed initial asset price A_0, the PM sets $\langle q(\beta) = 0, q(\gamma) < 1 \rangle$ for some ranges of parameters.

5.3 Multiple informed investors

In the real-world, there may be *several* investors with relevant private information. The mechanisms discussed above could be similarly implemented if the PM set rules such as, for case 1, "execute if and only if the price increase reflects the majority of informed investors buying the asset". In general, if the other $\mathcal{II}$'s all take the action that reports the true signal, no $\mathcal{II}$ prefers to deviate, *whether or not his deviation would change the policy.* No $\mathcal{II}$ gains from the policy itself, only from his investment profits; thus the identical IC constraints should guarantee this.

However, there might be a coordination failure among the $\mathcal{II}$s: they might be in an equilibrium that yields each of them lower payoffs, given the PM's announced mechanism. E.g., in case (i), the PM might set probabilities of execution such than any $\mathcal{II}$ gets non-negative profit from buying and inducing execution after a good signal, and a non-positive profit from buying and inducing execution after a bad signal. Nonetheless, all of the $\mathcal{II}$'s might choose "do nothing" after a good signal, inducing non-execution. In such a case no $\mathcal{II}$ would have an incentive to deviate to "buy", as this would not affect the PM's decision and thus the deviation would yield negative profits.

6 Conclusion

We have presented a simple binary-policy framework, intuitively depicting a complete set of relationships between private signals, asset values, and PM's welfare, and solving for the optimal policy mechanism, giving some intuition for the comparative statics. We found that listening is generically strictly welfare improving, and that the optimal mechanism involves sometimes/never executing after a *bad* signal, but always executes after a *good* signal. We compared our main model's results to the "buy/short" implementation, showing that allowing short-selling can *improve* the PM's welfare. We considered the robustness of our model to allowing multiple informed investors, and demonstrated the sensitivity of our results to our assumption that the initial asset prices are endogenously determined in light of the policy commitments.

Considering all three setups yields an additional insight (not presented above). Consider case vi, the *Paternalist* policy. The first best is achieved in our main model if shorting is not allowed; or, if shorting is allowed (section 5.1), if the good signal is the less likely one. This agrees with an intuitive notion of "aligned incentives": both the government and an *asset holder* want the policy executed if and only if it is *good*. If the A_0 were exogenous, as discussed in section 5.2, the first-best alignment would hold for *any* $P(\gamma)$. $\mathcal{II}$ would care only about the policy's impact on the *asset's fundamental value*; he would always want to buy after a good signal, inducing the "good" policy and making a profit on the asset, and would never want to buy after a bad signal. On the other hand, with an endogenous A_0, if the good signal is the more common one, shorting yields a greater "information advantage" than does buying, implying $q(\beta) > 0$ may be required to deter shorting after a good signal. A similar insight holds for case (iii), i.e., wherever the expected *direction* of the policy impact on the asset depends on the signal. A caution to policymakers: even if ex-post incentives are aligned, the incentive to buy or short sell an asset may not be.

For policies where voters do not have a strong issue identification, policy performance, rather than immediate public opinion, is what matters most.[34] Thus, after floating policy "trial balloons" politicians may listen to both polls and markets.[35] The ability to listen and conditionally commit, perhaps imprecisely, may be embodied in the political system. In a system with several branches of government, the framers of a constitution could either allow an executive (President or Premier)

[34]This is likely to hold for technical "hard issues" (Carmines and Stimson, 1980). Fiorina (1978), among others found some evidence for "retrospective voting"; however, there is debate over its explanatory power (see e.g., Fiorina et al., 2003).

[35]Listening to the market is not equivalent to a referendum; the commitments we describe allow the PM to use the market to *extract private information* about the potential results of policies. Unlike referendum voters, traders in the market are "voting" about what they think is profitable, but not necessarily voting for what benefits them as a private citizen.

to execute policy unilaterally, or require her to put a bill to the legislative branch.

While an executive cannot precisely set the probability of passage after good or bad market news, she can make the initial bill's language more or less palatable. For example, President Obama could have first submitted his health care bill mandating broad birth-control coverage, which presumably would have made it unlikely to pass. If it failed, but the market's reaction was favorable, he could then have submitted a similar bill without the birth-control provision. Constitutional and procedural rules determine the length of time a bill is considered, whether its sponsor can introduce unrelated "riders," and how much it can be adjusted throughout the process. A long deliberation permitting amendments and reconsiderations may allow legislators more time and flexibility to listen to the market; "fast-tracking" will limit this. The bill's precise timeline can also affect the conditional probabilities of execution. If the market's signal about the policy is revealed with a known hazard rate, the date that a bill is scheduled for a vote will determine the probability that the (good or bad) signal will have been revealed.

