Muddled Information[*]

Alex Frankel[†]   Navin Kartik[‡]

December 3, 2014

Abstract

We study a model of signaling in which agents are heterogeneous on two dimensions. An agent's natural action is her preferred action in the absence of signaling concerns. Her gaming ability parameterizes the cost of increasing actions. Equilibrium behavior muddles information about both dimensions. As incentives to take higher actions increase—due to higher stakes or more easily manipulated signaling technology—more information is revealed about gaming ability, and less about natural actions. Showing agents' actions to new observers may thus worsen information for existing observers. We discuss applications to credit scoring, school testing, and web search.

JEL Classification: D72; D82; D83

[†]University of Chicago Booth School of Business; afrankel@chicagobooth.edu.

[‡]Columbia University, Department of Economics; nkartik@columbia.edu.

Outline

1. Introduction

In many signaling environments, there is a concern that gaming behavior by agents can lead to "muddled" information. Educators worry that rich students have better access to SAT tutoring and test preparation than do poor students, and so the test may be a flawed measure of underlying student quality. Indeed, in March 2014, the College Board announced plans to redesign the SAT test, in part to "rein in the intense coaching and tutoring on how to take the test that often gave affluent students an advantage." (The New York Times, 2014) Google tries to prevent undesirable search engine optimization from contaminating the relevance of its organic search results. The Fair Isaac Corporation keeps its precise credit scoring formula secret to make it more difficult for consumers to game the algorithm.

In canonical signaling models (e.g. Spence, 1973), the only welfare cost from gaming, if any, is through an increase in costly effort.[1] Even though gaming may induce an inefficient rat race, it does not ultimately lead to a reduction in the information available to the market. In the absence of restrictions that may be viewed as artificial, there are typically equilibria that fully reveal agents' private information under standard conditions, viz., when the Spence-Mirrlees single-crossing assumption holds.

Motivated by the aforementioned applications, this paper studies how gaming can worsen market information. We introduce a new model of signaling in which agents have two-dimensional types. Both dimensions affect an agent's cost of sending a one-dimensional signal. The first dimension is an agent's *natural action*, which is the action taken in the absence of signaling concerns.[2] The second dimension is an agent's *gaming ability*, which parameterizes the costs of increasing actions beyond the natural level. In the credit scoring application, the signal is an agent's credit score; the natural action is the score the agent would obtain if this score would not be disseminated; and gaming ability determines how costly it is for an agent to increase her score. In the testing application, the natural action is the test score a student would receive without studying, and gaming ability captures how easily the student can increase her score by studying.

We assume that agents care about influencing a market's belief about their quality on

---

[1] By "gaming" we simply mean strategic behavior, which in a signaling context is an agent's choice of what signal to send. More broadly, concerns about gaming have been studied in other contexts; for example, Ederer et al. (2014) recently address it in a multi-tasking moral-hazard problem.

[2] Throughout the paper, we use the terms "signal" and "action" synonymously.

one of the two dimensions, which we refer to as the *dimension of interest*.[3] We are primarily motivated by situations in which the dimension of interest is the natural action. With credit scoring, the lending market may reward those perceived to have higher natural actions because they are lower risk; the market does not care about an agent's gaming ability because this trait merely reflects an agent's knowledge about how credit scores are computed and can be manipulated. Similarly for search engine optimization, where higher natural actions correspond to more relevant web pages. On the other hand, there are some applications where the dimension of interest is the gaming ability. For example, in the testing environment, gaming ability would not be of interest to the market if it solely represents "studying to the test", but gaming ability may correlate with the ability to study more broadly. In that case colleges or employers might value agents with higher gaming ability even more than those with higher natural actions. Or, in a standard job-market signaling model, gaming ability may be correlated with intelligence and work ethic, while the natural action—the amount of education that would be acquired if it were irrelevant to job search—may capture a dimension of preferences that is unrelated to job performance.

Our model delivers insights into how the combination of gaming ability and natural actions affects the information revealed to the market. The core of our analysis concerns comparative statics about the amount of equilibrium information on the dimension of interest as the costs and benefits of signaling vary.

In our formulation, detailed in Section 2, each dimension of an agent's type—natural action or gaming ability—satisfies a single-crossing property. Thus, the effects of heterogeneity on any one dimension alone are familiar. Indeed, if we were to assume homogeneity of natural actions and the dimension of interest to be gaming ability, then our model would be similar to a canonical signaling environment such as Spence (1973). If instead gaming ability were homogeneous and the dimension of interest were the natural action, then our model would share similarities with, for example, Kartik et al. (2007).[4] In both cases, full separation would be possible.

The novelty of this paper concerns the interaction of heterogeneity across the two dimensions of an agent's type. There will be *cross types*—one with higher natural action but lower

---

[3]More generally, some combination of beliefs on both dimensions could be relevant. We illustrate in Section 5 how our main themes can be extended to the more general case.

[4]Other signaling models with heterogeneity in natural actions include Bernheim (1994) and Bernheim and Severinov (2003). There is also a parallel in the literature on earnings management, wherein a market is assumed to observe a firm's reported earnings but not its "natural earnings" (e.g. Stein, 1989).

gaming ability than the other—whose ordering by marginal cost of signaling turns on the level of signaling. Our central assumption is that at low levels of signaling, the marginal cost depends more on natural actions, whereas at high levels of signaling, it depends more on gaming ability. To illustrate, suppose Anne has a natural credit score of 600 but limited gaming ability, while Bob has a lower natural score, 550, but a higher gaming ability. Then Anne's marginal cost of score improvement is likely to be lower than Bob's at scores just above 600, where Anne can start to address the most obvious flaws on her credit history while Bob has already made a lot of changes from his natural behavior. At higher scores around 700, though, Bob's marginal cost of improvement is lower—both Anne and Bob must engage in a lot of gaming to reach this level, and Bob is the one who knows more about how to game or is better at it.

Consequently, the market is faced with *muddled information*: equilibrium actions depend on both dimensions of an agent's type. While agents who take higher actions will tend to have both higher natural actions and higher gaming ability, any observed action will generally not reveal either dimension. Intermediate actions might come from an agent with a high natural action and a low gaming ability, an agent with a low natural action and a high gaming ability, or an agent who is in-between on both. Even though the market would like to evaluate an agent on his natural action (or gaming ability), the information revealed about this dimension of an agent's type is muddled with irrelevant information about his gaming ability (or natural action).

Our main results establish that when agents' incentives to take high actions increase—because the costs of signaling go down, for instance, or because the stakes in signaling go up—the muddled information reveals more about an agent's gaming ability and less about her natural action. The intuition is as follows: when incentives to raise actions are low and agents accordingly take relatively low actions, cost differences in signaling are largely driven by agents' natural actions; hence, equilibrium signals are more informative about natural actions than gaming ability. As the incentive to take higher actions increases, agents distort their actions more from the natural levels, in which case cost differences are more reflective of gaming ability than natural actions.

We believe this mechanism is compelling and robust. Moreover, our analysis clarifies that higher stakes do not always lead to less information—rather, they simultaneously generate less information on one dimension and more information on the other. We shed light on the properties of signaling costs required for these comparative statics. In a multidimensional

signaling model where our assumption on marginal costs does not hold, one would not have a strong reason to expect less information on natural actions at higher stakes.

After illustrating the core mechanics in Section 3, Section 4 establishes the comparative statics precisely in a canonical "two-by-two" setting and provides general results for small and large signaling stakes. Section 5 studies a tractable variant of our model, a linear-quadratic-normal specification: signaling benefits are linear in the market posterior; costs take a quadratic form; and the type distributions are bivariate normal. This specification affords a sharp equilibrium characterization and additional comparative-statics results.

In Section 6 we use our results to gain insights into applications. One issue we address is the value of giving agents more information about how to manipulate signals, for example by making the inner working of the signaling technology more transparent. A more transparent algorithm will lower the costs of signaling for all agents, increasing the incentives to take higher actions. Therefore, when the dimension of interest is the natural action, the market becomes less informed as the algorithm is made more transparent. This analysis explains why the College Board tried to keep past SAT questions secret for many years: it improves the informativeness of its test. Likewise, even today, Fair Isaac Corporation is willing to give only a vague discussion of how its credit scores are calculated (Weston, 2011, pp. 8–9), and Google maintains some mystery around the details of its ranking algorithm. (If the dimension of interest were gaming ability, our result would reverse, and one would want to make algorithms more transparent.)

It bears emphasis that it is not gaming *per se* that reduces information about natural actions; for example, if web sites were all equally prone to engage in search engine optimization, then their efforts could wash out and leave observers well informed. Rather, muddled information is driven by the fact that there is unobservable heterogeneity across agents in how prone they are to gaming.[5] This provides an explanation for why, in addition to announcing changes to the SAT itself in March 2014, the College Board also announced provision of free online test preparation to "level the playing field". Such a policy lowers the gaming cost for all agents, which, as noted above, by itself tends to further muddle information about natural actions. However, by disproportionately helping those with low intrinsic gaming ability (i.e. poor families), it reduces heterogeneity on this dimension, which should improve

---

[5]Indeed, in discussing an update to Google's ranking algorithm to combat undesirable search engine optimization, Google engineer Matt Cutts (2012) has blogged that "*a few sites* use techniques that don't benefit users ... [Google] can't divulge specific signals because we don't want to give people a way to game our search results and worsen the experience for users." (emphasis added)

market information about natural actions.

When the dimension of interest is the natural action, our aforementioned results imply a novel tradeoff when making the signal available to new observers. The new, or marginal, observers get information they didn't previously have. However, with more observers tracking her actions, an agent's stakes in signaling have grown. At higher stakes, the signal is less informative. So there is a negative informational externality on those inframarginal observers who already had access to the signal. In the context of credit scoring, allowing employers and insurance companies to use credit reports will improve information in those markets, but at a cost of reducing the information available in the loan market. Subsection 6.2 illustrates that the social value of information across markets can decline after the signal is made available to new markets.

While the literature on signaling with multidimensional types is limited, ours is not the first paper in which muddled information emerges in equilibrium.[6] Related phenomena arise in Austen-Smith and Fryer (2005), Bénabou and Tirole (2006), and Bagwell (2007). Indeed, the particular linear-quadratic-normal (LQN) specification that we discuss in Section 5 overlaps with Fischer and Verrecchia (2000) and Bénabou and Tirole (2006, Section II.B), as elaborated further in Section 5.[7,8] Our analysis emphasizes the general forces that determine the equilibrium amount of information, and the contrasting effects that changes in stakes have on information about different dimensions of an agent's type.

We should note that there have been arguments for incomplete revelation of information even when agents have one-dimensional types satisfying single crossing. Separation may be precluded if there are bounds on the signal space, in which case there can be bunching

---

[6]Early work on signaling with multidimensional types (Quinzii and Rochet, 1985; Engers, 1987) establishes the existence of fully separating equilibria under suitable "global ordering" or single-crossing assumptions. As already noted, our model satisfies single-crossing within each dimension but not globally. See Araujo et al. (2007) for a multidimensional-type model in which it is effectively as though single-crossing fails even within a single dimension, which leads to "counter-signaling" equilibria. Feltovich et al. (2002), Chung and Esö (2013), and Sadowski (2013) make related points in models with a single-dimensional type.

[7]The information muddling—or "signal extraction" problem in their terminology—in the basic version of Bénabou and Tirole (2006) is less of an equilibrium phenomenon and more a direct consequence of the assumptions about agents' preferences. Specifically, in their baseline analysis (Section II.A of their paper), an agent's preferences over combinations of market beliefs and actions are determined by a single-dimensional statistic of the two dimensions of her private information; therefore, at best the market can infer this single-dimensional statistic. In our model, every pair of types has distinct preferences over combinations of market beliefs and actions.

[8]In addition to Fischer and Verrecchia (2000), there are other models in the earnings-management accounting literature that feature muddled information (in contrast to the separating equilibrium of Stein (1989)), for example Dye and Sridhar (2008) and Beyer et al. (2014).

at the edges of the type space (Cho and Sobel, 1990). However, the signal space itself is often a choice object, e.g. for credit or test scores. There are also critiques of the focus on separating equilibria even when these exist (Mailath et al., 1993); recently, Daley and Green (2013) note that separating equilibria need not be strategically stable when the market exogenously receives sufficiently precise information about the agents. Another reason why the market may not be able to perfectly infer the agent's type is that the signaling technology may be inherently noisy (Matthews and Mirman, 1983), although this can again be a choice object (Rick, 2013).

2. The Model

Our model is a reduced-form signaling game. One or more agents independently take observable actions; for convenience, we often refer to "the" agent. The agent has two-dimensional private information—her *type*—that determines her cost of taking a single-dimensional action. The agent then receives a benefit that depends on the public belief about her type that is formed after observing her chosen action.

2.1. Types and signaling costs

The agent takes an action, $a \in A \equiv \mathbb{R}$. The agent's type, her private information, is $\theta = (\eta, \gamma)$, drawn from a cumulative distribution F with compact support $\Theta \subset \mathbb{R} \times \mathbb{R}_{++}$. We write Θ_η and Θ_γ for the projections of Θ onto dimension η and γ respectively. The first dimension of the agent's type, η, which we call her *natural action*, represents the agent's intrinsic ideal point, or the highest action that she can take at minimum cost. The second dimension, γ, which we call *gaming ability*, parameterizes the agent's cost of increasing her action above the natural level: a higher γ will represent lower cost. (It will be helpful to remember the mnemonics θ for *t*ype, η for *n*atural, and γ for *g*aming.) The cost for an agent of type $\theta = (\eta, \gamma)$ of taking action a is given by $C(a, \eta, \gamma)$, also written as $C(a, \theta)$. Using subscripts on functions to denote partial derivatives as usual, we make the following assumption on signaling costs:

Assumption 1. *The cost function $C : \mathbb{R} \times \mathbb{R} \times \mathbb{R}_{++} \to \mathbb{R}$ is differentiable, twice-differentiable except possibly when $a = \eta$, and satisfies:*

    1. For all γ and $a \leq \eta$, $C(a, \eta, \gamma) = 0$.

2. *For all γ and $a > \eta$, $C_{aa}(a, \eta, \gamma) > 0$.*

3. *For all γ and $a > \eta$, $C_{a\eta}(a, \eta, \gamma) < 0$ and $C_{a\gamma}(a, \eta, \gamma) < 0$.*

4. *For any $\underline{\eta} < \overline{\eta}$ and $\underline{\gamma} < \overline{\gamma}$, $\frac{C_a(\cdot, \overline{\eta}, \underline{\gamma})}{C_a(\cdot, \underline{\eta}, \overline{\gamma})}$ is strictly increasing on $[\overline{\eta}, \infty)$ and there exists $a^{or} > \overline{\eta}$ such that $C_a(a^{or}, \overline{\eta}, \underline{\gamma}) = C_a(a^{or}, \underline{\eta}, \overline{\gamma})$.*

Together, Parts 1 and 2 of Assumption 1 say that (i) the natural action $a = \eta$ is an agent's highest cost-minimizing action, with cost normalized to zero; (ii) the agent can costlessly take actions below her natural action ("free downward deviations"); (iii) the marginal cost of increasing her action is zero at her natural action; and (iv) the agent must pay an increasing and convex cost to take actions above this level.[9] Part 3 of the assumption stipulates that the marginal cost of increasing one's action is lower for agents with either higher natural actions or higher gaming ability. Consequently, $C(\cdot)$ satisfies decreasing differences (and hence a single-crossing property): if $\underline{a} < \overline{a}$ and $\underline{\theta} < \overline{\theta}$ in the component-wise order, then

$$C(\overline{a}, \underline{\theta}) - C(\underline{a}, \underline{\theta}) \geq C(\overline{a}, \overline{\theta}) - C(\underline{a}, \overline{\theta}),$$

where the inequality is strict if $\underline{a}$ is no smaller than $\underline{\theta}$'s natural action.

However, the key feature of our model is that owing to the multidimensionality of Θ, the single-crossing property of $C(\cdot)$ does not, in general, hold globally on $A \times \Theta$. The fourth part of Assumption 1 places structure on how $C(\cdot)$ behaves for pairs of *cross types*, where one type, $(\underline{\eta}, \overline{\gamma})$, has a strictly higher natural action but a strict lower gaming ability than the other, $(\overline{\eta}, \underline{\gamma})$. We write $(\underline{\eta}, \overline{\gamma}) <> (\overline{\eta}, \underline{\gamma})$ as shorthand for $\underline{\eta} < \overline{\eta}$ and $\overline{\gamma} > \underline{\gamma}$. At low actions, the type with the lower γ (and higher η) has a lower marginal cost of increasing its action. But this type's marginal cost grows faster than the other type's, and there is some cutoff action, a^{or}, at which the marginal-cost ordering of the two types reverses: at higher actions the type with the higher γ (and lower η) now has a lower marginal cost of increasing its action. We refer to the action a^{or} as the *order-reversing action*.

Assumption 1 implies the existence of another cutoff action, one at which the cross types share an equal signaling cost. We denote this action by a^{ce} and refer to it as the *cost-equalizing action*. For any action below a^{ce}, the type with lower γ (but higher η) bears a lower cost, whereas the relationship is reversed for higher actions.

---

[9]The assumption of free downward deviations is not crucial for our main points, but it simplifies the analysis and is appropriate for many applications. Section 5 considers a specification without free downward deviations and confirms our main insights.

Lemma 1. *For any $\underline{\eta} < \bar{\eta}$ and $\underline{\gamma} < \bar{\gamma}$, there exists $a^{ce} > a^{or}$ such that $C(a^{ce}, \underline{\eta}, \bar{\gamma}) = C(a^{ce}, \bar{\eta}, \underline{\gamma})$. Furthermore, for any $a > \underline{\eta}$, $\mathrm{sign}[C(a, \underline{\eta}, \bar{\gamma}) - C(a, \bar{\eta}, \underline{\gamma})] = \mathrm{sign}[a^{ce} - a]$.*

(The proof of this and all other formal results can be found in the Appendices.)

Figure 1 summarizes the above discussion by illustrating the implications of Assumption 1 when Θ consists of four types: a low type, $(\underline{\eta}, \underline{\gamma})$; two intermediate cross types, $(\underline{\eta}, \bar{\gamma})$ and $(\bar{\eta}, \underline{\gamma})$; and a high type $(\bar{\eta}, \bar{\gamma})$.
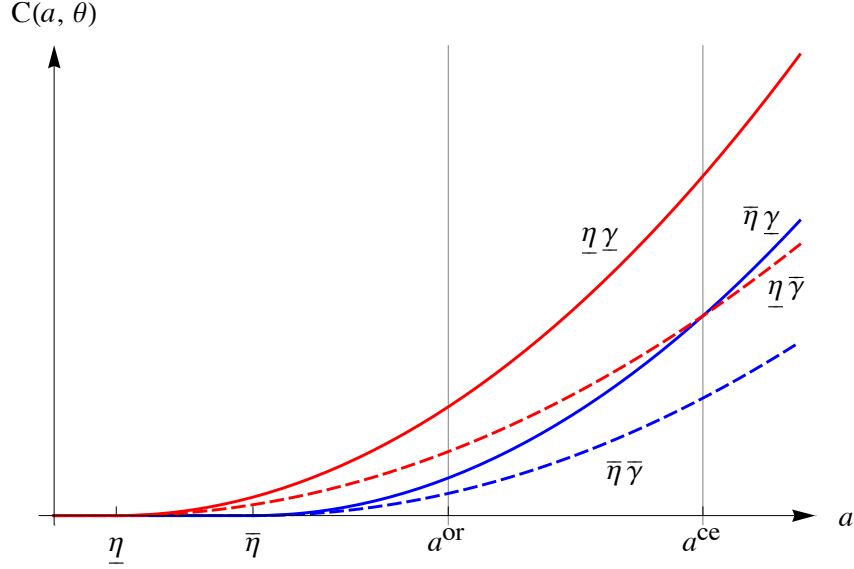

Figure 1 – Cost curves for $\Theta = \{\underline{\eta}, \bar{\eta}\} \times \{\underline{\gamma}, \bar{\gamma}\}$ with $\underline{\eta} < \bar{\eta}$ and $\underline{\gamma} < \bar{\gamma}$. The solid red curve is $C(\cdot, \underline{\eta}, \underline{\gamma})$ and the solid blue curve is $C(\cdot, \bar{\eta}, \underline{\gamma})$; the dashed red curve is $C(\cdot, \underline{\eta}, \bar{\gamma})$ and the dashed blue curve is $C(\cdot, \bar{\eta}, \bar{\gamma})$.

Example 1. The canonical functional form to keep in mind is $C(a, \eta, \gamma) = \frac{c(a,\eta)}{\gamma}$. In this case, the first three parts of Assumption 1 reduce to only requiring the analogous properties on $c(a, \eta)$ (with the second requirement of Part 3 automatically ensured). Since $\frac{C_a(a, \bar{\eta}, \underline{\gamma})}{C_a(a, \underline{\eta}, \bar{\gamma})} = \frac{c_a(a, \bar{\eta})}{c_a(a, \underline{\eta})} \frac{\bar{\gamma}}{\underline{\gamma}}$, a sufficient condition for the fourth part is that for any $\underline{\eta} < \bar{\eta}$, $\frac{c_a(\cdot, \bar{\eta})}{c_a(\cdot, \underline{\eta})}$ is strictly increasing on the relevant domain with $\lim_{a \to \infty} \frac{c_a(a, \bar{\eta})}{c_a(a, \underline{\eta})} = 1$. In particular, given any exponent $r > 1$, the cost function $C(a, \eta, \gamma) = \frac{(a-\eta)^r}{\gamma}$ for $a > \eta$ (and 0 for $a \leq \eta$) satisfies Assumption 1. This family will be our leading example. In this family, for any pair of cross types $(\underline{\eta}, \bar{\gamma}) <> (\bar{\eta}, \underline{\gamma})$, a^{or} and a^{ce} can be computed as

$$ a^{or} = \frac{\bar{\eta}\bar{\gamma}^k - \underline{\eta}\underline{\gamma}^k}{\bar{\gamma}^k - \underline{\gamma}^k} \quad \text{and} \quad a^{ce} = \frac{\bar{\eta}\bar{\gamma}^l - \underline{\eta}\underline{\gamma}^l}{\bar{\gamma}^l - \underline{\gamma}^l}, \quad \text{where } k = \frac{1}{r-1} \text{ and } l = \frac{1}{r}. \qquad \parallel $$

8

2.2. Beliefs, payoffs, and equilibrium

We suppose that there is one *dimension of interest* about the agent's type, $\tau \in \{\eta, \gamma\}$.[10] After observing the agent's action, an observer or "market" forms a posterior belief $\beta_\tau \in \Delta(\Theta_\tau)$ over the dimension of interest, where $\Delta(X)$ is the set of probability distributions on a set X. We assume the market evaluates the agent by the expected value of her type on dimension τ, which we denote $\hat{\tau}$: $\hat{\tau} \equiv \mathbb{E}_{\beta_\tau}[\tau]$. We refer to $\hat{\tau}$ as the *market belief* about the agent. The gross benefit from signaling for an agent who induces belief $\hat{\tau}$ is denoted $V(\hat{\tau}; s)$, where $s \in \mathbb{R}_{++}$ parameterizes the signaling *stakes*.[11] Notice that this benefit is independent of an agent's type. We maintain the following assumption about the benefit function.

Assumption 2. *The benefit function, $V(\hat{\tau}; s)$, is continuous and satisfies:*

1. *For any s, $V(\cdot; s)$ is strictly increasing.*

2. *$V(\cdot)$ has strictly increasing differences, i.e. for any $\hat{\tau}' > \hat{\tau}$, $V(\hat{\tau}'; \cdot) - V(\hat{\tau}; \cdot)$ is strictly increasing.*

3. *For any $\hat{\tau}' > \hat{\tau}$, $V(\hat{\tau}'; s) - V(\hat{\tau}; s) \to \infty$ as $s \to \infty$ and $V(\hat{\tau}'; s) - V(\hat{\tau}; s) \to 0$ as $s \to 0$.*

In other words, the agent prefers higher market beliefs; higher beliefs are more valuable when the stakes are higher; as stakes grow unboundedly, the benefit of inducing any higher belief grows unboundedly; and analogously as stakes vanish. A simple example that we will sometimes refer to is $V(\hat{\tau}; s) = sv(\hat{\tau})$ for some strictly increasing function $v(\cdot)$. We highlight that a higher value of stakes does *not* represent greater direct benefits from taking higher actions; rather, it implies greater rewards to inducing higher market beliefs.

Combining the benefits and costs of signaling, an agent of type $\theta = (\eta, \gamma)$ who plays action a yielding beliefs $\hat{\tau}$ on dimension τ has net (von-Neumann Morgenstern) payoff $V(\hat{\tau}; s) - C(a, \theta)$. This payoff function together with the prior distribution of types, F, induces a signaling game in the obvious way. We focus on Bayesian Nash equilibria—simply equilibria, hereafter—of this signaling game. Since the market does not take any actions, Bayesian Nash equilibria coincide with Perfect Bayesian equilibria in our setting. Given that

---

[10]We will abuse notation by using the same symbols to denote both the dimension of interest and the agent's type realization on the dimension of interest.

[11]More generally, one may be interested in functions $V(\hat{\eta}, \hat{\gamma}; s)$ that depend on beliefs about both dimensions of an agent's type; we study one such specification in Subsection 5.3.

the agent cares about the market belief on only one dimension of his type, equilibria cannot generally fully reveal both dimensions (cf. Stamland, 1999). We say that an equilibrium is *separating* if it fully reveals the agent's private information on the dimension of interest; an equilibrium is *pooling* if it reveals no information on the dimension of interest; and an equilibrium is *partially-pooling* if it is neither separating nor pooling. We say that two equilibria are *equivalent* if they share the same mapping from types to (distributions over) the posterior belief, β_τ, and the same mapping from types to (distributions over) signaling costs.

The assumption of free downward deviations implies that equilibrium beliefs must be monotone over on-path actions. More precisely, following the convention that $\sup \emptyset = -\infty$:

Lemma 2. *In any equilibrium, if $a' < a''$ are both on-path actions, then $\hat{\tau}(a') \leq \hat{\tau}(a'')$. Moreover, for any equilibrium, there is an equivalent equilibrium in which (i) if $a' < a''$ are both on-path actions, then $\hat{\tau}(a') < \hat{\tau}(a'')$, and (ii) if a is an off-path action,*

$$\hat{\tau}(a) = \max\{\min \Theta_\tau, \sup\{\hat{\tau}(a') : a' \text{ is on path and } a' < a\}\}.$$

The first statement of the lemma is straightforward. Property (i) of the second statement follows from the observation that if there are two on-path actions $a' < a''$ with $\hat{\tau}(a') = \hat{\tau}(a'')$, then one can shift any type's use of a'' to a' without altering either the market belief at a' or any incentives. We will refer to this property as *belief monotonicity*, and without loss of generality, restrict attention to equilibria that satisfy it. Part (ii) of the second statement of Lemma 2 assures that there would be no loss in also requiring weak monotonicity of beliefs off the equilibrium path. However, we will not impose this property because it is expositionally more convenient for some of our arguments to assume that off-path actions are assigned "worst beliefs" (i.e. a belief corresponding to the lowest type on the dimension of interest).

Remark 1. By free downward deviations, there is always a pooling equilibrium in which all types play $a = \min \Theta_\eta$.

Remark 2. If either the agent's natural action, η, or her gaming ability, γ, were known to the market—i.e., there is private information on only one dimension—and the dimension of interest is the dimension on which there is private information, then there would be a separating equilibrium. More generally, if there are no cross types in Θ then there is a separating equilibrium due to the single-crossing property.

