Muddled Information*

Alex Frankel[†]   Navin Kartik[‡]

March 26, 2017

Abstract

We study a model of signaling in which agents are heterogeneous on two dimensions. An agent's natural action is the action taken in the absence of signaling concerns. Her gaming ability parameterizes the cost of increasing the action. Equilibrium behavior muddles information across the 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. We explore a new externality: showing agents' actions to additional observers can worsen information for existing observers. Applications to credit scoring, school testing, and web search are discussed.

JEL Classification: D72; D82; D83

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

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

1. Introduction

In many signaling environments, there is a concern that agents' gaming can lead to "muddled" information. Google tries to prevent 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. 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." (New York Times, 2014)

In canonical signaling models (e.g., Spence, 1973), standard assumptions such as the Spence-Mirrlees single-crossing condition ensure the existence of separating equilibria: equilibria that fully reveal agents' private information. So the only welfare cost from gaming—i.e., strategic behavior—is through an increase in costly effort. Even though gaming may induce an inefficient rat race, it does not lead to a reduction in market information.

This paper studies how gaming can worsen market information. We develop a 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 (synonymous with the signal sent) in the absence of signaling concerns. 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 one of the two dimensions, which we refer to as the *dimension of interest*.[1] Situations abound in which the dimension of interest is the natural action. For example, people with higher natural credit scores default less often on loans; the credit market does not care about gaming ability because this trait merely reflects one's knowledge about how to manipulate credit scores.

---

[1] More generally, beliefs on both dimensions could be relevant; we study such a case in Subsection 4.3.

Similarly for search engine optimization, where higher natural actions correspond to more relevant web pages. Yet, there are contexts in which the dimension of interest is the gaming ability. In the testing environment, gaming ability would not be of interest to the market if it solely represents "studying to the test", but colleges or employers might value gaming ability if it correlates with the ability to study more broadly. Or, in a 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 for schooling that is unrelated to job performance.

We explore how the combination of heterogeneous gaming ability and natural actions interact in determining the market's information. 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).[2] In both cases, full separation would be possible.

With two dimensions of heterogeneity, the market is typically faced with *muddled information*. Even though the market would like to evaluate an agent on her natural action (or gaming ability), the information revealed about this dimension of an agent's type is muddled with irrelevant information about her gaming ability (or natural action). 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.

A key contribution of this paper is to identify a relationship between the signaling costs of *cross types*—pairs in which one has higher natural action but lower gaming ability than the other. Our central assumption, formalized in Assumption 1 (part 4), is that at low levels of signaling, differences in marginal cost are driven by differences in natural actions; at higher levels of signaling, they depend more on differences in gaming ability. In other words, as more gaming occurs, gaming ability becomes relatively more important in determining

---

[2] 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).

signaling costs.

The core of our analysis concerns comparative statics on equilibrium information on the dimension of interest. We establish that when agents' incentives to take high actions increase—because the stakes in signaling go up, for instance, or the costs of signaling go down—the muddled information reveals more about an agent's gaming ability and less about her natural action. Hence, as a search engine like Google becomes more popular and the stakes for web sites to game its algorithm grow, Google searches can become less informative—even after Google adjusts its algorithm to account for this extra gaming. Notwithstanding, our analysis clarifies that while higher stakes lead to less information on one dimension, they generate more information on the other.

Section 3 establishes these comparative statics globally in a canonical "two-by-two" setting and provides general results for small and large signaling stakes. Section 4 develops a linear-quadratic-elliptical specification: signaling benefits are linear in the market belief; costs are quadratic; and the types are jointly elliptically distributed. This specification affords a sharp equilibrium characterization and additional comparative-statics results.

In Subsection 5.1 we consider 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 evaluators often try to obscure the details of their evaluation metrics, such as the College Board keeping past SAT questions secret for many years: it improves the informativeness of its test.[3]

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. 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 disproportionately helps those with low intrinsic gaming ability (i.e., poor families). By reducing heterogeneity on gaming ability, it should improve market information about natural actions.

---

[3] There are, of course, alternative mechanisms by which opacity can improve welfare in other contexts.

In Subsection 5.2 we explore a novel tradeoff in making a signal available to new observers. With more observers tracking her actions, an agent's stakes in signaling grow. At higher stakes, the signal becomes less informative about the natural action. So there is a negative informational externality on those 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. The social value of information across markets can decline after the signal is made available to new markets.

Muddled information—information loss on the dimension of interest owing to other dimensions of private information—is not a new phenomenon in signaling environments.[4] See, among others, Austen-Smith and Fryer (2005), Bénabou and Tirole (2006), Esteban and Ray (2006) and Bagwell (2007) in the economics literature, and Dye and Sridhar (2008) and Beyer et al. (2014) in the earnings-management accounting literature. As already mentioned, our main contribution is developing the comparative statics of market information when it is muddled and uncovering the general forces underlying these comparative statics.

The closest antecedent in this respect is the innovative work of Fischer and Verrecchia (2000). They study a linear-quadratic-normal model (also seen in Bénabou and Tirole (2006)) that is related to our linear-quadratic-elliptical specification in Section 4 with the dimension of interest being the natural action. Their motivation is a manager's report of firm earnings given private information on both true earnings (analogous to our natural action) and her own objectives (analogous to our gaming ability). Among other things, Fischer and Verrechia show how "price efficiency"—the information on true earnings contained in reported earnings—changes with parameters. For reasons explained in Section 4, we use elliptical distributions with bounded support rather than normal distributions. The linear-quadratic-normal and -elliptical specifications are appealing in their tractability. In particular, they yield a scalar measure of information that one can combine with explicit equilibrium computation to deduce comparative statics. However, owing to the limitations of functional form assumptions and algebraic calculations, we believe that a proper understanding of the underlying forces requires a more general analysis based on more fundamental

---

[4] 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) makes a related point in a model with a single-dimensional type.

assumptions. We aspire for Sections 2 and 3 of the current paper to elucidate these more general forces.

We should note that there are 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 at the edges of the type space (Cho and Sobel, 1990). However, this does not seem relevant for applications such as school testing or credit scores where few people have perfect scores.[5] Indeed, if bunching at the edges is ever a problem, it may be possible to simply expand the signal space: a test can be made more difficult. On the other hand, there are critiques of the focus on separating equilibria even when these exist (Mailath et al., 1993); recently, Daley and Green (2014) 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

We study a reduced-form signaling game. An agent takes an observable action; we will sometimes refer to "agents" for expositional convenience. The agent has two-dimensional private information—her *type*—that determines her cost of taking a single-dimensional action. The agent chooses an action and then receives a benefit that depends on an observer's belief about her type.

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.[6] The second dimension, γ, which we call *gaming ability*, parameterizes the agent's cost of increasing her

---

[5] In 2014, less than 0.1% of students taking the SAT got a perfect 2400; the 99[th] percentile score was 2250 (College Board, 2014b). Also in 2014, only about 1% of the U.S. population had a perfect FICO credit score of 850; less than 20% of people had a score between 800 and 850 (Wall Street Journal, 2015).

[6] We will abuse notation by using the same symbols to denote both dimensions and realizations.

action above the natural level: a higher γ will represent lower cost. (It will be helpful to remember the mnemonics θ for *type*, η for *natural*, and γ for *gaming*.) 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 in the usual manner, 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 $\underline{\gamma} < \overline{\gamma}$, $C_a(\cdot, \overline{\eta}, \underline{\gamma})/C_a(\cdot, \underline{\eta}, \overline{\gamma})$ 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 incurs an increasing and convex cost to take actions above this level. 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) among ordered types: if $\underline{a} < \overline{a}$ and $\underline{\theta} < \overline{\theta}$ in the component-wise order, then $C(\overline{a}, \underline{\theta}) - C(\underline{a}, \underline{\theta}) \geq C(\overline{a}, \overline{\theta}) - C(\underline{a}, \overline{\theta})$, with a strict inequality so long as $\overline{a}$ is strictly larger than $\underline{\theta}$'s natural action.

The fourth part of Assumption 1 places structure on how $C(\cdot)$ behaves for pairs of *cross types*, where one type, $(\overline{\eta}, \underline{\gamma})$, has a strictly higher natural action but a strictly lower gaming ability than the other, $(\underline{\eta}, \overline{\gamma})$. At low actions, the type with the higher η (and lower γ) has a lower marginal cost of increasing its action. But this type's marginal cost grows faster than the other type's. 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* for the given pair.

6

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 actions above a^{ce}.

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

(All proofs are in the Supplementary Appendices unless otherwise noted.)

Figure 1 summarizes the implications of Assumption 1 when Θ consists of four types: a low type, $(\underline{\eta}, \underline{\gamma})$; two intermediate cross types, $(\underline{\eta}, \overline{\gamma})$ and $(\overline{\eta}, \underline{\gamma})$; and a high type, $(\overline{\eta}, \overline{\gamma})$. Subsection 2.3 elaborates on the economics of the Assumption.
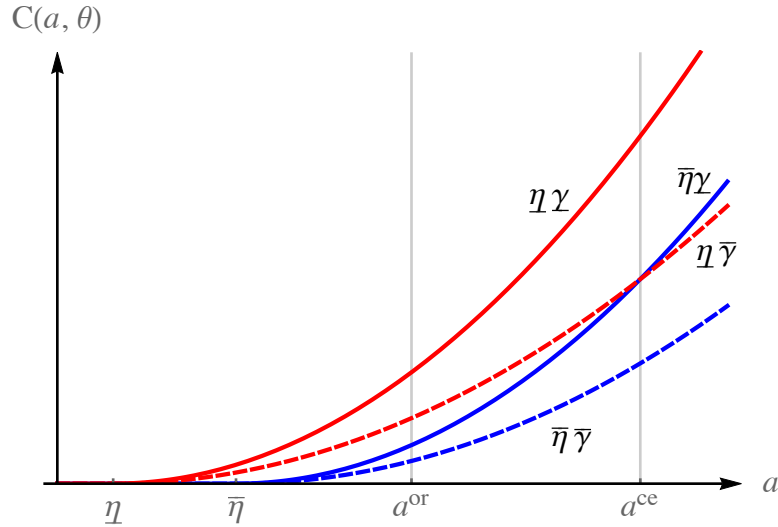

Figure 1 – Cost curves when $\Theta = \{\underline{\eta}, \overline{\eta}\} \times \{\underline{\gamma}, \overline{\gamma}\}$ with $\underline{\eta} < \overline{\eta}$ and $\underline{\gamma} < \overline{\gamma}$. The solid red curve is $C(\cdot, \underline{\eta}, \underline{\gamma})$ and the solid blue curve is $C(\cdot, \overline{\eta}, \underline{\gamma})$; the dashed red curve is $C(\cdot, \underline{\eta}, \overline{\gamma})$ and the dashed blue curve is $C(\cdot, \overline{\eta}, \overline{\gamma})$. The $C(\cdot, \overline{\eta}, \underline{\gamma})$ and $C(\cdot, \underline{\eta}, \overline{\gamma})$ curves have equal slopes at a^{or}.

Example 1. A canonical functional form is $C(a, \eta, \gamma) = c(a, \eta)/\gamma$. In this case the first three parts of Assumption 1 reduce to requiring the analogous properties on $c(a, \eta)$, with the second requirement of part 3 automatically ensured. Since $\frac{C_a(a, \overline{\eta}, \underline{\gamma})}{C_a(a, \underline{\eta}, \overline{\gamma})} = \frac{c_a(a, \overline{\eta})}{c_a(a, \underline{\eta})} \frac{\overline{\gamma}}{\underline{\gamma}}$, a sufficient condition for part 4 is that for any $\underline{\eta} < \overline{\eta}$, $\frac{c_a(\cdot, \overline{\eta})}{c_a(\cdot, \underline{\eta})}$ is strictly increasing on the relevant domain with $\lim\limits_{a \to \infty} \frac{c_a(a, \overline{\eta})}{c_a(a, \underline{\eta})} = 1$. In particular, given any exponent $r > 1$, the cost function $C(a, \eta, \gamma) = (\max\{a - \eta, 0\})^r / \gamma$ satisfies Assumption 1. This family will be our leading example.[7]   ∥

---

[7] In this family, for any pair of cross types $(\underline{\eta}, \overline{\gamma})$ and $(\overline{\eta}, \underline{\gamma})$, a^{or} and a^{ce} can be computed as $a^{\text{or}} = \frac{\overline{\eta}\overline{\gamma}^k - \underline{\eta}\underline{\gamma}^k}{\overline{\gamma}^k - \underline{\gamma}^k}$

2.2. Beliefs, payoffs, and equilibrium

There is one *dimension of interest* about the agent's type, $\tau \in \{\eta, \gamma\}$. 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 (measurable) set X. 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. Gross of costs, the value or 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*. 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: 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, and higher beliefs are more valuable when stakes are higher. The benefit of inducing any higher belief grows unboundedly as stakes grow unboundedly, and analogously as stakes vanish. An example that we will refer to is $V(\hat{\tau}; s) = sv(\hat{\tau})$ for some strictly increasing $v(\cdot)$. Note that higher stakes do *not* represent greater direct benefits from taking higher actions; rather, they capture 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 (weak) Perfect Bayesian equilibria—simply equilibria, hereafter—of this signaling game: every type of the agent chooses its action optimally given the market belief function $\hat{\tau}(a)$, and the market belief is derived from Bayes Rule on the equilibrium path (with no restrictions off path). Given that the agent cares about the market belief on only one dimension of her type, equilibria cannot generally fully

---

and $a^{\mathrm{ce}} = \frac{\overline{\eta}\overline{\gamma}^l - \underline{\eta}\underline{\gamma}^l}{\overline{\gamma}^l - \underline{\gamma}^l}$, where $k = 1/(r-1)$ and $l = 1/r$.

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, then*

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

The first statement of the lemma is straightforward. Part (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, we restrict attention to equilibria that satisfy it. Part (ii) assures that there would be no loss in also requiring weak monotonicity of beliefs off the equilibrium path.

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

Remark 2. If the agent has private information only on the dimension of interest, with the component of her type on the other dimension known to the market, then there is 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.

2.3. Discussion of the model

Assumptions. Two of our assumptions warrant additional discussion. The first is free downward deviations (part 1 of Assumption 1): an agent can costlessly take any action below her natural action. As noted above, free downward deviations ensures that equilibrium beliefs are monotonic in actions and that a pooling equilibrium always exists. These two properties are common features of signaling games. The fact that free downward deviations

guarantees the properties simplifies our analysis. In making the assumption, though, we are primarily motivated by applications. It is much easier to make web pages appear to be of lower than higher quality; it is obvious how to wreck one's credit score but not how to raise it; and it is virtually costless to get questions wrong on a test, whereas getting more questions right is difficult.

Notwithstanding, in some settings there may be a direct cost of deviating downwards. An accountant manipulating financial reports, as in Fischer and Verrecchia (2000), cannot easily make them look worse than they truly are; lowering one's credit score by failing to pay a bill on time may incur monetary costs; or, as in Kartik (2009), agents make dislike lying regardless of the direction in which they lie.[8]

The mechanism explored in our paper does not turn on the assumption of free downward deviations. Indeed, take any equilibrium strategy profile in which each type's action is weakly above her natural action. This strategy profile would remain an equilibrium even with costly downward deviations—unplayed lower actions would now be even less attractive. Section 4 studies a specification of the model and a class of equilibria in which actions are in fact never below natural actions. Hence, the results therein would be unchanged if downward deviations were made costly.