Adding legislator-specific favors and "pork", or unfavorable "poison" to a bill will also can directly affect the conditional probabilities. Suppose that there is uncertainty about some legislator's preferences, while some legislators' preferences can be identified, e.g., known "swing Democrats" always support the bill after a good signal, oppose it after a bad signal. Suppose that offering pork for a legislator's district raises the probability that she votes for a bill after *any* signal – as long as she is not already certain to vote for (against) the bill. Here, offering pork for a swing Democrat will increase $q(\beta)$ without affecting $q(\gamma)$.

Another such strategy is *committing political capital.* A government or party may come out strongly in support of a policy, putting their credibility on the line, and making it costly (but not impossible) to later vote against the policy if the market reveals a negative signal.[36] This might have *no effect* on the probability of execution after a *good* signal. The strength of the commitment will depend on the number of legislators asked to speak in favor of the policy. Suppose there is asymmetric information over legislators' true preferences, or there are other sources of randomness such as revealed public support. Thus, the likelihood that the bill is passed after a bad signal would

[36] This cost could come from a loss of reputation for managerial expertise as in Prendergast and Stole (1996) or a simple voter dislike or distrust of inconsistency and "flip-flopping", perhaps signaling untrustworthiness.

increase in the number of legislators committed to the policy.

The commitment might also take the form of a policy trial, perhaps one with a high probability of a type-I or type-II error. The PM could commit to follow the results of the trial if they are strongly significant in one direction or the other, which may be a small fraction of the time. Here, the PM might not actually expect to learn from the trial itself; she cares more about how the market reacts to the *announcement* of the trial. Where the trial's results are *not* significant, she can follow the signals generated by $\mathcal{II}$'s behavior.

For more technical policies and technocratic policymakers commitments might involve *explicit* randomization.[37] However, we think it more likely will occur through the less precise methods just described, exploiting randomness inherent in the political system. We have presented suggestive evidence that legislators and executives are already taking the market's reaction into account. Still, even if policymakers do not *explicitly* consider introducing randomness, we present a framework for considering the use of market signals to adapt policy.

We chose to focus on an indirect mechanism because an *explicit direct* mechanism – paying investors for insider information – may not be feasible. Such "handouts" would likely be politically unpopular. It also may be difficult to know *which* investors have the accurate information. The government might ask investors who claim to have inside information to pledge a large enough forfeit – in case they are proven wrong – to induce only truth-tellers to come forward. However, limited liability may make this impossible. Existing asset markets are already suited to deal with these commitment problems. Furthermore, if the government offered commitments to these "informants", it would still have to carefully monitor the informants' asset market positions, as outside investments could undermine their incentives truth telling. To the extent that the government designs explicit mechanisms, the considerations we discuss will still be relevant.

Our model could be extended in several ways. Future work might consider a policymaker who

[37]Wolfers and Zitzewitz (2006) note that if there are several public and private signals of a policy's efficacy, the difference in a conditional asset's value may be difficult to interpret. E.g., the US may only execute a carbon cap if there is severe flooding of the Eastern seaboard; hence the conditional expectation of GDP in the event of a carbon cap may also reflect the expectation of the effect this flooding. Thus, introducing exogenous randomness might have an additional benefit: it may help improve the interpretation of market signals even *without* manipulation. This is scope for future work.

seeks to influence *investor* behavior. Alternately, one could model an investor with an *inherent* interest, e.g., who owns an asset affected by the policy which cannot be sold without costs. The most valuable extensions may be empirical. Economists should seek to identify and measure the *ways* in which particular asset values will be differentially affected by policies, and how this relates to the policies' welfare consequences. Where a connection is found, economists should measure the extent to which information is concentrated in the hands of potential "manipulators." Armed with this knowledge, policymakers could set up a "listening process", bearing in mind the implementation concerns we highlight. In particular, they may need to limit the extent to which good or bad news feeds directly into policy, incorporating literal or approximate policy commitments.

References

Abramowicz, M. (2004). Information markets, administrative decisionmaking, and predictive cost-benefit analysis. *The University of Chicago Law Review*, 933–1020.

Arrow, K. J., R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. O. Ledyard, S. Levmore, R. Litan, P. Milgrom, F. D. Nelson, et al. (2008). The promise of prediction markets. *Science - New York then Washington 320*(5878), 877.

Bagehot, W. (1971). The only game in town.

Bernanke, B. S. and F. S. Mishkin (1997). Inflation targeting: a new framework for monetary policy? Technical report, National Bureau of Economic Research.