Remark 3. Consider our leading family of cost functions: for some $r > 1$, $C(a, \eta, \gamma) = \frac{(a-\eta)^r}{\gamma}$ for $a > \eta$ (and 0 for $a \leq \eta$). With the change of variables $e \equiv \frac{a-\eta}{\gamma^{1/r}}$, the agent's net payoff

function can rewritten as $V(\hat{\tau}; s) - e^r \cdot \mathbb{1}_{\{e>0\}}$. The problem is thus isormophic to one where the agent chooses "effort" e at a type-independent cost e^r (if $e > 0$, and 0 otherwise) and the market observes "output" $a = \eta + \gamma^{1/r}e$ before forming its expectation of $\tau \in \{\eta, \gamma\}$. Here, η can be interpreted as the agent's baseline talent while γ parameterizes her marginal product of effort.

2.3. Interpretations of heterogeneity

Before turning to the analysis, we pause to discuss some interpretations of heterogeneity among agents, which also allows us to note additional applications of the model.

Heterogeneity in natural actions is uncontroversial in most applications we have in mind, as it simply reflects that agents would be choosing different actions absent signaling concerns. This could be viewed as a reduced form for factors outside the model. Even in the context of education signaling (Spence, 1973), it is plausible that workers have different intrinsic preferences over education, for example owing to heterogenous effects on human capital.

There are a number of (non-exclusive) sources for heterogeneity in gaming ability. An obvious one is underlying skills: some students may simply be more facile at studying, web designers may vary in their skill with respect to search engine optimization, or agents may incur different lying costs in strategic information transmission (cf. fn. 27). But there are other sources.

Differences in information. Agents are likely to vary in their understanding of how to game a signal. This may be due to experience or other factors.[12] In his book "The Big Short", Michael Lewis (2011, p. 100) describes how Wall Street firms raced to find ways to manipulate credit rating agencies' algorithms during the subprime mortgage crisis (emphasis added): "The models...were riddled with...[gaming] opportunities. *The trick was finding them before others did*—finding, for example, that both Moody's and S&P favored floating-rate mortgages with low teaser rates over fixed-rate ones." Alternatively, many "white hat" or benign search engine optimization methods are considered best practices for web design—using appropriate keywords, linking to relevant sites, coding web sites in ways that are easily

---

[12]Sometimes, experience with a signaling technology is endogenous and could be considered as part of gaming; e.g. a student must decide how many practice tests to take. But in other cases experience may be thought of as more akin to a fixed type: web designers have different levels of experience prior to their role at their current job; a consumer who has applied for a loan before may have a better understanding of the credit scoring market.

parsed by search engines—but professional web designers would likely be more attuned to these techniques than amateurs.[13]

Information differences may be especially relevant for consumers trying to improve their credit scores, but who may not know what strategies to use. One should obviously make on-time payments. But, according to the Fair Isaac Corporation, payment history counts for only 35% of one's FICO score. The credit score also puts a 30% weight on credit utilization (having a higher credit limit and a lower balance on one's card is helpful), 15% on length of credit history, 10% on recent searches for credit, and 10% on types of credit used. The formula is presumably derived from correlations with loan defaults, but not in ways that are transparent to consumers. A large genre of books (e.g. Weston, 2011) and other products is devoted to explicating strategies to improve one's credit score.

Differences in preferences for gaming. Students may vary in how much they enjoy or dislike studying. Those who enjoy it more will face a lower cost of increasing their test scores above their natural levels. When gaming involves monetary costs, we can also interpret those with access to more money as having a lower disutility of spending money relative to the signaling benefits. The College Board worries that richer students can better afford private tutoring and test prep courses for its SAT test.[14] More profitable web sites have a lower relative cost of spending money on search engine optimization to improve their ranking.

Moreover, when there are somewhat unethical approaches to gaming signals, differences in "integrity" could affect agents' preferences towards gaming. Students can not just study for exams, but also find ways to cheat. Some consumers pay "credit doctors" to raise credit scores in ambiguously legal ways, such as reporting fake accounts with low usage and/or great repayment records (Businessweek, 2008). Similarly, colleges can and sometimes do engage in dubious activities to affect their US News & World Report (US News) rankings. For example, in 2008, Baylor University offered monetary rewards for already-admitted freshmen to retake the SAT, presumably to improve the average score of its incoming freshman class and thereby raise its US News ranking (The New York Times, 2008).

---

[13]Many of the these white-hat techniques are encouraged by search engines as a way of improving the quality of web sites to users. Incorporating this aspect into our model would be analogous to making education intrinsically productive in Spence (1973); while it would complicate the analysis, it would not fundamentally alter the insights.

[14]Our analysis will imply that this concern is only warranted when colleges or other interested parties are not able to directly observe all relevant aspects of students' economic backgrounds, at least not perfectly. Indeed, to the extent that observers do obtain some exogenous information about his dimension, they would wish to treat students with identical scores differently—an issue that is much debated in college admissions.

Differences in signaling benefits. One can also interpret an agent's gaming ability as parameterizing his *private stakes* or benefit from signaling rather than cost. After all, if for simplicity $C(a, \eta, \gamma) = c(a, \eta)/\gamma$, then the net payoff function $V(\hat{\tau}; s) - \frac{c(a,\eta)}{\gamma}$ is indistinguishable from the payoff function $\gamma V(\hat{\tau}; s) - c(a, \eta)$. Intuitively, it is identical to state that rich students have a lower cost of paying for coaching relative to their benefit from higher scores, or that these students have a higher benefit (in dollar terms) relative to the monetary cost of such coaching. In addition to such differences driven by wealth or amounts of money at stake, web site owners might simply differ on their interest in attracting more hits, or colleges may differ in the benefits they obtain from boosting their US News rankings.

2.4. Measuring information and welfare

The natural measure for *agent welfare* is the expected payoff across types, $\mathbb{E}[V(\hat{\tau}; s) - C(a, \theta)]$. We say that *allocative efficiency* is the expected benefit from signaling gross of signaling costs: $\mathbb{E}[V(\hat{\tau}; s)]$. Besides these standard quantities, our focus in this paper will be on the amount of information revealed on the dimension of interest about the agent's type, τ.

Recall that $\beta_\tau \in \Delta(\Theta_\tau)$ is the market posterior (the marginal distribution) over the dimension of interest, τ. From the ex-ante point of view, any equilibrium induces a probability distribution over β_τ, which is an element of $\Delta(\Delta(\Theta_\tau))$. In any equilibrium, the expectation over β_τ must be the prior distribution over τ. Equilibria may differ, however, in the distribution they induce over β_τ. A separating equilibrium is *fully informative* on the dimension of interest: after any on-path action, β_τ will be degenerate. A pooling equilibrium is *uninformative*: after any on-path action, β_τ is simply the prior over τ. To compare informativeness of equilibria in between these two extremes, we will use the canonical partial ordering of Blackwell (1951, 1953). We say that a distribution of beliefs or posteriors is *more informative* than another if the former is a mean-preserving spread of the latter.[15] An equilibrium e' is more informative about dimension τ than an equilibrium e'' if the distribution of β_τ under e' is more informative than that under e''.

We will also be interested in information specifically about the posterior mean of the dimension of interest, i.e., information about $\hat{\tau}$ rather than about the entire β_τ. An equilibrium e' is *more informative about $\hat{\tau}$* than e'' if the distribution of $\hat{\tau}$ under e' is a mean-preserving spread of the distribution under e''. In particular, an equilibrium is uninformative about

---

[15]Throughout this paper, we use the terminological convention that binary comparisons are always in the weak sense (e.g., "more informative" means "at least as informative as") unless explicitly indicated otherwise.

$\hat{\tau}$ if the distribution it induces over $\hat{\tau}$ is a point mass at the prior mean of τ. Note that an equilibrium can be uninformative about $\hat{\tau}$ even if the equilibrium is informative; on the other hand, an equilibrium is fully informative about $\hat{\tau}$ if and only if it is fully informative. In general, the partial order on equilibria generated by information about $\hat{\tau}$ is finer than that generated by Blackwell information: more informative implies more informative about $\hat{\tau}$, but more informative about $\hat{\tau}$ does not necessarily imply more informative.[16]

Comparing equilibria according to their informativeness is appealing—when possible—because of the fundamental connection between this statistical notion and economic quantities of interest. Consider, in particular, allocative efficiency $\mathbb{E}[V(\hat{\tau}; s)]$. If the benefit function $V(\cdot; s)$ is convex, then for fixed stakes a more informative distribution of beliefs on dimension τ implies higher allocative efficiency. Indeed, as $V(\cdot; s)$ depends only on $\hat{\tau}$ rather than β_{τ}, a more informative distribution about $\hat{\tau}$ implies higher allocative efficiency. If $V(\cdot; s)$ is concave, the opposite holds: allocative efficiency is maximized by pooling together all types and leaving the market belief at the prior. For a linear $V(\cdot; s)$, allocative efficiency is independent of the information about $\hat{\tau}$.

We are primarily motivated by situations where information has an allocative benefit, corresponding to a weakly convex benefit function. Consider, for instance, a market in which consumers (agents) bring differing service costs to a firm that provides them a product. Revealing information about consumer costs means that higher cost consumers will be offered higher prices. This information transfers surplus from high cost to low cost consumers but also improves the efficiency of the allocation (strictly, so long so long as demand is not perfectly inelastic). Appendix A provides an explicit example relating the demand curve for a product to the shape of a convex benefit function.

3. The Logic with Two Types

To illustrate the core mechanism and implications of muddled information, we begin with a simple two type example in which $\Theta = \{\theta_1, \theta_2\}$. If the two types are not cross types, then a separating equilibrium exists. Throughout this section, then, we consider the case of cross types: $\theta_1 = (\underline{\eta}, \overline{\gamma}) <> (\overline{\eta}, \underline{\gamma}) = \theta_2$. For ease of exposition, we refer to θ_1 as "the gamer" type

---

[16]If Θ_{τ} is binary, then the posterior mean $\hat{\tau}$ is a sufficient statistic for the posterior distribution β_{τ}. In this case, more informative about $\hat{\tau}$ does imply more informative, and uninformative about $\hat{\tau}$ implies uninformative.

(i.e., the one with the higher gaming ability), and θ_2 as "the natural."

3.1. Dimension of interest is the natural action

Let the dimension of interest be the natural action: $\tau = \eta$. When can agents' private information be fully revealed? In a separating equilibrium, the gamer type θ_1 plays some action a_1 while the natural type θ_2 plays $a_2 > a_1$.[17] The incentive constraints that each type is willing to play its own action over the other's are $C(a_2, \theta_1) - C(a_1, \theta_1) \geq V(\overline{\eta}; s) - V(\underline{\eta}; s) \geq C(a_2, \theta_2) - C(a_1, \theta_2)$. The first inequality says that the gamer type's incremental cost of playing a_2 rather than a_1 is higher than the incremental benefit, while the second inequality says that the incremental benefit is greater than the incremental cost for the natural type. Without loss, we can take $a_1 = \underline{\eta}$, as the gamer would never pay a positive cost to be revealed as the type with the lower natural action. Substituting $a_1 = \underline{\eta}$ and rewriting the incentive constraints yields

$$C(a_2, \theta_1) \geq V(\overline{\eta}; s) - V(\underline{\eta}; s) \geq C(a_2, \theta_2). \tag{1}$$

We find that a separating equilibrium, or equivalently an action a_2 satisfying (1), exists if and only stakes are sufficiently low.[18] The logic behind this result can be seen by considering the two cross types' costs, as illustrated in Figure 1. For small s, separation can be obtained through the natural actions. For moderate values of s, the standard signaling mechanism operates: the natural or "high" type distorts its action upwards to a point that is too costly for the gamer or "low" type to mimic. But as stakes increase the natural type cannot continue to separate by taking ever higher actions, because the single-crossing order reverses and eventually the gamer has lower costs. Formally, we can keep raising stakes without violating the inequalities in (1) until the first inequality binds, and so we choose action a_2 to make $C(a_2, \theta_1)$ as large as possible; the constraint is that we must choose a_2 small enough such that $C(a_2, \theta_1) \geq C(a_2, \theta_2)$. As $C(a_2, \theta_1) \geq C(a_2, \theta_2)$ if and only if a_2 is below the cost-equalizing action a^{ce} (Lemma 1), a separating equilibrium exists for stakes up to a cutoff s_η^*

---

[17]Due to free downward deviations, each type can be mixing over multiple actions in a separating equilibrium; however, such an equilibrium is equivalent to one in which the agent uses a pure strategy. While we set aside such details here to ease exposition, the full range of possibilities is covered by the formal results in Subsection 4.1 for a setting that subsumes the current one.

[18]We need not consider incentive constraints on taking actions other than a_1 or a_2, as we can set beliefs to be $\hat{\eta} = \underline{\eta}$ at all off-path actions.

that solves

$$V(\overline{\eta}; s_\eta^*) - V(\underline{\eta}; s_\eta^*) = C(a^{\mathrm{ce}}, \theta_1).$$

Action a^{ce} acts like an *endogenous* upper bound on actions insofar as separating the two types is concerned. At any action $a > a^{\mathrm{ce}}$ the signaling cost is lower for the gamer type, and hence a cannot be used as a credible signal of a high natural action by the natural type.[19]

While separation is impossible at stakes $s > s_\eta^*$, there are in fact informative equilibria even at arbitrarily high stakes. Recall that two types have the same incremental cost of taking action $a_2 = a^{\mathrm{ce}}$ rather than $a_1 = \underline{\eta}$. (The subsequent logic can be applied to any pair of actions $a_1 < a_2$ for which $C(a_2, \theta) - C(a_1, \theta)$ is the same for both types θ; for any a_1 in $[\underline{\eta}, a^{\mathrm{or}})$, there is a corresponding a_2 in $(a^{\mathrm{or}}, a^{\mathrm{ce}}]$ at which this holds.) At high stakes, both types can mix over this pair of actions: there is an equilibrium where the natural type plays a_2 with a suitably higher probability than the gamer type so that the incremental benefit $V(\hat{\eta}(a_2); s) - V(\hat{\eta}(a_1); s)$ is equal to the incremental cost $C(a_2, \theta) - C(a_1, \theta)$. This partially-pooling equilibrium is informative, as $\hat{\eta}(a_2) > \hat{\eta}(a_1)$. Thus, even though a^{ce} acts like an upper bound on separation, the current setting is very different from a canonical (one-dimensional) signaling model with a bounded signal space where at high stakes there would only be pooling equilibria.[20]

Observation 1. *In the two cross types setting, if the dimension of interest is the natural action, then there is a threshold $s_\eta^* > 0$ such that a separating equilibrium exists if and only if $s \le s_\eta^*$. At higher stakes, a partially-pooling equilibrium exists.*

Nevertheless, as $s \to \infty$, the informativeness of any equilibrium vanishes. In the equilibrium described above it must hold that $\hat{\eta}(a_2) \to \hat{\eta}(a_1)$ as $s \to \infty$ in order to maintain $V(\hat{\eta}(a_2); s) - V(\hat{\eta}(a_1); s) = C(a_2, \theta) - C(a_1, \theta)$. Otherwise, by Part 3 of Assumption 2,

---

[19]At sufficiently high stakes, there will exist pooling equilibria in which both types always play $a > a^{\mathrm{ce}}$. However, it can be shown that no action above a^{ce} can be used in an informative equilibrium.

[20]Some readers may note that the partially-informative equilibrium constructed here violates strategic-stability-based refinements. For example, the D1 criterion (Cho and Kreps, 1987) would imply that no actions above a^{or} can be taken in equilibrium, since at actions $a > a^{\mathrm{or}}$ it would be the natural type rather than the gamer type who would have a greater cost savings from deviating slightly downwards. Hence, under D1, at high enough stakes both types would pool at a^{or}—just as in a canonical model one-dimensional signaling model with maximum action a^{or}. Our approach in this paper is to set aside questions of plausible out-of-equilibrium beliefs and to instead focus on gaining what we view as more robust insights by considering local or global comparative statics given an arbitrary equilibrium. In any case, richer type spaces—including the 2×2 type space considered in Subsection 4.1—will have informative equilibria that satisfy the D1 criterion for all stakes.

the left-hand side goes off to infinity while the right-hand side stays fixed. Now take an arbitrary equilibrium containing any two on-path actions $a' < a''$ with different beliefs. To satisfy belief monotonicity, the gamer θ_1 must use a' while the natural θ_2 must use a'', hence it must hold that $a'' \leq a^{\mathrm{ce}}$ (cf. fn. 19). This implies the incentive constraints $V(\hat{\eta}(a''); s) - V(\hat{\eta}(a'); s) \leq C(a'', \theta_1) - C(a', \theta_1) \leq C(a^{\mathrm{ce}}, \theta_1)$. Again we see that if stakes are large, beliefs on the two actions, $\hat{\eta}(a')$ and $\hat{\eta}(a'')$, must be approximately equal—and thus both must be approximately equal to the prior. So at large stakes, any equilibrium must be approximately uninformative.

The preceding discussion suggests a general comparative static: as stakes increase, less information can be revealed in equilibrium about natural actions. We show formally in Subsection 4.1 that this is indeed true in a setting that embeds the current two cross types. It is instructive, however, to explain the basic logic in this simpler setting. Take some equilibrium at a given stakes, and consider a marginal reduction in stakes. We perturb the old equilibrium to construct a more informative equilibrium at the lower stakes. In particular, we alter mixing probabilities to shift gamer types left, from high actions to low ones, and shift natural types right, from low to high. As the beliefs on the dimension of interest are increasing in actions, this perturbation moves high natural action types from lower to higher beliefs, and moves low natural action types from higher to lower beliefs. This generates a mean-preserving spread of the distribution of posteriors, i.e., a more informative equilibrium.

To illustrate the approach, consider a partial-pooling equilibrium of the form depicted in Figure 2: both types (denoted in the figure by $\theta_1 = \underline{\eta}\bar{\gamma}$ and $\theta_2 = \bar{\eta}\underline{\gamma}$) put positive probability on the two on-path actions a_1 and a_2, with $a_1 < a_2$ and $\hat{\eta}(a_1) < \hat{\eta}(a_2)$.[21] As stakes decrease, we aim to perturb the equilibrium in a way that only moves gamer types left and natural types right.
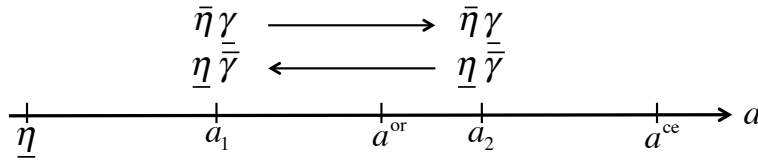

Figure 2 – Increasing informativeness as s decreases, when $\tau = \eta$.

---

[21] As depicted in the figure, it must hold that $\underline{\eta} \leq a_1 < a^{\mathrm{or}} < a_2 \leq a^{\mathrm{ce}}$ in order for both types θ to be indifferent because they must have the same cost difference $C(a_2, \theta) - C(a_1, \theta)$.

There are two relevant cases to consider. The first case is that both types are indifferent over actions a_1 and a_2, but no other incentive constraints are binding; the agents are not tempted to take any alternative actions. As stakes decrease, a_2 would become less attractive relative to a_1 (by Part 2 of Assumption 2) if we held strategies fixed. We can recover the indifference over the two actions, and thus the equilibrium, if we make a_2 relatively more attractive by decreasing beliefs at a_1 and increasing beliefs at a_2. Figure 2 shows exactly how to do this: move natural types right from a_1 to a_2, and/or move gamer types left from a_2 to a_1.

The second case is when there is another binding constraint: suppose $a_1 > \underline{\eta}$ and that the gamer type is originally indifferent not just over a_1 and a_2 at beliefs $\hat{\eta}(a_1) < \hat{\eta}(a_2)$, but also over its ideal point $a = \underline{\eta}$ at belief $\hat{\eta} = \underline{\eta}$.[22]A perturbation of the previous form would no longer yield an equilibrium. Lowering the stakes makes $a = \underline{\eta}$ more appealing than a_1, and then shifting types to reduce beliefs at a_1 would make $\underline{\eta}$ even more appealing; gamers would no longer be willing to play their equilibrium action a_1. Instead, we use a different perturbation as depicted in Figure 3. First shift some of the gamer types left from action a_1 to $\underline{\eta}$ (see the horizontal arrow). This shift holds beliefs at $\underline{\eta}$ at $\hat{\tau} = \underline{\eta}$, but increases the belief at a_1. After a small reduction in s, we can increase beliefs at a_1 in this manner until the gamer is again indifferent between playing $\underline{\eta}$ and a_1. But both types now prefer a_1 to a_2, since stakes have gone down while beliefs at a_1 have gone up. The next step is to increase beliefs at a_2 by moving gamer types left to action $\underline{\eta}$ (see the curved arrow of Figure 3). Beliefs at $\underline{\eta}$ are unchanged at $\hat{\eta} = \underline{\eta}$, and so this second shift can recover the indifference of both types between a_1 and a_2 without affecting the indifference of gamer types between a_1 and $\underline{\eta}$. These two steps of the perturbation only ever move gamer types left, from high-belief actions to low-belief actions, so the perturbation again improves informativeness.

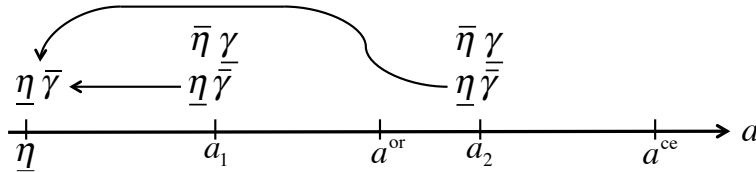**Figure 3** – Increasing informativeness with two steps as s decreases, when $\tau = \eta$.

Similar techniques (see the proof of Lemma 3) show that, starting from any equilibrium,

---

[22]Since the marginal cost of signaling is higher for the gamer on the domain $(\underline{\eta}, a^{\mathrm{or}})$, and because $a_1 < a^{\mathrm{or}}$, the gamer type prefers to deviate from a_1 to a lower action before the natural type does.

it is always possible to weakly improve informativeness as stakes decrease.

3.2. Dimension of interest is the gaming ability

Now suppose the dimension of interest is the gaming ability: $\tau = \gamma$. A separating equilibrium will exist if and only if the stakes are sufficiently large, by contrast to the previous subsection. In a separating equilibrium, without loss the natural type $\theta_2 = (\overline{\eta}, \underline{\gamma})$ plays $a_2 = \underline{\eta}$ and the gamer type $\theta_1 = (\underline{\eta}, \overline{\gamma})$ plays $a_1 > \overline{\eta}$. (If $a_1 \leq \overline{\eta}$, free downward deviations implies that the natural would simply mimic the gamer.) So a separating equilibrium exists if and only if there is an $a_1 > \overline{\eta}$ such that

$$C(a_1, \theta_2) \geq V(\overline{\gamma}; s) - V(\underline{\gamma}; s) \geq C(a_1, \theta_1). \tag{2}$$

The first inequality is the incentive constraint for the natural not to mimic the gamer, while the second inequality is the gamer's incentive constraint not to deviate to its natural action. Combining the implication of (2) that $C(a_1, \theta_2) \geq C(a_1, \theta_1)$ with $a_1 > \overline{\eta}$ shows that the gamer's action a_1 must in fact be above the cost-equalizing action a^{ce}. Furthermore, the value of separation must be high enough that the gamer is willing to bear the cost of taking an action sufficiently above a^{ce} to deter the natural from mimicking it. Formally, a separating equilibrium exists if and only if $s \geq s_\gamma^*$, where s_γ^* solves

$$V(\overline{\gamma}; s_\gamma^*) - V(\underline{\gamma}; s_\gamma^*) = C(a^{\mathrm{ce}}, \theta_1).$$

Next, we consider whether an informative equilibrium exists at low stakes. At s small enough such that $V(\overline{\gamma}; s) - V(\underline{\gamma}; s) \leq C(\overline{\eta}, \theta_1)$, it is clear that the gamer type would never take any action above $\overline{\eta}$ in equilibrium, and so the natural type could costlessly mimic any action taken by the gamer; hence, informative equilibria cannot exist. Indeed, under any informative equilibrium the gamer type must actually play some action above a^{or}, the order-reversing action. Otherwise, say that the highest action she plays is $a_1 \leq a^{\mathrm{or}}$; it follows that this action has the highest on-path market belief $\hat{\gamma}$. But because the natural type has uniformly lower marginal costs at actions below a^{or}, the natural type must uniquely play a_1, so a_1 has the lowest equilibrium belief as well; hence, the equilibrium is uninformative. No informative equilibrium exists at stakes below s_γ^{**}, where s_γ^{**} satisfies $V(\overline{\gamma}; s_\gamma^{**}) - V(\underline{\gamma}; s_\gamma^{**}) = C(a^{\mathrm{or}}, \theta_1)$.

For moderate stakes between s_γ^{**} and s_γ^*, partially-pooling equilibria do exist. For instance,

the two types can both mix over a pair of actions a_1 and a_2, with $a_1 < a^{\text{or}} < a_2$ and $C(a_2, \theta) - C(a_1, \theta)$ equal across the two types. At stakes just above s^*_γ, we would take a_1 and a_2 close to each other and make the belief difference $\hat{\gamma}(a_2) - \hat{\gamma}(a_1)$ close to zero; at stakes just below s^*_γ, we would take a_1 close to $\underline{\eta}$, a_2 close to a^{ce} and make the belief difference $\hat{\gamma}(a_2) - \hat{\gamma}(a_1)$ close to $\overline{\eta} - \underline{\gamma}$.

Observation 2. *In the two cross types setting, if the dimension of interest is gaming ability, then there are thresholds $s^{**}_\gamma < s^*_\gamma$ such that a separating equilibrium exists if and only if $s \geq s^*_\gamma$, and all equilibria are pooling when $s \leq s^{**}_\gamma$. At stakes $s \in (s^{**}_\gamma, s^*_\gamma)$ there exist partially-pooling equilibria.*

We establish exact comparative statics that for $\tau = \gamma$, information increases in stakes (rather than decreases as when $\tau = \eta$) in Subsection 4.1, which embeds the current cross-types setting. To see the intuition, suppose that at some stakes s, there is a partially-pooling equilibrium with both types mixing over two on-path actions, $a_1 < a_2$. Consider a marginal increase in stakes to $s' > s$. Here we want to perturb the equilibrium to be more informative by moving gamer types right, from low beliefs on gaming ability to high beliefs, and/or moving natural types left.

There are two important differences with the logic given in the previous subsection, where stakes decreased (recall Figure 2 and Figure 3). As stakes have increased, both types find their natural actions less attractive relative to actions a_1 and a_2. This difference simplifies matters, as we will be able to ignore any incentives to deviate down to one's natural action. However, because action a_2 has also gotten more attractive relative to a_1, maintaining in-difference by altering the behavior only across the two actions would require reducing $\hat{\gamma}(a_2)$ and increasing $\hat{\gamma}(a_1)$. That shift would involve moving gamers left and naturals right, which would induce a mean-preserving contraction in the distribution of market beliefs—a less informative equilibrium. Instead, as depicted by the arrow in Figure 4, we move gamers right from a_2 to some new higher action a_3. By moving gamers right, from a low belief $\hat{\gamma}(a_2)$ to a higher belief $\overline{\gamma}$, we produce a more informative equilibrium.

Specifically, to get the new strategies to be an equilibrium, we shift gamer types right to reduce beliefs at a_2 until we recover the indifference of both types between a_1 and a_2. The action a_3 is then set to make gamers indifferent between a_2 at the new beliefs, and a_3 at beliefs $\hat{\gamma} = \overline{\gamma}$. Because $a_2 > a^{\text{or}}$ and because gamers have lower marginal costs above a^{or}, if the gamers are indifferent between a_2 and a_3, the naturals strictly prefer to stick with a_2.
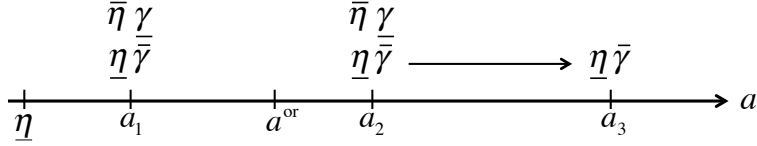
Figure 4 – Increasing informativeness as s increases, when $\tau = \gamma$.