The second assumption to highlight, our key assumption, is part 4 of Assumption 1: for any pair of cross types, the ratio of marginal costs of the high natural action type to the high gaming ability type is increasing in the action. The interpretation is that as more gaming occurs—i.e, as agents choose higher actions—cost differences become less driven by variation in natural actions, and more by variation in gaming ability. In the credit scoring example, suppose Anne has a natural credit score of 675 and low gaming ability, while Bob has a lower natural score, 600, but a higher gaming ability. If both agents aim for a credit score around 700, Anne's marginal cost of score improvement is lower than Bob's: Anne can 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 800, though, Bob's marginal cost of improvement is lower than Anne's: 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.

In many settings, over the relevant range of actions, we think our key assumption is likely to be valid. On the other hand, one can conceive of violations. For instance, no matter

---

[8] In the sender-receiver lying cost application, η represents a payoff-relevant state for the receiver, a is the sender's message, and the sender's lying cost is given by, say, $-(a-\eta)^2/(2\gamma)$. The parameter γ captures how much the sender dislikes lying. Kartik (2009) studies a related model without heterogeneity on γ.

the amount of studying (gaming), only those with sufficiently high natural ability (natural actions) might be able to attain the highest scores on an IQ test. It turns out, however, that the assumption is crucial for the comparative statics on how market information responds to changes in stakes. Indeed, we view one of our main contributions as identifying the role and importance of such an assumption. In the conclusion of the paper we discuss how our analysis helps shed light on situations in which the assumption may not be satisfied.

Applications. Our analysis considers either the natural action or the gaming ability as the possible dimension of interest.

Natural action as the dimension of interest. The main applications we are motivated by are school tests, web search, and credit scores. Heterogeneity of natural actions reflects that agents would take different actions absent signaling concerns. Students get different SAT scores prior to studying; web sites are more or less relevant for a given query; and even without a formal credit score, consumers differ in their propensity to pay bills on time. This natural action is of direct interest to the colleges admitting students, people searching the web, and banks offering loans.

There are a number of (non-exclusive) sources for heterogeneity in gaming ability in these applications. One is underlying skills: some students may simply be more facile at studying. Another is that agents could have heterogeneous understanding of how to game a signal due to differing experience or information. Some students may have access to tutors with better practice materials; professional web designers are more attuned than amateurs to search engine optimization techniques; and some consumers do not know strategies to improve one's credit score, as evidenced by the large genre of books on the subject.

Agents may also have different preferences for gaming: students vary in how much they enjoy or dislike studying. Those who enjoy it more face a lower cost of increasing their test scores. 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. In particular, the College Board worries that richer students can better afford private tutoring and test prep courses for its SAT test (College Board, 2014a).[9]

When there are unethical approaches to gaming, differences in "integrity" could affect agents' preferences towards gaming. In addition to studying for exams, students can find

---

[9] Our analysis will imply that a student's score should be interpreted differently based on the available information about her wealth—an issue that is much debated in college admissions.

ways to cheat. Web sites can engage in undesirable behaviors like "webspam" or "black hat SEO" to improve their search rankings. Colleges can and sometimes do engage in dubious activities to affect their US News & World Report (US News) rankings (e.g., New York Times, 2008).

One can also interpret an agent's gaming ability as parameterizing her *private benefits*: if $C(a, \eta, \gamma) = c(a, \eta)/\gamma$, then the net payoff function $V(\hat{\tau}; s) - c(a, \eta)/\gamma$ represents the same preferences as the payoff function $\gamma V(\hat{\tau}; s) - c(a, \eta)$. Intuitively, it is indistinguishable whether rich students have a lower cost of paying for coaching relative to their benefit from higher scores, or whether these students have a higher benefit (in dollar terms) relative to the monetary cost of such coaching. Alternatively, some web site owners are more interested in attracting hits than others. In related applications, managers value the market's evaluation of their firm's earnings differently (Fischer and Verrecchia, 2000), and individuals vary in how much they care about their social image (Bénabou and Tirole, 2006).

Gaming ability as the dimension of interest. There are contexts in which what we refer to as the gaming ability would be the dimension of interest. Indeed, Spence's (1973) framework of job-market signaling is precisely one in which the market values "gaming" ability, because the ability to "game" by completing undesirable schooling at lower cost is positively correlated with productivity. While the simplest version of Spence (1973) has a homogeneous natural action of acquiring no education, there may in fact be underlying preference variation over education that is irrelevant to employers. Similarly, pre-existing variation in SAT scores might arise from differences in socioeconomic status or high school quality while "gaming" ability correlates with a broader ability to study and learn new skills. (Colleges or employers could value a mix of both dimensions; see Subsection 4.3.)

Other applications for gaming ability as the dimension of interest emerge when gaming ability is reinterpreted as private benefits, as discussed above. For instance, better students may tend to have a stronger preference to attend better colleges, expecting to get more out of the experience. Colleges would then prefer to admit those students with higher private benefits. Esteban and Ray (2006) make a related point in the context of signaling quality for license procurement.

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 about the dimension of interest of 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 about τ*: after any on-path action, β_τ will be degenerate. A pooling equilibrium is *uninformative about τ*: 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.[10] An equilibrium e' is *more informative about τ* than an equilibrium e'' if the distribution of β_τ under e' is more informative than that under e''.

As the agent's signaling benefit depends only on the market's posterior mean on the dimension of interest, $\hat{\tau}$, we will also be interested in information specifically about $\hat{\tau}$ rather than about the entire distribution 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''.[11] An equilibrium is *uninformative about $\hat{\tau}$* if the distribution it induces over $\hat{\tau}$ is a point mass at the prior mean of τ; it is *fully informative about $\hat{\tau}$* if every on-path action reveals the agent's true mean on the dimension of interest. Note that an equilibrium can be uninformative about $\hat{\tau}$ even if the equilibrium is informative about τ. On the other hand, an equilibrium is fully informative about $\hat{\tau}$ if and only if it is fully informative about τ. In general, the partial order on equilibria generated by information about $\hat{\tau}$ is finer than that generated by information about τ: more informative about τ implies more informative about $\hat{\tau}$, but more informative about $\hat{\tau}$ does not necessarily imply more informative about τ.[12]

Comparing equilibria according to their informativeness is appealing because of the fundamental connection between this statistical notion and allocative efficiency $\mathbb{E}[V(\hat{\tau}; s)]$. If

---

[10] 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.

[11] Our notion of informativeness about $\hat{\tau}$ is the same as Ganuza and Penalva's (2010) *integral precision*. Informativeness about $\hat{\tau}$ is prior-dependent, unlike informativeness about τ.

[12] 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 about τ, and uninformative about $\hat{\tau}$ implies uninformative about τ.

the benefit function $V(\cdot; s)$ is convex, then for fixed stakes there is higher allocative efficiency when the beliefs about $\hat{\tau}$ are more informative (and hence also when the beliefs about τ are more informative). If $V(\cdot; s)$ is concave, the opposite holds: allocative efficiency is maximized by pooling 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. Appendix A provides an explicit example relating the demand curve for a product to the shape of a convex benefit function.

3. The Effect of Stakes on Muddled Information

3.1. A 2×2 setting

This section considers a 2×2 setting: $\Theta \subseteq \{\underline{\eta}, \overline{\eta}\} \times \{\underline{\gamma}, \overline{\gamma}\}$, with $\underline{\eta} < \overline{\eta}$ and $\underline{\gamma} < \overline{\gamma}$. We will be able to establish global comparative statics here on the informativeness of equilibria with respect to the stakes.

First, to develop intuition, consider a special case in which the prior's support is the two cross types, $(\underline{\eta}, \overline{\gamma})$ and $(\overline{\eta}, \underline{\gamma})$. Call $(\overline{\eta}, \underline{\gamma})$ "the natural type," as it has the higher natural action, and $(\underline{\eta}, \overline{\gamma})$ "the gamer." With only these cross types, the following observation suggests why information about the natural action decreases with stakes while information about the gaming ability increases.

Observation 1. *When $\Theta = \{(\underline{\eta}, \overline{\gamma}), (\overline{\eta}, \underline{\gamma})\}$:*

1. *If $\tau = \eta$, then there is a threshold $s_\eta^* > 0$ such that a separating equilibrium exists if and only if $s \leq s_\eta^*$; for all $s > s_\eta^*$, a partially-pooling equilibrium exists; and as $s \to \infty$, all equilibria are approximately uninformative.[13]*

2. *If $\tau = \gamma$, then there are thresholds $0 < s_\gamma^{**} < s_\gamma^*$ such that all equilibria are pooling if $s \leq s_\gamma^{**}$; and a separating equilibrium exists if and only if $s \geq s_\gamma^*$.*

---

[13] See Subsection 3.2 for a definition of approximately uninformative. It suffices in the present context that for any $\varepsilon > 0$, there is an $\overline{s} > 0$ such that when $s > \overline{s}$, in any equilibrium and for any on-path action a, $|\hat{\eta}(a) - \mathbb{E}[\eta]| < \varepsilon$.

The logic driving the observation is worth going through. First suppose the dimension of interest is the natural action ($\tau = \eta$), so that both types want to be thought of as the natural type. If stakes are low, one separating equilibrium has both types playing their natural actions, at zero cost. The natural type obviously prefers not to deviate downwards even though it could do so for free; and the gamer is unwilling to bear the cost of mimicking the natural type. On the other hand, there cannot be a separating equilibrium when stakes are high. The gamer would be willing to take any action below the cost-equalizing action, a^{ce}, to be thought of as the natural type; and the natural type cannot separate by taking an action above a^{ce}, as any such action would be less costly for the gamer.

When the dimension of interest is the gaming ability ($\tau = \gamma$), both types want to be thought of as the gamer. Separation now requires high stakes, as the gamer cannot separate by taking an action below a^{ce}. At high enough stakes, there will be an $a > a^{\mathrm{ce}}$ such that only the gamer would be willing to take a in order to be thought of as the gamer. On the other hand, at low enough stakes, the gamer would not be willing to take any action above $\bar{\eta}$; by free downward deviations, the natural type can costlessly mimic the gamer, and hence only pooling equilibria exist.

In addition to separating and pooling equilibria, there can be partially-pooling equilibria. One—but not necessarily the only—form of partial pooling is as follows. Pick any action $a_1 \in [\underline{\eta}, a^{\mathrm{or}})$, where a^{or} is the order-reversing action. There is a corresponding action $a_2 \in (a^{\mathrm{or}}, a^{\mathrm{ce}}]$ such that the gamer and the natural type bear the same incremental cost of moving from a_1 to a_2. At high enough stakes, regardless of the dimension of interest, there is a partially-pooling equilibrium in which both types mix over these two actions.[14] The two types can both be indifferent because they pay the same additional cost and receive the same additional signaling benefit when increasing their action from a_1 to a_2. If the dimension of interest is the natural action, these equilibria are the only informative ones at high enough stakes. But they become uninformative as stakes go to infinity: the belief at the lower action a_1 must converge to the belief at the higher action a_2 in order for the signaling benefit of increasing from a_1 to a_2 to remain constant (equal to the unchanging cost difference).

So, with only cross types, as summed up in Observation 1, when $\tau = \eta$ the market can

---

[14] When both types mix in this fashion, no action in (a_1, a_2) can be taken; however, there may be on-path actions above a_2 or below a_1. By Assumption 1 part 4, the gamer has a lower incremental cost than the natural type of moving to actions above a_2, and also a larger cost reduction of moving to actions below a_1. So any actions besides a_1 and a_2 can only be taken by the gamer. Since beliefs must be monotonic on path, the gamer can take an action below a_1 when $\tau = \eta$ and above a_2 when $\tau = \gamma$. When $\tau = \eta$, however, at sufficiently high stakes the gamer would no longer be willing to reveal itself by taking an action below a_1.

get full information at low stakes but approximately no information at high stakes. When $\tau = \gamma$, there is no information at low stakes but there can be full information at high stakes.

Armed with this intuition, let us turn to global comparative statics. We seek to show that information decreases in stakes when $\tau = \eta$ and increases in stakes when $\tau = \gamma$, not just for the case of two cross types but for the more general 2×2 type space. For any given stakes, there are typically multiple equilibria; these equilibria need not all be ranked by their (Blackwell) informativeness. We use the weak set order to compare equilibrium sets: equilibrium set $\mathcal{Q}$ is *more informative about* τ than equilibrium set $\mathcal{Q}'$, and $\mathcal{Q}'$ is *less informative about* τ than $\mathcal{Q}$, if (i) for any equilibrium $e \in \mathcal{Q}$ there exists $e' \in \mathcal{Q}'$ with e more informative about τ than e', and (ii) for any $e' \in \mathcal{Q}'$ there exists $e \in \mathcal{Q}$ with e more informative about τ than e'. Condition (i) is satisfied whenever the comparison is between all equilibria at different parameters, simply by taking $e' \in \mathcal{Q}'$ to be a pooling equilibrium (which always exists, as was noted in Remark 1). So it is only condition (ii) that has bite when comparing sets of equilibria across parameters: for any equilibrium in the less informative set, there is an equilibrium in the more informative set that is more informative.[15]

The following proposition is the main result of this subsection.

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

1. *When $\tau = \eta$, the set of equilibria under $\underline{s}$ is more informative about η than the set of equilibria under $\overline{s}$.*

2. *When $\tau = \gamma$, the set of equilibria under $\underline{s}$ is less informative about γ than the set of equilibria under $\overline{s}$.*

To prove Proposition 1, we first fix a dimension of interest, a level of stakes, and an arbitrary equilibrium. We then look at nearby equilibria as we perturb the stakes. We construct a path of new equilibria in which the belief distribution (continuously) becomes more informative as stakes move in the appropriate direction: lower stakes when $\tau = \eta$ and higher stakes when $\tau = \gamma$. Formally:

Lemma 3. *In the 2×2 setting, let $\mathcal{Q}(s)$ be the set of equilibria at stakes $s > 0$, and fix some equilibrium e_0 at stakes $s_0 > 0$.*

---

[15] There may exist no "most informative" equilibrium in an equilibrium set. If we were to extend the Blackwell (1951) partial ordering to a complete ordering, then our notion of equilibrium set $\mathcal{Q}$ being more informative than equilibrium set $\mathcal{Q}'$ (which contains a pooling equilibrium) would correspond to the most informative element of $\mathcal{Q}$ being more informative than that of $\mathcal{Q}'$.

1. *When $\tau = \eta$, there is a function $e_\eta(s)$ from stakes $s \in (0, s_0]$ to an equilibrium in $\mathcal{Q}(s)$ such that (i) $e_\eta(s_0) = e_0$; (ii) the distribution of β_η under $e_\eta(s)$ is continuous in s;[16] and (iii) $e_\eta(s'')$ is less informative than $e_\eta(s')$ for any $s'' > s'$.*

2. *When $\tau = \gamma$, there is a function $e_\gamma(s)$ from stakes $s \in [s_0, \infty)$ to an equilibrium in $\mathcal{Q}(s)$ such that (i) $e_\gamma(s_0) = e_0$; (ii) the distribution of β_γ under $e_\gamma(s)$ is continuous in s; and (iii) $e_\gamma(s'')$ is more informative than $e_\gamma(s')$ for any $s'' > s'$.*

Lemma 3 implies Proposition 1. The lemma's proof is involved because, even in this 2×2 setting, an equilibrium can have many different combinations of binding incentive constraints. We provide a sketch of the proof in Appendix C. The proof confirms that, starting at any such combination, a suitable perturbed equilibrium can be found as the stakes go up (when $\tau = \gamma$) or down (when $\tau = \eta$). The same basic logic applies in some form for each case: to increase information about τ as the stakes vary, we shift mixing probabilities of high-τ types from low actions with low beliefs to high actions with high beliefs, and/or shift mixing probabilities of low-τ types from high actions to low actions. The main cases begin from an equilibrium akin to the partially-pooling ones discussed in the context of Observation 1: the two cross types are both indifferent between the same pair of on-path actions $a_1 \in [\underline{\eta}, a^{\mathrm{or}})$ and $a_2 \in (a^{\mathrm{or}}, a^{\mathrm{ce}}]$. The low type, $(\underline{\eta}, \underline{\gamma})$, takes an action no larger than a_1. The high type, $(\overline{\eta}, \overline{\gamma})$, takes an action no smaller than a_2.