Breinlich, H. (2011). Heterogeneous firm-level responses to trade liberalization: A test using stock price reactions.

Carmines, E. G. and J. A. Stimson (1980). The two faces of issue voting. *The American Political Science Review*, 78–91.

Fiorina, M., S. Abrams, and J. Pope (2003). The 2000 us presidential election: Can retrospective voting be saved? *British Journal of Political Science 33*(2), 163–187.

Fiorina, M. P. (1978). Economic retrospective voting in american national elections: A micro-analysis. *American Journal of Political Science*, 426–443.

Friedman, J. N. (2009). The incidence of the medicare prescription drug benefit: using asset prices to assess its impact on drug makers. *Harvard University*.

Hahn, R. and P. Tetlock (2003). Using information markets to improve public decision making. *Harvard Journal of Law and Public Policy, Summer*, 04–18.

Hanson, R. (2013). Shall we vote on values, but bet on beliefs? *Journal of Political Philosophy*.

Hanson, R. and R. Oprea (2009). A manipulator can aid prediction market accuracy. *Economica 76*(302), 304–314.

Kau, J. B., J. S. Linck, and P. H. Rubin (2008). Do managers listen to the market? *Journal of Corporate Finance 14*(4), 347–362.

Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets 3*, 205–258.

Milani, F. (2010). Public option and private profits. *Applied health economics and health policy 8*(3), 155–165.

Prendergast, C. and L. Stole (1996). Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning. *Journal of Political Economy*, 1105–1134.

Spence, M. (1973). Job Market Signaling. *Quarterly Journal of Economics 87*(3), 355–374.

Sumner, S. and A. L. Jackson (2008). Using prediction markets to guide global warming policy.

Wolfers, J. and E. Zitzewitz (2006). Prediction markets in theory and practice. Technical report, national bureau of economic research.