4. The Effect of Stakes on Muddled Information

While the previous section aimed to provide intuition for why higher stakes tends to decrease equilibrium information when the dimension of interest is the natural action, and conversely when the market is interested in gaming ability, it is obviously artificial to only consider two types that are cross types. With this in mind, we turn to our main results.

4.1. A 2×2 type space

We first consider the minimal non-trivial setting where Θ is a product set, which we refer to as the 2×2 setting: $\Theta = \{\underline{\eta}, \overline{\eta}\} \times \{\underline{\gamma}, \overline{\gamma}\}$ with $\underline{\eta} < \overline{\eta}$ and $\underline{\gamma} < \overline{\gamma}$. Here, we are able to establish global comparative statics on the informativeness of equilibria with respect to the stakes.

For any given stakes, there are typically multiple equilibria; moreover, these equilibria need not be ranked according to (Blackwell) informativeness. We seek to compare the informativeness of the equilibrium sets at different stakes. We do so using the weak set order: equilibrium set $\mathcal{E}$ is *more informative* than equilibrium set $\mathcal{E}'$, and $\mathcal{E}'$ is *less informative* than $\mathcal{E}$, if (i) for any equilibrium $e \in \mathcal{E}$ there exists $e' \in \mathcal{E}'$ with e more informative than e', and (ii) for any $e' \in \mathcal{E}'$ there exists $e \in \mathcal{E}$ with e more informative than e'. In our setting, condition (i) is satisfied for any pair of equilibrium sets by taking $e' \in \mathcal{E}'$ to be a pooling equilibrium. So it is only condition (ii) that in fact has bite: for any equilibrium in the less informative set, there is an equilibrium in the more informative set that is more informative. Note that there may exist no "most informative" equilibrium in either set because the Blackwell (1951) informativeness criterion is a partial order.[23]

---

[23]If we were to extend the Blackwell partial ordering to a complete ordering on informativeness of equilibria—for example, by taking any convex function $U(\beta_\tau)$ and comparing equilibria by $\mathbb{E}[U(\cdot)]$—then our notion of equilibrium set $\mathcal{E}$ being more informative than equilibrium set $\mathcal{E}'$ would correspond to the most informative element of $\mathcal{E}$ being more informative than that of $\mathcal{E}'$.

We now formalize the sense in which information decreases (resp., increases) in the stakes when the dimension of interest is the natural action (resp., gaming ability).

Proposition 1. *In the 2×2 setting, consider stakes $\underline{s} < \bar{s}$:*

1. *If the dimension of interest is the natural action, the set of equilibria under $\underline{s}$ is more informative than the set of equilibria under $\bar{s}$.*

2. *If the dimension of interest is gaming ability, the set of equilibria under $\underline{s}$ is less informative than the set of equilibria under $\bar{s}$.*

The first part of Proposition 1 is a consequence of the following lemma.

Lemma 3. *Consider the 2×2 setting with $\tau = \eta$. Let $\mathcal{E}(s)$ be the set of equilibria at stakes s, and fix some equilibrium e_0 at stakes $s_0 > 0$. There exists a function $e(s)$ from stakes $s > 0$ to equilibria in $\mathcal{E}(s)$ that satisfies (i) $e(s_0) = e_0$; (ii) the distribution of β_η under $e(s)$ is continuous in s;[24] and (iii) $e(s'')$ is less informative than $e(s')$ for any $s'' > s'$.*

In words, Lemma 3 says that starting from any equilibrium at some stakes, say s_0, we can find a path of equilibria that is continuous in information—i.e. in the distribution of the market posterior—and decreases information as stakes increase from s_0, and increases information as stakes decrease from s_0. This result is substantially stronger than the first part of Proposition 1 because of the continuity in information that is assured from the arbitrary starting point.

Similarly, the second part of Proposition 1 is a consequence of the following lemma.

Lemma 4. *Consider the 2×2 setting with $\tau = \gamma$. Let $\mathcal{E}(s)$ be the set of equilibria at stakes s, and fix some equilibrium e_0 at stakes $s_0 > 0$. There exists a function $e(s)$ from stakes $s \in [s_0, \infty)$ to equilibria in $\mathcal{E}(s)$ that satisfies (i) $e(s_0) = e_0$; (ii) the distribution of β_γ under $e(s)$ is continuous in s; and (iii) $e(s'')$ is more informative than $e(s')$ for any $s'' > s'$.*

This lemma is a slightly weaker counterpart of Lemma 3: starting from any equilibrium at some stakes, it only assures a continuous increase in information when stakes increase, but not necessarily a continuous decrease in information when stakes decrease.

---

[24]That is, for a sequence $s \to s^*$, the corresponding distributions under s converge weakly to that under s^*.

The proofs of Lemma 3 and Lemma 4 are constructive, building on and generalizing the arguments illustrated in Section 3. We note that we would conjecture the conclusion of Proposition 1 to also hold when attention is restricted to equilibria satisfying stability-based refinements such as D1 or divinity (Cho and Kreps, 1987; Banks and Sobel, 1987).[25]

4.2. General type spaces

For more general type spaces, $\Theta \subset \mathbb{R} \times \mathbb{R}_{++}$, we are unable to get exact results on comparative statics on information as stakes vary. Instead, to extend the intuition that observers tend to be more informed about the natural action at low stakes and more informed about gaming ability at high stakes, we focus on the limiting cases when stakes get arbitrarily small or large. For these limits, we provide results on either the existence of separating equilibria, or on conditions guaranteeing that equilibria become approximately uninformative about $\hat{\tau}$. Formally, we say that equilibria are approximately uninformative about $\hat{\tau}$ at high (low) stakes if for any sequence of equilibria e_s at stakes s, it holds that as $s \to \infty$ ($s \to 0$), the distribution of $\hat{\tau}$ under e_s approaches the uninformative distribution, a point mass at $\mathbb{E}[\tau]$.

4.2.1. Dimension of interest is the natural action

Let the dimension of interest be the natural action: $\tau = \eta$. We first generalize the idea of Observation 1 by establishing conditions under which full information can only be revealed at low stakes.

Proposition 2. *Assume $\tau = \eta$. If $|\Theta_\eta| < \infty$, then for sufficiently small stakes, an equilibrium exists that is fully informative on η. If Θ has any cross types, then a fully-informative equilibrium does not exist for sufficiently large stakes.*

If there were no cross types, then standard arguments based on the single-crossing property imply that a fully-informative equilibrium would exist for any stakes.

We next seek to formalize the idea that for larger s, the information on η becomes limited. The first step is to take an arbitrary pair of cross types and show that the equilibrium belief difference between the high-η type and the low-η type is bounded above by a function that decreases in s.

---

[25]See Bagwell (2007) for a thorough analysis of equilibria satisfying the intuitive criterion in a model with 2×2 types; he does not analyze comparative statics of informativeness.

Lemma 5. *Assume $\tau = \eta$ and fix any two cross types, $\theta_1 = (\underline{\eta}, \overline{\gamma}) <> \theta_2 = (\overline{\eta}, \underline{\gamma})$, with the corresponding cost-equalizing action, a^{ce}. Across all type spaces Θ containing $\{\theta_1, \theta_2\}$ and across all equilibria, it holds that if $\hat{\eta}_i$ is some belief that θ_i induces in equilibrium $(i = 1, 2)$, then*

$$V(\hat{\eta}_2; s) - V(\hat{\eta}_1; s) \leq C(a^{ce}, \theta_1) = C(a^{ce}, \theta_2).$$

The important consequence of Lemma 5 (and Part 3 of Assumption 2) is that for any pair of cross types, $\theta_1 <> \theta_2$, as stakes get arbitrarily large, the type with lower natural action and higher gaming ability (θ_1) cannot induce a belief much lower than any belief induced by the type with higher natural action and lower gaming ability (θ_2). Depending on the type space and the parameters, there may be equilibria in which θ_1 induces a strictly higher belief than θ_2, because θ_1 may (partially-)pool with types with even higher natural actions than θ_2; this possibility does not arise in a 2×2 setting. In the limit of $s \to \infty$, any type with strictly higher gaming ability than another type must induce a weakly higher belief about its natural action.

This monotonicity of beliefs in the limit provides an upper bound on how informative an equilibrium can be about natural actions at very large stakes: any limiting distribution of beliefs on η must be "ironed" so that the set of γ types consistent with a belief $\hat{\eta}$ is weakly increasing in $\hat{\eta}$ (in the sense of the strong set order). In certain salient cases any limiting distribution is necessarily uninformative on the posterior mean $\hat{\eta}$; in these cases any equilibrium is approximately uninformative about $\hat{\eta}$ at large stakes.

Proposition 3. *Assume $\tau = \eta$. If the marginal distribution of γ has a continuous distribution and if $\mathbb{E}[\eta|\gamma]$ is non-increasing in γ over the support of Θ, then equilibria are approximately uninformative about $\hat{\eta}$ for large stakes.*

A simple condition which implies that $\mathbb{E}[\eta|\gamma]$ is non-increasing in γ is that the natural action and gaming ability are independent. Note that equilibria becoming uninformative about $\hat{\eta}$ does not imply their becoming uninformative about η. In the limit, observers learn nothing about the mean of the agent's natural action—as this moment is what the agent cares about—but may still learn something about other belief moments, such as the variance.

Two observations help explain why some conditions are needed in Proposition 3. First, if the distribution of γ were not continuous, then, for example, a mass of types with the lowest γ and a low η (or the highest γ and a high η) may separate from other types even in the limit of $s \to \infty$, revealing information about their η. Second, even with a continuous distribution

of γ, if the expectation $\mathbb{E}[\eta|\gamma]$ were increasing—e.g. η and γ are positively correlated—then types with higher γ may be able to signal their higher average η by taking higher actions.

4.2.2. Dimension of interest is the gaming ability

Let the dimension of interest be the gaming ability: $\tau = \gamma$. The mirror image of Proposition 2 now holds:

Proposition 4. *Assume $\tau = \gamma$. If $|\Theta_\gamma| < \infty$, then for sufficiently large stakes, an equilibrium exists that is fully informative on γ. If Θ has any cross types, then a fully-informative equilibrium does not exist for sufficiently small stakes.*

We can also find bounds on belief differences between a pair of cross types $\theta_1 = (\underline{\eta}, \overline{\gamma}) <> \theta_2 = (\overline{\eta}, \underline{\gamma})$, similarly to Subsection 4.2.1. In particular, we will show that at sufficiently low s, the type θ_2 with lower γ but higher η can at worst replicate the beliefs of θ_1, and potentially can do even better.

Lemma 6. *Assume $\tau = \gamma$. Fix any pair of cross types in Θ, $\theta_1 <> \theta_2$, with order-reversing action a^{or}. If $V(\max \Theta_\gamma; s) - V(\min \Theta_\gamma; s) \leq C(a^{or}, \theta_1)$, then in any equilibrium it holds that $\hat{\gamma}_1 \leq \hat{\gamma}_2$ for any beliefs $\hat{\gamma}_1$ and $\hat{\gamma}_2$ that θ_1 and θ_2 respectively induce.*

By compactness of Θ_γ and Assumption 2, $V(\max \Theta_\gamma; s) - V(\min \Theta_\gamma; s) \to 0$ as $s \to 0$. Since $a^{\text{or}} > \overline{\eta} > \underline{\eta}$ for any pair of cross types, $(\underline{\eta}, \overline{\gamma}) <> (\overline{\eta}, \underline{\gamma})$, Lemma 6 implies that for small enough stakes, a low-γ cross type must induce a weakly higher belief about its gaming ability than its high-γ counterpart. In the limit as $s \to 0$, we see that the set of η types compatible with some belief $\hat{\gamma}$ must be increasing in $\hat{\gamma}$; as before, only "ironed" limiting distributions are possible. When such ironing is only compatible with an equilibrium that is uninformative about $\hat{\gamma}$, it follows that the informativeness of equilibria about the mean of the gaming ability must vanish as stakes gets small. In particular:

Proposition 5. *Assume $\tau = \gamma$. If the marginal distribution of η has a continuous distribution and $\mathbb{E}[\gamma|\eta]$ is non-increasing in η over the support of Θ, then equilibria are approximately uninformative about $\hat{\gamma}$ for small stakes.*

5. A Linear-Quadratic-Normal Specification

As an antidote to the somewhat abstract approach thus far of providing comparative statics without actually characterizing equilibria, we now turn to a tractable specification—more precisely, a variant—of our model. In this section alone, we take the agent's payoff function to be

$$s\gamma\hat{\tau} - \frac{(a-\eta)^2}{2}, \tag{3}$$

for some $s > 0$. We also take $\eta, \gamma \sim \mathcal{N}(\mu_\eta, \mu_\gamma, \sigma_\eta^2, \sigma_\gamma^2, \rho)$, i.e. the two variables are jointly distributed accordingly to a bivariate normal distribution with means μ_η and μ_γ, variances $\sigma_\eta^2 > 0$ and $\sigma_\gamma^2 > 0$, and correlation coefficient ρ (equivalently, covariance $\rho\sigma_\eta\sigma_\gamma$). For obvious reasons, we refer to this setting as the *linear-quadratic-normal*, or just *LQN*, specification.[26]

Some remarks are in order to relate this specification with our baseline model. First, on a technical note, the type space here is not compact. Second, for $\gamma > 0$, (3) is simply a positive affine transformation of $s\hat{\tau} - \frac{(a-\eta)^2}{2\gamma}$. This expression satisfies Assumption 1 and Assumption 2, except for free downward downward deviations: the agent must now incur a cost to take actions that are lower than her natural action. Third, when $\gamma < 0$, (3) implies that the agent prefers lower market beliefs. As discussed in Subsection 2.3, one can interpret γ as representing the agent's heterogeneous *private stakes* in signaling. While negative private stakes are of intrinsic interest in some applications,[27] the main reason we allow them here is to leverage normal distributions; the probability of $\gamma < 0$ can be specified to be arbitrarily small if desired. Note that for $\gamma < 0$, private stakes are not isomorphic to our canonical gaming-ability formulation in which the cost function is $\frac{(a-\eta)^2}{2\gamma}$.

In the LQN specification, we focus on *linear equilibria*: equilibria in which an agent of type (η, γ) takes action $a = l_\eta \eta + l_\gamma \gamma + k$, for some constants l_η, l_γ, and k.[28] In such an equilibrium, the market posterior, $\hat{\tau}(a)$, is a linear function of the agent's action; moreover, ex ante, $\hat{\tau}$ is normally distributed with mean μ_τ and some variance $\text{Var}(\hat{\tau}) \in [0, \sigma_\tau^2]$. It can be verified

---

[26]For related specifications, Tamura (2014) characterizes the optimal information disclosure policy an agent would use if he could commit to a mapping from his type to costless signals.

[27]Consider a sender-receiver "costly talk" application (Kartik et al., 2007). Here, η represents a payoff-relevant state for the receiver, a is the sender's message, and $s|\gamma|$ captures the marginal rate of substitution for the sender between the receiver's beliefs and her own lying costs, i.e. the intensity of the sender's bias. The sign of γ captures the direction of the sender's bias, i.e. whether the sender prefers to induce higher or lower beliefs in the receiver.

[28]In the LQN specification, any smooth pure-strategy equilibrium—one in which the agent uses pure strategy that is continuous and the market posterior is twice differentiable—must be linear. See Proposition 11 in Appendix F.

that in the class of linear equilibria, informativeness about τ coincides with informativeness about $\hat{\tau}$, so that an equilibrium with higher $\mathrm{Var}(\hat{\tau})$ is (Blackwell) more informative on the dimension of interest for a fixed prior distribution of τ. As $\mathrm{Var}(\tau|a)$ is independent of a in linear equilibria (see the proof of Lemma 9), a higher $\mathrm{Var}(\hat{\tau})$ is also equivalent to a lower $\mathrm{Var}(\tau|a)$ for all a.

Our analysis in this section is related to that of Bénabou and Tirole (2006, Section II.B), and even more so, Fischer and Verrecchia (2000). Bénabou and Tirole focus on the comparative statics of equilibrium *actions*, while we emphasize comparative statics of equilibrium *informativeness*, which also appears in Fischer and Verrecchia (2000, Corollary 3).[29] Both these papers assume (their analogs of) η and γ are uncorrelated and focus on $\tau = \eta$. We also study the effects of correlation and treat both dimensions of interest, $\tau = \eta$ and $\tau = \gamma$. These are important for our applications. For example, students from a higher socioeconomic class may benefit more from increasing the market belief (higher γ) and may also tend to be intrinsically better prepared for college (higher η).

We begin by characterizing linear equilibria. Given any linear market posterior

$$\hat{\tau}(a) = La + K, \tag{4}$$

the agent's optimal action is given by

$$a = \eta + sL\gamma. \tag{5}$$

Under this linear strategy, the market's posterior beliefs will be normally distributed with mean linear in the agent's action. Equilibria are determined by solving for a fixed point: values of L and K under which the market's induced beliefs have mean equal to that hypothesized. While we relegate the details to Lemma 9 in Appendix F, it is useful to note that the equilibrium value of L is determined as:

$$L = \frac{\mathcal{L}(s, L, \tau)\sigma_\tau^2 + \mathcal{L}(s, L, \neg\tau)\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2 L^2 \sigma_\gamma^2 + 2sL\rho\sigma_\gamma\sigma_\eta}, \tag{6}$$

where $\neg\tau$ refers to the dimension other than the dimension of interest (e.g. $\neg\tau = \gamma$ when $\tau = \eta$), and $\mathcal{L}(s, \tilde{L}, \eta) \equiv 1$ and $\mathcal{L}(s, \tilde{L}, \gamma) \equiv s\tilde{L}$.

---

[29]In ongoing work using a related model, Ali and Bénabou (2014) also discuss comparative statics of equilibrium informativeness.

By Equation 4, the equilibrium constant L measures the responsiveness of the market belief $\hat{\tau}$ to the agent's action; the market belief on the dimension of interest is independent of the agent's action if and only if $L = 0$. We focus below on equilibria in which the solution to Equation 6 is non-negative, which we refer to as *non-decreasing equilibria*, to ensure that the market belief on the dimension of interest is non-decreasing in the agent's action.[30] We say that an equilibrium is *increasing* if the solution to Equation 6 is strictly positive.

5.1. Dimension of interest is the natural action

Assume $\tau = \eta$. A routine computation using Equations 4–6 shows that $\mathrm{Var}(\hat{\eta}) = L^2 s \rho \sigma_\eta \sigma_\gamma + L \sigma_\eta^2$. To streamline the exposition, the following proposition restricts attention to non-negative correlation, $\rho \geq 0$; we comment on negative correlation subsequently.

Proposition 6. *In the LQN specification, assume $\tau = \eta$, $\rho \geq 0$, and restrict attention to linear non-decreasing equilibria.*

1. *There is a unique equilibrium; it is increasing.*

2. *As $s \to 0$, $\mathrm{Var}(\hat{\eta}) \to \sigma_\eta^2$.*

3. *As $s \to \infty$, $\mathrm{Var}(\hat{\eta}) \to \rho^2 \sigma_\eta^2$.*

4. *The following comparative statics hold:*

    (a) $\dfrac{\mathrm{d}}{\mathrm{d}s} \mathrm{Var}(\hat{\eta}) < 0.$

    (b) $\dfrac{\mathrm{d}}{\mathrm{d}\mu_\gamma} \mathrm{Var}(\hat{\eta}) = 0.$

    (c) $\dfrac{\mathrm{d}}{\mathrm{d}\sigma_\gamma} \mathrm{Var}(\hat{\eta}) < 0.$

    (d) $\dfrac{\mathrm{d}}{\mathrm{d}\rho} \mathrm{Var}(\hat{\eta}) > 0.$

For the case of $\rho = 0$, the content of Proposition 6 is largely contained in Fischer and Verrecchia (2000). Part 1 of Proposition 6 is self-explanatory. Part 2 says that as stakes vanish, the unique equilibrium is approximately fully informative about η; this intuitive

---

[30] In our general model, this property is guaranteed by free downward deviations. As the LQN specification violates free downward deviations, there are parameters for which it has equilibria with $\hat{\tau}(a)$ decreasing.

result is a counterpart to the first portion of Proposition 2. Part 3 implies that as stakes grow unboundedly, the equilibrium becomes uninformative about η when $\rho = 0$, which is consistent with Proposition 3. Part 3 also implies that for any $\rho > 0$, there is some information revealed about η even in the limit of unbounded stakes; as explained after Proposition 3, the intuition is that when $\mathbb{E}[\eta|\gamma]$ is increasing in γ, higher γ types can signal their higher average η by taking higher actions. As this intuition suggests, the expression in Part 3 verifies that the equilibrium becomes arbitrarily close to fully informative as $\rho \to 1$ and $s \to \infty$.

Part 4 of Proposition 6 provides comparative statics on equilibrium informativeness about η for any value of stakes. Part 4a confirms our fundamental theme that higher stakes reduce information about natural actions. The other three parts address comparative statics that we have not touched on so far. Part 4b notes that the ex-ante mean of private stakes has no effect on equilibrium informativeness about η; rather, changes in μ_γ only shift the agent's action and the market belief function by a constant. This observation underscores that muddled information is not due to gaming *per se*, but rather heterogeneity in gaming ability.

Part 4c of the proposition says that greater ex-ante uncertainty about γ reduces equilibrium information about η. To understand why, notice from Equation 3 that s simply scales γ; thus, an increase in s is isomorphic to increasing both the ex-ante mean μ_γ and the ex-ante variance σ_γ^2 (albeit differently). Since we have just seen that μ_γ has no effect on equilibrium informativeness, the directional effect of increasing σ_γ^2 must be same as that of increasing s. Finally, Part 4d says that increasing an already non-negative correlation between η and γ leads to more equilibrium information about η. An intuition is that a greater non-negative correlation reduces the amount of heterogeneity in γ conditional on any η: under bivariate normality, $\text{Var}(\gamma|\eta) = \sigma_\gamma^2(1 - \rho^2)$. Indeed, when $\rho = 1$, the type space is effectively one-dimensional and the equilibrium fully reveals all private information.

Although Proposition 6 is stated for $\rho \geq 0$, the key points also extend to $\rho < 0$. The complication is that when $\rho < 0$ and stakes are intermediate, there can be multiple (linear increasing) equilibria. Nevertheless, the comparative statics in s, μ_γ, and σ_η all generalize subject to the caveat of focusing on the appropriate equilibria.[31]

---

[31]Three points bear clarification. First, when there are multiple equilibria, there are generically three of them. In these cases, the comparative statics hold for the equilibria with the highest and lowest solutions to Equation 6. Second, for $\rho < 0$, $\text{Var}(\hat{\eta}) \to 0$ as $s \to \infty$ (for large enough s, there is unique equilibrium), as is consistent with the "ironing" logic of Proposition 3. Third, comparative statics on ρ when $\rho < 0$ are not clear-cut.

5.2. Dimension of interest is private stakes

Assume $\tau = \gamma$. A routine computation using Equations 4–6 shows that $\mathrm{Var}(\hat{\gamma}) = L^2 s \sigma_\gamma^2 + L\rho\sigma_\eta\sigma_\gamma$. As before, we streamline our exposition by focusing on $\rho \geq 0$ in the following proposition.

Proposition 7. *In the LQN specification, assume $\tau = \gamma$, $\rho \geq 0$, and restrict attention to linear non-decreasing equilibria.*

1. *An equilibrium exists.*

   (a) *If $\rho = 0$, then (i) there is an equilibrium that is uninformative about γ, and (ii) there is also an increasing equilibrium (which is informative about γ) if and only if $s > \sigma_\eta^2/\sigma_\gamma^2$; the increasing equilibrium is unique when it exists.*

   (b) *If $\rho > 0$, then there is a unique equilibrium; it is increasing.*

2. *As $s \to 0$, $\mathrm{Var}(\hat{\gamma}) \to \rho^2\sigma_\gamma^2$.*

3. *As $s \to \infty$ and in the sequence of increasing equilibria, $\mathrm{Var}(\hat{\gamma}) \to \sigma_\gamma^2$.*

4. *The following comparative statics hold for an increasing equilibrium:*

   (a) $\dfrac{\mathrm{d}}{\mathrm{d}s}\mathrm{Var}(\hat{\gamma}) > 0.$

   (b) $\dfrac{\mathrm{d}}{\mathrm{d}\mu_\eta}\mathrm{Var}(\hat{\gamma}) = 0.$

   (c) $\dfrac{\mathrm{d}}{\mathrm{d}\sigma_\eta}\mathrm{Var}(\hat{\gamma}) < 0.$

   (d) $\dfrac{\mathrm{d}}{\mathrm{d}\rho}\mathrm{Var}(\hat{\gamma}) > 0.$

Part 1 of Proposition 7 notes that there is an equilibrium in which the agent plays $a = \eta$ if and only if $\rho = 0$, because only in the absence of correlation would the market belief about γ then be independent of the agent's action. There is a unique increasing equilibrium for any $\rho > 0$ or for sufficiently high stakes when $\rho = 0$.

To interpret Part 2 of Proposition 7, observe that were $s = 0$ the agent would play $a = \eta$; hence the market would fully learn η and form a belief $\hat{\gamma} = \mathbb{E}[\gamma|\eta] = \mu_\gamma + \sigma_\gamma\rho\left(\frac{\eta-\mu_\eta}{\sigma_\eta}\right)$, which implies $\mathrm{Var}(\hat{\gamma}) = \rho^2\sigma_\gamma^2$. Consequently, so long as the correlation in η and γ is small, there is

little information revealed about γ when s vanishes.[32] Part 3 of Proposition 7 says that as stakes grow unboundedly, any sequence of increasing equilibria approaches full information about γ; this is a counterpart to the first portion of Proposition 4. Part 4a confirms our fundamental theme that, even away from limiting stakes, higher stakes increase information about γ. The remaining comparative statics in Part 4 are analogous to those discussed in the context of Proposition 6. Note, however, that a change in σ_η^2 cannot be viewed through the lens of a change in s (unlike a change in σ_γ^2). Instead, the intuition for Part 4c is that greater ex-ante uncertainty about η makes the market belief about γ less responsive to the agent's action for any given agent's strategy, which in turn causes the agent's action to become less responsive to her private stakes, thereby generating less information about γ.

5.3. Mixed dimensions of interest

The tractability of the LQN specification makes it possible to study a number of additional questions within this setting. Here, we only take up one: what if the agent cares about the market's belief about both η and γ? To this end, generalize the agent's signaling benefit in the LQN specification to $s\gamma[g\hat{\gamma} + (1-g)\hat{\eta}]$, where $g \in [0,1]$ is a commonly known parameter; call this the *LQN specification with mixed dimensions of interest*.

When the agent uses a linear strategy, the market updates on each dimension as given in Lemma 9. Given any pair of linear market beliefs $\hat{\gamma}(a) = L_\gamma a + K_\gamma$ and $\hat{\eta}(a) = L_\eta a + K_\eta$, the agent's best response is to play $a = \eta + s(gL_\gamma + (1-g)L_\eta)\gamma$. Routine substitutions imply that an equilibrium is now characterized by a pair of constants, (L_η, L_γ), that simultaneously solve

$$L_\eta = \frac{\sigma_\eta^2 + s(gL_\gamma + (1-g)L_\eta)\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2(gL_\gamma + (1-g)L_\eta)^2\sigma_\gamma^2 + 2s(gL_\gamma + (1-g)L_\eta)\rho\sigma_\gamma\sigma_\eta},$$

$$L_\gamma = \frac{s(gL_\gamma + (1-g)L_\eta)\sigma_\gamma^2 + \rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2(gL_\gamma + (1-g)L_\eta)^2\sigma_\gamma^2 + 2s(gL_\gamma + (1-g)L_\eta)\rho\sigma_\gamma\sigma_\eta}.$$

Call an equilibrium where $L_\eta > 0$ and $L_\gamma > 0$ an *increasing equilibrium*. Consider, for simplicity, $\rho = 0$. Proposition 12 in the Supplementary Appendix shows that for any

---

[32]It is worth noting that by continuity with $s = 0$, $\mathrm{Var}(\hat{\gamma}) \to \rho^2\sigma_\gamma^2$ as $s \to 0$ even when $\rho < 0$. Consequently, for any $\rho < 0$, information about γ does not vanish as stakes vanish; to the contrary, information about γ in the limit can be arbitrarily high when $\rho \approx -1$. This discrepancy with Proposition 5 owes to the failure of free downward deviations in the LQN specification.