Although Proposition 1 is stated for the entire set of equilibria, we conjecture that its conclusion would also hold were attention restricted to equilibria satisfying stability-based refinements such as D1 or divinity (Cho and Kreps, 1987; Banks and Sobel, 1987).[17]

3.2. General type spaces

For more general type spaces, $\Theta \subset \mathbb{R} \times \mathbb{R}_{++}$, we are unable to get global comparative statics on information as stakes vary. Instead, to extend the theme that observers tend to be more informed about the natural action at low stakes and more informed about gaming ability at high stakes, we generalize Observation 1. As stakes get arbitrarily small or large, we provide results on the existence of separating equilibria, as well as conditions guaranteeing that equilibria become approximately uninformative about $\hat{\tau}$. Formally, we say that at high (resp., low) stakes, equilibria are approximately uninformative about $\hat{\tau}$ if for any sequence

---

[16] I.e., for a sequence $s \to s^*$, the corresponding distributions under s converge weakly to that under s^*.

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

of equilibria e_s at stakes $s > 0$, it holds that as $s \to \infty$ (resp., $s \to 0$), the distribution of $\hat{\tau}$ under e_s weakly converges to the uninformative distribution, a point mass at $\mathbb{E}[\tau]$.

First consider the natural action as the dimension of interest.

Proposition 2. *Assume $\tau = \eta$.*

1. *If $|\Theta_\eta| < \infty$, then at low stakes there is a fully informative equilibrium about η.*

2. *If Θ has any cross types, then at high stakes there is no fully informative equilibrium about η.*

3. *If the marginal distribution of γ is continuous and if $\mathbb{E}[\eta|\gamma]$ is non-increasing in γ, then at high stakes equilibria are approximately uninformative about $\hat{\eta}$.*

(Owing to their centrality, the proofs of Proposition 2 and Proposition 3 are in Appendix D rather than the Supplementary Appendices.)

Parts 1 and 2 of Proposition 2 are relatively straightforward given our discussion in the 2×2 setting. Regarding part 2, recall that if there were no cross types, then standard arguments based on the single-crossing property imply that a separating equilibrium would exist at any level of stakes.

Part 3 of Proposition 2 is a consequence of Lemma 4 in Appendix D, which states that for any pair of cross types, as stakes get large, the type with higher gaming ability must induce a belief not much lower, and possibly strictly higher, than any belief induced by the other type. In the limit as $s \to \infty$, any type with strictly higher gaming ability than another type induces 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). Under the hypotheses of part 3 of Proposition 2, any limiting distribution is necessarily uninformative about the posterior mean $\hat{\eta}$.

Two observations help explain the hypotheses in part 3 of Proposition 2. First, if the distribution of γ were not continuous, then a mass of types with the lowest γ and a low η (or the highest γ and a high η) might separate from other types even in the limit as $s \to \infty$, revealing information about their η. Second, even with a continuous distribution of γ, if the expectation $\mathbb{E}[\eta|\gamma]$ were strictly increasing in γ—e.g., because of positive correlation between

18

η and γ —then types with higher γ might be able to signal their higher average η by taking higher actions.

Turning to gaming ability as the dimension of interest:

Proposition 3. *Assume $\tau = \gamma$.*

1. *If $|\Theta_\gamma| < \infty$, then at high stakes there is a fully informative equilibrium about γ.*

2. *If Θ has any cross types, then at low stakes there is no fully informative equilibrium about γ.*

3. *If the marginal distribution of η is continuous and if $\mathbb{E}[\gamma|\eta]$ is non-increasing in η, then at low stakes equilibria are approximately uninformative about $\hat\gamma$.*

The logic is more or less a mirror image of that Proposition 2, with Lemma 5 in Appendix D playing an analogous role to Lemma 4.

4. A Linear-Quadratic-Elliptical Specification

This section studies a specification of our general framework that permits an explicit equilibrium characterization and additional comparative statics. We specialize to a linear benefit $V(\hat\tau; s) = s\hat\tau$ and a quadratic cost $C(a, \eta, \gamma) = (\max\{a - \eta, 0\})^2/(2\gamma)$. Given stakes $s > 0$, the agent's payoff is thus

$$s\hat\tau - \frac{(\max\{a - \eta, 0\})^2}{2\gamma}. \tag{1}$$

Furthermore, the agent's type $\theta = (\eta, \gamma)$ is drawn from an *elliptical distribution*: a distribution in which there is a constant probability density on each concentric ellipse about a mean. We refer to this specification of preferences and type distribution as the *linear-quadratic-elliptical*, or *LQE*, specification.

Formally, an (absolutely continuous) elliptical distribution $\mathcal{E}(\mu, \Sigma, g)$ over a two-dimensional realization $x = (x_1, x_2)$ is defined by $\mu = (\mu_1, \mu_2) \in \mathbb{R}^2$, $\Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$ a positive definite matrix, and $g(\cdot) : \mathbb{R}_+ \to \mathbb{R}_+$ a measurable function called the density generator. The probability density of this distribution is $f(x) = k|\Sigma|^{-1/2}g((x - \mu)\Sigma^{-1}(x - \mu)')$,

with $k = 1/(\pi \int g(t)\mathrm{d}t) \in \mathbb{R}_{++}$ a constant of integration.[18] We take θ to be drawn from $\mathcal{E}(\mu_\theta, \Sigma_\theta, g_\theta)$, where $\mu_\theta = (\mu_\eta, \mu_\gamma)$ and $\Sigma_\theta = \begin{pmatrix} \sigma_\eta^2 & \rho\sigma_\eta\sigma_\gamma \\ \rho\sigma_\eta\sigma_\gamma & \sigma_\gamma^2 \end{pmatrix}$ with $\sigma_\eta > 0$, $\sigma_\gamma > 0$, and $\rho \in (-1, 1)$. Our maintained assumption of a compact support corresponds to a requirement that g_θ has compact support. Without loss, let the support of g_θ be contained in $[0, 1]$; then for $i \in \{\eta, \gamma\}$, the support of the marginal distribution of i is $[\mu_i - \sigma_i, \mu_i + \sigma_i]$. In order to guarantee our maintained assumption that $\gamma > 0$ for all types, assume $\mu_\gamma > \sigma_\gamma$.

In an elliptical distribution with a given density generator g, the marginal distribution of component $i = 1, 2$ depends only on μ_i and σ_i. (See Gómez et al. (2003) for an accessible introduction to this and other properties of elliptical distributions.) The vector of means is μ. The covariance matrix is $\alpha\Sigma$ for some constant $\alpha > 0$ which depends only on g. The correlation coefficient between the two components is therefore $\sigma_{12}/(\sigma_1\sigma_2)$. The coefficient of determination in a linear regression of one component on the other, what is commonly referred to as the R^2, is equal to the square of this correlation coefficient: $R^2 \equiv \sigma_{12}^2/(\sigma_1^2\sigma_2^2)$.

Elliptical distributions are a generalization of joint normal distributions. Normality corresponds to the density generator $g(t) = \exp\{(-1/2)t\}$. We cannot use normal distributions (and have ruled them out by requiring g_θ to have compact support) because they would entail types with $\gamma < 0$; an agent with $\gamma < 0$ and objective (1) would obtain direct benefits rather than incurring costs from taking higher actions. A simple example of an elliptical distribution with compact support is a uniform distribution over the interior of an ellipse, which corresponds to $g(t) = \mathbb{1}_{\{t \leq 1\}}$. Elliptical distributions preserve many useful properties of joint normal distributions. Crucially, when (η, γ) is elliptically distributed and the action a is any linear function of η and γ, it holds that $\mathbb{E}[\tau|a]$ is a linear function of a. So a linear strategy in the agent's type will imply a linear market belief. (Quadratic costs and linear benefits will ensure a linear strategy is optimal given a linear market belief.)

Our analysis in this section is related to Fischer and Verrecchia (2000), Bénabou and Tirole (2006, Section II.B), and Gesche (2016). Fischer and Verrecchia (2000) and Bénabou and Tirole (2006) have previously studied related specifications to what we use; they take their type distribution to be bivariate independent normal and focus on the dimension of interest being (their analog of) $\tau = \eta$.[19] Bénabou and Tirole study equilibrium *actions*.

---

[18] We assume $g(\cdot)$ is Lebesgue integrable with $\int g(t)\mathrm{d}t \in \mathbb{R}_{++}$. The notation $|\Sigma|$ refers to the determinant of Σ and $(x - \mu)'$ refers to the transpose of $(x - \mu)$. Vectors x and μ are row vectors prior to transposition.

[19] Recall from Subsection 2.3 that γ can be reinterpreted as parameterizing private benefits rather than gaming ability. Specifically, when $\gamma > 0$, the objective (1) is equivalent to $s\gamma\hat{\tau} - (\max\{a - \eta, 0\})^2/2$. This

We emphasize comparative statics of equilibrium *informativeness*, which were studied in Fischer and Verrecchia's Corollary 3, and which play a role in Ali and Bénabou (2016). Our analysis adds broader dimensions of interest (allowing for $\tau = \gamma$ and even a mixture) as well as correlation of types. Permitting correlation is important for our applications. For example, types are correlated when students from a higher socioeconomic class can more easily pay for effective test preparation (higher γ) and also tend to be better prepared for college (higher η). We owe the idea of using elliptical distributions to Gesche (2016); his equilibrium characterization is related to ours for the case of $\tau = \eta$, but he makes somewhat different assumptions than we do and he does not focus on market information.

Consistent with the common practice in models with normal distributions, we will focus on *linear equilibria* in our LQE specification: equilibria in which an agent of type (η, γ) takes action $a = l_\eta \eta + l_\gamma \gamma + b$, for some constants l_η, l_γ, and b. In any such equilibrium, the market belief, $\hat{\tau}(a)$, is a linear function of the agent's action.[20] Moreover, (i) the vector (τ, a) is elliptically distributed with the same density generator g_θ as (η, γ), and (ii) the ex-ante distribution of posteriors β_τ about τ given a is determined entirely by $R^2_{\tau a}$, the R^2 between τ and a. Fixing g_θ and fixing the prior distribution of τ, a higher $R^2_{\tau a}$ implies an equilibrium that is more informative about the market belief $\hat{\tau}$. An $R^2_{\tau a}$ of 1 implies a fully informative equilibrium about $\hat{\tau}$ (and hence also about τ). An $R^2_{\tau a}$ of 0 is uninformative about $\hat{\tau}$.[21] See Lemma 6 and Lemma 8 in Supplementary Appendix SA.3 for details.

We begin by characterizing linear equilibria. As in our general analysis earlier, free downward deviations ensures that any equilibrium must have a market posterior $\hat{\tau}(\cdot)$ that is non-decreasing in a on the equilibrium path, and that there is a pooling equilibrium in which $\hat{\tau}(\cdot)$ is constant. Thus, a linear equilibrium is informative about $\hat{\tau}$ if and only if the

---

latter objective is still meaningful when some agents have $\gamma < 0$, under the interpretation that some agents prefer lower market beliefs to higher. If the cost function is then modified from $(\max\{a - \eta, 0\})^2 / 2$ to $(a - \eta)^2 / 2$, i.e., to let downward deviations from the natural action be symmetrically costly to upward deviations, one recovers the objective function analyzed by Fischer and Verrecchia (2000), Bénabou and Tirole (2006), and Gesche (2016). In that specification, agents with $\gamma < 0$ will take actions below their ideal point η at a positive cost in order to reduce market beliefs.

[20] Strictly speaking, the market belief is only pinned down at on-path actions. It is without loss for our purposes to stipulate a globally linear market belief.

[21] $R^2_{\tau a} = 0$ implies that τ and a are uncorrelated: $\hat{\tau}$, the posterior mean conditional on a, is constant with respect to a. However, while the equilibrium is uninformative about $\hat{\tau}$, it is still informative about τ. The support of τ depends on a, for instance. Indeed, when two random variables are jointly elliptically distributed, they can be independent only if the distribution is joint normal (Kelker, 1970).

market belief is increasing in the agent's action. An *increasing* linear equilibrium has

$$\hat{\tau}(a) = La + K \tag{2}$$

for some $L > 0$. Given a market belief of the form (2), the agent's optimal action is unique:

$$a = \eta + sL\gamma. \tag{3}$$

When the agent plays a linear strategy, the market's posterior beliefs will be elliptically distributed with mean linear in the agent's action. Increasing linear equilibria are determined by solving for a fixed point: values of $L > 0$ 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 SA.3, it is useful to note that an equilibrium value of $L > 0$ is determined as:

$$L = \frac{\mathcal{L}(s, L, \tau)\sigma_\tau^2 + \mathcal{L}(s, L, \neg\tau)\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2 L^2 \sigma_\gamma^2 + 2sL\rho\sigma_\gamma\sigma_\eta}, \tag{4}$$

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 $\mathcal{L}(s, \tilde{L}, \gamma) \equiv s\tilde{L}$. By Equation 2, the equilibrium constant L measures the responsiveness of the market belief $\hat{\tau}$ to the agent's action.

Remark 3. By Equation 3, the agent takes an action above her natural action in any increasing linear equilibrium. Consequently, such equilibria are unaffected by relaxing free downward deviations: making it costly for the agent to take actions $a < \eta$ (e.g., $C(a, \eta, \gamma) = -(a - \eta)^2/(2\gamma)$) would only make some deviations even less attractive.

Remark 4. Fixing a joint distribution over η and γ with $\rho \geq 0$, an equilibrium will be less informative about $\hat{\eta}$ and more informative about $\hat{\gamma}$ when the coefficient sL in Equation 3 is larger; see Lemma 7 in Appendix SA.3. Fixing the marginal distribution of τ and varying ρ or $\sigma_{\neg\tau}$, though, sL is no longer a sufficient statistic for information. One will need to look at $R_{\tau a}^2$, with explicit formulas given in Equations SA.6 and SA.7 of Appendix SA.3.

4.1. Dimension of interest is the natural action

Assume $\tau = \eta$. As described above, informativeness about $\hat{\eta}$ is captured by the one-dimensional value $R_{\eta a}^2 \in [0, 1]$.

Proposition 4. *In the LQE specification, assume $\tau = \eta$ and $\rho \geq 0$.*

1. *There is a unique increasing linear equilibrium.*

2. *In that equilibrium: (a) As $s \to 0$, $R_{\eta a}^2 \to 1$; (b) As $s \to \infty$, $R_{\eta a}^2 \to \rho^2$.*

3. *Furthermore: (a)* $\dfrac{\mathrm{d}}{\mathrm{d}s} R_{\eta a}^2 < 0$; *(b)* $\dfrac{\mathrm{d}}{\mathrm{d}\mu_\gamma} R_{\eta a}^2 = 0$; *(c)* $\dfrac{\mathrm{d}}{\mathrm{d}\sigma_\gamma} R_{\eta a}^2 < 0$; *(d)* $\dfrac{\mathrm{d}}{\mathrm{d}\rho} R_{\eta a}^2 > 0$.

For the case of $\rho = 0$, Fischer and Verrecchia (2000) have obtained similar conclusions to Proposition 4 using their specification.

Part 1 of Proposition 4 is self-explanatory. Part 2(a) says that as stakes vanish, the increasing linear equilibrium becomes fully informative about $\hat{\eta}$; this result is a counterpart to part 1 of Proposition 2. Part 2(b) says that as stakes grow unboundedly, the equilibrium becomes uninformative about $\hat{\eta}$ when $\rho = 0$, which is consistent with part 3 of Proposition 2. However, for any $\rho > 0$, there is some information revealed about $\hat{\eta}$ even in the limit of unbounded stakes; as explained after Proposition 2, 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. The unbounded-stakes limit becomes fully informative as $\rho \to 1$.