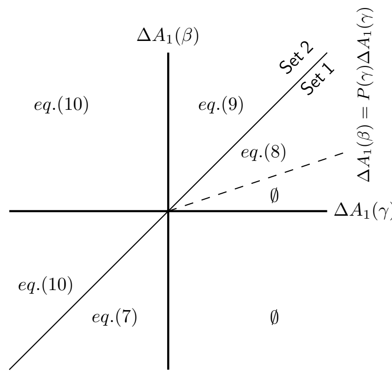

Figure 1: Binding constraints without short-selling

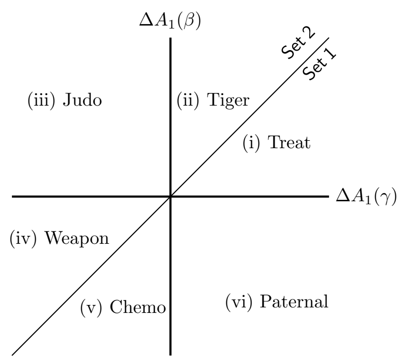

Figure 2: Taxonomy of incentive alignment between PM and an *asset-owner*

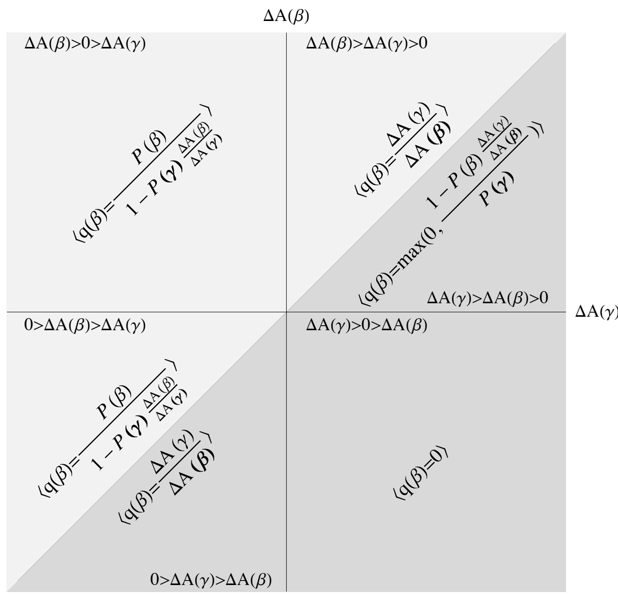

Figure 3: Without short-selling

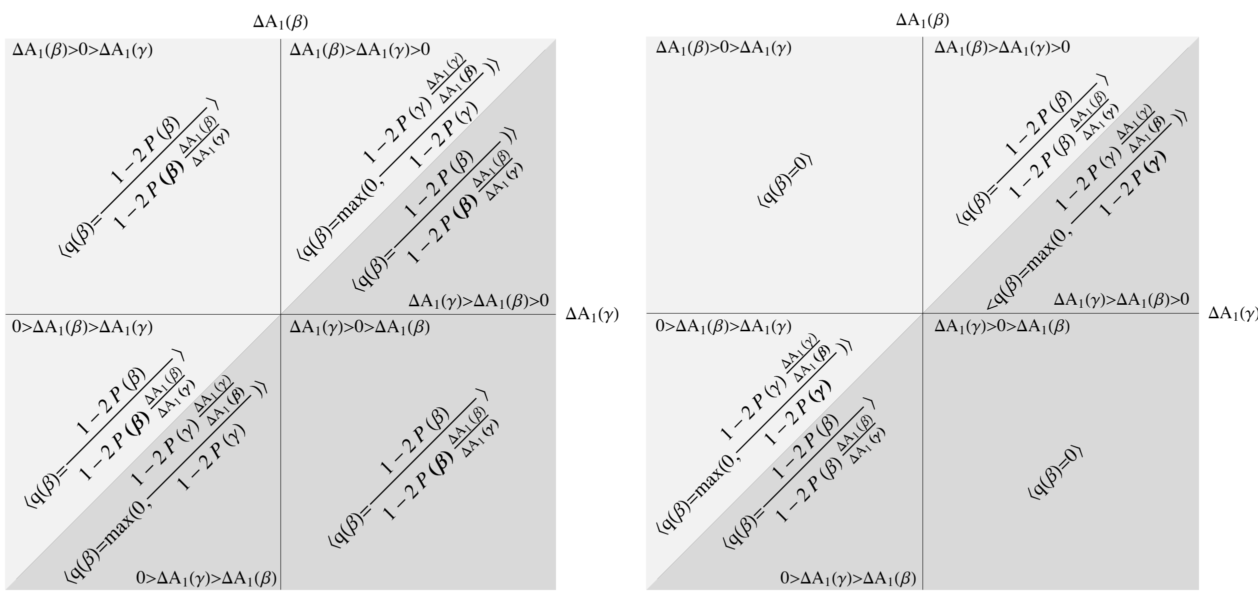

$P(\gamma) > \frac{1}{2} > P(\beta)$ $\qquad$ $P(\gamma) < \frac{1}{2} < P(\beta)$

Figure 4: Allowing short-selling

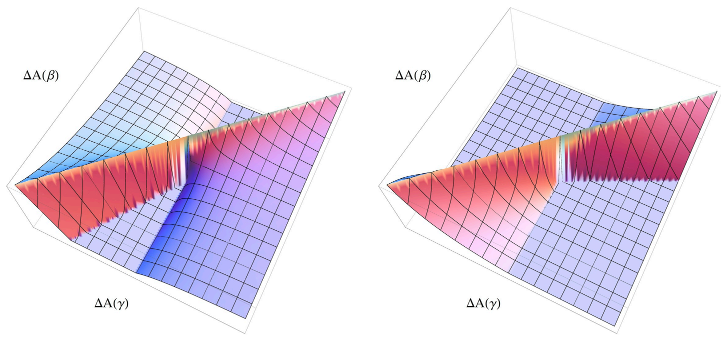

$P(\gamma) = \frac{3}{4} > P(\beta) = \frac{1}{4}$ $P(\gamma) = \frac{1}{4} < P(\beta) = \frac{3}{4}$

Figure 5: With short-selling, suppressing "do nothing"

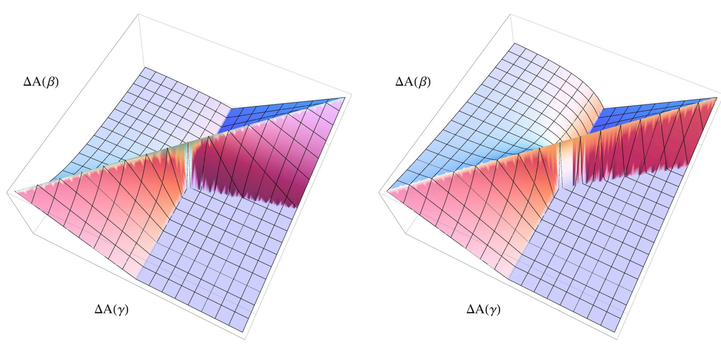

$P(\gamma) = \frac{3}{4} > P(\beta) = \frac{1}{4}$ $P(\gamma) = \frac{1}{4} < P(\beta) = \frac{3}{4}$

Figure 6: Without short-selling