$g \in (0, 1)$, an increasing equilibrium exists, and that in such equilibria, (i) $\text{Var}(\hat{\eta})$ is decreasing in stakes and ranges from σ_η^2 (as $s \to 0$) to 0 (as $s \to \infty$), and (ii) $\text{Var}(\hat{\gamma})$ is increasing in stakes and ranges from 0 (as $s \to 0$) to σ_γ^2 (as $s \to \infty$). Thus, our main themes about how changes in stakes affect the market's information about both η and γ extend to this setting where the agent cares about the market belief on both dimensions of her type.

6. Applications

6.1. Manipulability and information provision

While our main comparative statics results are couched in terms of the signaling stakes, the results have other interpretations. Consider adding a parameter to the model, the *manipulability* of the signal, $M > 0$. Formally, suppose that an agent's payoff function is now given by

$$V(\hat{\tau}; s) - \frac{C(a, \theta)}{M}. \tag{7}$$

Manipulability parameterizes the ease with which agents can affect their signaling action: higher manipulability reduces the signaling costs for all types of agents. As this payoff function is isomorphic to $MV(\hat{\tau}; s) - C(a, \theta)$, it follows that increasing M is, from an agent's perspective, identical to increasing the marginal benefit of inducing a higher market belief, which is the same effect that an increase in stakes has. Thus, our previous analysis already implies the relevant comparative statics on information due to changes in M. When the dimension of interest is the natural action, increasing manipulability reduces information. When the dimension of interest is the gaming ability, increasing manipulability increases information.

In some cases manipulability might be a fundamental property of the signaling technology. Google's PageRank algorithm for web site ranking is considered harder to game than some earlier search engines based primarily on keyword density. Different types of tests might cover material that is easier or harder to study for; it is often thought that tests which measure "aptitude" rather than "achievement" are particularly difficult to study for.[33]

---

[33]The College Board's "SAT" was originally an acronym for Scholastic Aptitude Test, but the name was changed in 1994 to remove the connotation that it measured something innate and unchangeable. The College Board now insists that "SAT" is no longer an acronym for anything. Without endorsing the distinction between aptitude and achievement, Malcolm Gladwell (2001) writes in the New Yorker:

Manipulability may also depend on the agents' collective knowledge about the signaling technology. One aspect of this is the agents' familiarity—how much signaling experience they have had. A designer trying to increase market information about natural actions may benefit from changing the signaling technology (e.g., moving to a new test format, or switching out the credit scoring, credit rating, or search engine ranking algorithm) even if the new one is inherently no less manipulable than the old one. Simply by virtue of being new, agents may know less about how to game it, at least in the short run. This leads to a constant battle between designers and agents; it is reported that Google tweaks its search algorithm as often as 600 times a year partly to mitigate undesirable search engine optimization.

Designers could also make agents more or less knowledgeable by controlling how much information is given out about the workings of the signaling technology. Google and Fair Isaac keep their precise algorithms secret; past SAT questions were kept hidden until the 1980s; and US News sometimes changes its ranking algorithm with an explicit announcement that the algorithm will not be revealed until after the rankings are published (Morse, 2010). Our model gives a straightforward reason for such practices: a more opaque signaling technology may be less manipulable or more costly to game than a transparent one. When the dimension of interest is the natural action, less manipulable implies more information.

It is important, to recognize, however that revealing information about the workings of an algorithm might not increase gaming ability *uniformly* across agents. As discussed earlier, information differences may have been a source of the original heterogeneity on gaming ability. In that case, revealing information need not act like an increase in M in expression (7); rather, it may serve to effectively increase gaming ability for all types, but moreso for the low-γ types than the high-γ ones. In general, this has an ambiguous effect on information. Consider expression (7) with the specification $C(a, \eta, \gamma) = c(a, \eta)/\gamma$. On one hand, increasing all gaming abilities by a constant factor is identical to increasing M, which would lead to a decrease in information when the dimension of interest is natural actions. On the other hand, raising γ differentially to reduce its heterogeneity can increase information about natural actions. This is always true at the extreme when heterogeneity is eliminated altogether because that would allow for a separating equilibrium with full information; even

---

[After World War II, the SAT] was just beginning to go into widespread use. Unlike existing academic exams, it was intended to measure innate ability—not what a student had learned but what a student was capable of learning—and it stated clearly in the instructions that cramming or last-minute reviewing was pointless.

away from the extreme, we found in the LQN specification that reducing the variance of γ improves information no matter how the mean of γ is altered. So, in some cases, it might make sense to reveal past SAT questions to everyone if the College Board is concerned that some people already have access to them.

Similarly, subsidizing direct monetary costs of gaming would effectively increase the manipulability of the mechanism by raising the gaming ability of all agents. This would seem to be a bad policy for a designer interested in increasing information about agents' natural actions. But if heterogeneity on gaming ability had been driven by wealth differences in the first place, then such a policy would increase ability more for the low-γ poor students than the high-γ rich ones. This is our interpretation of the College Board's recent move to produce and publicize free test prep material. Rich students were already paying for this sort of material, and so the first order effect of the subsidy would be to level the playing field across agents and reduce heterogeneity.[34]

6.2. Informational externalities across markets

Consider a designer who can choose whether or not to reveal agents' actions to observers in different segments of a market or different markets. There is a direct informational benefit of revealing actions to observers who would otherwise lack access to this information about agents' types. However, the previous analysis shows there is likely to also be an indirect informational effect. As the signal is used in more markets, agents have stronger incentives to signal that they are of higher quality: signaling stakes increase. When the dimension of interest is the gaming ability, this extra incentive makes the signal more informative. But when the dimension of interest is the natural action, the information value of the signal is reduced as the stakes go up. In other words, preexisting observers are now less informed. In this latter case, a designer who is considering revealing the action to additional markets must trade off the benefits to the marginal observers with the negative informational externality imposed on the inframarginal observers.

---

[34]The College Board is teaming up with an online education service, Khan Academy, to produce the free test prep. The following excerpt from the College Board's press release (2014) is telling:

> "For too long, there's been a well-known imbalance between students who could afford test-prep courses and those who couldn't," said Sal Khan, founder and executive director of Khan Academy. "We're thrilled to collaborate closely with the College Board to level the playing field by making truly world-class test-prep materials freely available to all students."

To make the point concrete, consider markets in which credit reports may be used. Lenders use credit scores to determine how much credit to offer borrowers and at what terms. Landlords run credit checks to screen prospective tenants. A growing number of employers check credit reports during the hiring process to assess risks like employee theft or general trustworthiness; surveys suggest that only 19% of employers used credit checks in 1996 but 60% did in 2010. Automobile and homeowners insurance companies are now also commonly using information from credit reports to help determine insurance rates.

All these uses of credit reports are permitted in the United States on the federal level by the Fair Credit Reporting Act; some other uses, such as for targeted marketing, are federally prohibited. Moreover, states may pass additional regulations. In recent years, a number of states have considered or passed legislation restricting the use of credit history in insurance underwriting. As of November 2013, 10 states, including California and Illinois, limit or ban employers from using credit history in hiring. Many of the arguments in favor of these laws are based on some notion of "fairness"—it is not fair to deny someone employment or to raise their insurance rates based on missed mortgage payments or high credit card debt.

We suggest a very different sort of concern. Making credit information available to markets such as insurance and employment could alter consumers' gaming behavior in a way that dilutes the information contained in credit scores, thereby harming the efficiency of the loan market.

To formalize this informational externality in the simplest possible manner, we focus now on a setting where information is more socially valuable in some markets than in others. Let the dimension of interest be η, and let

$$V(\hat{\eta}, s) = v(\hat{\eta}) + w(\mathbb{E}[\eta]) + (s - 1)\big(w(\hat{\eta}) - w(\mathbb{E}[\eta])\big) \text{ for } s \in \{1, 2\}$$

for some strictly increasing functions $v(\cdot)$ and $w(\cdot)$. The designer chooses $s \in \{1, 2\}$ to maximize allocative efficiency $\mathbb{E}[V]$.

The interpretation of the value function V and the stakes s is that there is some benchmark first market that always observes the signal—e.g., the loan market in the credit scoring application. Showing the signal to this market corresponds to baseline stakes of $s = 1$, and gives agents a payoff $v(\hat{\eta})$. Agents also participate in a second market, the marginal one, in which actions may or may not be observed. If the signal is not observed ($s = 1$), the agent gets a payoff $w(\mathbb{E}[\eta])$ in the second market; if it is observed ($s = 2$), the agent gets a payoff $w(\hat{\eta})$ in the second market.

To make the analysis as stark as possible, let $v(\cdot)$ be strictly convex and $w(\cdot)$ be linear, say, $w(\hat{\eta}) = \alpha\hat{\eta}$. (Taking w to be approximately linear but with a small amount of convexity would give similar results; a concave w would only strengthen the point.) So there is a positive social value of information in the first market: more information implies higher allocative efficiency. For example, in the loan market, it more efficient to make (larger) loans to lower-financial-risk consumers. In the second market, though, any information just redistributes surplus across types without affecting allocations. For example, in the market for auto or homeowners insurance, lower-financial-risk consumers receive lower rates because this is correlated with them filing fewer claims, but only a small fraction of people are on the margin of car ownership based on their insurance rates. In this case, allowing an agent to signal that she is a lower financial-risk type may provide a significant positive private value but little social value.

By construction, there is no direct effect on allocative efficiency in the second market from allowing it to observe agents' actions. There is only an indirect effect on allocative efficiency of the first market. Our previous analysis implies that the higher stakes lead to a signal with lower information, and therefore lower allocative efficiency. We have argued for this comparative static result in a number of ways throughout the paper; a precise statement obtains, for example, in the 2×2 model using Proposition 1 in Subsection 4.1:

Corollary 1. *Consider the 2×2 setting with dimension of interest the natural action. Let $V(\hat{\eta}, s) = v(\hat{\eta}) + \alpha\mathbb{E}[\eta] + (s-1)\alpha(\hat{\eta} - \mathbb{E}[\eta])$ on $s \in \{1, 2\}$, with $\alpha > 0$ and $v(\cdot)$ strictly convex. For any equilibrium under $s = 2$, there exists an equilibrium under $s = 1$ with higher $\mathbb{E}[V]$.*

Thus, to support the efficiency of the "loan" market where the information is socially valuable (strictly convex v), the action should be hidden from the "auto insurance" market in which the social value is small (linear or approximately linear w). There is, of course, a converse: if the dimension of interest were gaming ability rather than natural actions, then revealing one's action to a second market would generate a positive rather than a negative informational externality to the first market. The designer might want to show the information to the purely redistributive second market just to ramp up the stakes for the agents and thereby increase gaming.

In Appendix B, we further analyze informational externalities in settings where information is equally valuable across markets: all markets have the same convex $v(\cdot)$ function. There is now an allocative benefit as well as a cost when providing information to new markets. In that setting we can study whether allocative efficiency is maximized by showing

signals to as many observers as possible, or whether there is an interior optimum after which the benefit to new observers is outweighed by the cost to existing observers.

7. Conclusion

This paper has proposed a model of signaling with two-dimensional types and one-dimensional actions. Equilibria typically confound the two dimensions of an agent's type: there is muddled information. In a nutshell, we find that when stakes increase, there is less equilibrium information about the agent's natural action and more information about her gaming ability. We have argued that this simple but robust finding provides insight into a number of applications and yields a novel tradeoff regarding how information, e.g. credit scores, should be regulated across markets. We close by noting some additional issues.

Our comparative statics on information stem from costs of signaling being driven by the natural action at low actions and the gaming ability at high actions. Specifically, we assumed that for any given pair of cross types, the ratio of marginal costs for the type with higher natural action relative to the type with higher gaming ability is monotonically increasing in the action (Assumption 1 part 4). We believe this assumption is a good approximation for many applications, including those we have highlighted.

However, our framework also provides insight into settings where this marginal cost ratio is not everywhere increasing, or where signaling costs depend on more than two dimensions of a type. While obtaining global comparative statics is generally infeasible, it is possible to derive limiting results on information as stakes get large or small. Essentially, at low (resp., high) stakes the market learns about the dimension that determines marginal costs at low (resp., high) actions. Suppose, for instance, that musicians differ on natural talent and on quality of coaching. For untrained musicians, performance quality—or the cost of performance improvement—is determined by natural talent. As musicians begin to train and reach higher levels of performance, the quality of coaching becomes important. Finally, those at the highest levels have absorbed all coaching lessons, and performance improvements are again driven by differences in talent. We would then expect a non-monotonic relationship between stakes and information about natural talent: very low and very high stakes generate precise information about natural talent.

A simplification of our model is that an agent's action is directly observed by the market. In practice, markets sometimes only observe a noisy signal of an agent's action. For example,

in test taking, agents choose effort which stochastically translates into the test score. Such noisy signaling would complicate the analysis, but, we believe, not fundamentally change our main points. That said, it would be interesting to explore whether greater noise can sometimes to lead to more equilibrium information.

Finally, our analysis of information yields direct implications for the allocative efficiency of equilibria. Other quantities may also be of interest. One might study agent welfare or, as in Bénabou and Tirole (2006), the average action taken by agents. In ongoing work, we are exploring some of the broader implications of multidimensional heterogeneity and muddled information on the equilibria of signaling games.

Bibliography

ALI, S. N. AND R. BÉNABOU (2014): "Image Versus Information," Personal communication.

ARAUJO, A., D. GOTTLIEB, AND H. MOREIRA (2007): "A model of mixed signals with applications to countersignalling," *RAND Journal of Economics*, 38, 1020–1043.

AUSTEN-SMITH, D. AND R. G. FRYER (2005): "An Economic Analysis of "Acting White"," *Quarterly Journal of Economics*, 120, 551–583.

BAGWELL, K. (2007): "Signalling and entry deterrence: A multidimensional analysis," *The RAND Journal of Economics*, 38, 670–697.

BANKS, J. S. AND J. SOBEL (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55, 647–661.

BÉNABOU, R. AND J. TIROLE (2006): "Incentives and Prosocial Behavior," *American Economic Review*, 96, 1652–1678.

BERNHEIM, B. D. (1994): "A Theory of Conformity," *Journal of Political Economy*, 102, 841–877.

BERNHEIM, B. D. AND S. SEVERINOV (2003): "Bequests as Signals: An Explanation for the Equal Division Puzzle," *Journal of Political Economy*, 111, 733–764.

BEYER, A., I. GUTTMAN, AND I. MARINOVIC (2014): "Earnings Management and Earnings Quality: Theory and Evidence," Unpublished.

BLACKWELL, D. (1951): "Comparison of Experiments," in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, vol. 1, 93–102.

——— (1953): "Equivalent Comparisons of Experiments," *The Annals of Mathematical Statistics*, 24, 265–272.

BUSINESSWEEK (2008): "Credit Scores: Not-So-Magic Numbers," http://www.businessweek.com/stories/2008-02-06/credit-scores-not-so-magic-numbers.

CHO, I.-K. AND D. KREPS (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179–221.

CHO, I.-K. AND J. SOBEL (1990): "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory*, 50, 381–418.

CHUNG, K.-S. AND P. ESÖ (2013): "Persuasion and learning by countersignaling," *Economics Letters*, 121, 487–491.

COLLEGE BOARD (2014): "The College Board Announces Bold Plans to Expand Access to Opportunity; Redesign of the SAT," https://www.collegeboard.org/releases/2014/expand-opportunity-redesign-sat.

CUTTS, M. (2012): "Another step to reward high-quality sites," Google Webmaster Central Blog; http://googlewebmastercentral.blogspot.com/2012/04/another-step-to-reward-high-quality.html.

DALEY, B. AND B. GREEN (2013): "Market signaling with grades," Forthcoming in *Journal of Economic Theory*.

DYE, R. A. AND S. S. SRIDHAR (2008): "A positive theory of flexibility in accounting standards," *Journal of Accounting and Economics*, 46, 312–333.

EDERER, F., R. HOLDEN, AND M. MEYER (2014): "Gaming and Strategic Opacity in Incentive Provision," Unpublished.

ENGERS, M. (1987): "Signalling with Many Signals," *Econometrica*, 55, 663–674.

FELTOVICH, N., R. HARBAUGH, AND T. TO (2002): "Too Cool for School? Signaling and Countersignaling," *RAND Journal of Economics*, 33, 630–649.

FISCHER, P. E. AND R. E. VERRECCHIA (2000): "Reporting bias," *The Accounting Review*, 75, 229–245.

GLADWELL, M. (2001): "Examined Life: What Stanley H. Kaplan taught us about the S.A.T." The New Yorker; http://www.newyorker.com/archive/2001/12/17/011217crat_atlarge.

KARTIK, N., M. OTTAVIANI, AND F. SQUINTANI (2007): "Credulity, Lies, and Costly Talk," *Journal of Economic Theory*, 134, 93–116.

LEWIS, M. (2011): *The Big Short: Inside the Doomsday Machine*, WW Norton & Company.

MAILATH, G., M. OKUNO-FUJIWARA, AND A. POSTLEWAITE (1993): "Belief-Based Refinements in Signalling Games," *Journal of Economic Theory*, 60, 241–276.

MATTHEWS, S. A. AND L. J. MIRMAN (1983): "Equilibrium Limit Pricing: The Effects of Private Information and Stochastic Demand," *Econometrica*, 51, 981–96.

MORSE, R. (2010): "U.S. News Takes Steps to Stop Law Schools From Manipulating the Rankings," US News Blog; http://www.usnews.com/education/blogs/college-rankings-blog/2010/05/20/us-news-takes-steps-to-stop-law-schools-from-manipulating-the-rankings.

QUINZII, M. AND J.-C. ROCHET (1985): "Multidimensional signalling," *Journal of Mathematical Economics*, 14, 261–284.

RICK, A. (2013): "The Benefits of Miscommunication," Unpublished.

SADOWSKI, P. (2013): "Overeagerness," Unpublished.

SPENCE, M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 87, 355–374.

STAMLAND, T. (1999): "Partially informative signaling," *Journal of Economic Theory*, 89, 148–161.

STEIN, J. C. (1989): "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics*, 104, 655–669.

TAMURA, W. (2014): "A Theory of Multidimensional Information Disclosure," Unpublished.

THE NEW YORK TIMES (2008): "Baylor Rewards Freshmen Who Retake SAT," http://www.nytimes.com/2008/10/15/education/15baylor.html.

——— (2014): "A New SAT Aims to Realign With Schoolwork," http://www.nytimes.com/2014/03/06/education/major-changes-in-sat-announced-by-college-board.html.

WESTON, L. P. (2011): *Your Credit Score: How to Improve the 3-digit Number that Shapes Your Financial Future*, FT Press, 4th ed.

Appendices (For Online Publication Only)

A. Relating the Signaling Benefit to Consumer Demand

Here we illustrate how the benefit function $V(\cdot; s)$ can be derived from the demand curve for a product by consumers with heterogeneous costs. In the example below, allocative efficiency corresponds to consumer surplus. The allocative benefit of information on consumer costs corresponds to the convexity of the induced benefit function.

Example 2. A consumer chooses an action a that is observed by a competitive market of firms before they offer the consumer a price p. The consumer then (mechanically) purchases a quantity given by a demand curve $D(p, s)$, where s captures the importance of this market to the consumer. The firms' expected cost of transacting with or serving the consumer depends on the consumer's type on the dimension of interest, τ, with higher types having lower cost. Specifically, expected service cost is $k - \alpha\tau$, with the most costly consumer having $\tau = 0$. As the market is competitive, the price offered to a consumer will equal the expected cost: $p = k - \alpha\hat{\tau}$.[35] A consumer who is thought to be type $\hat{\tau}$ receives gross consumer surplus

$$V(\hat{\tau}; s) = \int_{k-\alpha\hat{\tau}}^{k} D(k - \alpha t, s) dt.$$

Allocative efficiency—the expected consumer surplus—measures the full social value of information in this example, because firms receive no surplus. Agent (consumer) welfare, which is allocative efficiency minus expected total costs of signaling, is also total welfare here.

If the consumer's demand curve were completely inelastic—$D(p, s)$ independent of p—then $V(\cdot, s)$ would be linear. Additional information in the market would just transfer surplus from low to high types: a consumer of higher type would see lower prices and a consumer of lower type would see higher prices, but purchases would be unaffected. With a downward sloping demand curve, however, information becomes socially valuable. A consumer with higher cost of service purchases less, while a consumer with lower cost of service purchases more, leading to more efficient allocations. This induces a convex value function $V(\cdot; s)$. ‖

B. Informational Externalities across Homogeneous Markets

In this section we consider a class of settings in which there is a genuine tradeoff between providing socially-valuable information to additional observers and the negative information externality this imposes on existing observers. Throughout this section, let the dimension of interest be the natural action, η, and let

$$V(\hat{\eta}; s) = sv(\hat{\eta}) \text{ for } v(\hat{\eta}) = w(\hat{\eta}) - w(\mathbb{E}[\eta]),$$

where $w(\cdot)$, and therefore $v(\cdot)$, is strictly increasing, convex, and twice differentiable. The designer chooses stakes $s \geq 0$, and is interested in maximizing allocative efficiency $\mathbb{E}[V(\hat{\eta}; s)]$. For any given s, we focus on equilibria that maximize the allocative efficiency.

---

[35]Thus, one can think of a market of consumers each having identical continuous demand curves, where the cost of service is per-unit transaction; or, the demand curve may represent the stochastic demand for a single unit of the good, where this demand is realized after the consumer's action is chosen.

The interpretation of the value function V and the stakes s is that there are a number of markets in which the agent participates. The agent's value as a function of beliefs is $w(\cdot)$ in each market, scaled by the size of the market. In some of these markets, the agent's action is observed, in which case the beliefs on the agent's type are $\hat{\eta}$. In other markets, the agent's action is hidden, and beliefs remain at the prior $\mathbb{E}[\eta]$. Showing the signal to a given market turns it from uninformed to informed, leading to a net agent benefit of $v(\hat{\eta}) = w(\hat{\eta}) - w(\mathbb{E}[\eta])$. (We omit the constant payoff that the agent receives in all of the other uninformed markets.) The stakes s then corresponds to the mass of markets in which the signal is observable. The total value of information aggregated across markets, $\mathbb{E}[V(\hat{\eta}; s)]$, is the value per observation of $\mathbb{E}[v]$ times the mass of observers s: $s\mathbb{E}[v]$.

The designer chooses the mass of markets s in which the signal is observable. (Equivalently we can think of the designer as choosing the probability that a given market observes the signal.) The key assumption here is that the function determining the value of information, $w(\cdot)$, is homogeneous across markets. In the consumer pricing example (Example 2), the corresponding assumption would be that (up to linear scalings) the demand curve is identical across markets, and that beliefs on η yield the same prices across markets.

When a designer is choosing s, it is clear that if the market doesn't become uninformative in the limit as the stakes s go to infinity—if $\mathbb{E}[w(\hat{\eta})]$ stays bounded away from $w(\mathbb{E}[\eta])$, or equivalently $\mathbb{E}[v(\hat{\eta})]$ is bounded away from 0—then the designer who only cares about maximizing the social value of information (the allocative efficiency) should increase s to infinity.[36] The tradeoff between the marginal benefit and inframarginal cost of adding observers is interesting when the market becomes uninformed in the limit. Should we show the information to as many people as possible, taking s to infinity, or does the tradeoff lead us to an interior solution?

We analyze this question for two simple settings in which signals become uninformative in the limit. The first, as a proof of concept, shows that information is maximized at an interior solution to s when the type space is composed of a pair of cross types. The second shows that with a richer type space in which η and γ are independently distributed, we should in fact take s as large as possible: the direct effect of increasing the information for marginal observers dominates the indirect effect of worsening the information for inframarginal observers. We consider this latter result to be the takeaway message from this section.

It is important to emphasize that our focus below will be solely on the social value of information, i.e., the allocative efficiency. If the designer's welfare criterion also includes signaling costs, then he may not want to take s to infinity even if that maximizes information. Rather, the negative informational externality due to information muddling would instead be one factor that could lead the designer to reveal the signal to fewer observers than he or she otherwise might.

Two Cross Types. When there are only two cross types, information is maximized at an interior value of the stakes s:

Proposition 8. *Assume $\Theta = \{\theta_1, \theta_2\}$ with $\theta_1 = (\underline{\eta}, \overline{\gamma}) <> (\overline{\eta}, \underline{\gamma}) = \theta_2$ and $V(\hat{\eta}; s) = sv(\hat{\eta})$ with $v(\hat{\eta}) = w(\hat{\eta}) - w(\mathbb{E}[\eta])$ for some strictly convex w. The allocative efficiency $\mathbb{E}[V]$ is maximized over choice of s and choice of equilibrium at some finite $s \in (0, \infty)$.*

---

[36]This result applies, for example, when Θ contains an extreme type—one below or above all other types in both natural action and gaming ability—that has positive probability. The extreme type(s) can separate from the other types even in the limit, and the positive mass hypothesis keeps $v(\hat{\eta}) - v(\mathbb{E}[\eta])$ bounded away from zero.

(The proof is provided later in the section.) Recall from Subsection 3.1 that in a two-cross-types setting, there is an informative equilibrium for any $s > 0$ but the informativeness vanishes as $s \to \infty$. In the current context, this implies that $\mathbb{E}[v] > 0$ for any $s > 0$ and $\mathbb{E}[v] \to 0$ as $s \to \infty$. The question is what happens to $\mathbb{E}[V] = s\mathbb{E}[v]$ as $s \to \infty$. We establish in the proof of Proposition 8 that $s\mathbb{E}[v]$ goes to 0 at a rate faster than $\frac{1}{s}$ as $s \to \infty$. Allocative efficiency is positive at any positive s, and since it goes to 0 as $s \to 0$ or as $s \to \infty$, it has an interior optimum.

Independent η and γ. Now consider the more natural case when η and γ are independently distributed. The simplest specification with independent types that gets an uninformative limiting equilibrium ($\mathbb{E}[v] \to 0$ as $s \to \infty$) is when η has a binary distribution while γ is continuously distributed; see Proposition 3. In this benchmark setting, we find that under general conditions the designer would increase s without bound if he could, because even though allocative efficiency increases at a less than linear rate with s, it still grows without bound.

Proposition 9. *Assume η and γ are independent, $\Theta_\eta = \{\underline{\eta}, \overline{\eta}\}$, and γ is continuously distributed. Let $V(\hat{\eta}; s) = sv(\hat{\eta})$ with $v(\hat{\eta}) = w(\hat{\eta}) - w(\mathbb{E}[\eta])$ for some strictly convex w, and let $C(a, \eta, \gamma) = \frac{(a-\eta)^r}{\gamma}$ on $a \geq \eta$ for some $r > 1$. As $s \to \infty$, there exists a sequence of equilibria with allocative efficiency $\mathbb{E}[V] \to \infty$.*

(The proof is provided later in the section.) We prove Proposition 9 by constructing a class of two-action equilibria—all types take actions at either $a = \underline{\eta}$, or at some higher stake-dependent action $a = \tilde{a}$—in which $\mathbb{E}[v]$ goes to 0 at a limiting rate of $s^{\frac{-2}{r+1}}$ in s, and hence allocative efficiency $s\mathbb{E}[v]$ increases at a limiting rate of $s^{\frac{r-1}{r+1}}$.[37] In fact, the proof is more general than independent types, and establishes some conditions under which allocative efficiency increases linearly in the stakes when the distribution of gaming abilities for high-η agents is in some sense above that of low-η agents.