Part 3 of Proposition 4 provides comparative statics for interior stakes. Part 3(a) confirms our fundamental theme that higher stakes reduce information about the natural action. The other parts address comparative statics that we have not touched on so far. Part 3(b) notes that the ex-ante mean of the gaming ability has no effect on equilibrium informativeness about $\hat{\eta}$; rather, changes in μ_γ only shift the agent's action and the market belief function by a constant. Part 3(c) says that greater ex-ante uncertainty about γ reduces equilibrium information about $\hat{\eta}$. Together, parts 3(b) and 3(c) underscore that loss of information about the natural action is not due to gaming *per se*, but rather heterogeneity in gaming ability. Finally, part 3(d) says that increasing an already non-negative correlation between η and γ leads to more equilibrium information about $\hat{\eta}$. An intuition is that a greater non-negative correlation reduces the amount of heterogeneity in γ conditional on any η: for elliptical distributions, $\mathrm{Var}(\gamma|\eta) = \alpha\sigma_\gamma^2(1 - \rho^2)$. At the limit when $\rho = 1$, the type space is effectively one-dimensional and the equilibrium fully reveals all private information.

Although Proposition 4 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 increasing linear equilibria. Nevertheless, the comparative statics in s, μ_γ, and σ_η all generalize subject to the caveat of focusing on the appropriate equilibria.[22]

---

[22] Two points bear clarification about $\rho < 0$. First, $R_{\eta a}^2 \to 0$ as $s \to \infty$, as is consistent with the "ironing" logic discussed in the context of part 3 of Proposition 2. Second, comparative statics on ρ are not clear-cut.

4.2. Dimension of interest is the gaming ability

Assume $\tau = \gamma$. Informativeness about $\hat\gamma$ is captured by $R^2_{\gamma a}$.

Proposition 5. *In the LQE specification, assume $\tau = \gamma$ and $\rho \geq 0$.*

1. *(a) When $\rho = 0$, there is an increasing linear equilibrium if and only if $s > \sigma^2_\eta/\sigma^2_\gamma$; the increasing linear equilibrium is unique when it exists.*

   (b) When $\rho > 0$, there is a unique increasing linear equilibrium.

2. *In the increasing linear equilibrium: (a) When $\rho > 0$ and as $s \to 0$, $R^2_{\gamma a} \to \rho^2$; (b) As $s \to \infty$, $R^2_{\gamma a} \to 1$.*

3. *Furthermore: (a) $\dfrac{\mathrm{d}}{\mathrm{d}s} R^2_{\gamma a} > 0$; (b) $\dfrac{\mathrm{d}}{\mathrm{d}\mu_\eta} R^2_{\gamma a} = 0$; (c) $\dfrac{\mathrm{d}}{\mathrm{d}\sigma_\eta} R^2_{\gamma a} < 0$; (d) $\dfrac{\mathrm{d}}{\mathrm{d}\rho} R^2_{\gamma a} > 0$.*

Part 1 of Proposition 5 says that an increasing linear equilibrium exists only when either $\rho = 0$ and stakes are sufficiently large, or when $\rho > 0$. When $\rho = 0$ (but not when $\rho > 0$), for any level of stakes there is an equilibrium in which the agent plays $a = \eta$: the agent takes her natural action at no cost, and the market learns nothing about $\hat\gamma$.[23]

To interpret Part 2(a) of Proposition 5, observe that as $s \to 0$ the agent's play in an increasing linear equilibrium must converge to $a = \eta$; at the limit, the market learns η, which implies $R^2_{\gamma a} = R^2_{\gamma \eta} = \rho^2$. Part 2(b) says that as stakes grow unboundedly, the equilibrium becomes fully informative about $\hat\gamma$; this is a counterpart to part 1 of Proposition 3. Part 3(a) confirms our fundamental theme that, even away from limiting stakes, higher stakes increase information about gaming ability. The remaining comparative statics in part 3 are analogous to those discussed in the context of Proposition 4.

4.3. Mixed dimensions of interest

The tractability of the LQE specification makes it possible to study a number of additional questions. Supplementary Appendix SA.3.4 studies an extension in which the agent cares about the market's belief about both η and γ. Specifically, we consider a signaling benefit

---

[23] As $\rho \to 0_+$, the unique increasing linear equilibrium converges to the increasing linear equilibrium of $\rho = 0$ if $s > \sigma^2_\eta/\sigma^2_\gamma$, while it otherwise converges to the equilibrium where the agent plays $a = \eta$. When $\rho = 0$, the equilibrium in which $a = \eta$ is uninformative about $\hat\gamma$, but it is not uninformative about γ; see fn. 21. Note that this equilibrium exists irrespective of free downward deviations. There also exists a fully uninformative equilibrium (just as when $\tau = \eta$) supported using free downward deviations, in which the agent pools on a sufficiently low action independent of type.

$s[\kappa\hat{\gamma}+(1-\kappa)\hat{\eta}]$, where $\kappa \in (0,1)$. We show that even with these mixed dimensions of interest, higher stakes reduce market information about the agent's natural action and increase market information about the agent's gaming ability.

5. Applications

5.1. Manipulability and information provision

Our main results are couched in terms of changes in the signaling stakes, which affect the benefits of signaling. Now let us add a parameter to the model that affects the costs of signaling: the *manipulability* of the signal, $M > 0$. Consider an agent's payoff function of the form

$$V(\hat{\tau}; s) - \frac{C(a,\theta)}{M}. \tag{5}$$

Higher manipulability scales down the signaling costs for all agent types. For an agent, this payoff function is isomorphic to $MV(\hat{\tau}; s) - C(a,\theta)$, with M satisfying all the conditions on s in Assumption 2. Increasing manipulability thus has the same equilibrium effect as raising the stakes. When the dimension of interest is the natural action, it reduces market information; when the dimension of interest is the gaming ability, it increases information.

In some cases manipulability is a property of the signaling technology. Google's PageRank algorithm for web site ranking is considered harder (i.e., costlier) to game than earlier search engines that were based primarily on keyword density. Different types of tests might cover material that is easier or harder to study for; for instance, it is often thought to be harder to study for tests that measure "aptitude" rather than "achievement."[24]

Manipulability may also depend on agents' knowledge about the signaling technology. One aspect of this is familiarity or experience with the technology. A designer trying to increase market information about natural actions may benefit from altering a test format or by replacing a credit scoring or search engine ranking algorithm with a new one. Even if the new signaling technology is inherently no less manipulable, agents may initially find it more difficult to game. A constant battle between designers and agents could result. Indeed, it is reported that Google tweaks its search algorithm as often as 600 times a year partly to mitigate undesirable search engine optimization (http://moz.com/google-algorithm-change).

---

[24] "SAT" was originally an acronym for Scholastic Aptitude Test, but the College Board changed the name in 1994 to remove the connotation that it measured something innate and unchangeable. It is no longer an acronym for anything.

Designers can also control how much information is given out about the workings of the signaling technology. Google keeps its precise algorithm 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 informative. Weston (2011, pp. 8–9) writes that prior to legislation requiring the Fair Isaac Corporation to reveal some details about its FICO algorithm, "The company said it worried that consumers [...] would try to 'game the system' if they knew more. Fair Isaac feared that its formulas would lose their predictive abilities if consumers started changing their behavior to boost their scores."

It is important to recognize, however, that revealing information about the workings of an algorithm need not scale down gaming costs *uniformly* across agents as posited in (5). If gaming costs were reduced differentially across types, then the impact on market information could be ambiguous. To make this point precise, recall the cost specification of $C(a, \theta) = c(a, \eta)/\gamma$, in which case (5) becomes $V(\hat{\tau}; s) - \frac{c(a,\eta)}{M\gamma}$. Increasing M is now isomorphic to uniformly scaling up the gaming abilities of all agent types. But, as discussed earlier, knowledge of or experience with the signaling technology may be a source of heterogeneity on gaming ability. In that case, revealing information could effectively increase gaming ability more for the low-γ types than the high-γ ones. At the extreme, revealing all information about the algorithm could raise all gaming abilities to the same high level, eliminating heterogeneity on γ. There would then be a separating equilibrium with full information on η. Away from the extreme, Proposition 4 (part 3) for the LQE specification suggests a simple rule of thumb: the amount of information about η decreases with the variance of γ, while it is independent of the mean of γ. Any intervention lowering the variance of gaming ability should tend to improve market information about the natural action. So, if some people are better able to study for the SAT than others because they have greater access to past SAT questions, the College Board may prefer to reveal these questions to everyone.

Similarly, subsidizing direct monetary costs of gaming could effectively raise the gaming ability of all agents. This would seem to be a bad policy to increase information about natural actions. But if heterogeneity on gaming ability had been driven by wealth differences in the first place, then such a policy might increase gaming ability more for low-γ types than high-γ types, with the effect of increasing information about natural actions. This is our

interpretation of the College Board's recent move to "level the playing field" between rich (high γ) and poor (low γ) students by producing and publicizing free test prep material (College Board, 2014a). The intervention does not just transfer surplus from rich to poor; it improves the information content of SAT scores.

5.2. Informational externalities across markets

Suppose a designer can choose whether or not to reveal agents' actions to different observers. The direct effect of revealing the actions to a new observer is that the new observer is more informed. With more observers, however, agents have stronger incentives to signal: signaling stakes increase. Our analysis therefore implies an indirect informational effect on preexisting observers. When the dimension of interest is the gaming ability, an agent's action becomes more informative; but when the dimension of interest is the natural action, the action becomes less informative. In this latter case, a designer must trade off the benefits of revealing the action to marginal observers with the negative informational externality imposed on inframarginal observers.

For concreteness, consider different markets in which credit reports are useful. Lenders use credit scores to determine how much credit to offer borrowers, and at what terms. Employers check credit reports during the hiring process to assess risks like employee theft or general trustworthiness. Automobile and homeowners insurance companies now also commonly use information from credit reports to help determine insurance rates. In recent years, a number of states have considered or passed legislation restricting the use of credit history in insurance underwriting and in employment. Many of the arguments in favor of these laws are based on some notion of "fairness"—it is not fair to deny someone employment based on their high credit card debt, or to raise their insurance rates because of missed mortgage payments.

We suggest a different 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 efficiency in the loan market.

To formalize the argument, let the dimension of interest be the natural action, η, and let

$$V(\hat{\eta}; s) = v(\hat{\eta}) + w(\mathbb{E}[\eta]) + (s-1)\big(w(\hat{\eta}) - w(\mathbb{E}[\eta])\big) \text{ for } s \in \{1, 2\}, \tag{6}$$

with $v(\cdot)$ and $w(\cdot)$ strictly increasing. The designer chooses $s \in \{1, 2\}$ to maximize allocative efficiency, $\mathbb{E}[V]$. The interpretation is that there is some benchmark first market which always observes the signal—e.g., the loan market in the credit scoring application. Showing

27

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—e.g., automobile insurance—in which the signal 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 the signal is observed ($s = 2$), the agent gets a payoff $w(\hat{\eta})$ in the second market.

The key condition is that *information is more socially valuable in the first market.* In particular, let $v(\cdot)$ be strictly convex and let $w(\cdot)$ be linear. (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.) There is a positive social value of information in the first market in that more information implies higher allocative efficiency. For example, in the loan market, it efficient to make more, or larger, loans to lower-financial-risk consumers. In the second market, though, any information just redistributes surplus across types. For example, in the market for auto insurance, lower-financial-risk consumers receive lower rates because this is correlated with their filing fewer claims, but (almost) no people are on the margin of car ownership based on their insurance rates. So information about financial risk leads to monetary transfers in the auto market but does not affect allocations.

By construction, there is no direct effect on allocative efficiency in the second market from observing agents' actions. The only effect on allocative efficiency is the indirect effect on the first market. Our analysis implies that higher stakes lead to a less informative signal, 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 setting using Proposition 1:

Corollary 1. *Consider the 2×2 setting with $\tau = \eta$. Let $V(\cdot)$ be of the form (6), with $v(\cdot)$ strictly convex and $w(\cdot)$ linear. For any equilibrium under $s = 2$, there exists an equilibrium under $s = 1$ with weakly higher allocative efficiency, $\mathbb{E}[V]$.*

In other words, to support the efficiency of the "loan" market in which information about natural actions is socially valuable, the signal should be hidden from the "auto insurance" market in which the social value is negligible. There is, of course, a converse: if the dimension of interest were the gaming ability rather than the natural action, then revealing the signal to a second market would generate a positive informational externality on the first market. The designer might want to show the agent's action to the purely redistributive second market just to ramp up stakes, which would increase gaming and information about gaming ability.

In [Appendix B](), we analyze informational externalities when information is equally valuable across markets: all markets have the same convex $v(\cdot)$ function. There is now a direct allocative benefit as well as an indirect cost when a new market observes signals. We study whether allocative efficiency is maximized by showing signals to as many markets as possible, or whether there is an interior optimum after which the benefit to new markets is outweighed by the cost to existing markets. Of course, a designer interested in agent welfare—allocative efficiency minus signaling costs—would also consider the benefits that lower stakes confer through less gaming.

6. Conclusion

This paper has studied 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 such as 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]()). As discussed in [Subsection 2.3](), 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.

Finally, our analysis of information yields direct implications for the allocative efficiency of equilibria. Other quantities are also of of interest. A promising avenue for future research is to explore the broader implications of multidimensional heterogeneity and muddled information in signaling games.

Bibliography

ALI, S. N. AND R. BÉNABOU (2016): "Image Versus Information: Changing Societal Norms and Optimal Privacy," Unpublished.

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.

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.

COLLEGE BOARD (2014a): "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.

——— (2014b): "SAT Percentile Ranks for Males, Females, and Total Group," https://secure-media.collegeboard.org/digitalServices/pdf/sat/sat-percentile-ranks-composite-crit-reading-math-writing-2014.pdf.

DALEY, B. AND B. GREEN (2014): "Market Signaling with Grades," *Journal of Economic Theory*, 151, 114–145.

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

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

ESTEBAN, J. AND D. RAY (2006): "Inequality, Lobbying, and Resource Allocation," *American Economic Review*, 96, 257–279.

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.

GANUZA, J.-J. AND J. S. PENALVA (2010): "Signal Orderings Based on Dispersion and the Supply of Private Information in Auctionsased on dispersion and the supply of private information in auctions," *Econometrica*, 78, 1007–1030.

GESCHE, T. (2016): "De-biasing Strategic Communication," Unpublished.

GÓMEZ, E., M. A. GÓMEZ-VILLEGAS, AND J. M. MARÍN (2003): "A Survey on Continuous Elliptical Vector Distributions," *Revista matemática complutense*, 16, 345–361.

KARTIK, N. (2009): "Strategic Communication with Lying Costs," *Review of Economic Studies*, 76, 1359–1395.

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

KELKER, D. (1970): "Distribution Theory of Spherical Distributions and a Location-Scale Parameter Generalization," *Sankhya: The Indian Journal of Statistics, Series A*, 32, 419–430.

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.

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.

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

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

WALL STREET JOURNAL (2015): "How to Perfect Your Credit Score," http://www.wsj.com/articles/how-to-perfect-your-credit-score-1420832973.

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

Appendices

A. Relating the Signaling Benefit to Consumer Demand

Here we illustrate how a benefit function $V(\cdot; s)$ may be derived from demand curves in a competitive market in which consumers have heterogeneous costs of service. The convexity of V relates to the efficiency improvement from firms' learning about consumer costs.

Example 2. A consumer chooses an action a that is observed by a competitive market of firms. Firms then offer the consumer a price p, and the consumer (mechanically) purchases a quantity given by the demand curve $D(p, s)$. The stakes parameter s is a public and exogenous demand-shifter: $\frac{\partial D(p,s)}{\partial s} > 0$. 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, the service cost per unit is $k - \alpha\tau$, with the most costly consumer having $\tau = \underline{\tau}$. The market is competitive, so the price offered to a consumer will equal the expected cost: $p = k - \alpha\hat{\tau}$. A consumer who is thought to be type $\hat{\tau}$ receives gross consumer surplus $V(\hat{\tau}; s) = \int_{k-\alpha\hat{\tau}}^{k-\alpha\underline{\tau}} D(k - \alpha t, s)dt$. 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 costs of signaling, is total welfare.

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 face a lower price and a consumer of lower type would face a higher price, but purchases would be unaffected. With a downward sloping demand curve, information becomes socially valuable. A consumer with higher cost of service purchases less, as is efficient, while a consumer with lower cost of service purchases more. This induces a convex benefit function $V(\cdot; s)$. ‖

B. More on Informational Externalities Across Markets

Revealing the agent's action to a new observer makes that observer more informed about the agent's type. But, as we argued in Subsection 5.2, it can have an informational externality on others who already observed the agent's action. When the dimension of interest is the natural action, existing observers tend to become less informed when "the stakes are raised" by revealing the action to new observers.

Continuing from Subsection 5.2, suppose an agent participates in a number of markets. Information about the agent's action may be revealed, or not, to each such market. Contrasting with the environment of Corollary 1, however, information is socially valuable in all markets (i.e., strictly convex benefit functions); indeed, markets are *homogeneous* insofar as information is equally valuable in each market. There is no longer a way to reveal the agent's action to (more convex) high-value markets before revealing it to (less convex) low-value markets. Is the social value of information—allocative efficiency—maximized by showing signals to as many observers as possible, or is there is an interior optimum?

Throughout this section, let the dimension of interest be the natural action, η. Let

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

with $w(\cdot)$, and therefore $v(\cdot)$, 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 with the highest allocative efficiency.

The interpretation of the benefit function V and the stakes s is that there is a large 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 a mass s of these markets, the agent's action is observed, in which case the beliefs on the agent's type are $\hat{\eta}$. In all 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 the other uninformed markets.) The total value of information aggregated across markets is the value per observation, $\mathbb{E}[v]$, times the mass of observers, s: $\mathbb{E}[V(\hat{\eta}; s)] = s \cdot \mathbb{E}[v]$.

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 of Appendix A), 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.

It is clear that if the market doesn't become uninformative in the limit as the stakes s gets unboundedly large—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 cares about maximizing allocative efficiency should increase s without bound.[25] The tradeoff between the marginal benefit

---

[25] This result applies, for example, when Θ contains an extreme type—one below or above all other types

and inframarginal cost of adding observers is only interesting when each observer becomes uninformed in the limit.

We analyze two simple settings. The first, as a proof of concept, shows that information is maximized at an interior s when the type space is a single pair of cross types. The second illustrates a richer type space in which η and γ are independently distributed. There we find that it is optimal to take $s \to \infty$: the gain from increasing information for marginal observers dominates the loss from worsening information for inframarginal observers.

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

Proposition 6. *Assume $\Theta = \{\theta_1, \theta_2\}$ where θ_1 and θ_2 are cross types, 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 $s\mathbb{E}[v]$ is maximized over choice of s and choice of equilibrium at some finite $s > 0$.*

Recall from Observation 1 part 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, therefore, $\mathbb{E}[v] > 0$ at all $s > 0$ and $\mathbb{E}[v] \to 0$ as $s \to \infty$. We establish in the proof of Proposition 6 (in Supplementary Appendix SA.4.2) that $s\mathbb{E}[v] \to 0$ as $s \to \infty$.

Independent η and γ. Now consider the 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 2 part 3. Here we find that 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 7. *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 $C(a, \eta, \gamma) = (\max\{a - \eta, 0\})^r / \gamma$ for some $r > 1$. As $s \to \infty$, there exists a sequence of equilibria with allocative efficiency $s\mathbb{E}[v] \to \infty$.*

We prove Proposition 7 in Supplementary Appendix SA.4.2 by constructing a class of two-action equilibria—all types take actions at either $a = \underline{\eta}$, or at some higher stake-dependent

---

in both η and γ—that occurs with 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 0.

action—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 $s^{\frac{r-1}{r+1}}$.[26] 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.

Our analysis of homogeneous markets suggests 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 the social value of information is maximized by showing signals to as many observers as possible. We stress, however, that our discussion ignores signaling costs. A designer who takes agents' signaling costs into account might prefer to hide information from some observers.

## C. Proof Sketch of Lemma 3 in Subsection 3.1

The proof of Lemma 3 is in Supplementary Appendix SA.2.2. Here we provide a sketch.

Part 1 of Lemma 3. Let $\tau = \eta$, and fix some equilibrium at some $s_0 > 0$. It suffices to show that as s varies in a neighborhood of s_0, we can perturb the equilibrium to be (weakly) more informative as s decreases and less informative as s increases.

For small changes in stakes, in many cases one can maintain the distribution of beliefs β_τ as stakes vary simply by continuously moving actions while not changing the (mixed) play of each type across these actions. For instance, as stakes go down, one may be able to go "left to right" decreasing each subsequent on-path action in a manner that keeps all of the relevant incentive constraints binding.

This approach fails when there is a pair of actions $a_1 < a_2$ such that two different types are both mixing over these actions. In that case it is impossible to move either of the actions without breaking one of the types' indifference. When this occurs, we hold a_1 and a_2 fixed as stakes vary; instead, we alter the mixing probabilities such that as stakes go down, (i) incentive constraints remain binding, and (ii) high natural action types are shifted to higher actions, and low natural action types are shifted to lower actions. Point (ii) ensures that the equilibrium becomes more informative about η as stakes are reduced: beliefs in the initial equilibrium were monotonic, and the perturbation spreads beliefs out further.

---

[26] We do not suggest this rate is optimal; there are examples in which $\mathbb{E}[V]$ increases at a faster rate.
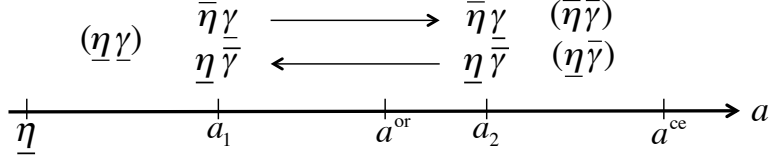
(a) When the gamer type $\underline{\eta}\overline{\gamma}$ does not have an incentive to take an action below a_1.
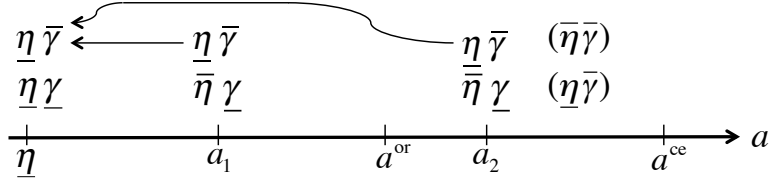
(b) When the gamer type $\underline{\eta}\overline{\gamma}$ has an incentive to take an action below a_1.

Figure 2 – 2×2 setting, $\tau = \eta$: Perturbing an equilibrium to become more informative as s decreases.

More specifically, the above situation involves the gamer $(\underline{\eta}, \overline{\gamma})$ and the natural type $(\overline{\eta}, \underline{\gamma})$ both indifferent between actions $a_1 \in [\underline{\eta}, a^{\mathrm{or}})$ and $a_2 \in (a^{\mathrm{or}}, a^{\mathrm{ce}}]$. Start from an equilibrium at stakes s_0 in which the gamer and the natural type both mix over actions a_1 and a_2. There are two qualitative ways in which the mixing probabilities may need to be shifted while holding a_1 and a_2 fixed. First, suppose that at stakes s_0 the gamer strictly prefers a_1 over any action below a_1. A reduction in stakes makes a_1, which has a lower market belief, relatively more appealing than a_2, which has a higher belief. To recover indifferences, we adjust the mixing probabilities so that the gamer shifts some mass from a_2 to a_1 while the natural type shifts from a_1 to a_2. Since $\tau = \eta$, this increases the belief at a_2 and reduces the belief at a_1, making a_2 relatively more appealing once more. See Figure 2 panel (a); in the figure, parentheses indicate that there may be a mass of such types taking actions in some range. Second, suppose that at stakes s_0 the gamer is indifferent between a_1 and an action below a_1; in equilibrium, this lower action will be $\underline{\eta}$. A reduction in stakes makes $a = \underline{\eta}$ relatively more attractive than a_1, and also makes a_1 more attractive than a_2. We can adjust mixing probabilities by having the gamer move some mass from a_2 to $\underline{\eta}$, improving beliefs at a_2, and then move from mass from a_1 to $\underline{\eta}$, improving beliefs at a_1. Once again, this change recovers indifferences and increases information. See Figure 2 panel (b). *Q.E.D.*

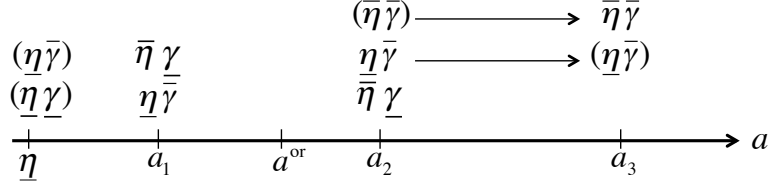**Part 2 of Lemma 3.** Let $\tau = \eta$, and fix some equilibrium at some $s_0 > 0$. We argue that

$$(\bar{\eta}\,\bar{\gamma}) \longrightarrow \bar{\eta}\,\bar{\gamma}$$

$$(\bar{\eta}\,\bar{\gamma}) \quad \bar{\eta}\,\underline{\gamma} \qquad \underline{\eta}\,\bar{\gamma} \longrightarrow (\underline{\eta}\,\bar{\gamma})$$

$$(\underline{\eta}\,\underline{\gamma}) \quad \underline{\eta}\,\bar{\bar{\gamma}} \qquad \bar{\bar{\eta}}\,\underline{\gamma}$$

Figure 3 – 2×2 setting, $\tau = \gamma$: Perturbing an equilibrium to become more informative as s increases.

as stakes increase in a neighborhood, we can perturb the equilibrium to (weakly) increase information. Analogously to the proof sketch of Lemma 3 part 1, in many cases one can maintain the distribution of beliefs β_τ as stakes increase by going "left to right" and increasing actions while not changing the behavior of each type across these actions.

This approach fails when the gamer $(\underline{\eta}, \bar{\gamma})$ and the natural type $(\bar{\eta}, \underline{\gamma})$ are both indifferent over on-path actions $a_1 \in [\underline{\eta}, a^{\mathrm{or}})$ and $a_2 \in (a^{\mathrm{or}}, a^{\mathrm{ce}}]$. When this occurs, we hold a_1 and a_2 fixed as stakes vary; instead, we alter the mixing probabilities such that as stakes increase, (i) incentive constraints remain binding, and (ii) high gaming ability types are shifted to higher actions, and low gaming ability types are shifted to lower actions. Point (ii) ensures that the equilibrium becomes more informative about γ as stakes are increased. More specifically, an increase in stakes makes action a_2 (with higher belief) relatively more appealing than a_1 (with lower belief). So, we shift a high gaming ability type—either the gamer, or the type $(\bar{\eta}, \bar{\gamma})$—who had previously been playing a_2 to an action a_3 above a_2. This change reduces the belief at a_2, recovering the indifference between a_1 and a_2. See Figure 3, where parentheses indicate that there may be a mass of such types taking the specified action. *Q.E.D.*

D. Additional Results and Proofs for Subsection 3.2

Lemma 4. *Assume $\tau = \eta$. Fix any two cross types, $\theta_1 = (\underline{\eta}, \bar{\gamma})$ and $\theta_2 = (\bar{\eta}, \underline{\gamma})$ with $\underline{\eta} < \bar{\eta}$ and $\underline{\gamma} < \bar{\gamma}$, 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 $V(\hat{\eta}_2; s) - V(\hat{\eta}_1; s) \leq C(a^{ce}, \theta_1) = C(a^{ce}, \theta_2)$.*

Proof of Lemma 4. Fix any type space containing θ_1 and θ_2, and any 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 $a_1 = \underline{\eta}$.   Q.E.D.

Proof of Proposition 2. <u>Part 1</u>: Θ_η is finite by hypothesis. We claim that for small enough $s > 0$ there is an equilibrium in which every $\theta = (\eta, \gamma)$ takes its natural action, $a = \eta$; any 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. 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 $s > 0$ is small enough. The proposition's hypotheses about Θ imply there is an $\varepsilon > 0$ such that $C(a, \eta, \gamma) > \varepsilon$ for all $(\eta, \gamma) \in \Theta$ 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.

<u>Part 2</u>: The result follows from Lemma 4 and part 3 of Assumption 2.

<u>Part 3</u>: 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 4 and part 3 of Assumption 2.

Now fix an arbitrary sequence of equilibria as $s \to \infty$. Given the hypothesis that the marginal distribution of γ is continuous, 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 $\mathbb{E}[\eta|\theta \in S_\Delta] = \mathbb{E}[\hat{\eta}^*(\theta)|\theta \in S_\Delta] \geq \mathbb{E}[\eta] + \Delta.$ Q.E.D.

Lemma 5. *Assume $\tau = \gamma$. Fix any two cross types, $\theta_1 = (\underline{\eta}, \overline{\gamma})$ and $\theta_2 = (\overline{\eta}, \underline{\gamma})$ for $\underline{\eta} < \overline{\eta}$ and $\underline{\gamma} < \overline{\gamma}$, with the corresponding cost-equalizing action a^{or}. Suppose $\{\theta_1, \theta_2\} \subseteq \Theta$ and $V(\max \Theta_\gamma; s) - V(\min \Theta_\gamma; s) \leq C(a^{or}, \theta_1)$. Then across all equilibria, it holds that if $\hat{\gamma}_i$ is some belief that θ_i induces in equilibrium $(i = 1, 2)$, then $\hat{\gamma}_1 \leq \hat{\gamma}_2$.*

Proof of Lemma 5. 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 left-hand side of this inequality is the incremental cost while the right-hand side 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 3. <u>Part 1</u>: 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 $\{\eta_{\max}\} \times \Theta_\gamma$ 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.

Part 2: The result follows from Lemma 5 and part 3 of Assumption 2.

Part 3: The argument is analogous to that provided for Proposition 2 part 3, switching γ and η, taking $s \to 0$ rather than $s \to \infty$, and using Lemma 5 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. $Q.E.D.$

Supplementary Appendices for Online Publication

SA.1. 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 $a^{\text{or}} > \overline{\eta}$ 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 and attains a strictly positive maximum at $a = a^{\text{or}}$. Moreover, the integral is zero at a unique point (on the domain $a > \underline{\eta}$), $a^{\text{ce}} > a^{\text{or}}$, because Assumption 1 implies there is some $\varepsilon > 0$ such that the integrand is 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

$$\hat{\tau}(a') = \hat{\tau}(a''). \tag{SA.1}$$

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

$$C(a', \theta) = C(a'', \theta) = 0. \tag{SA.2}$$

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 SA.1, the induced belief at a' does not change; hence, by Equation SA.2, 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.*

SA.2. Proofs for Section 3

SA.2.1. Proof of Observation 1

Throughout this proof, we refer to the gamer $(\underline{\eta}, \overline{\gamma})$ as θ_1 and the natural type $(\overline{\eta}, \underline{\gamma})$ as θ_2.