From the analysis of homogeneous markets we see that, while it is possible to construct examples in which the social value of information is maximized by obscuring information from some observers, in more plausible settings it seems likely that the social value of information is maximized by showing signals to as many observers as possible. In other words, even though signals do become less informative as the number of observers grows, the inframarginal cost is outweighed by the marginal benefit. We emphasize that our discussion ignores signaling costs; a designer who takes those into account might prefer to hide information from some observers.

B.1. Proof of Proposition 8

Denote the prior mean of η by $P \equiv \mathbb{E}[\eta]$. Without loss, we will assume that no actions below $\underline{\eta}$ are played in an equilibrium.[38]

Step 1: For sufficiently large s, any informative equilibrium has two on-path actions. This follows because the two types can simultaneously be indifferent only over a single pair of actions. Suppose, to contradiction, that an informative equilibrium involves three (or more) actions. Then one of these actions

---

[37]We do not suggest this rate is optimal; indeed, there are examples in which $\mathbb{E}[V]$ increases at a faster rate.

[38]If they are, they must all have the same market belief (as these actions are costless to both types), in which case there is an outcome-equivalent equilibrium that collapses all such actions to single one at $\underline{\eta}$.

is only played by a single type, and so has degenerate beliefs of $\hat{\eta} = \underline{\eta}$ or $\hat{\eta} = \bar{\eta}$. If one action induces beliefs $\hat{\eta}_1 = \underline{\eta}$, then there must be another action inducing beliefs $\hat{\eta}_2 > \mathbb{E}[\eta]$. If one action induces beliefs $\hat{\eta}_2 = \bar{\eta}$, then there must be another action inducing beliefs $\hat{\eta}_1 < \mathbb{E}[\eta]$. In either case, the benefit of increasing from the low to the high belief, $sv(\hat{\eta}_2) - sv(\hat{\eta}_1) \to \infty$ as $s \to \infty$. But by Lemma 5, $sv(\hat{\eta}_2) - sv(\hat{\eta}_1) \leq C(a^{\text{ce}}, \theta_1)$ for any s; a contradiction.

Step 2: A program that bounds allocative efficiency in two-action equilibria. Fix some stakes s. Say that in a two-action equilibrium with actions $a_1 < a_2$, a probability mass of q_i agents take action a_i and induce corresponding belief $\hat{\eta}_i$, for $i = 1, 2$. Using the identities $q_2\hat{\eta}_2 + q_1\hat{\eta}_1 = P$ and $q_1 + q_2 = 1$, we can solve for q_1 and q_2 as

$$q_1 = \frac{\hat{\eta}_2 - P}{\hat{\eta}_2 - \hat{\eta}_1}, \ q_2 = \frac{P - \hat{\eta}_1}{\hat{\eta}_2 - \hat{\eta}_1}.$$

Thus, allocative efficiency is

$$s\mathbb{E}[v(\hat{\eta})] = s\mathbb{E}[w(\hat{\eta}) - w(P)] = s\big(q_2 w(\hat{\eta}_2) + q_1 w(\hat{\eta}_1) - w(P)\big)$$
$$= s\left(\frac{P - \hat{\eta}_1}{\hat{\eta}_2 - \hat{\eta}_1}w(\hat{\eta}_2) + \frac{\hat{\eta}_2 - P}{\hat{\eta}_2 - \hat{\eta}_1}w(\hat{\eta}_1) - w(P)\right).$$

Lemma 5 implies that $sw(\hat{\eta}_2) - sw(\hat{\eta}_1) \leq C(a^{\text{ce}}, \theta_1)$. The following program therefore gives us an upper bound on allocative efficiency across all two-action equilibria:

$$\max_{\hat{\eta}_2, \hat{\eta}_1} \ s\left(\frac{P - \hat{\eta}_1}{\hat{\eta}_2 - \hat{\eta}_1}w(\hat{\eta}_2) + \frac{\hat{\eta}_2 - P}{\hat{\eta}_2 - \hat{\eta}_1}w(\hat{\eta}_1) - w(P)\right) \tag{8}$$

subject to $\hat{\eta}_1 \leq P \leq \hat{\eta}_2$ and

$$w(\hat{\eta}_2) - w(\hat{\eta}_1) \leq \frac{C(a^{\text{ce}}, \theta_1)}{s}. \tag{9}$$

We will show that the value of the above program tends to zero as $s \to \infty$; this establishes the desired conclusion that allocative efficiency is maximized at some interior s, because for any $s > 0$ there is an equilibrium with strictly positive allocative efficiency while allocative efficiency is obviously zero in any equilibrium when $s = 0$.

Step 3: The value of the program asymptotes to 0. In the solution to program (8), constraint (9) must be satisfied with equality; otherwise there is a mean-preserving spread of $\hat{\eta}$ that strictly increases allocative efficiency. Notice also that as $s \to \infty$, the constraints jointly imply $\hat{\eta}_2, \hat{\eta}_1 \to P$, and hence, for s large, $w(\hat{\eta}_2) - w(\hat{\eta}_1) \simeq w'(P) \cdot (\hat{\eta}_2 - \hat{\eta}_1)$. Thus, for any $\alpha < w'(P)$,

$$\alpha(\hat{\eta}_2 - \hat{\eta}_1) \leq w(\hat{\eta}_2) - w(\hat{\eta}_1) \leq \frac{C(a^{\text{ce}}, \theta_1)}{s}$$

at any solution at large enough s.

Fixing any $k > \frac{C(a^{\text{ce}}, \theta_1)}{w'(P)}$ and noting that relaxing constraint (9) can only increase the value of the program, it follows that the value of the program (8) for large enough s is no larger than the value of the following

program:

$$\max_{\hat{\eta}_2,\hat{\eta}_1} s\left(\frac{P-\hat{\eta}_1}{\hat{\eta}_2-\hat{\eta}_1}w(\hat{\eta}_2) + \frac{\hat{\eta}_2-P}{\hat{\eta}_2-\hat{\eta}_1}w(\hat{\eta}_1) - w(P)\right)$$

subject to $\hat{\eta}_1 \le P \le \hat{\eta}_2$ and

$$\hat{\eta}_2 - \hat{\eta}_1 \le \frac{k}{s}. \tag{10}$$

Since (10) will again be satisfied with equality, we substitute $\hat{\eta}_2 - \hat{\eta}_1 = \frac{k}{s}$ into the objective to simplify the above program to:

$$\max_{\hat{\eta}_2\in[P,P+\frac{k}{s}]} \varphi(\hat{\eta}_2;s) \equiv s\left(\frac{P-\hat{\eta}_2+\frac{k}{s}}{\frac{k}{s}}w(\hat{\eta}_2) + \frac{\hat{\eta}_2-P}{\frac{k}{s}}w\left(\hat{\eta}_2 - \frac{k}{s}\right) - w(P)\right). \tag{11}$$

Claim 1. *The function φ from (11) satisfies $s\varphi\left(P+\frac{kx}{s};s\right) \overset{s\to\infty}{\longrightarrow} \frac{k^2}{2}x(1-x)w''(P)$ for $x \in [0,1]$.*

Proof. Using the definition of $\varphi(\cdot)$,

$$s\varphi\left(P+\frac{kx}{s};s\right) = s^2\left[(1-x)w\left(P+\frac{kx}{s}\right) + xw\left(P-(1-x)\frac{k}{s}\right) - w(P)\right].$$

The second-order Taylor expansion, $w(P+\varepsilon) = w(P) + \varepsilon w'(P) + \frac{\varepsilon^2}{2}w''(P) + o(\varepsilon^3)$, yields

$$s\varphi\left(P+\frac{kx}{s};s\right) = s^2\left[(1-x)\left(w(P) + \frac{kx}{s}w'(P) + \frac{k^2x^2}{2s^2}w''(P) + o(s^{-3})\right)\right.$$

$$\left. + x\left(w(P) - \frac{k(1-x)}{s}w'(P) + \frac{k^2(1-x)^2}{2s^2}w''(P) + o(s^{-3})\right) - w(P)\right]$$

$$= s^2\left((1-x)\frac{k^2x^2}{2s^2}w''(P) + x\frac{k^2(1-x)^2}{2s^2}w''(P) + o(s^{-3})\right)$$

$$= \frac{k^2}{2}x(1-x)w''(P) + o(s^{-1})$$

$$\to \frac{k^2}{2}x(1-x)w''(P) \text{ as } s\to\infty. \hspace{2cm} Q.E.D.$$

Claim 1 implies that $s\max_{\hat{\eta}_2\in[P,P+\frac{k}{s}]} \varphi(\hat{\eta}_2;s)$ asymptotes to a constant. Hence, allocative efficiency, which was shown to be bounded above for large s by $\max_{\hat{\eta}_2\in[P,P+\frac{k}{s}]} \varphi(\hat{\eta}_2;s)$, asymptotes to 0.

B.2. Proof of Proposition 9

We prove a result that is slightly more general than Proposition 9. We maintain the assumption of binary η and continuous γ, but we generalize from independent types. (Allowing for discrete distributions of γ or mixed distributions with atoms would not materially change any conclusions but would complicate notation.)

From Proposition 8, we have an example in which the distribution of high type gaming abilities is below the low type gaming abilities, and in which allocative efficiency has an interior maximum in s. To guarantee

A-5

that there do exist equilibria where allocative efficiency diverges to infinity, we will invoke a condition that the distribution of high type gaming abilities is in some sense not below that of the low types.

Condition 1. *The natural actions are in $\Theta_\eta = \{\underline{\eta}, \overline{\eta}\}$, with probability mass $p \in (0,1)$ on $\overline{\eta}$ and $1 - p$ on $\underline{\eta} < \overline{\eta}$. Conditional on $\eta \in \Theta_\eta$, gaming ability γ is continuously distributed with cdf G_η, pdf g_η, and compact support in $\mathbb{R}_{++}$. Moreover, either*

1. $\max \operatorname{Supp} G_{\underline{\eta}} \leq \min \operatorname{Supp} G_{\overline{\eta}}$*; or*

2. *there exists γ^* such that $0 < G_{\overline{\eta}}(\gamma^*) < G_{\underline{\eta}}(\gamma^*) < 1$; or*

3. *there exists γ^* such that $0 < G_{\overline{\eta}}(\gamma^*) = G_{\underline{\eta}}(\gamma^*) < 1$, and $\lim_{\gamma \uparrow \gamma^*} g_{\overline{\eta}}(\gamma) > 0$.*

Part 1 says that all high-η types have gaming ability above that of any low types; it implies the type space is ordered by single-crossing as there are no cross types. Part 2 says that the support of gaming ability overlaps for the two values of η and that the gaming-ability distribution for high-η types is strictly above that of low-η types at some point (in the sense that the value of the cdf is below). Part 3 requires that the distributions are equal at some point where the left-neighborhood contains some high η types. When γ and η are independent and γ is continuously distributed as in Proposition 9, Condition 1 part 3 is satisfied.

Proposition 10 (Generalization of Proposition 9). *Assume the joint distribution of types satisfies Condition 1. Let $\tau = \eta$; let $V(\hat{\eta}; s) = sv(\hat{\eta})$ with $v(\hat{\eta}) = (w(\hat{\eta}) - w(\mathbb{E}[\eta]))$ for some strictly convex w; and let $C(a, \eta, \gamma) = \frac{(a-\eta)^r}{\gamma}$ on $a \geq \eta$, for some $r > 1$. As $s \to \infty$, there exists a sequence of equilibria with allocative efficiency $s\mathbb{E}[v(\hat{\eta})] \to \infty$. In particular, if Condition 1 parts 1 or 2 are satisfied, then there exists a sequence of equilibria with $s\mathbb{E}[v(\hat{\eta})]$ increasing at a linear rate in s. Under Condition 1 part 3, there exists a sequence of equilibria with $s\mathbb{E}[v(\hat{\eta})]$ increasing at rate $s^{\frac{r-1}{r+1}}$ (or faster).*

If there are no cross types—that is, Condition 1 part 1 holds—then for all s there exists a separating equilibrium on η. The result is then trivial because $\mathbb{E}[v(\hat{\eta})] = pv(\overline{\eta}) + (1-p)v(\underline{\eta})$ in any separating equilibrium, and hence allocative efficiency $s\mathbb{E}[v(\hat{\eta})]$ increases linearly in s along a sequence of separating equilibria.

The more interesting cases are when parts 2 or 3 of Condition 1 hold. In these cases, we can construct a two-action equilibrium of the following form, illustrated in Figure 5.

Lemma 7. *Suppose $\tau = \eta$, Condition 1 holds with either part 2 or 3 satisfied, and $C(a, \eta, \gamma) = \frac{(a-\eta)^r}{\gamma}$ on $a \geq \eta$, for $r > 1$. Given any γ^* satisfying $0 < G_{\underline{\eta}}(\gamma^*) \leq G_{\overline{\eta}}(\gamma^*) < 1$, there exists an action $\tilde{a} > \underline{\eta}$ and a gaming ability $\tilde{\gamma} < \gamma^*$ such that there is a two-action equilibrium in which agents of type (η, γ) play action $a_\eta(\gamma)$ defined by*

$$a_{\underline{\eta}}(\gamma) = \begin{cases} \underline{\eta} & \text{if } \gamma \leq \gamma^* \\ \tilde{a} & \text{if } \gamma > \gamma^*, \end{cases}$$

$$a_{\overline{\eta}}(\gamma) = \begin{cases} \underline{\eta} & \text{if } \gamma \leq \tilde{\gamma} \\ \tilde{a} & \text{if } \gamma > \tilde{\gamma}. \end{cases}$$
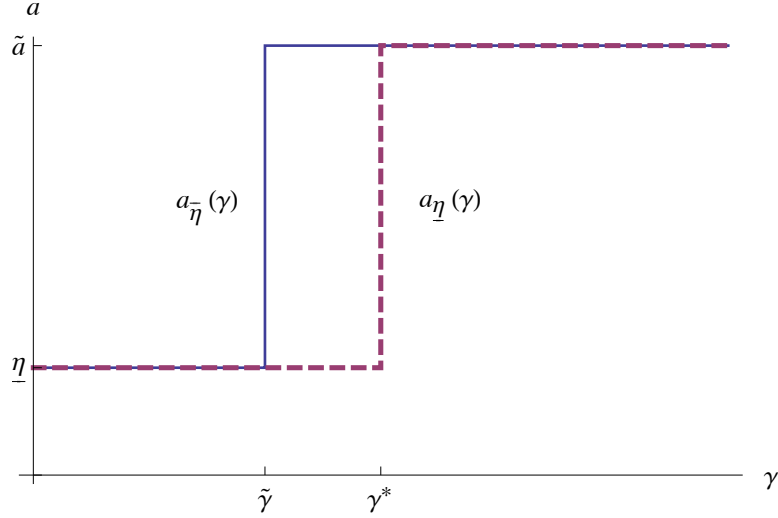
Figure 5 – Strategies under the equilibrium of Lemma 7.

The following conditions are sufficient for the strategy in Lemma 7 to constitute an equilibrium:[39]

$$\begin{cases} sv(\hat\eta(\overline\eta)) = sv(\hat\eta(\tilde a)) - \frac{(\tilde a - \overline\eta)^r}{\tilde\gamma} & \text{if } \tilde a > \overline\eta \\ \tilde\gamma = 0 & \text{if } \tilde a \le \overline\eta, \end{cases} \tag{12}$$

$$sv(\hat\eta(\underline\eta)) = sv(\hat\eta(\tilde a)) - \frac{(\tilde a - \underline\eta)^r}{\gamma^*}. \tag{13}$$

Equation 13 requires type $(\underline\eta, \gamma^*)$ to be indifferent between actions $\underline\eta$ and $\tilde a$. Condition 12 requires indifference for type $(\overline\eta, \tilde\gamma)$ if $\tilde a > \overline\eta$; on the other hand, if $\tilde a \le \overline\eta$, then all high-$\eta$ types strictly prefer $\tilde a$ (because both actions are costless but $\tilde a$ induces a higher belief, as confirmed below) and we take $\tilde\gamma = 0$. Note that deviations to off-path actions can be deterred by assigning any off-path action the belief $\hat\eta = \underline\eta$.

Given the strategy of Lemma 7, the observer's posterior probability $\Pi(a)$ that the agent's type is $\eta = \overline\eta$ conditional on action $a \in \{\underline\eta, \tilde a\}$ is given by

$$\Pi(\underline\eta) = \frac{pG_{\overline\eta}(\tilde\gamma)}{pG_{\overline\eta}(\tilde\gamma) + (1-p)G_{\underline\eta}(\gamma^*)}, \tag{14}$$

$$\Pi(\tilde a) = \frac{p(1 - G_{\overline\eta}(\tilde\gamma))}{p(1 - G_{\overline\eta}(\tilde\gamma)) + (1-p)(1 - G_{\underline\eta}(\gamma^*))}. \tag{15}$$

For $a \in \{\underline\eta, \tilde a\}$, $\hat\eta(a) = \Pi(a)\overline\eta + (1 - \Pi(a))\underline\eta$; that $\hat\eta(\tilde a) > \hat\eta(\underline\eta)$, or equivalently that $\Pi(\tilde a) > \Pi(\underline\eta)$, follows from $G_{\overline\eta}(\gamma^*) \le G_{\underline\eta}(\gamma^*)$ and $\tilde\gamma < \gamma^*$.

If $G_{\overline\eta}(\gamma^*) < G_{\underline\eta}(\gamma^*)$, as it is under Condition 1 part 2, it is straightforward to compute that even if $\tilde\gamma \to \gamma^*$ as $s \to \infty$ (and *a fortiori* if $\tilde\gamma \not\to \gamma^*$), $\hat\eta(\tilde a)$ remains bounded away from $\hat\eta(\underline\eta)$. In this case, $\mathbb{E}[v]$ stays bounded away from 0, and hence allocative efficiency $s\mathbb{E}[v] \to \infty$ at a linear rate as $s \to \infty$. This proves Proposition 10 for Condition 1 part 2.

---

[39]Modulo inessential multiplicities in $\tilde\gamma$, the conditions are also necessary.

For Condition 1 part 3, the conclusion of Proposition 10 follows from the following two claims.

Claim 2. *In a sequence of equilibria of Lemma 7, if $\gamma^* - \tilde\gamma \to 0$ at a rate of $f(s)$, then $\mathbb{E}[v] \to 0$ at a rate of $(f(s))^2$.*

Claim 3. *Assume Condition 1 holds with part 3 satisfied. In a sequence of equilibria of Lemma 7, $\gamma^* - \tilde\gamma \to 0$ at a rate of $s^{-\frac{1}{r+1}}$.*

Accordingly, it remains to prove Lemma 7, Claim 2, and Claim 3.

Proof of Lemma 7. The strategy has two free parameters: $\tilde a$ and $\tilde\gamma$, with the constraints that $\tilde a > \underline\eta$ and $\tilde\gamma < \gamma^*$, and the equilibrium conditions (12) and (13). To ease notation going forward, let

$$\tilde\eta(\tilde\gamma) \equiv \frac{p(1 - G_{\overline\eta}(\tilde\gamma))}{p(1 - G_{\overline\eta}(\tilde\gamma)) + (1 - p)(1 - G_{\underline\eta}(\gamma^*))}\overline\eta + \frac{(1 - p)(1 - G_{\underline\eta}(\gamma^*))}{p(1 - G_{\overline\eta}(\tilde\gamma)) + (1 - p)(1 - G_{\underline\eta}(\gamma^*))}\underline\eta$$

denote the observer's belief about η when action $\tilde a$ is observed given a value of $\tilde\gamma$.

<u>Case 1:</u> $s[v(\tilde\eta(0)) - v(\underline\eta)] \le \frac{(\overline\eta - \underline\eta)^r}{\gamma^*}$.

For this case we look for a solution with $\tilde\gamma = 0$ and $\tilde a \in (\underline\eta, \overline\eta]$. The induced on-path beliefs will be $\hat\eta(\underline\eta) = \underline\eta$ and $\hat\eta(\tilde a) = \tilde\eta(0)$. It suffices to check that there exists $\tilde a \in (\underline\eta, \overline\eta]$ to satisfy the $(\underline\eta, \gamma^*)$ indifference at these beliefs, condition (13). A solution exists precisely under the hypothesis of the case being considered.

<u>Case 2:</u> $s[v(\tilde\eta(0)) - v(\underline\eta)] > \frac{(\overline\eta - \underline\eta)^r}{\gamma^*}$.

For this case, we look for solutions with $\tilde a > \overline\eta$ and $\tilde\gamma > 0$. Supposing that $\tilde a > \overline\eta$, combine the indifferences in (12) and (13) to get $\frac{(\tilde a - \overline\eta)^r}{\tilde\gamma} = \frac{(\tilde a - \underline\eta)^r}{\gamma^*}$. Fixing any $\tilde\gamma \in (0, \gamma^*)$, this equality uniquely pins down a corresponding $\tilde a$, which we write as a function

$$\tilde a(\tilde\gamma) = \overline\eta + \frac{\overline\eta - \underline\eta}{(\gamma^*/\tilde\gamma)^{\frac{1}{r}} - 1}. \tag{16}$$

The function $\tilde a(\tilde\gamma)$ is continuous and increasing in $\tilde\gamma$, with $\tilde a(\tilde\gamma) \to \overline\eta$ as $\tilde\gamma \to 0^+$ and $\tilde a(\tilde\gamma) \to \infty$ as $\tilde\gamma \to \gamma^*$. We seek to find a $\tilde\gamma > 0$ which, along with the corresponding action $\tilde a = \tilde a(\tilde\gamma)$, constitutes an equilibrium. It is enough to find $\tilde\gamma$ satisfying the $(\underline\eta, \gamma^*)$ indifference, condition (12), as the $(\overline\eta, \tilde\gamma)$ indifference is assured by $\tilde a = \tilde a(\tilde\gamma)$. Rearranging (12) gives

$$s\big(v(\hat\eta(\tilde a)) - v(\hat\eta(\underline\eta))\big) = \frac{(\tilde a(\tilde\gamma) - \underline\eta)^r}{\gamma^*}. \tag{17}$$

For any sufficiently small but positive $\tilde\gamma$, $\hat\eta(\tilde a) = \tilde\eta(0)$ and $\hat\eta(\underline\eta) = \underline\eta$ because $\min \operatorname{Supp} G_{\overline\eta} > 0$ and $G_{\underline\eta}(\gamma^*) > 0$. Hence, as $\tilde\gamma \to 0^+$, the LHS of (17) goes to $s(v(\tilde\eta(0)) - v(\underline\eta))$, while the RHS goes to $\frac{(\overline\eta - \underline\eta)^r}{\gamma^*}$. From the hypothesis of Case being considered, it follows that for sufficiently small $\tilde\gamma$, the LHS of (17) is greater than the RHS. On the other hand, as $\tilde\gamma \to \gamma^*$, the LHS of (17) converges to a constant while the RHS diverges to ∞ (because $\tilde a(\tilde\gamma) \to \infty$). Hence, for sufficiently large $\tilde\gamma$, the LHS of (17) is less than the RHS. By continuity, there exists $\tilde\gamma > 0$ for which (17) holds; this value of $\tilde\gamma$ together with $\tilde a(\tilde\gamma)$ constitutes an equilibrium. *Q.E.D.*

Proof of Claim 2. To ease notation, let $G^* \equiv G_{\underline\eta}(\gamma^*) = G_{\overline\eta}(\gamma^*) \in (0, 1)$ and $g^* \equiv \lim_{\gamma\uparrow\gamma^*} g(\gamma) > 0$.

For $\tilde{\gamma} = \gamma^* - \varepsilon$ with $\varepsilon > 0$ small, $G_{\overline{\eta}}(\tilde{\gamma})$ is approximately linear in ε; to a first order Taylor approximation, $G_{\overline{\eta}}(\tilde{\gamma}) \simeq G^* - \varepsilon g^*$. As $\varepsilon \to 0^+$, this yields beliefs at the two actions approaching $\mathbb{E}[\eta]$ at a linear rate, with

$$\hat{\eta}(\underline{\eta}) \simeq \mathbb{E}[\eta] - \frac{p(1-p)g^*}{G^*}(\overline{\eta} - \underline{\eta})\varepsilon,$$

$$\hat{\eta}(\tilde{a}) \simeq \mathbb{E}[\eta] + \frac{p(1-p)g^*}{1 - G^*}(\overline{\eta} - \underline{\eta})\varepsilon.$$

Since $\mathbb{E}[v] = (pG_{\overline{\eta}}(\tilde{\gamma}) + (1-p)G^*)v(\hat{\eta}(\underline{\eta})) + (p(1 - G_{\overline{\eta}}(\tilde{\gamma})) + (1-p)(1 - G^*))v(\hat{\eta}(\tilde{a}))$, substituting in the approximations for $\hat{\eta}(\cdot)$ at small ε gives $\mathbb{E}[v]$ approximately quadratic in ε. *Q.E.D.*

Proof of Claim 3. First, it is routine to verify using conditions (12) and (13) that $\tilde{\gamma} \uparrow \gamma^*$ and $\tilde{a} = \tilde{a}(\tilde{\gamma}) \to \infty$ as $s \to \infty$.

Second, while $\tilde{a}(\tilde{\gamma}) \to \infty$ as $\tilde{\gamma} \to \gamma^*$, it holds that $\varepsilon \cdot \tilde{a}(\gamma^* - \varepsilon)$ goes to a constant as $\varepsilon \to 0^+$. In particular, Equation 16 implies

$$\varepsilon \cdot \tilde{a}(\gamma^* - \varepsilon) = \overline{\eta}\varepsilon + (\overline{\eta} - \underline{\eta})\frac{\varepsilon}{\left(\frac{\gamma^*}{\gamma^* - \varepsilon}\right)^{\frac{1}{r}} - 1} \overset{\varepsilon \to 0^+}{\longrightarrow} (\overline{\eta} - \underline{\eta})\gamma^* r$$

where the limit follows from L'Hopital's Rule. In other words, for large s, $\tilde{a}$ is of order $\frac{1}{\varepsilon}$.

Finally, we observe that for $\tilde{\gamma} = \gamma^* - \varepsilon$, for ε small we have $v(\hat{\eta}(\tilde{a})) - v(\hat{\eta}(\underline{\eta}))$ approximately linear in ε (following the proof of Claim 2), and therefore Equation 13 implies $(\tilde{a} - \underline{\eta})^r$—and therefore $\tilde{a}^r$—diverges to ∞ at a rate that is linear in $s \cdot \varepsilon$. Applying the previous result, we have ε^{-r} approximately linear in $s \cdot \varepsilon$, or equivalently $\varepsilon = \gamma^* - \tilde{\gamma}$ is of order $s^{\frac{-1}{r+1}}$. *Q.E.D.*

C. Proofs for Section 2

Proof of Lemma 1. For any $a > \underline{\eta}$,

$$C(a, \underline{\eta}, \overline{\gamma}) - C(a, \overline{\eta}, \underline{\gamma}) = \int_{\underline{\eta}}^{a} C_a(\hat{a}, \underline{\eta}, \overline{\gamma})\left[1 - \frac{C_a(\hat{a}, \overline{\eta}, \underline{\gamma})}{C_a(\hat{a}, \underline{\eta}, \overline{\gamma})}\right] d\hat{a}.$$

Since $C_a(\cdot) > 0$ on the relevant region, Part 4 of Assumption 1 implies that the sign of the intergrand above is the same as that of $a^{\text{or}} - \hat{a}$, where a^{or} is the order-reversing action for the given cross types. This implies that on the domain $a > \underline{\eta}$, the integral is strictly quasi-concave with a maximum at $a = a^{\text{or}}$. Furthermore, on the domain $a > \underline{\eta}$, the integral is strictly positive for $a \leq a^{\text{or}}$ and can only change sign at most once. To see that it must change sign at some point, observe that Assumption 1 also implies that there is some $\varepsilon > 0$ such that the integrand is strictly less than $-\varepsilon$ for all $a > a^{\text{or}} + \varepsilon$. *Q.E.D.*

Proof of Lemma 2. Fix an equilibrium in which $a' < a''$ are two on-path actions with $\hat{\tau}(a') \geq \hat{\tau}(a'')$. Part 1 of Assumption 2 implies that

$$\hat{\tau}(a') = \hat{\tau}(a''), \tag{18}$$

for otherwise, any type would strictly prefer a'' to a'. Furthermore, any type $\theta = (\eta, \gamma)$ that uses a'' must have $\eta \geq a''$, and hence

$$C(a', \theta) = C(a'', \theta) = 0, \tag{19}$$

for otherwise, Assumption 1 implies θ would strictly prefer a' to a''. Now consider a new strategy in which the agent behaves identically to the given equilibrium except for playing a' whenever she was to play a''. By Equation 18, the induced belief at a' does not change; hence, by Equation 19, this new strategy also constitutes an equilibrium in which a'' is off path.