<u>Part 1</u>: In a separating equilibrium, the gamer θ_1 plays some action a_1 while the natural type θ_2 plays $a_2 > a_1$.[27] 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'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

$$C(a_2, \theta_1) \geq V(\overline{\eta}; s) - V(\underline{\eta}; s) \geq C(a_2, \theta_2). \tag{SA.3}$$

---

[27] 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.

From (SA.3), 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 (SA.3) 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

$$H(a) := C(a, \theta_1) - C(a, \theta_2) \tag{SA.4}$$

is continuous. Let s_η^* solve $V(\overline{\eta}; s_\eta^*) - V(\underline{\eta}; s_\eta^*) = C(a^{\text{ce}}, \theta_1)$. 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 $s > s_\eta^*$. We show that a partially pooling equilibrium exists. 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.

Finally, we show that equilibria become uninformative as $s \to \infty$. First, note that the belief $\hat{\eta}(a) \in [\underline{\eta}, \overline{\eta}]$ at an equilibrium action a is strictly above $\underline{\eta}$ if and only if the action is played in equilibrium (with some probability) by a natural type θ_2, and the belief is strictly below $\overline{\eta}$ if and only if the action is played by a gamer θ_1. So the lowest belief at an equilibrium action must be achieved at some action played by θ_1, and the highest equilibrium at some action played by θ_2. Call these minimum and maximum beliefs $\hat{\eta}_1 \leq \hat{\eta}_2$. It suffices to show that $\hat{\eta}_2 - \hat{\eta}_1 \to 0$ as $s \to \infty$. Towards contradiction, suppose there exists a sequence of equilibria with $\hat{\eta}_2 - \hat{\eta}_1$ approaching $\varepsilon > 0$ as $s \to \infty$. There must then be a subsequence with $\hat{\eta}_2 \to \hat{\eta}_2^{\text{lim}}$ and $\hat{\eta}_1 \to \hat{\eta}_1^{\text{lim}}$, with $\hat{\eta}_2^{\text{lim}} > \hat{\eta}_1^{\text{lim}}$. Apply Lemma 4 from Appendix D, it holds that for each s, $V(\hat{\eta}_2; s) - V(\hat{\eta}_1; s) \leq C(a^{\text{ce}}, \theta_1)$. Passing to the limit, it must be that

$$\lim_{s \to \infty} V(\hat{\eta}_2^{\text{lim}}; s) - V(\hat{\eta}_1^{\text{lim}}; s) \leq C(a^{\text{ce}}, \theta_1).$$

But Assumption 2 part 3 implies that the left-hand side (LHS) above diverges to infinity, a contradiction.

<u>Part 2</u>: First, let $s_\gamma^{**} > 0$ solve $V(\overline{\gamma}; s_\gamma^{**}) - V(\underline{\gamma}; s_\gamma^{**}) = C(a^{\text{or}}, \theta_1)$. We show that no informative equilibrium exists at $s \leq s_\gamma^{**}$. The value s_γ^{**} is defined so that at any equilibrium under $s \leq s_\gamma^{**}$, θ_1 would never play any action above a^{or}; it would strictly prefer to receive the worst belief $\hat{\gamma} = \underline{\gamma}$ at action $a = \underline{\eta}$. Towards contradiction, suppose there is an informative equilibrium at $s \leq s_\gamma^{**}$. 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 the fact 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 $\hat{\gamma}(a') < \hat{\gamma}(a'')$.

Next, let s_γ^* solve $V(\overline{\gamma}; s_\gamma^*) - V(\underline{\gamma}; s_\gamma^*) = C(a^{\text{ce}}, \theta_1)$; note that $s_\gamma^* > s_\gamma^{**}$ because $a^{\text{ce}} > a^{\text{or}}$. We show that there is a separating equilibrium when $s > s_\gamma^*$. In a separating equilibrium, without loss the natural type $\theta_2 = (\overline{\eta}, \underline{\gamma})$ plays $a_2 = \underline{\eta}$ and the gamer $\theta_1 = (\underline{\eta}, \overline{\gamma})$ plays $a_1 > \overline{\eta}$. (If $a_1 \leq \overline{\eta}$, free downward deviations implies that the natural type would mimic the gamer.) So a separating equilibrium exists if and only if there is an $a_1 > \overline{\eta}$ such that

$$C(a_1, \theta_2) \geq V(\overline{\gamma}; s) - V(\underline{\gamma}; s) \geq C(a_1, \theta_1). \tag{SA.5}$$

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. Define $a_1(s)$ by setting the second inequality of (SA.5) to hold with equality; $a_1(s)$ is continuous and strictly increasing. 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 (SA.4). By Lemma 1, $H(a_1(s)) \leq 0$ for $a_1(s) > \overline{\eta}$ if and only if $a_1(s) \geq a^{\mathrm{ce}}$; and $a_1(s) \geq a^{\mathrm{ce}}$ if and only if $s \geq s_\gamma^*$ as defined above.

SA.2.2. Proof of Lemma 3

We first establish some straightforward preliminaries for the analysis of equilibria of the 2×2 setting. Throughout this section, in addition to referring to $(\underline{\eta}, \overline{\gamma})$ as the gamer and $(\overline{\eta}, \gamma)$ 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 refer to $(\underline{\eta}, \gamma)$ as the low type and $(\overline{\eta}, \overline{\gamma})$ as the high type.

Claim 1. *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.[28] Consequently, an upper bound on the number of equilibrium actions is $|\Theta| + 2\binom{|\Theta|}{2}$.   *Q.E.D.*

Claim 2. *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^*.[29]*

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 setting specifically, we 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 (Blackwell) information. 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 1, there are finitely many actions in the support and each has strictly positive probability.)

---

[28] 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.

[29] 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 3. *In the 2×2 setting, 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 in the reverse way would lead to less informative rather than more informative beliefs.)

Proof of Claim 3. One posterior distribution of beliefs is Blackwell more informative than another if and only if, for any continuous and convex function over beliefs U, $\mathbb{E}[U(\beta_\tau)]$ 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 $\mathbb{E}[U(\hat{\tau})]$, 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. The belief $\hat{\tau}(a)$ at action a is then

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

and $\mathbb{E}[U(\hat{\tau})]$ is given by $\sum_a U(\hat{\tau}(a)) \sum_\theta f(\theta) p_\theta(a)$ where the 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 $\frac{d}{dp_{\theta'}(a_h)}\mathbb{E}[U(\hat{\tau})] - \frac{d}{dp_{\theta'}(a_l)}\mathbb{E}[U(\hat{\tau})]$. When $\tau(\theta') = \overline{\tau}$, we can evaluate these derivatives and simplify to get

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

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 $\frac{d}{dp_{\theta'}(a_l)}\mathbb{E}[U(\hat{\tau})] - \frac{d}{dp_{\theta'}(a_h)}\mathbb{E}[U(\hat{\tau})]$. 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. $\hspace{2cm}$ *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 1.)

Claim 4. *In the 2×2 setting, any equilibrium is equivalent to one that falls into one of two cases: (1) for any pair of actions weakly above $\underline{\eta}$, at most a single type is willing to play both actions; 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. Moreover:*

1. *In case (1): Consider 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.*

2. *In case (2): It holds that $\underline{\eta} \le a_1 < a^{or} < a_2 \le a^{ce}$. The gamer $(\underline{\eta}, \overline{\gamma})$ 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.*

Proof of Claim 4. The first assertion of the claim (before the enumerated items) is trivial. So we prove the two enumerated items. 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 $[a_1, a_3]$; i.e., 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 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, with $\underline{\eta} \le a_1 < a_2$, that two types are both willing to play. This cannot hold for any pair of types that are single-crossing ordered, and so it must be that the two types are the cross types, $(\underline{\eta}, \overline{\gamma})$ and $(\overline{\eta}, \underline{\gamma})$.[30] It follows that $C(a_2, \underline{\eta}, \overline{\gamma}) - C(a_1, \underline{\eta}, \overline{\gamma}) = C(a_2, \overline{\eta}, \underline{\gamma}) - C(a_1, \overline{\eta}, \underline{\gamma})$, and hence that $\underline{\eta} \le a_1 < a^{or} < a_2 \le a^{ce}$. The same logic also implies that there cannot be an action weakly above $\underline{\eta}$ other than a_1 and a_2 that both types are willing to play.

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 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 $\hat{\gamma} = \underline{\gamma}$. $\qquad$ Q.E.D.

We now proceed to prove the two parts of Lemma 3. 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

---

[30] We can rule out that $(\overline{\eta}, \underline{\gamma})$ and $(\overline{\eta}, \overline{\gamma})$ are both indifferent over a pair of actions in $[\underline{\eta}, \overline{\eta}]$, 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 the lower action.

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 case of only two cross types.

Proof of Lemma 3 part 1. 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 2) guarantees that this local construction around any given equilibrium e_0 at any stakes s_0 extends to a path of equilibria on $s \in (0, s_0)$ that are less informative at higher stakes. In fact, our argument will imply something slightly stronger than claimed in the lemma: we also show that as stakes increase, the equilibrium can be perturbed to increase information, implying that we can extend to a global path of equilibria on $s \in (0, \infty)$.[31]

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 illustrated in Figure 2. As stakes decrease, we move types with $\eta = \underline{\eta}$ down from actions with high beliefs to ones with low beliefs, and/or move types with $\eta = \overline{\eta}$ up from actions with low beliefs 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 3, increase information. As stakes increase, we can do the reverse moves to decrease information.

Using Claim 4, 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 $\underline{\eta} \leq a_1 < a^{\mathrm{or}} < a_2 \leq \overline{\eta}$ such that the gamer $(\underline{\eta}, \overline{\gamma})$ and the natural type $(\overline{\eta}, \gamma)$ are both willing to play a_1 and a_2. No actions in (a_1, a_2) are played in equilibrium. Natural types $(\overline{\eta}, \gamma)$ are not willing to play any action strictly below a_1 or above a_2, low types $(\underline{\eta}, \gamma)$ 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.*

(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 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.

---

[31] 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.

(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 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 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 gamers play a_2 with positive probability. In that case, we follow the logic of Figure 2 panel (a) and move either natural types up from a_1 to a_2, and/or gamers 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 3, 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 type more than for the gamer, and because the natural type is indifferent, the gamer 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 gamers play a_1 with positive probability. In that case, we effectively reverse the direction of Figure 2 panel (a): 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 gamers 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 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 gamers 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.[32]

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 2 panel (b): 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 gamers 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 gamers 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 gamers 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, we can move gamers directly from a_0 to a_2, essentially reversing the direction of Figure 2 panel (b)); 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 gamers 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 gamers 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 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 gamers play a_2. In that case we move gamers 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.

---

[32] 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 $\overline{\eta}$.

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. When there are two actions above a_2, the lower one a_3 is played by the gamer 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 3 part 2. We proceed similarly to the proof of part 1 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 part 1 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 4, 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 $\underline{\eta} \leq a_1 < a^{\mathrm{or}} < a_2 \leq \overline{\eta}$ such that the gamer $(\underline{\eta}, \overline{\gamma})$ 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 gamers 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 Case 1 in Lemma 3 part 1. 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 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 3. (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 gamers 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 $a_1 < a^{\text{or}}$, sliding a_1 up increases the cost of taking a_1 more for gamers than for natural types; and because $a_2 > a^{\text{or}}$, sliding a_2 up increases the cost of taking a_2 less for gamers than for natural types. 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.*

SA.3. Additional Results and Proofs for Section 4

Define the following notation: for any $d \in \{\eta, \gamma\}$, let $\neg d$ denote η if $d = \gamma$ and γ if $d = \eta$.

SA.3.1. Characterizing linear equilibria

Lemma 6. *Consider the LQE specification and fix $\tau \in \{\eta, \gamma\}$. Suppose agents of type (η, γ) take action $a = l_\eta \eta + l_\gamma \gamma + b$ for some constants $l_\eta \neq 0$, $l_\gamma \neq 0$, and b. Then:*

1. *The vector (τ, a) is distributed according to $\mathcal{E}\left((\mu_\tau, \mu_a), \begin{pmatrix} \sigma_\tau^2 & \sigma_{\tau a} \\ \sigma_{\tau a} & \sigma_a^2 \end{pmatrix}, g_{\tau a}\right)$ with*

$$\mu_a = l_\eta \mu_\eta + l_\gamma \mu_\gamma + b,$$
$$\sigma_a^2 = 2 l_\eta l_\gamma \rho \sigma_\eta \sigma_\gamma + l_\eta^2 \sigma_\eta^2 + l_\gamma^2 \sigma_\gamma^2,$$
$$\sigma_{\tau a} = l_\tau \sigma_\tau^2 + l_{\neg \tau} \rho \sigma_\eta \sigma_\gamma,$$
$$g_{\tau a}(\cdot) = g_\theta(\cdot).$$

2. $\hat{\tau}(a) \equiv \mathbb{E}[\tau | a] = \mu_\tau + \dfrac{\sigma_{\tau a}}{\sigma_a^2}(a - \mu_a).$

3. *(τ, a) has an R^2 of*

$$R_{\tau a}^2 = \frac{\left(l_\tau \sigma_\tau^2 + l_{\neg \tau} \rho \sigma_\eta \sigma_\gamma\right)^2}{\sigma_\tau^2 \left(2 l_\eta l_\gamma \rho \sigma_\eta \sigma_\gamma + l_\eta^2 \sigma_\eta^2 + l_\gamma^2 \sigma_\gamma^2\right)}.$$

Proof of Lemma 6. Parts 1 and 2 follow from the properties of elliptical distributions covered in Gómez et al. (2003); respectively, see Theorems 5 and 8 of that paper. Part 3 is immediate from Part 1 and that $R_{\tau a}^2 \equiv \sigma_{\tau a}^2 / (\sigma_\tau^2 \sigma_a^2)$. *Q.E.D.*

For the strategy $a = \eta + sL\gamma$ (Equation 3) with $L > 0$, a simplification of the expressions from Lemma 6 part 3 yields

$$R_{\eta a}^2 = \frac{(\sigma_\eta + sL\rho \sigma_\gamma)^2}{2sL\rho \sigma_\eta \sigma_\gamma + \sigma_\eta^2 + s^2 L^2 \sigma_\gamma^2}, \tag{SA.6}$$

$$R_{\gamma a}^2 = \frac{(\rho \sigma_\eta + sL\sigma_\gamma)^2}{2sL\rho \sigma_\eta \sigma_\gamma + \sigma_\eta^2 + s^2 L^2 \sigma_\gamma^2}. \tag{SA.7}$$

Lemma 7. *Hold fixed $\sigma_\eta > 0$, $\sigma_\gamma > 0$, and $\rho \geq 0$. If $sL > 0$, then $R_{\eta a}^2$ in (SA.6) is decreasing in sL while $R_{\gamma a}^2$ in (SA.7) is increasing in sL.*

Proof of Lemma 7. Both results follow by straightforwardly taking derivatives. The derivative of $R_{\eta a}^2$ with respect to sL can be simplified to

$$\frac{-2sL\sigma_\eta \sigma_\gamma^2 (1 - \rho^2)(\sigma_\eta + sL\rho \sigma_\gamma)}{\left(2sL\rho \sigma_\eta \sigma_\gamma + \sigma_\eta^2 + s^2 L^2 \sigma_\gamma^2\right)^2} < 0.$$

The derivative of $R_{\gamma a}^2$ with respect to sL can be simplified to

$$\frac{2\sigma_\eta^2 \sigma_\gamma (1 - \rho^2)(\rho \sigma_\eta + sL\sigma_\gamma)}{\left(2sL\rho \sigma_\eta \sigma_\gamma + \sigma_\eta^2 + s^2 L^2 \sigma_\gamma^2\right)^2} > 0. \qquad Q.E.D.$$

For the next result, recall that we use β_τ to denote the induced posterior beliefs on τ after observing the realization of a. These beliefs are determined in equilibrium by the joint distribution of a and τ.

Lemma 8. *Fix $\tau \in \{\eta, \gamma\}$, as well as σ_τ, μ_τ, and $g_{\tau a}$. If the equilibrium distribution of (τ, a) is*

$$\mathcal{E}\left((\mu_\tau, \mu_a), \begin{pmatrix} \sigma_\tau^2 & \sigma_{\tau a} \\ \sigma_{\tau a} & \sigma_a^2 \end{pmatrix}, g_{\tau a}\right),$$

then the ex ante distribution of β_τ (prior to the realization of the action) depends only on $R_{\tau a}^2 \equiv \dfrac{\sigma_{\tau a}^2}{\sigma_\tau^2 \sigma_a^2}$. Beliefs are more informative about $\hat{\tau}$ if $R_{\tau a}^2 \in [0, 1]$ is larger, with beliefs being uninformative about $\hat{\tau}$ if $R_{\tau a}^2 = 0$ and fully informative about $\hat{\tau}$ (and τ) if $R_{\tau a}^2 = 1$.

Proof of Lemma 8. For any joint distribution of τ and a, the ex ante distribution of β_τ is clearly preserved under affine transformations of the action. Specifically, subtract μ_a from all actions; multiply actions by $\frac{1}{\sigma_a}$; and then multiply actions by -1 if $\sigma_{\tau a} < 0$. The transformation normalizes $\mu_a = 0$, $\sigma_a^2 = 1$, and $\sigma_{\tau a} \geq 0$ while preserving $R_{\tau a}^2$. Since $\sigma_{\tau a} \in [0, \sigma_\tau]$ is the only remaining parameter under the normalization, it is immediate that the ex ante distribution of β_τ depends only on $R_{\tau a}^2$. Moreover, there is a bijection from $\sigma_{\tau a} \in [0, \sigma_\tau]$ to $R_{\tau a}^2 \in [0, 1]$ in which higher $\sigma_{\tau a}$ corresponds to higher $R_{\tau a}^2$.

Given the normalizations above of $\sigma_a^2 = 1$ and $\mu_a = 0$, Theorem 8 of Gómez et al. (2003) implies that $\hat{\tau}(a) = \mu_\tau + a\sigma_{\tau a}$. Since the marginal distribution of a is independent of $\sigma_{\tau a}$, changes in $\sigma_{\tau a}$ simply correspond to scale shifts of the distribution of $\hat{\tau}$; hence, for any $0 \leq \underline{\sigma}_{\tau a} < \overline{\sigma}_{\tau a} \leq \sigma_\tau$, the distribution of $\hat{\tau}$ when $\sigma_{\tau a} = \overline{\sigma}_{\tau a}$ is a mean-preserving spread of that when $\sigma_{\tau a} = \underline{\sigma}_{\tau a}$.

Finally, beliefs are uninformative about $\hat{\tau}$ when $R_{\tau a}^2 = 0$ because that corresponds to the case where $\sigma_{\tau a} = 0$, in which case $\hat{\tau}(a)$ is constant across actions. Beliefs are fully informative about $\hat{\tau}$ when $R_{\tau a}^2 = 1$ because that corresponds to the case where $\sigma_{\tau a} = \sigma_\tau$, in which case the two variables are perfectly correlated and the conditional distribution of τ given a is degenerate at a single point. *Q.E.D.*

Lemma 9. *Consider the LQE specification. Given linear market beliefs $\hat{\tau}(a) = La + K$ with $L > 0$ (Equation 2), the agent's unique optimal action is given by $a = \eta + sL\gamma$ (Equation 3).*

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, an increasing linear equilibrium for dimension of interest $\tau \in \{\eta, \gamma\}$ is characterized by (2) and (3), where $L > 0$ solves

$$L = \frac{\mathcal{L}(s, L, \tau)\sigma_\tau^2 + \mathcal{L}(s, L, \neg\tau)\rho\sigma_\eta\sigma_\gamma}{\sigma_\eta^2 + s^2L^2\sigma_\gamma^2 + 2sL\rho\sigma_\gamma\sigma_\eta}. \tag{SA.8}$$

Proof of Lemma 9. First consider the agent's best response to any linear market belief, $\hat{\tau}(a) = La + K$ with $L > 0$. For any η, γ, the agent solves

$$\max_{a \in \mathbb{R}} s(La + K) - \frac{(\max\{a - \eta, 0\})^2}{2\gamma}.$$

This maximand is globally weakly concave in a, and strictly concave on $a \geq \eta$. Since $L > 0$, it is suboptimal for the agent to choose $a \leq \eta$, and so the first order condition yields the unique optimum $a = \eta + sL\gamma$.

SA-14

Now consider the market posterior given a strategy $a = \eta + sL\gamma$. Plugging $l_\eta = 1$ and $l_\gamma = sL$ into Lemma 6 parts 1 and 2 gives the stated expression.

By Theorem 8 of Gómez et al. (2003), it follows that for any observed action a, the marginal distribution of the market posterior on dimension $d \in \{\eta, \gamma\}$ is elliptically distributed with mean $\mu_{d|a}$ given by

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

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

Q.E.D.

SA.3.2. Proof of Proposition 4

Part 1 of Proposition 4 is subsumed in Lemma 10 below; part 2(a) and 2(b) of Proposition 4 are immediate consequences of expression (SA.11) and Lemma 11 below; and part 3 of Proposition 4 is Lemma 12 below. Given $\tau = \eta$, it will be convenient to rewrite Equation SA.8 as

$$f_\eta(L, s, \sigma_\eta, \sigma_\gamma, \rho) := s^2\sigma_\gamma^2 L^3 + 2s\rho\sigma_\gamma\sigma_\eta L^2 + (\sigma_\eta^2 - s\rho\sigma_\eta\sigma_\gamma)L - \sigma_\eta^2 = 0. \tag{SA.9}$$

Lemma 10. *If $\rho \geq 0$ then there is unique solution to Equation SA.9 on the non-negative domain; this solution satisfies $L \in (0, 1)$ and $\frac{\partial f_\eta}{\partial L} > 0$ at the solution.*

Proof of Lemma 10. 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 SA.9 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 SA.9 that $L \to 1$.*

Proof of Lemma 11. The result for $s \to 0$ is immediate from Equation SA.9, so we only prove the limits as $s \to \infty$. Assume $\rho \geq 0$. Divide Equation SA.9 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 SA.9 by dividing by Ls as

$$s\sigma_\gamma^2 L^2 + 2L\rho\sigma_\gamma\sigma_\eta - \rho\sigma_\eta\sigma_\gamma - \frac{\sigma_\eta^2}{Ls}(1 - L) = 0. \tag{SA.10}$$

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 (SA.10) above vanish, which means the LHS of (SA.10) goes to something strictly negative, a contradiction. Hence, $Ls \to \infty$. This implies the LHS of (SA.10) 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) $\frac{d}{d\mu_\gamma}R_{\eta a}^2 = 0$, (ii) $\frac{d}{d\sigma_\gamma}R_{\eta a}^2 < 0$, (iii) $\frac{d}{ds}R_{\eta a}^2 < 0$, and (iv) $\frac{d}{d\rho}R_{\eta a}^2 > 0$.*

Proof of Lemma 12. First observe that Equation SA.9 implies

$$2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2 = \frac{\sigma_\eta}{L}(\sigma_\eta + sL\rho\sigma_\gamma).$$

We now rewrite Equation SA.6 and plug in the above identity to get

$$
\begin{aligned}
R_{\eta a}^2 &= \frac{(\sigma_\eta + sL\rho\sigma_\gamma)^2}{2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2} \\
&= \frac{(\sigma_\eta + sL\rho\sigma_\gamma)^2}{\frac{\sigma_\eta}{L}(\sigma_\eta + sL\rho\sigma_\gamma)} \\
&= L + sL^2\rho\frac{\sigma_\gamma}{\sigma_\eta}.
\end{aligned}
\tag{SA.11}
$$

For the first part, note from (SA.11) that $R_{\eta a}^2$ depends neither directly nor indirectly (through the solution L to Equation SA.9) 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 $x = \sigma_\gamma, s, \rho$,

$$\frac{\mathrm{d}}{\mathrm{d}x}R_{\eta a}^2 = \frac{\partial R_{\eta a}^2}{\partial L}\frac{\partial L}{\partial x} + \frac{\partial R_{\eta a}^2}{\partial x} = \frac{1}{\frac{\partial f_\eta}{\partial L}}\left(\frac{\partial R_{\eta a}^2}{\partial x}\frac{\partial f_\eta}{\partial L} - \frac{\partial R_{\eta a}^2}{\partial L}\frac{\partial f_\eta}{\partial x}\right). \tag{SA.12}$$

Each partial derivative in (SA.12) can be computed from either (SA.9) or (SA.11), and some manipulations then yield parts (ii), (iii), and (iv) of the result. In more detail:

$$\frac{\partial R_{\eta a}^2}{\partial L} = 1 + 2sL\rho\frac{\sigma_\gamma}{\sigma_\eta}, \ \frac{\partial R_{\eta a}^2}{\partial \sigma_\gamma} = \frac{sL^2\rho}{\sigma_\eta}, \ \frac{\partial R_{\eta a}^2}{\partial s} = L^2\rho\frac{\sigma_\gamma}{\sigma_\eta}, \ \frac{\partial R_{\eta a}^2}{\partial \rho} = sL^2\frac{\sigma_\gamma}{\sigma_\eta},$$

and

$$
\begin{aligned}
\frac{\partial f_\eta}{\partial L} &= \sigma_\eta^2 + 3L^2s^2\sigma_\gamma^2 + (4L-1)\rho s\sigma_\gamma\sigma_\eta, \\
\frac{\partial f_\eta}{\partial \sigma_\gamma} &= Ls\left(2L^2s\sigma_\gamma + (2L-1)\rho\sigma_\eta\right), \\
\frac{\partial f_\eta}{\partial s} &= L\sigma_\gamma\left(2L^2s\sigma_\gamma + (2L-1)\rho\sigma_\eta\right), \\
\frac{\partial f_\eta}{\partial \rho} &= L(2L-1)s\sigma_\gamma\sigma_\eta.
\end{aligned}
$$

Since $\frac{\partial f_\eta}{\partial L} > 0$ at the solution to L (Lemma 10), we drop the $\frac{1}{\frac{\partial f_\eta}{\partial L}}$ term in (SA.12) and compute

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\sigma_\gamma}R_{\eta a}^2 &\propto \frac{Ls}{\sigma_\eta}\left(-L^3\rho s^2\sigma_\gamma^2 + (1-L)\rho\sigma_\eta^2 + Ls\sigma_\gamma\sigma_\eta\left(\rho^2 - 2L\right)\right) \\
&= \frac{Ls}{\sigma_\eta}\left(\rho\left(2L^2\rho s\sigma_\gamma\sigma_\eta + L\left(\sigma_\eta^2 - \rho s\sigma_\gamma\sigma_\eta\right) - \sigma_\eta^2\right) - (L-1)\rho\sigma_\eta^2 + Ls\sigma_\gamma\sigma_\eta\left(\rho^2 - 2L\right)\right) \\
&= -2L^3\left(1 - \rho^2\right)s^2\sigma_\gamma < 0,
\end{aligned}
$$

where the first line obtains from plugging in the formulae for partial derivatives into (SA.12) and some algebraic manipulation, the second line obtains from substituting in for $-L^3\rho s^2\sigma_\gamma^2$ from Equation SA.9, and

SA-16

third line is algebraic simplification.

Analogous steps prove $\frac{\mathrm{d}R_{\eta a}^2}{\mathrm{d}s} < 0$. Finally, plugging in partial derivatives into (SA.12) and simplifying,

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

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

SA.3.3. Proof of Proposition 5

Part 1 of Proposition 5 is subsumed in Lemma 13 below; Part 2(a) and 2(b) of Proposition 5 are immediate consequences of expression (SA.14) and Lemma 14 below; and Part 3 of Proposition 5 is Lemma 15 below. Given $\tau = \gamma$, it will be convenient to rewrite Equation SA.8 as

$$f_\gamma(L, s, \sigma_\gamma, \sigma_\eta, \rho) := s^2\sigma_\gamma^2 L^3 + 2s\rho\sigma_\gamma\sigma_\eta L^2 + (\sigma_\eta^2 - s\sigma_\gamma^2)L - \rho\sigma_\eta\sigma_\gamma = 0. \tag{SA.13}$$

Lemma 13. *If $\rho > 0$ there is a unique solution to Equation SA.13 on the non-negative domain; it is strictly positive. If $\rho = 0$, (i) $L = 0$ is always a solution, and (ii) there is a 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 SA.13, it holds that $\frac{\partial f_\gamma}{\partial L} > 0$.*

Proof of Lemma 13. Assume $\rho \geq 0$. First, there is at most one strictly positive solution to Equation SA.13 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 SA.13.

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 SA.13. 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 SA.13 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 SA.13 that $L \to \rho\sigma_\gamma$.*[33]

Proof of Lemma 14. The proof is omitted as the argument is analogous to that in the proof of Lemma 11, but applied to Equation SA.13.   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 SA.13): (i) $\frac{\mathrm{d}}{\mathrm{d}\mu_\eta}R_{\gamma a}^2 = 0$, (ii) $\frac{\mathrm{d}}{\mathrm{d}\sigma_\eta}R_{\gamma a}^2 < 0$, (iii) $\frac{\mathrm{d}}{\mathrm{d}s}R_{\gamma a}^2 > 0$, and (iv) $\frac{\mathrm{d}}{\mathrm{d}\rho}R_{\gamma a}^2 > 0$.*

---

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

Proof of Lemma 15. First observe that Equation SA.13 implies

$$2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2 = \frac{\sigma_\gamma}{L}(\rho\sigma_\eta + sL\sigma_\gamma).$$

We now rewrite Equation SA.7 and plug in the above identity to get

$$R_{\gamma a}^2 = \frac{(\rho\sigma_\eta + sL\sigma_\gamma)^2}{2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2} = \frac{(\rho\sigma_\eta + sL\sigma_\gamma)^2}{\frac{\sigma_\gamma}{L}(\rho\sigma_\eta + sL\sigma_\gamma)}$$

$$= sL^2 + L\rho\frac{\sigma_\eta}{\sigma_\gamma}. \tag{SA.14}$$

The rest of the argument is analogous to that in the proof of Lemma 12, but applied to expressions (SA.14) and (SA.13). The algebraic details are available on request. *Q.E.D.*

SA.3.4. Mixed dimensions of interest

Here we provide formal results for the LQE specification with mixed dimensions of interest discussed in Subsection 4.3. 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 $\hat\eta(a) = L_\eta a + K_\eta$ with $\kappa L_\gamma + (1-\kappa)L_\eta > 0$, the agent's unique best response is to play $a = \eta + s\left(\kappa L_\gamma + (1-\kappa)L_\eta\right)\gamma$. Using the shorthand $L \equiv \kappa L_\gamma + (1-\kappa)L_\eta$, we have $a = \eta + sL\gamma$. Routine substitutions imply that an equilibrium is now characterized by a pair of constants, (L_η, L_γ), that simultaneously solve

$$L_\eta = \frac{\sigma_\eta^2 + sL\rho\sigma_\eta\sigma_\gamma}{2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2}, \tag{SA.15}$$

$$L_\gamma = \frac{sL\sigma_\gamma^2 + \rho\sigma_\eta\sigma_\gamma}{2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2}, \tag{SA.16}$$

An equilibrium is said to be increasing if both $L_\eta > 0$ and $L_\gamma > 0$. From Equation SA.6,

$$R_{\eta a}^2 = \frac{(\sigma_\eta + sL\rho\sigma_\gamma)^2}{2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2} = L_\eta Ls\rho\frac{\sigma_\gamma}{\sigma_\eta} + L_\eta, \tag{SA.17}$$

where the second equality can be verified by substituting for L_η using Equation SA.15. Similarly, Equation SA.7 gives

$$R_{\gamma a}^2 = \frac{(\rho\sigma_\eta + sL\sigma_\gamma)^2}{2sL\rho\sigma_\eta\sigma_\gamma + \sigma_\eta^2 + s^2L^2\sigma_\gamma^2} = L_\gamma\rho\frac{\sigma_\eta}{\sigma_\gamma} + L_\gamma sL, \tag{SA.18}$$

where the second equality can be verified by substituting for L_γ using Equation SA.16.