Finally, note that in any equilibrium, assigning any off-path action $\tilde{a}$ the belief specified in item (ii) of the lemma preserves the property that no type has a strict incentive to use $\tilde{a}$. *Q.E.D.*

D. Proofs for Section 3

Proof of Observation 1. From (1), there is a separating equilibrium with $a_2 = \overline{\eta}$ (and $a_1 = \underline{\eta}$) if $s < s_\eta^{**}$, where $V(\overline{\eta}; s_\eta^{**}) - V(\underline{\eta}; s_\eta^{**}) = C(a_2, \theta_1)$. So restrict attention to $s \geq s_\eta^{**}$. For the existence of a separating equilibrium, there is no loss in assuming the second inequality in (1) holds with equality. This implicitly defines a strictly increasing and continuous function, $a_2(s)$, whose range is $[\overline{\eta}, \infty)$ for $s \in [s_\eta^{**}, \infty)$. A separating equilibrium exists at stakes s if and only if $H(a_2(s)) \geq 0$, where

$$H(a) := C(a, \theta_1) - C(a, \theta_2) \tag{20}$$

is continuous. Since $\text{sign}[H(a)] = \text{sign}[a^{\text{ce}} - a]$ by Lemma 1, it follows from the definition of s_η^* that a separating equilibrium exists if and only if $s \leq s_\eta^*$.

Now suppose $s > s_\eta^*$ and consider a strategy where each type mixes (possibly degenerately) between actions $\underline{\eta}$ and a^{ce}. By choosing the mixing probabilities suitably, we can induce via Bayes rule any $\hat{\eta}(a^{\text{ce}})$ and $\hat{\eta}(\underline{\eta})$ such that $\overline{\eta} \geq \hat{\eta}(a_2) > \mathbb{E}[\eta] > \hat{\eta}(a_1) \geq \underline{\eta}$. By the definition of s_η^*, there is a pair $\hat{\eta}(a^{\text{ce}})$ and $\hat{\eta}(\underline{\eta})$ satisfying these inequalities such that $V(\hat{\eta}(a^{\text{ce}}); s) - V(\hat{\eta}(\underline{\eta}); s) = C(a^{\text{ce}}, \theta_1) = C(a^{\text{ce}}, \theta_2)$; the corresponding mixing probabilities define an equilibrium strategy when belief $\underline{\eta}$ is assigned to any off-path actions. *Q.E.D.*

Proof of Observation 2. Define $a_1(s)$ by setting the second inequality of (2) to hold with equality; $a_1(s)$ is strictly increasing and unbounded above. It is straightforward that a separating equilibrium exists at stakes s if and only $a_1(s) > \overline{\eta}$ and $H(a_1(s)) \leq 0$, where $H(\cdot)$ was defined in (20). By continuity of $H(\cdot)$ and Lemma 1, there is separating equilibrium if and only if $s \geq s_\gamma^*$, where $V(\overline{\gamma}; s_\gamma^*) - V(\underline{\gamma}; s_\gamma^*) = C(a^{\text{ce}}, \theta_1)$.

Next, suppose $s \leq s_\gamma^{**}$ where s_γ^{**} solves $V(\overline{\gamma}; s_\gamma^*) - V(\underline{\gamma}; s_\gamma^*) = C(a^{\text{or}}, \theta_1)$. To contradiction, consider any informative equilibrium. Let a' and a'' be two on path actions such that $\hat{\gamma}(a') < \hat{\gamma}(a'')$. It follows that $a' < a'' \leq a^{\text{or}}$, by belief monotonicity and that θ_1 will not play any $a > a^{\text{or}}$. Since $C_a(a, \theta_2) < C_a(a, \theta_1)$ for $a \in (\underline{\eta}, a^{\text{or}})$, it further follows that type θ_2 will not play a'. But this implies $\hat{\gamma}(a') = \overline{\gamma}$, contradicting $\hat{\gamma}(a') < \hat{\gamma}(a'')$.

Finally, consider $s \in (s_\gamma^{**}, s_\gamma^*)$. Let the two types mix over two actions, $a_2 > a_1$. By choosing the mixing probabilities suitably, we can induce via Bayes rule any beliefs $\hat{\gamma}(a_2)$ and $\hat{\gamma}(a_1)$ such that $\hat{\gamma}(a_2) - \hat{\gamma}(a_1) \in [0, \overline{\gamma} - \underline{\gamma}]$. By Lemma 1, for any $x \in (0, C(a^{\text{ce}}, \theta_1)]$, we can find a_1 and a_2 such that $\underline{\eta} \leq a_1 < a^{\text{or}} < a_2 \leq a^{\text{ce}}$ and $C(a_2, \theta_1) - C(a_1, \theta_1) = C(a_2, \theta_2) - C(a_1, \theta_2) = x$. It follows that by assigning off-path actions the belief $\underline{\gamma}$, there is an informative equilibrium with two on-path actions so long as $V(\overline{\gamma}; s) - V(\underline{\gamma}; s) \leq C(a^{\text{ce}}, \theta_1)$ and $V(\overline{\gamma}; s) - V(\underline{\gamma}; s) > C(a^{\text{or}}, \theta_1)$. These two inequalities are assured by the hypothesis that $s \in (s_\gamma^{**}, s_\gamma^*)$. *Q.E.D.*

E. Proofs for Section 4

E.1. Proofs for Subsection 4.1

We first establish some straightforward preliminaries for the analysis of equilibria of the 2×2 model. Throughout this section, in addition to referring to $(\underline{\eta}, \overline{\gamma})$ as the gamer type and $(\overline{\eta}, \underline{\gamma})$ as the natural type (with a^{or} and a^{ce} being the order-reversing and cost-equalizing actions with respect to these cross types), we also refer to $(\underline{\eta}, \underline{\gamma})$ as the low type and $(\overline{\eta}, \overline{\gamma})$ as the high type.

Claim 4. *For any finite type space Θ, up to equivalence, there is a finite upper bound on the number of actions used in equilibrium over all s and all equilibria.*

Proof. First, for any type θ, there can only be a single action (up to equivalence) which is played by θ alone in a given equilibrium; otherwise that type would be playing two actions with the same beliefs but different costs. Second, for any pair of types, there can be at most two distinct actions that both types are both willing to play.[40] Consequently, an upper bound on the number of equilibrium actions is $|\Theta| + 2\binom{|\Theta|}{2}$. *Q.E.D.*

Claim 5. *Fix any finite type space Θ and let $s_n \to s^* > 0$ be a sequence of stakes. If $e_n \to e^*$, where each e_n is an equilibrium (strategy profile) at stakes s_n with corresponding distribution of market beliefs $\delta_n \in \Delta(\min \Theta_\tau, \max \Theta_\tau)$, then (i) there exists δ^* such that $\delta_n \to \delta^*$, and (ii) e^* is an equilibrium at s^*.[41]*

Proof. The claim follows from standard upper-hemicontinuity arguments, with one caveat. We need to show that if as $s_n \to s^*$ there are two sequences $a_n \to a^*$ and $a'_n \to a^*$, where a_n and a'_n are each on-path actions in e_n, then the respective equilibrium beliefs $\hat{\tau}(a_n)$ and $\hat{\tau}(a'_n)$ converge to the same limit. This ensures that the belief at a^* under e^* is equal to the limiting belief along both a_n and a'_n, and therefore that the payoff of a^* under e^* is equal to the limit of the payoffs along any sequence of actions approaching a^*, whereafter routine arguments apply.

Suppose, to contradiction, that $\hat{\tau}(a_n) \to h$ and $\hat{\tau}(a'_n) \to l$ with $h > l$. For any $\theta \in \Theta$, $C(a_n, \theta) \to C(a^*, \theta)$ and $C(a'_n, \theta) \to C(a^*, \theta)$; hence, for any θ and sufficiently large n, $V(\hat{\tau}(a_n); s_n) - C(a_n, \theta) > V(\hat{\tau}(a'_n); s_n) - C(a'_n, \theta)$, which contradicts a'_n being on path. *Q.E.D.*

Now, for the 2×2 model specifically, we formally establish that moving high-τ types from actions with low beliefs to ones with high beliefs, or moving low-τ types from high to low beliefs, increases information in the sense of Blackwell. By "moving" a type θ from action a to a' we mean marginally altering the mixed strategy to slightly reduce the probability that θ takes a, and to correspondingly increase the probability that it takes a'. (As established in Claim 4, there are finitely many actions in the support and each has strictly positive probability.) We informally applied this result in the analysis of cross types in Section 3 (see Figure 2 through Figure 4), and we will again apply it in the 2×2 model which embeds the cross types.

---

[40]Given any market belief function, $\hat{\tau}(a)$, type θ is said to be willing to play action a' if a' is optimal for θ.

[41]Convergence of probability distributions is in the sense of weak convergence. A sequence of equilibria converges if the corresponding mixed strategies of each type converge. Note that for any sequence of equilibria e_n, there is an equivalent subsequence that converges. This is because as $s_n \to s^*$, up to equivalence, equilibrium actions are contained in compact set that is bounded below by $\min \Theta_\eta > -\infty$ and above by $\tilde{a} < \infty$ satisfying $V(\max \Theta_\tau; s^*) = C(\tilde{a}, \max \Theta_\eta, \max \Theta_\gamma)$.

Claim 6. *In the 2×2 model, consider information on dimension τ, where $\Theta_\tau = \{\underline{\tau}, \overline{\tau}\}$. Take some distribution of types over actions, with two actions a_l and a_h in the support inducing respective beliefs $\hat{\tau}_l < \hat{\tau}_h$. If we move either a type with $\tau = \overline{\tau}$ from a_l to a_h, or move a type with $\tau = \underline{\tau}$ from a_h to a_l, then the posterior beliefs become more informative about the dimension of interest.*

Moving types across actions in the reverse way would lead to less informative rather than more informative beliefs.

Proof of Claim 6. One posterior distribution of beliefs is Blackwell more informative than another if and only if, for any continuous and convex function over beliefs, U, it holds that $\mathbb{E}[U(\beta_\tau)]$ is higher under the first distribution than the second. Moreover, in the 2×2 setting, beliefs β_τ about the dimension of interest are fully captured by the expectation $\hat{\tau}$. So moving types increases Blackwell information if and only if for any continuous and convex function $U : [\underline{\tau}, \overline{\tau}] \to \mathbb{R}$, the move yields an increase in $\mathbb{E}[U(\hat{\tau})]$.

To calculate $\mathbb{E}[U(\hat{\tau})]$, we first establish some notation. Let $f(\theta)$ be the probability of type θ under the prior distribution F, let $p_\theta(a)$ indicate the probability that an agent of type θ chooses action a under a specified strategy, and let $\tau(\theta)$ indicate the component of θ on the dimension of interest. Then the belief $\hat{\tau}(a)$ at action a is

$$\hat{\tau}(a) = \frac{\sum_\theta \tau(\theta) f(\theta) p_\theta(a)}{\sum_\theta f(\theta) p_\theta(a)},$$

and $\mathbb{E}[U(\hat{\tau})]$ is given by $\sum_a U(\hat{\tau}(a)) \sum_\theta f(\theta) p_\theta(a)$ where we sum over all actions a in the support.

The effect on $\mathbb{E}[U(\hat{\tau})]$ of a marginal move from a_l to a_h of a type θ' is given by $\frac{\mathrm{d}}{\mathrm{dp}_{\theta'}(a_h)} \mathbb{E}[U(\hat{\tau})] - \frac{\mathrm{d}}{\mathrm{dp}_{\theta'}(a_l)} \mathbb{E}[U(\hat{\tau})]$. When $\tau(\theta') = \overline{\tau}$, we can evaluate these derivatives and simplify to get

$$\frac{\mathrm{d}}{\mathrm{dp}_{\theta'}(a_h)} \mathbb{E}[U(\hat{\tau})] - \frac{\mathrm{d}}{\mathrm{dp}_{\theta'}(a_l)} \mathbb{E}[U(\hat{\tau})] = f(\theta') \cdot \left[\left(U(\tau_h) + (\overline{\tau} - \tau_h) U'(\tau_h) \right) - \left(U(\tau_l) + (\overline{\tau} - \tau_l) U'(\tau_l) \right) \right].$$

Convexity of U combined with $\tau_l < \tau_h \leq \overline{\tau}$ guarantees that the bracketed term is nonnegative (positive under strict convexity), so as required the move increases $\mathbb{E}[U(\hat{\tau})]$.

Likewise, the effect on $\mathbb{E}[U(\hat{\tau})]$ of a marginal move from a_h to a_l of a type θ' is given by $\frac{\mathrm{d}}{\mathrm{dp}_{\theta'}(a_l)} \mathbb{E}[U(\hat{\tau})] - \frac{\mathrm{d}}{\mathrm{dp}_{\theta'}(a_h)} \mathbb{E}[U(\hat{\tau})]$. With $\tau(\theta') = \underline{\tau}$ the expression evaluates to

$$f(\theta') \cdot \left[\left(U(\tau_l) - (\tau_l - \underline{\tau}) U'(\tau_l) \right) - \left(U(\tau_h) - (\tau_h - \underline{\tau}) U'(\tau_h) \right) \right]$$

which is nonnegative (positive) under $\underline{\tau} \leq \tau_l < \tau_h$ and (strict) convexity of U. *Q.E.D.*

To clarify terminology, say that a type plays an action if its strategy assigns positive probability to this action. An action is an equilibrium action if some type of positive measure plays this action. (Recall Claim 4.)

Claim 7. *In the 2×2 model, all equilibria are equivalent to one in which either (1) for any pair of actions weakly above $\underline{\eta}$, at most a single type is willing to play both of these; or (2) there exist actions a_1 and a_2 with $\underline{\eta} \leq a_1 < a_2$ such that at least two types are willing to play both a_1 and a_2. Under these cases:*

Case (1). *Take any three actions $a_1 < a_2 < a_3$, with $a_1 \geq \underline{\eta}$. If types θ_1 and θ_2 are both willing to play a_3, if θ_1 is willing to play a_1, and if θ_2 is willing to play a_2, then it cannot be the case that any type plays a_2.*

A-12

Case (2). *It holds that $\underline{\eta} \leq a_1 < a^{or} < a_2 \leq \overline{\eta}$. The gamer type $(\underline{\eta}, \overline{\gamma})$ and the natural type $(\overline{\eta}, \underline{\gamma})$ are both willing to play a_1 and a_2. No other type is willing to play both of these actions, and for any other pair of actions weakly above $\underline{\eta}$ at most a single type is willing to play both.*

Moreover, no action in (a_1, a_2) is played in equilibrium. The natural type $(\overline{\eta}, \underline{\gamma})$ only plays actions a_1 and/or a_2 in equilibrium, and is not willing to play any action a with $a < a_1$ or $a > a_2$. The low type $(\underline{\eta}, \underline{\gamma})$ is not willing to play any actions a with $a > a_1$, and the high type $(\overline{\eta}, \overline{\gamma})$ is not willing to play any action a with $a < a_2$.

Proof of Claim 7. Up to equivalence, we can take all equilibrium actions to be weakly above $\underline{\eta}$.

Case (1). Take three actions $a_1 < a_2 < a_3$, with types θ_1 and θ_2 both willing to play a_3, with θ_1 willing to play a_1, and θ_2 willing to play a_2. By assumption of Case (1), θ_2 is not willing to play a_1 and θ_1 is not willing to play a_2. So it must be that the types are not single-crossing ordered over the range of $[a_1, a_3]$; that is, the two types θ_1 and θ_2 must be the gamer $(\underline{\eta}, \overline{\gamma})$ and the natural $(\overline{\eta}, \underline{\gamma})$—not necessarily in that order—and it must be that $a_1 < a^{or} < a_3$. The low type $(\underline{\eta}, \underline{\gamma})$ can only take actions up to a_1, and the high type $(\overline{\eta}, \overline{\gamma})$ can only take actions down to a_3. So we see that only type θ_2 can be willing to play a_2.

Now suppose that type θ_2, which is the only one willing to play a_2, does play a_2 with positive probability. If it did, then the beliefs at a_2 would reveal the type of θ_2. But if θ_2 has a high type on the dimension of interest, $\tau = \overline{\tau}$, then a_2 would be at least as appealing to θ_1 as a_3, so θ_1 would be attracted to a_2. On the other hand, if θ_2 has a low type on the dimension of interest ($\tau = \underline{\tau}$), then a_1 would be at least as appealing to θ_2 as a_2, so θ_2 would be attracted to a_1. Either case yields a contradiction, since θ_1 is not willing to play a_2 and θ_2 is not willing to play to a_1.

Case (2). Take some pair of actions a_1 and a_2 that two types are both willing to play. This cannot hold for any pair of actions that are single-crossing ordered, and so it must be that the two types are $(\underline{\eta}, \overline{\gamma})$ and $(\overline{\eta}, \underline{\gamma})$. [42]

Both players are willing to play both actions means that $C(a_2, \underline{\eta}, \overline{\gamma}) - C(a_1, \underline{\eta}, \overline{\gamma}) = C(a_2, \overline{\eta}, \underline{\gamma}) - C(a_1, \overline{\eta}, \underline{\gamma})$. Since these cross types are single-crossing-ordered on the range of $a \geq a^{or}$ (where $(\underline{\eta}, \overline{\gamma})$ has lower marginal costs) and on the range of $a \leq a^{or}$ (where $(\overline{\eta}, \underline{\gamma})$ has lower marginal costs), it must be that one of these actions, a_1 must be strictly below a^{or} and the other, a_2, must be strictly above. There can only be a single pair of such actions; otherwise the two types would both be willing to play two distinct actions above a^{or} or below it, which we established cannot happen.

By assumption, we take the lower action a_1 to be weakly greater than $\underline{\eta}$. And the higher action a_2 must be weakly below a^{ce} because for any $\tilde{a} > a^{ce}$ and $a_1 < \tilde{a}$, it holds that $C(\tilde{a}, \underline{\eta}, \overline{\gamma}) - C(a_1, \underline{\eta}, \overline{\gamma}) < C(\tilde{a}, \overline{\eta}, \underline{\gamma}) - C(a_1, \overline{\eta}, \underline{\gamma})$.

By single crossing in the region above a^{or}, $(\overline{\eta}, \underline{\gamma})$ cannot be willing to take any actions above a_2 or else $(\underline{\eta}, \overline{\gamma})$ would strictly prefer that action to a_2; and by single-crossing in the region below a^{or}, the type $(\overline{\eta}, \underline{\gamma})$ cannot be willing to take any actions below a_1 or else $(\underline{\eta}, \overline{\gamma})$ would strictly prefer that action to a_1. Additionally, the high type $(\overline{\eta}, \overline{\gamma})$ is unwilling to play any action below a_2, and the low type is unwilling

---

[42] We can rule out that $(\overline{\eta}, \underline{\gamma})$ and $(\overline{\eta}, \overline{\gamma})$ are both indifferent over a pair of actions above $\underline{\eta}$ but below $\overline{\eta}$ because those actions would have the same costs and the same beliefs on the dimension of interest. So, up to equivalence, the two actions could be rolled into a single action. (We take our notion of equivalence to say that if two actions have the same costs for all types playing them and the same beliefs on the dimension of interest, but different beliefs on the opposite dimension, then it is "equivalent" to put them together into a single action.)

to play any action above a_1; if one of these types were willing to play such an action, then another type currently playing a_1 or a_2 would strictly prefer to deviate to that action.

Finally, $(\underline{\eta}, \overline{\gamma})$ is unwilling to take any action in (a_1, a_2) because if this type were willing to take such an action, then $(\overline{\eta}, \underline{\gamma})$ would strictly prefer it to a_1 and a_2. So only $(\overline{\eta}, \underline{\gamma})$ can possibly be willing to take an action in (a_1, a_2), but in equilibrium this type does not play any such actions because doing so would break the equilibrium. In particular, taking such an action in equilibrium would reveal her type. Under $\tau = \eta$ this would mean she strictly preferred the intermediate action with $\hat{\eta} = \overline{\eta}$ to a_2; and under $\tau = \gamma$ she would strictly prefer a_1 under $\hat{\gamma}(a_1)$ to the intermediate action under beliefs $\hat{\gamma} = \underline{\gamma}$.   Q.E.D.

We now proceed to prove the main lemmas. Throughout, we maintain the assumption that there is a positive measure of both natural and gamer types; otherwise, the type space is fully ordered and we can straightforwardly maintain the equilibrium information level of any equilibrium e_0 at stakes s_0 as stakes vary by sliding actions up and down. We allow for there to be a zero measure of high or low types, so that we subsume the Section 3 case of only two cross types.

Proof of Lemma 3. Starting from any given equilibrium at some stakes, we show that as stakes decrease the equilibrium can be continuously perturbed in a manner that increases information. We will give local arguments, which show the existence of a path of equilibria nearby the starting point. The upper-hemicontinuity of the equilibrium set (Claim 5) guarantees that this local construction around any given equilibrium e_0 at any stakes s_0 extends to a global path of equilibria on $s \in (0, \infty)$.[43]

There are two kinds of perturbations involved. One slides the location of actions up or down without changing the distribution of types across actions, which has no effect on information. The other follows steps similar to those discussed in Section 3. As stakes decrease, we move low types with $\eta = \underline{\eta}$ down from actions with high beliefs to ones with low beliefs, and/or move high types with $\eta = \overline{\eta}$ up from actions with low to high beliefs. (Recall that due to free downward deviations, higher equilibrium actions have strictly higher beliefs.) These moves spread beliefs out and, as formalized in Claim 6, increase information. As stakes increase, we can do the reverse moves to decrease information.

Using Claim 7, we can categorize all possible equilibria into a number of exhaustive cases (up to equivalence), and then address these cases separately.

Case 1. For any pair of distinct actions weakly above $\underline{\eta}$, at most a single type is willing to play both actions.

Case 2. There exist actions a_1 and a_2 satisfying $\underline{\eta} \leq a_1 < a^{\text{or}} < a_2 \leq \overline{\eta}$ such that the gamer type $(\underline{\eta}, \overline{\gamma})$ and the natural type $(\overline{\eta}, \underline{\gamma})$ are both willing to play a_1 and a_2. No actions in (a_1, a_2) are played in equilibrium. Natural types $(\overline{\eta}, \underline{\gamma})$ are not willing to play any action strictly below a_1 or above a_2, low types $(\underline{\eta}, \underline{\gamma})$ are not willing to play any action above a_1, and high types $(\overline{\eta}, \overline{\gamma})$ are not willing to play any action below a_2. If there is an equilibrium action a_0 strictly below a_1, it can only be played by types with $\eta = \underline{\eta}$, and so would have the worst possible beliefs; hence it must be that $a_0 = \underline{\eta}$.

We can divide this case into five subcases:

(a) *The actions a_1 and a_2 are played in equilibrium. Either $a_1 = \underline{\eta}$; or, $a_1 > \underline{\eta}$, and no type playing a_1 is willing to play $a = \underline{\eta}$ or any equilibrium any action in $(\underline{\eta}, a_1)$, and no type playing an equilibrium action below a_1 is willing to play a_1.*

---

[43]Our local arguments cover different cases separately, but extending to a global path may require patching together different cases as one leads in to another.

A-14

(b) *Either $a_1 = \underline{\eta}$ or a_2 is not played in equilibrium.* Up to equivalence, a_2 must be played in equilibrium; otherwise we could assign it low beliefs so that the natural and gamer type would strictly prefer a_1 to a_2. So it must be that a_2 is played, but a_1 is not played; moreover, up to equivalence, a_1 has the lowest possible beliefs $\hat{\eta} = \underline{\eta}$. For the gamer to be indifferent over a_1 and $\underline{\eta}$, then, it must be that $a_1 = \underline{\eta}$. Because low types play an action at least $\underline{\eta}$ and at most a_1, low types play a_1, and hence this case is only possible if the measure of low types is zero.

(c) *It holds that $a_1 > \underline{\eta}$, and there is some type of positive measure that plays both actions $a_0 = \underline{\eta}$ and a_1.* Such a type has $\eta = \underline{\eta}$, and can be a gamer or a low type. Beliefs $\hat{\eta}$ at a_0 are at $\underline{\eta}$. Beliefs at a_1 are in $(\underline{\eta}, \overline{\eta})$; the natural type plays a_1.

(d) *It holds that $a_1 > \underline{\eta}$, and there is some type of positive measure that plays action $a_0 = \underline{\eta}$, and that does not play a_1 but is willing to play a_1.* Such a type must be the low type—if it were the gamer, then only the natural type would play a_1, and the gamer would prefer a_1 over a_2. Beliefs at a_1 are in $(\underline{\eta}, \overline{\eta})$, and so the natural and gamer types both play a_1.

(e) *It holds that $a_1 > \underline{\eta}$, and there is some type of positive measure playing a_1 that is willing to play $a_0 = \underline{\eta}$ but does not play this action.* Such a type has $\eta = \underline{\eta}$, and can be a gamer or a low type. Beliefs at a_1 are in $(\underline{\eta}, \overline{\eta})$, and so the natural type must play a_1.

In all cases, we assume without loss that all equilibrium actions are weakly above $\underline{\eta}$.

Case 1. Suppose e_0 is a Case 1 equilibrium at stakes s_0: no two types are both willing to play the same pair of actions. We will show that as s varies locally, we can slide actions marginally up or down to maintain all indifferences without moving types across actions, and therefore without affecting the distribution of posterior beliefs. We work from left to right, the lowest equilibrium action to the highest. For all such actions we will check indifferences "to the left" – seeing whether any type that is willing to play the given action is also willing to play a lower action. Without loss, it suffices to check only indifferences to lower actions that are played in equilibrium with positive probability, and to action $a = \underline{\eta}$; by free downward deviations, other off-path actions can be taken to have sufficiently low beliefs that any agent willing to deviate to one of those would also deviate to a lower equilibrium action or to $a = \underline{\eta}$.

Base case: Start with the lowest equilibrium action, i.e., the lowest action played with positive probability by any type. If this action is $\underline{\eta}$, then keep it at $\underline{\eta}$ and move on to the next step. Otherwise, check if any type playing this action—in particular, the relevant one would just be the low type $(\underline{\eta}, \gamma)$—is willing to play $a = \underline{\eta}$ as well at the equilibrium beliefs. If not, then as we locally vary s no agent type wants to deviate down to $\underline{\eta}$, and so again keep this action fixed and move on to the next step. So suppose that there is a type playing this action which is willing to play $\underline{\eta}$; by assumption of Case 1 there can only be a single indifferent type. As s varies, slide this lowest equilibrium action up or down to keep this type indifferent at the given beliefs. In particular, when stakes s increase then the appeal of the current action relative to $a = \underline{\eta}$ increases, as there is now a larger benefit of taking a higher action to get higher beliefs, and so we slide the action up to recover the indifference by raising the costs of taking this. When stakes s decrease then the current action becomes less appealing relative to $a = \underline{\eta}$, and so we lower the action to recover the indifference by lowering costs. All the while we maintain the probability that each type chooses this action as it shifts around and therefore keep fixed beliefs at this action. Hence, as we locally vary s, no types currently playing this action want to deviate down to $a = \underline{\eta}$.

Inductive step: move on to the next-highest equilibrium action. Look at all types willing to play this action (whether they play it in equilibrium or not). If none of these types are willing to play a lower equilibrium

A-15

action or action $a = \underline{\eta}$, then keep this action fixed as we locally vary s; no types currently playing this action become attracted to a lower action, and no types playing a lower action become attracted to this one. If there is such a binding indifference, then again there can only be a single indifferent type; this follows from the assumption that no two types are indifferent over the same pair of actions, combined with the characterization of Case (1) from Claim 7 that if two types are willing to play two different lower actions a_1 and a_2, then a_2 cannot be an equilibrium action. Proceed as above, sliding the action up or down to maintain the indifference to lower actions without changing beliefs or moving types across actions. (Here the indifferences are affected both by the direct change in the stakes, and also by potentially having moved the lower actions up or down in previous steps.)

Continue to proceed by induction for each next-higher equilibrium action until we are done. (Recall that there are only finitely many equilibrium actions.) This gives us a new equilibrium at the locally perturbed stakes: no type playing one action strictly wants to shift to any new action, because every previously optimal action remains optimal. This new equilibrium induces the same distribution over beliefs, and so information has not changed as we varied s.

Case 2. First we will move types across actions and/or slide locations of actions in order to maintain the appropriate indifferences across all actions actions at or below a_2. We treat each subcase separately and show that for an increase in s these moves will decrease information, and for a decrease in s these moves will increase information. Following that, without treating each subcase separately, we will slide around actions to maintain appropriate indifferences over actions above a_2 in a way that does not additionally change information.

Subcase (a). For a marginal decrease in s, types which were previously willing to play both a_1 and a_2 become more attracted to a_1. In this subcase there are no relevant binding incentive constraints attracting types playing a_1 to actions below a_1. Consider two possibilities. First, beliefs $\hat{\eta}$ at either a_1 or a_2 are in the range $(\underline{\eta}, \overline{\eta})$, so either natural types play a_1 or gamer types play a_2 with positive probability. In that case, we follow the cross-type logic of Figure 2 and move either natural types up from a_1 to a_2, and/or gamer types down from a_2 to a_1, to increase beliefs at a_2 and decrease beliefs at a_1 until we recover the appropriate indifference of gamers and naturals between actions a_1 and a_2. By Claim 6, these moves increase information. The second possibility is that only types with $\eta = \underline{\eta}$ play a_1, and only types with $\eta = \overline{\eta}$ play a_2—we already have full separation. In that case the natural types at a_2 become more attracted to a_1 (which they were previously indifferent to); we can slide a_2 down to lower the cost of a_2 and recover the indifference of the natural type across a_1 and a_2. Because $a_2 > a^{\mathrm{or}}$, sliding a_2 down lowers the cost for the natural more than for the gamer type, and because the natural type is indifferent, the gamer type now strictly prefers a_1, the action it was playing, to a_2, the action it was not playing. In any event, this slide does not change information.

For a marginal increase in s, types which were previously willing to play both a_1 and a_2 become more attracted to a_2. Again, consider two possibilities. First, either natural types play a_2 or gamer types play a_1 with positive probability. In that case, we effectively reverse the direction of Figure 2: move gamers up from a_1 to a_2, and/or move naturals down from a_2 to a_1, to decrease beliefs at a_2 and increase beliefs at a_1 until we recover the indifferences. Such a move decreases information. Second, no natural types play a_2 and no gamer types play a_1 (this can occur if enough high types play a_2, and enough low types play a_1, that beliefs at a_2 are above beliefs at a_1). In that case, slide a_2 up without moving types across actions until we recover the indifference of gamers across a_1 and a_2. Because a_2 was above a^{or}, sliding a_2 up imposes a higher cost increase on naturals than on gamers, and so naturals are no longer

A-16

indifferent between a_1 and a_2; they now strictly prefer a_1, which they are already playing. In any event, this slide does not change information.

Subcase (b). For a marginal decrease in s, the argument proceeds as in subcase (a). Here we are in the "first possibility" where all natural and gamer types play a_2, and so moving gamer types down increases beliefs at a_2 while holding fixed beliefs of $\hat{\eta} = \underline{\eta}$ at a_1.

For a marginal increase in s, we simply keep a_2 as it is: natural and gamer types now strictly prefer a_2 to a_1, and they were not previously playing a_1.[44]

Subcase (c). For a marginal decrease in s, types which were previously willing to play both a_1 and a_2 become more attracted to a_1, and the indifferent type between a_0 and a_1 becomes more attracted to a_0. Consider two possibilities. First, gamers play a_2 with a positive probability. In that case we follow the logic of Figure 3: move the indifferent type (which may be gamers or low types, but in either case have $\eta = \underline{\eta}$) left from a_1 to a_0 to increase beliefs at a_1 and recover the indifference between a_0 and a_1. Then move gamer types left from a_2 to a_0 to increase beliefs at a_2 and recover the indifference between a_2 and a_1. Both moves increase information. The second possibility is that no gamer types play a_2. In that case, we again start by moving the indifferent type from a_1 to a_0 to recover that indifference and increase information. Then we slide a_2 down to recover the indifference of natural types between a_1 and a_2; this leaves gamer types now strictly preferring a_1 over a_2, and does not affect information.

For a marginal increase in s, types which were previously willing to play both a_1 and a_2 become more attracted to a_2, and the indifferent type between a_0 and a_1 becomes more attracted to a_1. Now consider two possibilities. The first possibility is that gamers play a_1 with positive probability. In that case we do two steps: first, move the indifferent type (with $\eta = \underline{\eta}$) up from a_0 to a_1 to lower beliefs at a_1 and recover the indifference of that type across a_1 and a_2. This decreases information, and also makes a_1 less attractive relative to a_2. Second, move the indifferent type right from a_0 to a_1 while moving gamers right from a_1 to a_2 at exactly the same rate (or, if the indifferent type was the gamer type, we can move gamers directly from a_0 to a_2, essentially reversing the direction of Figure 3); this keeps beliefs at a_1 fixed, and also at a_0, because beliefs were already at $\hat{\eta} = \underline{\eta}$. But it decreases beliefs at a_2, and so we can do this until we recover the indifference of naturals and gamers across a_1 and a_2. The second possibility is that gamers do not play a_1. In that case the indifferent type must be the low type, and gamers must not be willing to play a_0. Here, we first move the low type up from a_0 to a_1 to lower beliefs at a_1 and recover that indifference, just as before. This decreases information while further increasing the attractiveness of a_2 relative to a_1. Next, we slide action a_2 up to recover the indifference of natural types across a_1 and a_2; the gamer types now strictly prefer a_2 to a_1, since their costs of a_2 increase by less than those of the natural types. But they were not previously playing any action below a_2, so this does not affect their behavior.

Subcase (d). For a decrease in s, the low type which is indifferent over $a_0 = \underline{\eta}$ and a_1 becomes more attracted to a_0 relative to a_1, since the benefit of higher beliefs has gone down while costs have not changed. Likewise the gamer and natural type become more attracted to a_1 relative to a_2. Consider two possibilities. The first is that gamer types do not play a_2, so that beliefs at a_2 are at $\hat{\eta} = \overline{\eta}$. In that case we can move natural types up from a_1 to a_2 to decrease beliefs at a_1 until we recover the indifference of the natural and gamer types between a_1 and a_2. This makes a_1 less appealing relative

---

[44]The argument of the "first possibility" of subcase (a) would fail because moving natural types down would discretely rather than continuously increase equilibrium beliefs at a_1 from $\underline{\eta}$ to $\overline{\eta}$.

A-17

to a_0, so we do not have to worry about low types becoming attracted from a_0 to a_1. The second possibility is that gamer types play a_2. In that case we move gamer types down from a_2 to a_1 to reduce beliefs at a_1 and increase beliefs at a_2, recovering the indifference of naturals and gamers between a_1 and a_2 while increasing information. This again makes a_1 less appealing to the low types relative to a_0, so they continue to play a_0.

For a increase in s, proceed as in subcase (c).

Subcase (e). For a decrease in s, proceed as in subcase (c).

For an increase in s, proceed as in subcase (a). The increase in s makes a_1 more appealing relative to a_0, and any moves of types across actions do not decrease beliefs at a_1, and so the previously indifferent type now strictly prefers a_1 to a_0.

As mentioned earlier, the above analysis corrects all of the incentives across actions less than or equal to a_2. We now turn to actions above a_2. If the equilibrium has no actions above a_2 (i.e., high types only play a_2) then of course we are done. It may also be the case that the equilibrium has one action above a_2 (call it a_3), or two actions above a_2 (a_3 and a_4). When there is one action above a_2, it is played by the high type and possibly by the gamer type. When there are two actions above a_2, the lower one a_3 is played by the gamer type and the high type while the higher one a_4 is played only by the high type (the gamer and high type cannot both be indifferent over a pair of actions a_3 and a_4). In any case, as we vary s, we can slide actions around as in Case 1 to recover the appropriate equilibrium indifferences . Moving again from left to right, if a type that plays a_3 had been willing to play a_2, then after the above perturbations simply slide up or down as appropriate a_3 to maintain this indifference; if no such type had been indifferent, then keep a_3 where it was. Then do the same for a_4, maintaining any relevant indifferences with lower actions. There are no complications here because the two relevant types are ordered by single-crossing.   *Q.E.D.*

Proof of Lemma 4. We proceed similarly to the proof of Lemma 3 above, with three key differences. First, fixing an equilibrium e_0 at stakes s_0, we generate a path of equilibria only for $s \geq s_0$. Second, with the dimension of interest equal to gaming ability rather than natural action, the direction of the information effect is reversed: as stakes increase, we find perturbations that increase information. We do this by only ever moving low-γ types down, from actions with high beliefs to ones with low beliefs, and moving high-γ types up from low beliefs to high. Third, while in Lemma 3 all of the perturbations moved equilibria continuously, here we sometimes use a trick of looking for continuous perturbations about an equilibrium that induces the same distribution of beliefs as equilibrium e_0, but has different strategies for some types. This gives us a path of equilibria over $s \in [s_0, \infty)$ in which distributions of beliefs vary continuously with the stakes, but strategies may jump discretely at s_0.

Using Claim 7, we can categorize all possible equilibria into a number of exhaustive cases (up to equivalence), and then address these cases separately.

Case 1. For any pair of distinct actions weakly above $\underline{\eta}$, at most a single type is willing to play both of these.

Case 2. There exist actions a_1 and a_2 satisfying $\underline{\eta} \leq a_1 < a^{\text{or}} < a_2 \leq \overline{\eta}$ such that the gamer type $(\underline{\eta}, \overline{\gamma})$ and the natural type $(\overline{\eta}, \underline{\gamma})$ are both willing to play a_1 and a_2. No actions in (a_1, a_2) are played in equilibrium. Natural types $(\overline{\eta}, \underline{\gamma})$ are not willing to play any action strictly below a_1 or above a_2, low types $(\underline{\eta}, \underline{\gamma})$ are not willing to play any action above a_1, and high types $(\overline{\eta}, \overline{\gamma})$ are not willing to play any action below a_2.

Any equilibrium action above a_2 can only be played by the gamer and high type, and so it must have belief of $\hat{\gamma} = \overline{\gamma}$; thus, there can be at most one such equilibrium action, a_3. We now divide this case into subcases on two separate dimensions: each equilibrium is in one category (i)-(iii) characterizing its higher actions above a_2, and in one category (a)-(c) characterizing its lower actions below a_1.

(i) *Some type plays both a_2 and $a_3 > a_2$. This type can be the high type or the gamer.*

(ii) *No type plays an action above a_2.*

(iii) *The action a_3 is played only by the high type, and the high type does not play a_2. Here the gamer must play a_2; otherwise no $\gamma = \overline{\gamma}$ types would be playing a_2, so it would have to have beliefs $\hat{\gamma} = \underline{\gamma}$ and so a_1 would be preferred by all types to a_2.*

(a) *Either there are no equilibrium actions below a_1; or, there are lower actions, but no type playing a lower action is willing to play a_1. Because only low and gamer types can be willing to play actions below a_1, and because gamer types are willing to play a_1, it means that only a low type could be playing a lower action without being willing to play a_1. Up to equivalence, this low type would play $a = \underline{\eta}$.*

(b) *Some type plays an equilibrium action below a_1, call it a_0, and is also willing to play a_1. Moreover, natural types play a_2 in equilibrium. The indifferent type between a_0 and a_1 can be a low or a gamer type. There can only be one such indifference. In this case a_1 must be an equilibrium action.*

(c) *Some type plays an equilibrium action below a_1, call it a_0, and is also willing to play a_1. Moreover, natural types do not play a_2 in equilibrium. The indifferent type between a_0 and a_1 can be a low or a gamer type. There can only be one such indifference. In this case a_1 is an equilibrium action. Note that because natural types do not play a_2, beliefs are $\hat{\gamma} = \overline{\gamma}$ at a_2 so we must be in subcase (ii) as well, where there are no actions above a_2.*

In all cases, we assume without loss that all equilibrium actions are weakly above $\underline{\eta}$.

Case 1. This case proceeds exactly as in the proof of Lemma 3. As stakes increase, we can perturb the equilibrium by sliding actions around in a way which has no impact on the distribution of beliefs.

Case 2. First, prior to varying the stakes, we tweak the equilibrium in the following way which does not affect the information. In subcase (i), do nothing. In subcase (ii), do not change any strategies, but define a_3 to be the action such that the highest type playing a_2 (the high type, if such types have positive measure; the gamer otherwise) would be just indifferent between a_2 at the current beliefs under e_0 and action a_3 under beliefs $\hat{\gamma} = \overline{\gamma}$. Finally, in subcase (iii), slide a_3 down by a discrete amount until the gamer type is just indifferent to playing a_3. At this level the high type still strictly prefers a_3 to a_2 at the beliefs under equilibrium e_0. The key to these tweaks is that now the highest type playing a_2 (either gamer or high types) has become indifferent to deviating up to a_3, if it was not already indifferent.

Starting from this tweaked equilibrium, we now look for continuous perturbations that increase information as stakes s increase. As stakes increase, types playing lower actions may become more attracted to higher actions at higher beliefs.

Now, consider any equilibrium actions strictly below a_1. Only low and gamer types can play actions below a_1, and there are either zero, one, or two of these actions. If there are no such actions, then types playing a_1

only become less inclined to deviate downwards as s increases, so we can move on. If there is one such action, fix that action and move on. If there are two, then as stakes increase, fix the lowest action, and marginally slide the second-lowest action up as necessary in order to make sure that types playing the lowest action do not now want to deviate to the second one.

As stakes increase, types playing lower actions may become more attracted to a_1, but the reverse does not hold. Moreover, if in the previous step we slid a lower action up to maintain indifferences, that only makes the lower action even less appealing to a type playing a_1. So we now turn to maintaining indifferences between the lower actions and a_1, and also across higher actions. We treat subcases (a)-(c) separately.

Subcase (a). Here there are no relevant indifferences between lower actions and action a_1: as we marginally increase stakes, any types playing actions below a_1 do not become attracted to a_1. So we can move on to indifferences between actions a_1 and higher.

The increase in s makes makes a_2 more appealing relative to a_1. So to recover the indifference of naturals and gamers between a_1 and a_2, we move the highest type playing a_2 (either gamers or high types) up to a_3, as in Figure 4. (The equilibrium tweak from before guarantees that this type was previously indifferent to playing a_2 or a_3.) This move reduces beliefs at a_2 as desired, and makes the equilibrium more informative. Finally we slide a_3 as necessary to maintain the indifference of the type being moved between a_2 and a_3.

Subcase (b). As stakes increase, and as we potentially slide a_0 upwards to prevent deviations to a_0 from lower actions, action a_1 has become more attractive to the previously indifferent type. To recover the indifference between a_0 and a_1, we move natural types down from a_2 to a_1 to reduce beliefs at a_1. This move increases information, since it moves a low-γ type from a high belief action to a low one.

The increase in s makes makes a_2 more appealing relative to a_1, and the above move of natural types from a_2 to a_1—increasing beliefs at a_2 and decreasing beliefs at a_1—does the same. So to recover the indifference of naturals and gamers between a_1 and a_2, we proceed as in subcase (a) and move gamers or high types up from a_2 to a_3, then slide a_3 as necessary. These moves make the equilibrium more informative.

Subcase (c). As stakes increase, and as we potentially slide a_0 upwards to prevent deviations to a_0 from lower actions, action a_1 has become more attractive to the previously indifferent type. To recover the indifference between a_0 and a_1, we slide action a_1 up without moving types across actions. This has no effect on information. Next, slide action a_2 up to keep gamer types indifferent between a_2 and a_1; an increase in stakes has made a_2 relatively more attractive, and sliding a_1 up does the same, so we have to increase the costs of a_2 to keep the natural and gamer types from deviating to that. Notice that because $a_1 < a^{\mathrm{or}}$, sliding a_1 up increases the cost of taking a_1 more for gamer types than for naturals; and because $a_2 > a^{\mathrm{or}}$, sliding a_2 up increases the cost of taking a_2 less for gamer types than for naturals. So if gamers have been made indifferent between a_1 and a_2, naturals now strictly prefer a_1; but by assumption of subcase (c), the naturals had not been playing a_2, so we maintain the equilibrium. (In this subcase there is no need to address action a_3 because no types had been playing any action above a_2.) *Q.E.D.*

E.2. Proofs for Subsection 4.2

Proof of Proposition 2. Consider the claim about small stakes. By hypothesis, Θ_η is finite. We claim that for small enough s there is an equilibrium in which every type $\theta = (\eta, \gamma)$ takes its natural action, $a = \eta$; any

A-20

off-path action $a \notin \Theta_\eta$ is assigned the belief $\hat{\tau} = \min \Theta_\eta$. Clearly, no type has a profitable deviation to any off-path action nor to any action below its natural action. So it suffices to show that there is no incentive for any type to deviate to any on-path action above its natural action (an "upward deviation") when s is small enough. The proposition's hypotheses about Θ imply that there is some $\varepsilon > 0$ such that $C(a, \eta, \gamma) > \varepsilon$ for any $(\eta, \gamma) \in \Theta$ and $a \in \Theta_\eta \cap (\eta, \infty)$. For any type, the gain from deviating to any action is bounded above by $V(\max \Theta_\eta; s) - V(\min \Theta_\eta; s)$, which, by Assumption 2, tends to 0 as $s \to 0$. It follows that for small enough $s > 0$, the cost of any upward deviation outweighs the benefit for all types.

The claim about large stakes follows from Lemma 5 and Part 3 of Assumption 2. $\hfill$ Q.E.D.

Proof of Lemma 5. Fix an arbitrary type space containing $\theta_1 <> \theta_2$ and an arbitrary equilibrium in which each θ_i $(i = 1, 2)$ uses action a_i inducing belief $\hat{\eta}_i$. If $\hat{\eta}_2 \leq \hat{\eta}_1$ then the result is trivially true, so suppose $\hat{\eta}_2 > \hat{\eta}_1$. By belief monotonicity, $a_1 < a_2$. Incentive compatibility implies

$$C(a_2, \theta_2) - C(a_1, \theta_2) \leq V(\hat{\eta}_2; s) - V(\hat{\eta}_1; s) \leq C(a_2, \theta_1) - C(a_1, \theta_1).$$

Hence, $V(\hat{\eta}_2; s) - V(\hat{\eta}_1; s)$ is bounded above by the maximum of $C(a_2, \theta_1) - C(a_1, \theta_1)$ subject to $a_2 \geq a_1$ and $C(a_2, \theta_1) - C(a_1, \theta_1) \geq C(a_2, \theta_2) - C(a_1, \theta_2)$. Lemma 1 implies that the constraint is violated if $a_2 > a^{\mathrm{ce}}$; hence the maximum is obtained when $a_2 = a^{\mathrm{ce}}$ and $a_1 = \underline{\eta}$. $\hfill$ Q.E.D.

Proof of Proposition 3. Given an equilibrium, let $\hat{\eta}(\theta)$ denote a belief induced by type θ. Given a sequence of equilibria as $s \to \infty$, let $\hat{\eta}^*(\theta)$ denote any limit point of such beliefs as $s \to \infty$ (passing to sub-sequence if necessary). We first claim that in any sequence of equilibria, it holds for any $\theta' = (\eta', \gamma')$ and $\theta'' = (\eta'', \gamma'')$ with $\gamma'' > \gamma'$ that $\hat{\eta}^*(\theta'') \geq \hat{\eta}^*(\theta')$; in words, in the limit any type with a higher gaming ability induces a weakly higher belief about its natural action. If $\eta'' \geq \eta'$, the claim follows from the fact that θ'' and θ' are ordered by single-crossing and hence θ'' must induce a weakly larger belief than θ' in any equilibrium; if $\eta'' < \eta'$, the claim follows from Lemma 5 and Part 3 of Assumption 2.

Now fix an arbitrary sequence of equilibria as $s \to \infty$. It suffices to prove that for any type $\theta = (\eta, \gamma)$ with $\gamma \in (\min \Theta_\gamma, \max \Theta_\gamma)$, $\hat{\eta}^*(\theta) = \mathbb{E}[\eta]$. To contradiction, suppose there is a type θ' with $\gamma' \in (\min \Theta_\gamma, \max \Theta_\gamma)$ and $\hat{\eta}^*(\theta') = \mathbb{E}[\eta] + \Delta$ for $\Delta > 0$. (A symmetric argument applies if $\Delta < 0$.) Let $S_\Delta = \{\theta : \hat{\eta}^*(\theta) \geq \mathbb{E}[\eta] + \Delta\}$. The claim in the previous paragraph establishes that there exists some $\tilde{\gamma} < \max \Theta_\gamma$ such that S_Δ contains all types with gaming ability strictly above $\tilde{\gamma}$ and no types with gaming ability strictly below $\tilde{\gamma}$. In other words, modulo "boundary types" with $\gamma = \tilde{\gamma}$—a set that has probability zero, by the hypothesis of a continuous marginal distribution of γ—we can take $S_\Delta = \{\theta'' = (\eta'', \gamma'') | \gamma'' \geq \tilde{\gamma}\}$; note that S_Δ has positive probability. Hence, $\mathbb{E}[\eta | \theta \in S_\Delta] = \mathbb{E}[\eta | \gamma \geq \tilde{\gamma}] \leq \mathbb{E}[\eta]$, where the inequality owes to the hypothesis that $\mathbb{E}[\eta | \gamma]$ is non-increasing in γ. But this contradicts the Bayesian consistency requirement that $\mathbb{E}[\eta | \theta \in S_\Delta] = \mathbb{E}[\hat{\eta}^*(\theta) | \theta \in S_\Delta] \geq \mathbb{E}[\eta] + \Delta$. $\hfill$ Q.E.D.

Proof of Proposition 4. Let $\eta_{\max} := \sup \Theta_\eta$ and $\eta_{\min} := \inf \Theta_\eta$ and order the values of $\gamma \in \Theta_\gamma$ as $\gamma_1 < \gamma_2 < \cdots < \gamma_N$. Let a_i^{ce} be the cost-equalizing action between types $(\eta_{\max}, \gamma_i)$ and $(\eta_{\min}, \gamma_{i+1})$. At actions $a'' > a_i^{ce}$ and $a' < a''$, it holds that $C(a'', (\eta', \gamma_{i+1})) - C(a', (\eta', \gamma_{i+1})) < C(a'', (\eta'', \gamma_i)) - C(a', (\eta'', \gamma_i))$ for any η', η'' in Θ_η.

Now define $a_i(s)$ as follows. Set $a_1(s) = \eta_{\max}$. For $i \geq 1$, given $a_i(s)$, inductively define $a_{i+1}(s)$ to be the action such that $C(a_{i+1}(s), (\eta_{\max}, \gamma_i)) - C(a_i(s), (\eta_{\max}, \gamma_i)) = V(\gamma_{i+1}; s) - V(\gamma_i; s)$. To help interpret, observe that this would be the sequence of least-cost separating actions at stakes s if the type space were $\{\eta_{\max}\} \times \Theta_\gamma$

A-21

rather than Θ. Assumption 2 implies that for any $i \geq 2$, $a_i(s) \to \infty$ as $s \to \infty$. Hence, there exists $\tilde{s}$ such that for any $s > \tilde{s}$, $a_{i+1}(s) > a_i^{ce}$ for each $i = 1, ..., N-1$.

We claim that for any $s > \tilde{s}$, there is an equilibrium in which (i) any type $(\eta, \gamma_1) \in \Theta$ takes action η and (ii) any type $(\eta, \gamma_i) \in \Theta$ with $i > 1$ takes action $a_i(s)$. Plainly, this strategy is separating on γ. To see that we have an equilibrium when $s > \tilde{s}$, first consider local incentive constraints among the on-path actions. Plainly, no type wants to deviate upwards, because $(\eta_{\max}, \gamma_i)$ is by construction indifferent between playing a_i and a_{i+1}, while all other (η, γ_i) types prefer a_i to a_{i+1}. No type wants to deviate downwards because $a_{i+1}(s) > a_i^{ce}$, hence the indifference of $(\eta_{\max}, \gamma_i)$ between a_i and a_{i+1} implies that for any η, type (η, γ_{i+1}) prefers a_{i+1} to a_i. Standard arguments using the single-crossing property on dimension γ then imply global incentive compatibility among on-path actions. Finally, off-path actions can be deterred by assigning them the lowest belief, γ_1. $Q.E.D.$

Proof of Lemma 6. Suppose there is an equilibrium with a pair of actions a_1 and a_2, used respectively by θ_1 and θ_2, yielding beliefs $\hat{\gamma}_1 > \hat{\gamma}_2$. By belief monotonicity, $a_1 > a_2$. That neither type strictly prefers the other action over its own implies $a_1 > a^{or}$, for otherwise $C(a_1, \theta_2) - C(a_2, \theta_2) < C(a_1, \theta_1) - C(a_2, \theta_1)$. That type θ_1 is willing to play action a_1 rather than its natural action implies $C(a_1, \theta_1) \leq V(\max \Theta_\gamma; s) - V(\min \Theta_\gamma; s)$, because the LHS of the inequality is the incremental cost while the RHS is an upper bound on the incremental benefit. It follows from the strict monotonicity of $C(\cdot; \theta_1)$ on (a^{or}, a_1) that $C(a^{or}, \theta_1) < V(\max \Theta_\gamma; s) - V(\min \Theta_\gamma; s)$. $Q.E.D.$

Proof of Proposition 5. We omit the proof because the argument is analogous to that provided for Proposition 3, switching γ and η, taking $s \to 0$ rather than $s \to \infty$, and using Lemma 6 to conclude that in the limit of vanishing stakes, any type with a higher natural action induces a weakly higher belief about its gaming ability. $Q.E.D.$

F. Proofs for Section 5

F.1. Smooth pure-strategy equilibria are linear

Definition 1. A *smooth pure-strategy equilibrium* is one in which the agent plays a pure strategy $a^*(\eta, \gamma)$ that is continuous in both arguments and the market posterior $\hat{\tau}(a)$ is twice differentiable.

Lemma 8. *In any smooth pure-strategy equilibrium, for every γ and a, there is some η such that $a^*(\eta, \gamma) = a$.*

Proof. Since the agent solves

$$\max_{a \in \mathbb{R}} \left[s\gamma\hat{\tau}(a) - \frac{(a-\eta)^2}{2} \right], \tag{21}$$

differentiability of $\hat{\tau}(\cdot)$ yields the first-order condition

$$a^*(\eta, \gamma) = \eta + s\gamma\hat{\tau}'(a^*(\eta, \gamma)). \tag{22}$$

Now fix an arbitrary γ. We show that $a^{\max} := \sup_\eta a^*(\eta, \gamma) = \infty$. Suppose, per contra, that $a^{\max} < \infty$. Standard arguments establish that $a^*(\cdot, \gamma)$ is non-decreasing. Hence, for all sufficiently large η, $a^*(\eta, \gamma)$ is arbitrarily close to $a^{\max}$. But then, since $\hat{\tau}(\cdot)$ is twice differentiable (hence $\hat{\tau}'(a) \to \hat{\tau}'(a^{\max})$ as $a \to a^{\max}$),

A-22

Equation 22 cannot be satisfied for type (η, γ) once η is large enough, a contradiction. An analogous argument shows that $\inf_\eta a^*(\eta, \gamma) = -\infty$. Since $a^*(\cdot, \gamma)$ is continuous, the lemma's conclusion follows. $\quad$ Q.E.D.