Proposition 8. *Consider the LQE specification with mixed dimensions of interest and $\rho = 0$. Then, for any $\kappa \in (0,1)$, an increasing linear equilibrium exists; moreover:*

1. $\forall\varepsilon > 0$, $\exists\tilde{s} > 0$ such that if $s < \tilde{s}$ then in any such equilibrium, $R_{\eta a}^2 \in [1-\varepsilon, 1]$ and $R_{\gamma a}^2 \in [0, \varepsilon]$.

2. $\forall\varepsilon > 0$, $\exists\hat{s} > 0$ such that if $s > \hat{s}$ then in any such equilibrium, $R_{\eta a}^2 \in [0, \varepsilon]$ and $R_{\gamma a}^2 \in [1-\varepsilon, 1]$.

3. Any such equilibrium has the following local comparative statics: $\frac{d}{ds}R_{\eta a}^2 < 0$ and $\frac{d}{ds}R_{\gamma a}^2 > 0$.

Remark 5. Subject to technical qualifiers, the above result can be generalized to $\rho \geq 0$; the only modifications are that the limiting $R^2_{\eta a}$ as $s \to \infty$ becomes ρ^2 (cf. Proposition 4 part 2(b)) while the limiting $R^2_{\gamma a}$ as $s \to 0$ becomes ρ^2 (cf. Proposition 5 part 2(a)). ‖

Proof of Proposition 8. Assume $\rho = 0$ and $\kappa \in (0,1)$. With $\rho = 0$, Equations SA.15 and SA.16 can be rewritten as

$$f_{\kappa,\eta}(L_\eta, L_\gamma, s, \sigma_\eta, \sigma_\gamma) := L_\eta \sigma_\eta^2 + s^2 L_\eta L^2 \sigma_\gamma^2 - \sigma_\eta^2 = 0, \tag{SA.19}$$

$$f_{\kappa,\gamma}(L_\eta, L_\gamma, s, \sigma_\eta, \sigma_\gamma) := L_\gamma \sigma_\eta^2 + s^2 L_\gamma L^2 \sigma_\gamma^2 - sL\sigma_\gamma^2 = 0, \tag{SA.20}$$

and (SA.17) and (SA.18) simplify to

$$R^2_{\eta a} = L_\eta, \tag{SA.21}$$

$$R^2_{\gamma a} = sL_\gamma L. \tag{SA.22}$$

It is straightforward from Equation SA.19 and Equation SA.20 that there is a positive solution (i.e. both $L_\eta > 0$ and $L_\gamma > 0$);[34] moreover, in any solution, $L_\eta \in (0,1)$.

Manipulating (SA.19) and (SA.20) 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 SA.21 and Equation SA.22. To prove part 3, first note that by the implicit function theorem,

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

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

$$\frac{\partial L_\eta}{\partial s} = -\frac{2s\sigma_\gamma^2 L_\eta L^2}{s\sigma_\gamma^2 (3sL^2 - \kappa) + \sigma_\eta^2} \quad \text{and} \quad \frac{\partial L_\gamma}{\partial s} = \frac{\sigma_\gamma^2 L (2sL_\gamma L - 1)}{s\sigma_\gamma^2 (3sL^2 - \kappa) + \sigma_\eta^2}. \tag{SA.23}$$

Using (SA.22) and (SA.23), some algebra yields

$$\frac{\mathrm{d}}{\mathrm{d}s} R^2_{\gamma a} = L_\gamma L + s\frac{\partial L_\gamma}{\partial s} L + s\frac{\partial L}{\partial s} L_\gamma$$

$$= \frac{L\left[-s\sigma_\gamma^2 L_\gamma (2sL^2 - \kappa) + s\sigma_\gamma^2 (-\kappa L_\eta + sL_\gamma L^2 + L_\eta) + \sigma_\eta^2 L_\gamma\right]}{s\sigma_\gamma^2 (3sL^2 - \kappa) + \sigma_\eta^2}. \tag{SA.24}$$

---

[34] Note that this uses $\kappa < 1$ (cf. Proposition 5).

We manipulate (SA.24) as follows:

$$\text{Numerator of (SA.24)} = L\left[-s^2L^2\sigma_\gamma^2 L_\gamma + \sigma_\eta^2 L_\gamma + sL\sigma_\gamma^2\right] \quad \text{using } L = \kappa L_\gamma + (1-\kappa)L_\eta$$

$$= 2\sigma_\eta^2 L_\gamma L \quad \text{as Equation SA.20 implies } -s^2 L_\gamma L^2\sigma_\gamma^2 + sL\sigma_\gamma^2 = L_\gamma\sigma_\eta^2$$

$$> 0.$$

$$\text{Denominator of (SA.24)} = s\sigma_\gamma^2\left(sL^2 - \frac{L}{L_\gamma}\right) + \sigma_\eta^2 + s\sigma_\gamma^2\left(\frac{L}{L_\gamma} - \kappa + 2sL^2\right)$$

$$= s\sigma_\gamma^2\left(\frac{L}{L_\gamma} - \kappa + 2sL^2\right) \quad \text{as Equation SA.20 implies } s\sigma_\gamma^2\left(sL^2 - \frac{L}{L_\gamma}\right) + \sigma_\eta^2 = 0$$

$$= \frac{s\sigma_\gamma^2}{L_\gamma}\left((1-\kappa)L_\eta + 2L_\gamma sL^2\right) > 0.$$

Consequently, $\frac{\mathrm{d}}{\mathrm{d}s}R_{\gamma a}^2 > 0$. Finally, observe from (SA.21) and (SA.23) that

$$\frac{\mathrm{d}}{\mathrm{d}s}R_{\eta a}^2 = \frac{\mathrm{d}}{\mathrm{d}s}L_\eta \propto -s\sigma_\gamma^2\left(3sL^2 - \kappa\right) - \sigma_\eta^2 < 0. \qquad\qquad Q.E.D.$$

SA.4. Proofs for Subsection 5.2

The proof of Corollary 1 follows immediately from the results of Proposition 1.

SA.4.1. Proof of Proposition 6

Let $\theta_1 = (\underline{\eta}, \overline{\gamma})$ and $\theta_2 = (\overline{\eta}, \underline{\gamma})$, with $\underline{\eta} < \overline{\eta}$ and $\underline{\gamma} < \overline{\gamma}$. 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.[35]

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 is only played by a single type, and so has degenerate beliefs of $\hat\eta = \underline{\eta}$ or $\hat\eta = \overline{\eta}$. 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 $\hat\eta_2 = \overline{\eta}$, 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 4, $sv(\hat\eta_2) - sv(\hat\eta_1) \leq C(a^{\mathrm{ce}}, \theta_1)$ 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

$$q_1 = \frac{\hat\eta_2 - P}{\hat\eta_2 - \hat\eta_1}, \; q_2 = \frac{P - \hat\eta_1}{\hat\eta_2 - \hat\eta_1}.$$

---

[35] 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 $\underline{\eta}$.

SA-20

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

$$\max_{\hat{\eta}_2, \hat{\eta}_1} \; s\Big(\frac{P - \hat{\eta}_1}{\hat{\eta}_2 - \hat{\eta}_1} w(\hat{\eta}_2) + \frac{\hat{\eta}_2 - P}{\hat{\eta}_2 - \hat{\eta}_1} w(\hat{\eta}_1) - w(P)\Big) \tag{SA.25}$$

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

$$w(\hat{\eta}_2) - w(\hat{\eta}_1) \leq \frac{C(a^{\mathrm{ce}}, \theta_1)}{s}. \tag{SA.26}$$

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 (SA.25), constraint (SA.26) 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 $\alpha < w'(P)$,

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

at any solution at large enough s.

Fixing any $k > \frac{C(a^{\mathrm{ce}}, \theta_1)}{w'(P)}$ and noting that relaxing constraint (SA.26) can only increase the value of the program, it follows that the value of the program (SA.25) for large enough s is no larger than the value of the following program:

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

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

$$\hat{\eta}_2 - \hat{\eta}_1 \leq \frac{k}{s}. \tag{SA.27}$$

Since (SA.27) 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:

$$\max_{\hat{\eta}_2 \in [P, P + \frac{k}{s}]} \varphi(\hat{\eta}_2; s) \equiv s\left(\frac{P - \hat{\eta}_2 + \frac{k}{s}}{\frac{k}{s}} w(\hat{\eta}_2) + \frac{\hat{\eta}_2 - P}{\frac{k}{s}} w\left(\hat{\eta}_2 - \frac{k}{s}\right) - w(P)\right). \tag{SA.28}$$

Claim 5. *The function φ from* (SA.28) *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 of Claim 5. 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

$$\begin{aligned}
s\varphi\left(P + \frac{kx}{s}; s\right) &= s^2\Bigg[(1-x)\left(w(P) + \frac{kx}{s}w'(P) + \frac{k^2x^2}{2s^2}w''(P) + o(s^{-3})\right) \\
&\quad + x\left(w(P) - \frac{k(1-x)}{s}w'(P) + \frac{k^2(1-x)^2}{2s^2}w''(P) + o(s^{-3})\right) - w(P)\Bigg] \\
&= s^2\left((1-x)\frac{k^2x^2}{2s^2}w''(P) + x\frac{k^2(1-x)^2}{2s^2}w''(P) + o(s^{-3})\right) \\
&= \frac{k^2}{2}x(1-x)w''(P) + o(s^{-1}) \\
&\to \frac{k^2}{2}x(1-x)w''(P) \text{ as } s \to \infty. \hspace{4cm} Q.E.D.
\end{aligned}$$

Claim 5 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.

SA.4.2. Proof of Proposition 7

We prove a result that is slightly more general than Proposition 7. 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 6, 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 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. $\max \operatorname{Supp} G_{\underline{\eta}} \le \min \operatorname{Supp} G_{\overline{\eta}}$; 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 7, Condition 1 part 3 is satisfied.

Proposition 9 (Generalization of Proposition 7)**.** *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—i.e., 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 4.

Lemma 16. *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*

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

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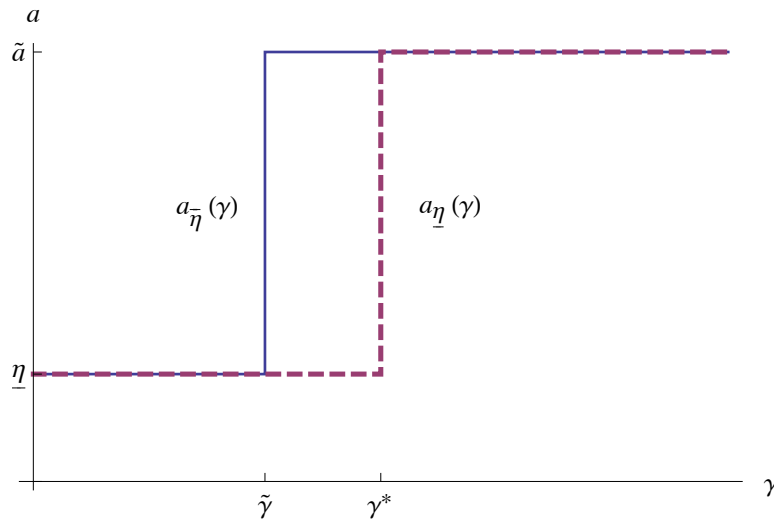

Figure 4 – Strategies under the equilibrium of Lemma 16.

The following conditions are sufficient for the strategy in Lemma 16 to constitute an equilibrium:[36]

$$\begin{cases} sv(\hat{\eta}(\overline{\eta})) = sv(\hat{\eta}(\tilde{a})) - \frac{(\tilde{a}-\overline{\eta})^r}{\tilde{\gamma}} & \text{if } \tilde{a} > \overline{\eta} \\ \tilde{\gamma} = 0 & \text{if } \tilde{a} \le \overline{\eta}, \end{cases} \tag{SA.29}$$

$$sv(\hat{\eta}(\underline{\eta})) = sv(\hat{\eta}(\tilde{a})) - \frac{(\tilde{a}-\underline{\eta})^r}{\gamma^*}. \tag{SA.30}$$

Equation SA.30 requires type $(\underline{\eta}, \gamma^*)$ to be indifferent between actions $\underline{\eta}$ and $\tilde{a}$. Condition SA.29 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 16, 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

$$\Pi(\underline{\eta}) = \frac{pG_{\overline{\eta}}(\tilde{\gamma})}{pG_{\overline{\eta}}(\tilde{\gamma}) + (1-p)G_{\underline{\eta}}(\gamma^*)}, \tag{SA.31}$$

$$\Pi(\tilde{a}) = \frac{p(1 - G_{\overline{\eta}}(\tilde{\gamma}))}{p(1 - G_{\overline{\eta}}(\tilde{\gamma})) + (1-p)(1 - G_{\underline{\eta}}(\gamma^*))}. \tag{SA.32}$$

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 9 for Condition 1 part 2.

For Condition 1 part 3, the conclusion of Proposition 9 follows from the following two claims.

Claim 6. *In a sequence of equilibria of Lemma 16, 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 7. *Assume Condition 1 holds with part 3 satisfied. In a sequence of equilibria of Lemma 16, $\gamma^* - \tilde{\gamma} \to 0$ at a rate of $s^{-\frac{1}{r+1}}$.*

Accordingly, it remains to prove Lemma 16, Claim 6, and Claim 7.

Proof of Lemma 16. 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 (SA.29) and (SA.30). 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>: $s[v(\tilde{\eta}(0)) - v(\underline{\eta})] \le \frac{(\overline{\eta}-\underline{\eta})^r}{\gamma^*}$.

---

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

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 (SA.30). 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 (SA.29) and (SA.30) 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

$$\tilde{a}(\tilde{\gamma}) = \overline{\eta} + \frac{\overline{\eta} - \underline{\eta}}{(\gamma^*/\tilde{\gamma})^{\frac{1}{r}} - 1}. \tag{SA.33}$$

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 (SA.29), as the $(\overline{\eta}, \tilde{\gamma})$ indifference is assured by $\tilde{a} = \tilde{a}(\tilde{\gamma})$. Rearranging (SA.29) gives

$$s\big(v(\hat{\eta}(\tilde{a})) - v(\hat{\eta}(\underline{\eta}))\big) = \frac{(\tilde{a}(\tilde{\gamma}) - \underline{\eta})^r}{\gamma^*}. \tag{SA.34}$$

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 (SA.34) goes to $s(v(\tilde{\eta}(0)) - v(\underline{\eta}))$, while the right-hand side (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 (SA.34) is greater than the RHS. On the other hand, as $\tilde{\gamma} \to \gamma^*$, the LHS of (SA.34) 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 (SA.34) is less than the RHS. By continuity, there exists $\tilde{\gamma} > 0$ for which (SA.34) holds; this value of $\tilde{\gamma}$ together with $\tilde{a}(\tilde{\gamma})$ constitutes an equilibrium. *Q.E.D.*

Proof of Claim 6. 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 7. First, it is routine to verify using conditions (SA.29) and (SA.30) 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 SA.33 implies

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

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 6), and therefore Equation SA.30 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.*