Proposition 11. *Any smooth pure-strategy equilibrium is linear.*

Proof. It suffices to show that the market posterior must be linear as this implies that the agent's strategy must be linear (Lemma 9). Since $\hat{\tau}(\cdot)$ is twice differentiable, (21) implies the second-order condition

$$s\gamma\hat{\tau}''(a^*(\eta, \gamma)) \leq 1. \tag{23}$$

Pick any $a \in \mathbb{R}$. If $\hat{\tau}''(a) > 0$ then for large enough $\gamma > 0$, (23) will be violated for a type (η, γ) choosing a (existence of such a type is assured by Lemma 8); analogously if $\hat{\tau}''(a) < 0$. Hence, $\hat{\tau}''(a) = 0$ for all a, which implies $\hat{\tau}(\cdot)$ is a linear function. $\quad$ Q.E.D.

F.2. Characterizing linear equilibria

To compactly state the following result characterizing linear equilibria in the LQN specification, define the following notation: for any $d \in \{\eta, \gamma\}$, let $\neg d$ denote η if $d = \gamma$ and γ if $d = \eta$.

Lemma 9. *Assume $\eta, \gamma \sim \mathcal{N}(\mu_\eta, \mu_\gamma, \sigma_\eta^2, \sigma_\gamma^2, \rho)$. Given linear market beliefs $\hat{\tau}(a) = La + K$ (Equation 4), the corresponding agent optimal action is given by $a = \eta + sL\gamma$ (Equation 5).*

Conversely, given a linear strategy of the form $a = \eta + s\tilde{L}\gamma$, and denoting $\mathcal{L}(s, \tilde{L}, \eta) \equiv 1$ and $\mathcal{L}(s, \tilde{L}, \gamma) \equiv s\tilde{L}$, the market belief on dimension $d \in \{\eta, \gamma\}$ is linear in the agent's action, with slope coefficient

$$\frac{\mathcal{L}(s, \tilde{L}, d)\sigma_d^2 + \mathcal{L}(s, \tilde{L}, \neg d)\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2\tilde{L}^2\sigma_\gamma^2 + 2s\tilde{L}\rho\sigma_\gamma\sigma_\eta}.$$

Therefore, a linear equilibrium for dimension of interest $\tau \in \{\eta, \gamma\}$ is characterized by (4) and (5), where L solves

$$L = \frac{\mathcal{L}(s, L, \tau)\sigma_\tau^2 + \mathcal{L}(s, L, \neg\tau)\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2L^2\sigma_\gamma^2 + 2sL\rho\sigma_\gamma\sigma_\eta}. \tag{24}$$

Proof. First consider the market posterior given a linear strategy. Let $\eta, \gamma \sim \mathcal{N}(\mu_\eta, \mu_\gamma, \sigma_\eta^2, \sigma_\gamma^2, \rho)$ and $a = l_\eta\eta + l_\gamma\gamma + k$. It is a standard result that for any $d \in \{\eta, \gamma\}$, the joint distribution of a and d is

$$\mathcal{N}\left(l_\eta\mu_\eta + l_\gamma\mu_\gamma + k, \mu_d, l_\eta^2\sigma_\eta^2 + l_\gamma^2\sigma_\gamma^2 + 2l_\eta l_\gamma\rho\sigma_\eta\sigma_\gamma, \sigma_d^2, \frac{l_d\sigma_d^2 + l_{\neg d}\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta\sigma_\gamma}\right).$$

Routine manipulations then yield that for any observed action, a, the marginal distribution of the market posterior on dimension $d \in \{\eta, \gamma\}$ is normally distributed with mean $\mu_{d|a}$ and variance $\sigma_{d|a}^2$ given by

$$\mu_{d|a} = \mu_d + \frac{l_d\sigma_d^2 + l_{\neg d}\rho\sigma_\eta\sigma_\gamma}{l_\eta^2\sigma_\eta^2 + l_\gamma^2\sigma_\gamma^2 + 2l_\eta l_\gamma\rho\sigma_\gamma\sigma_\eta}(a - k - l_\eta\mu_\eta - l_\gamma\mu_\gamma),$$

$$\sigma_{d|a}^2 = \frac{l_d^2(1 - \rho^2)\sigma_\eta^2\sigma_\gamma^2}{l_\eta^2\sigma_\eta^2 + l_\gamma^2\sigma_\gamma^2 + 2l_\eta l_\gamma\rho\sigma_\gamma\sigma_\eta}.$$

Plainly, $\mu_{d|a}$ is linear in a with slope coefficient as stated in the lemma (with obvious notational adjustment).

A-23

Now consider the agent's best response to any linear market belief, $\hat{\tau}(a) = La + K$. For any η, γ, the agent solves

$$\max_{a \in \mathbb{R}} s\gamma(La + K) - \frac{(a-\eta)^2}{2}.$$

The maximand is concave in a; the first order condition yields the solution $a = \eta + sL\gamma$.   *Q.E.D.*

F.3. Proof of Proposition 6

Part 1 of Proposition 6 is subsumed in Lemma 10 below; Part 2 and Part 3 of Proposition 6 are immediate consequences of expression (27) and Lemma 11 below; and Part 4 of Proposition 6 is Lemma 12 below. Given $\tau = \eta$, it will be convenient to rewrite Equation 24 as

$$f_\eta(L, s, \sigma_\eta, \sigma_\gamma, \rho) := s^2 \sigma_\gamma^2 L^3 + 2s\rho\sigma_\gamma\sigma_\eta L^2 + (\sigma_\eta^2 - s\rho\sigma_\eta\sigma_\gamma)L - \sigma_\eta^2 = 0. \tag{25}$$

Lemma 10. *If $\rho \geq 0$ then there is unique solution to Equation 25 on the non-negative domain; this solution satisfies $L \in (0,1)$ and $\frac{\partial f_\eta}{\partial L} > 0$ at the solution.*

Proof. Assume $\rho \geq 0$. Differentiation shows that $f_\eta(\cdot)$ is strictly convex in L on the domain $L > 0$. The result follows from the observation that $f_\eta(0, \cdot) < 0 < f_\eta(1, \cdot)$.   *Q.E.D.*

Lemma 11. *Let $\rho \geq 0$. As $s \to \infty$, it holds for the solution L to Equation 25 that $L \to 0$ and $L^2 s \to \rho\frac{\sigma_\eta}{\sigma_\gamma}$. As $s \to 0$, it holds for the solution L to Equation 25 that $L \to 1$.*

Proof. The result for $s \to 0$ is immediate from Equation 25, so we only prove the limits as $s \to \infty$. Assume $\rho \geq 0$. Divide Equation 25 by s^2 to get

$$\sigma_\gamma^2 L^3 - \frac{\sigma_\eta^2(1-L)}{s^2} + \rho\sigma_\eta\sigma_\gamma\frac{L}{s}(2L-1) = 0.$$

Let $s \to \infty$. Since $L \in (0,1)$, all terms on the LHS above except the first one vanish as $s \to \infty$, so it must be that $L \to 0$. Next rewrite Equation 25 by dividing by Ls as

$$s\sigma_\gamma^2 L^2 + 2L\rho\sigma_\gamma\sigma_\eta - \rho\sigma_\eta\sigma_\gamma - \frac{\sigma_\eta^2}{Ls}(1-L) = 0. \tag{26}$$

Suppose, to contradiction, that Ls converges to something finite (which must be non-negative). Then $L^2 s \to 0$, and hence the first two terms on the LHS of (26) above vanish, which means the LHS of (26) goes to something strictly negative, a contradiction. Hence, $Ls \to \infty$. This implies the LHS of (26) goes to $s\sigma_\gamma^2 L^2 - \rho\sigma_\eta\sigma_\gamma$, which implies $L^2 s \to \rho\frac{\sigma_\eta}{\sigma_\gamma}$.   *Q.E.D.*

Lemma 12. *If $\rho \geq 0$, then (i) $\frac{d}{d\mu_\gamma}\text{Var}(\hat{\eta}) = 0$, (ii) $\frac{d}{d\sigma_\gamma}\text{Var}(\hat{\eta}) < 0$, (iii) $\frac{d}{ds}\text{Var}(\hat{\eta}) < 0$, and (iv) $\frac{d}{d\rho}\text{Var}(\hat{\eta}) > 0$.*

Proof. Since $\hat{\eta}(a) = La + K$, we compute

$$\text{Var}(\hat{\eta}) = \text{Var}(L(\eta + sL\gamma) + K) = L\left(s^2 L^3 \sigma_\gamma^2 + L\sigma_\eta^2 + 2sL^2\rho\sigma_\eta\sigma_\gamma\right)$$
$$= L^2 s\rho\sigma_\eta\sigma_\gamma + L\sigma_\eta^2, \tag{27}$$

A-24

where L is the solution to Equation 25 and the third equality uses Equation 25.

For the first part, note from (27) that $\text{Var}(\hat{\eta})$ depends neither directly nor indirectly (through the solution L to Equation 25) on μ_γ. For the other parts, we use the chain rule and the implicit function theorem (which Lemma 10 ensures we can apply), to derive that for any $x = \sigma_\gamma, s, \rho$,

$$\frac{\mathrm{d}}{\mathrm{dx}}\text{Var}(\hat{\eta}) = \frac{\partial\text{Var}(\hat{\eta})}{\partial L}\frac{\partial L}{\partial x} + \frac{\partial\text{Var}(\hat{\eta})}{\partial x} = \frac{1}{\frac{\partial f_\eta}{\partial L}}\left(\frac{\partial\text{Var}(\hat{\eta})}{\partial x}\frac{\partial f_\eta}{\partial L} - \frac{\partial\text{Var}(\hat{\eta})}{\partial L}\frac{\partial f_\eta}{\partial x}\right). \tag{28}$$

Each partial derivative in (28) can be computed from either (25) or (27), and some manipulations then yield parts (ii), (iii), and (iv) of the result. In more detail:

$$\frac{\partial\text{Var}(\hat{\eta})}{\partial L} = 2Ls\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2, \frac{\partial\text{Var}(\hat{\eta})}{\partial\sigma_\gamma} = L^2s\rho\sigma_\eta, \frac{\partial\text{Var}(\hat{\eta})}{\partial s} = L^2\rho\sigma_\eta\sigma_\gamma, \frac{\partial\text{Var}(\hat{\eta})}{\partial\rho} = L^2s\sigma_\eta\sigma_\gamma,$$

and

$$\frac{\partial f_\eta}{\partial L} = \sigma_\eta^2 + 3L^2s^2\sigma_\gamma^2 + (4L-1)\rho s\sigma_\gamma\sigma_\eta,$$

$$\frac{\partial f_\eta}{\partial\sigma_\gamma} = Ls\left(2L^2s\sigma_\gamma + (2L-1)\rho\sigma_\eta\right),$$

$$\frac{\partial f_\eta}{\partial s} = L\sigma_\gamma\left(2L^2s\sigma_\gamma + (2L-1)\rho\sigma_\eta\right),$$

$$\frac{\partial f_\eta}{\partial\rho} = L(2L-1)s\sigma_\gamma\sigma_\eta.$$

Since $\frac{\partial f_\eta}{\partial L} > 0$ at the solution L (Lemma 10), we drop the $\frac{1}{\frac{\partial f_\eta}{\partial L}}$ term in (28) and compute

$$\frac{\mathrm{d}}{\mathrm{d}\sigma_\gamma}\text{Var}(\hat{\eta}) \propto Ls\sigma_\eta\left(-L^3\rho s^2\sigma_\gamma^2 + (1-L)\rho\sigma_\eta^2 + Ls\sigma_\gamma\sigma_\eta\left(\rho^2 - 2L\right)\right)$$

$$= Ls\sigma_\eta\left(\rho\left(2L^2\rho s\sigma_\gamma\sigma_\eta + L\left(\sigma_\eta^2 - \rho s\sigma_\gamma\sigma_\eta\right) - \sigma_\eta^2\right) - (L-1)\rho\sigma_\eta^2 + Ls\sigma_\gamma\sigma_\eta\left(\rho^2 - 2L\right)\right)$$

$$= -2L^3\left(1 - \rho^2\right)s^2\sigma_\gamma\sigma_\eta^2 < 0,$$

where the first line obtains from plugging in the formulae for partial derivatives into (28) and some algebraic manipulation, the second line obtains from substituting in for $-L^3\rho s^2\sigma_\gamma^2$ from Equation 25, and third line is algebraic simplification.

Analogous steps prove $\frac{\mathrm{d}\text{Var}(\hat{\eta})}{\mathrm{ds}} < 0$. Finally, plugging in partial derivatives into (28) and simplifying,

$$\frac{\mathrm{d}\text{Var}(\hat{\eta})}{\mathrm{d}\rho} \propto Ls\sigma_\gamma\sigma_\eta\left(3L^3s^2\sigma_\gamma^2 + (1-L)\sigma_\eta^2 + L\rho s\sigma_\gamma\sigma_\eta\right) > 0,$$

where the inequality uses the fact that $L < 1$, as was established in Lemma 10. Q.E.D.

F.4. Proof of Proposition 7

Part 1 of Proposition 7 is subsumed in Lemma 13 below; Part 2 and Part 3 of Proposition 7 are immediate consequences of expression (30) and Lemma 11 below; and Part 4 of Proposition 7 is Lemma 15 below. Given

A-25

$\tau = \gamma$, it will be convenient to rewrite Equation 24 as

$$f_\gamma(L, s, \sigma_\gamma, \sigma_\eta, \rho) := s^2\sigma_\gamma^2 L^3 + 2s\rho\sigma_\gamma\sigma_\eta L^2 + (\sigma_\eta^2 - s\sigma_\gamma^2)L - \rho\sigma_\eta\sigma_\gamma = 0. \tag{29}$$

Lemma 13. *If $\rho > 0$ there is a unique solution to Equation 29 on the non-negative domain; it is strictly positive. If $\rho = 0$, (i) $L = 0$ is always a solution, and (ii) there is a strictly positive solution if and only if $\sigma_\eta^2 < s\sigma_\gamma^2$; when a positive solution exists, it is unique. For any $\rho \geq 0$ and at any solution $L > 0$ to Equation 29, it holds that $\frac{\partial f_\gamma}{\partial L} > 0$.*

Proof. Assume $\rho \geq 0$. First, there is at most one strictly positive solution to Equation 29 because $f_\gamma(\cdot)$ intersects 0 from below at any strictly positive solution:

$$\left.\frac{\partial f_\gamma}{\partial L}\right|_{f_\gamma=0} = \sigma_\eta^2 + 3L^2 s^2 \sigma_\gamma^2 + 4s(L\rho)\sigma_\gamma\sigma_\eta - s\sigma_\gamma^2 > \sigma_\eta^2 + s^2\sigma_\gamma^2 L^2 + 2Ls\rho\sigma_\gamma\sigma_\eta - s\sigma_\gamma^2$$

$$= s\sigma_\gamma^2 + \frac{\rho\sigma_\eta\sigma_\gamma}{L} - s\sigma_\gamma^2 \geq 0,$$

where the second equality uses Equation 29.

Next, observe that because $f_\gamma(0, \cdot) \leq 0$ while $f_\gamma(L, \cdot) \to \infty$ as $L \to \infty$, there is always at least one non-negative solution to Equation 29. If $\rho > 0$, then $f_\gamma(0, \cdot) < 0$, so any non-negative solution is strictly positive. If $\rho = 0$, $f_\gamma(0, \cdot) = 0$, so $L = 0$ is always a solution; that there is a strictly positive solution if and only if $\sigma_\eta^2 < s\sigma_\gamma^2$ follows from the observations that $\frac{\partial f_\gamma(0,\cdot)}{\partial L} = \sigma_\eta^2 - s\sigma_\gamma^2$ and $\frac{\partial f_\gamma}{\partial L} > 0$ for all $L > 0$ if $\sigma_\eta^2 \geq s\sigma_\gamma^2$.   Q.E.D.

Lemma 14. *Let $\rho \geq 0$. As $s \to \infty$, it holds for the unique strictly positive solution L to Equation 29 that $L \to 0$ and $L^2 s \to 1$. If $\rho > 0$ then as $s \to 0$, it holds for the unique strictly positive solution L to Equation 29 that $L \to \rho\sigma_\gamma$.*[45]

Proof. The proof is omitted as the argument is analogous to that in the proof of Lemma 11, but applied to Equation 29.   Q.E.D.

Lemma 15. *If $\rho \geq 0$, then in an increasing equilibrium (when it exists, i.e. when there is a strictly positive solution to Equation 29): (i) $\frac{d}{d\mu_\eta}\mathrm{Var}(\hat{\gamma}) = 0$, (ii) $\frac{d}{d\sigma_\eta}\mathrm{Var}(\hat{\gamma}) < 0$, (iii) $\frac{d}{ds}\mathrm{Var}(\hat{\gamma}) > 0$, and (iv) $\frac{d}{d\rho}\mathrm{Var}(\hat{\gamma}) > 0$.*

Proof. Since $\hat{\gamma}(a) = La + K$, we compute

$$\mathrm{Var}(\hat{\gamma}) = \mathrm{Var}(L(\eta + sL\gamma) + K) = s^2 L^4 \sigma_\gamma^2 + L^2 \sigma_\eta^2 + 2L^3 s\rho\sigma_\eta\sigma_\gamma$$

$$= L^2 s\sigma_\gamma^2 + L\rho\sigma_\eta\sigma_\gamma, \tag{30}$$

where L is the solution to Equation 29 and the third equality uses Equation 29.

The rest of the argument is analogous to that in the proof of Lemma 12, but applied to expressions (30) and (29). The algebraic details are available on request.   Q.E.D.

---

[45]When $\rho = 0$, Lemma 13 assures that there is a strictly positive solution if and only if s is large enough.

F.5. Mixed dimensions of interest

Here we provide formal results for the LQN specification with mixed dimensions of interest discussed in Subsection 5.3. As discussed in the text, a linear equilibrium under mixed dimensions of interest is characterized by a solution (L_η, L_γ) to the following simultaneous equations:

$$L_\eta = \frac{\sigma_\eta^2 + sL\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2L^2\sigma_\gamma^2 + 2sL\rho\sigma_\gamma\sigma_\eta}, \tag{31}$$

$$L_\gamma = \frac{sL\sigma_\gamma^2 + \rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2L^2\sigma_\gamma^2 + 2sL\rho\sigma_\gamma\sigma_\eta}, \tag{32}$$

where we use the shorthand $L := gL_\gamma + (1-g)L_\eta$. An equilibrium is said to be increasing if both $L_\eta > 0$ and $L_\gamma > 0$. We compute

$$\mathrm{Var}(\hat\eta) = \mathrm{Var}(L_\eta(\eta + sL\gamma)) = L_\eta\Big(s^2L_\eta L^2\sigma_\gamma^2 + L_\eta\sigma_\eta^2 + 2sL_\eta L\rho\sigma_\eta\sigma_\gamma\Big) = L_\eta Ls\rho\sigma_\eta\sigma_\gamma + L_\eta\sigma_\eta^2, \tag{33}$$

where the third equality uses Equation 31. Similarly,

$$\mathrm{Var}(\hat\gamma) = \mathrm{Var}(L_\gamma(\eta + sL\gamma)) = L_\gamma\Big(s^2L_\gamma L^2\sigma_\gamma^2 + L_\gamma\sigma_\eta^2 + 2sL_\gamma L\rho\sigma_\eta\sigma_\gamma\Big) = L_\gamma\rho\sigma_\eta\sigma_\gamma + L_\gamma sL\sigma_\gamma^2, \tag{34}$$

where the third equality uses Equation 32.

Proposition 12. *Consider the LQN specification with mixed dimensions of interest and $\rho = 0$, and restrict attention to increasing equilibria. Then, for any $g \in (0,1)$, an equilibrium exists; moreover:*

1. *$\forall \varepsilon > 0$, $\exists \tilde{s} > 0$ such that if $s < \tilde{s}$ then in any equilibrium, $\mathrm{Var}(\hat\eta) \in [\sigma_\eta^2 - \varepsilon, \sigma_\eta]$ and $\mathrm{Var}(\hat\gamma) \in [0, \varepsilon]$.*

2. *$\forall \varepsilon > 0$, $\exists \hat{s} > 0$ such that if $s > \hat{s}$ then in any equilibrium, $\mathrm{Var}(\hat\eta) \in [0, \varepsilon]$ and $\mathrm{Var}(\hat\gamma) \in [\sigma_\gamma^2 - \varepsilon, \sigma_\gamma^2]$.*

3. *Any equilibrium has the following local comparative statics: $\frac{\mathrm{d}}{\mathrm{d}s}\mathrm{Var}(\hat\eta) < 0$ and $\frac{\mathrm{d}}{\mathrm{d}s}\mathrm{Var}(\hat\gamma) > 0$.*

Remark 4. Subject to technical qualifiers, the above result can be generalized to $\rho \geq 0$; the only modifications are that the limiting variance of $\hat\eta$ as $s \to \infty$ becomes $\rho^2\sigma_\eta^2$ (cf. Proposition 6 Part 3) while the limiting variance of $\hat\gamma$ as $s \to 0$ becomes $\rho^2\sigma_\gamma^2$ (cf. Proposition 7 Part 2). ‖

Proof of Proposition 12. Assume $\rho = 0$ and $g \in (0,1)$. Equation 31 and Equation 32 can be rewritten as

$$f_{g,\eta}(L_\eta, L_\gamma, s, \sigma_\eta, \sigma_\gamma) := L_\eta\sigma_\eta^2 + s^2L_\eta L^2\sigma_\gamma^2 - \sigma_\eta^2 = 0, \tag{35}$$

$$f_{g,\gamma}(L_\eta, L_\gamma, s, \sigma_\eta, \sigma_\gamma) := L_\gamma\sigma_\eta^2 + s^2L_\gamma L^2\sigma_\gamma^2 - sL\sigma_\gamma^2 = 0, \tag{36}$$

and (33) and (34) simplify to

$$\mathrm{Var}(\hat\eta) = L_\eta\sigma_\eta^2, \tag{37}$$

$$\mathrm{Var}(\hat\gamma) = sL_\gamma L\sigma_\gamma^2. \tag{38}$$

It is straightforward from Equation 35 and Equation 36 that there is a positive solution (i.e. both $L_\eta > 0$ and $L_\gamma > 0$);[46] moreover, in any solution, $L_\eta \in (0,1)$.

---

[46]Note that this uses $g < 1$ (cf. Proposition 7).

Manipulating (35) and (36) along similar lines to the proof of Lemma 11, it can be established that:

1. $\forall \varepsilon > 0$, $\exists \tilde{s} > 0$ such that if $s < \tilde{s}$ then in any positive solution, $L_\eta \in (1 - \varepsilon, 1)$ and $L_\gamma < \varepsilon$.

2. $\forall \varepsilon > 0$, $\exists \hat{s} > 0$ such that if $s > \hat{s}$ then in any positive solution, $L_\eta < \varepsilon$, $L_\gamma < \varepsilon$, and $sL_\gamma L \in (1 - \varepsilon, 1)$.

Part 1 and part 2 of the proposition follow applying these two facts to Equation 37 and Equation 38. To prove part 3, first note that by the implicit function theorem,

$$
\frac{\partial L_\eta}{\partial s} = -\frac{\det \begin{bmatrix} \frac{\partial f_{g,\eta}}{\partial s} & \frac{\partial f_{g,\eta}}{\partial L_\gamma} \\ \frac{\partial f_{g,\gamma}}{\partial s} & \frac{\partial f_{g,\gamma}}{\partial L_\gamma} \end{bmatrix}}{\det \begin{bmatrix} \frac{\partial f_{g,\eta}}{\partial L_\eta} & \frac{\partial f_{g,\eta}}{\partial L_\gamma} \\ \frac{\partial f_{g,\gamma}}{\partial L_\eta} & \frac{\partial f_{g,\gamma}}{\partial L_\gamma} \end{bmatrix}} \quad \text{and} \quad \frac{\partial L_\gamma}{\partial s} = -\frac{\det \begin{bmatrix} \frac{\partial f_{g,\eta}}{\partial L_\eta} & \frac{\partial f_{g,\eta}}{\partial s} \\ \frac{\partial f_{g,\gamma}}{\partial L_\eta} & \frac{\partial f_{g,\gamma}}{\partial s} \end{bmatrix}}{\det \begin{bmatrix} \frac{\partial f_{g,\eta}}{\partial L_\eta} & \frac{\partial f_{g,\eta}}{\partial L_\gamma} \\ \frac{\partial f_{g,\gamma}}{\partial L_\eta} & \frac{\partial f_{g,\gamma}}{\partial L_\gamma} \end{bmatrix}}.
$$

Computing all the relevant partial derivatives of $f_{g,\eta}$ and $f_{g,\gamma}$ and simplifying yields

$$
\frac{\partial L_\eta}{\partial s} = -\frac{2s\sigma_\gamma^2 L_\eta L^2}{s\sigma_\gamma^2 (3sL^2 - g) + \sigma_\eta^2} \quad \text{and} \quad \frac{\partial L_\gamma}{\partial s} = \frac{\sigma_\gamma^2 L (2sL_\gamma L - 1)}{s\sigma_\gamma^2 (3sL^2 - g) + \sigma_\eta^2}. \tag{39}
$$

Using (38) and (39), some algebra yields

$$
\frac{d}{ds}\text{Var}(\hat{\gamma}) = \sigma_\gamma^2 \left[L_\gamma L + s\frac{\partial L_\gamma}{\partial s} L + s\frac{\partial L}{\partial s} L_\gamma \right]
$$
$$
= \frac{\sigma_\gamma^2 L \left[-s\sigma_\gamma^2 L_\gamma (2sL^2 - g) + s\sigma_\gamma^2 (-gL_\eta + sL_\gamma L^2 + L_\eta) + \sigma_\eta^2 L_\gamma \right]}{s\sigma_\gamma^2 (3sL^2 - g) + \sigma_\eta^2}. \tag{40}
$$

We manipulate (40) as follows:

$$
\text{Numerator of (40)} = \sigma_\gamma^2 L \left[-L^2 s^2 \sigma_\gamma^2 L_\gamma + \sigma_\eta^2 L_\gamma + Ls\sigma_\gamma^2 \right] \quad \text{using } L = gL_\gamma + (1 - g)L_\eta
$$
$$
= 2\sigma_\eta^2 L_\gamma \sigma_\gamma^2 L \quad \text{as Equation 36 implies } -s^2 L_\gamma L^2 \sigma_\gamma^2 + sL\sigma_\gamma^2 = L_\gamma \sigma_\eta^2
$$
$$
> 0.
$$
$$
\text{Denominator of (40)} = s\sigma_\gamma^2 \left(sL^2 - \frac{L}{L_\gamma} \right) + \sigma_\eta^2 + s\sigma_\gamma^2 \left(\frac{L}{L_\gamma} - g + 2sL^2 \right)
$$
$$
= s\sigma_\gamma^2 \left(\frac{L}{L_\gamma} - g + 2sL^2 \right) \quad \text{as Equation 36 implies } s\sigma_\gamma^2 \left(sL^2 - \frac{L}{L_\gamma} \right) + \sigma_\eta^2 = 0
$$
$$
= \frac{s\sigma_\gamma^2}{L_\gamma} \left((1 - g)L_\eta + 2L_\gamma sL^2 \right) > 0.
$$

Consequently, $\frac{d}{ds}\text{Var}(\hat{\gamma}) > 0$. Finally, observe from (39) that

$$
\frac{d}{ds}\text{Var}(\hat{\eta}) \propto \frac{d}{ds}L_\eta \propto -s\sigma_\gamma^2 (3sL^2 - g) - \sigma_\eta^2 < 0. \qquad \text{Q.E.D.}
$$