How Wasteful is Signaling?[*]

Alex Frankel[†]   Navin Kartik[‡]

January 20, 2026

Abstract

Signaling is wasteful. But how wasteful? We study the fraction of surplus dissipated in a separating equilibrium. For isoelastic environments, this waste ratio has a simple formula: $\beta/(\beta+\sigma)$, where β is the benefit elasticity (reward to higher perception) and σ is the elasticity of higher types' relative cost advantage. The ratio is constant across types and independent of other parameters, including convexity of cost in the signal. A constant waste ratio characterizes the isoelastic class. In winner-take-all signaling tournaments with N candidates, exactly $(N-1)/N$ of the surplus dissipates—the same as in Tullock contests.

1. Introduction

Signaling is wasteful. In the canonical Spence (1973) model and its innumerable applications and descendants, agents take costly actions to distinguish themselves from lower types. The resulting separating equilibrium reveals information but necessarily dissipates surplus—a fundamental source of inefficiency under asymmetric information.[1]

But how wasteful is signaling, and what does it depend on? Despite more than 50 years of research, the literature offers no simple answer.

A natural intuition suggests that the magnitude of waste should depend on the *difficulty* of signaling. If signaling costs are highly convex in the action (a "hard" test), agents encounter high marginal costs quickly, which ought to limit total expenditure. This reasoning suggests that policies that make signaling more difficult—via exam difficulty, advertising costs, or certification requirements—could reduce waste. However, notice that while such policies reduce the level of signaling, they also increase the cost of lower signals.

Reducing signaling *stakes* instead—scaling down the benefits of being thought of as a higher type—also reduces the level of signaling, and does indeed reduce signaling costs. But it also lowers signaling benefits. For both difficulty and stakes, then, what is the overall effect on the *waste ratio*, i.e., the proportion of surplus burned through signaling?

Our paper studies the classic continuum-type signaling model used in economics, presented formally in Section 2, and focuses on the essentially unique (fully-)separating equilibrium. Our main result, Theorem 1 in Section 3, has two parts. First, under a standard multiplicative cost structure, the proportion of surplus burned by any type—the waste ratio—is invariant to both difficulty and stakes. Importantly, difficulty captures not just the scale of costs, but also the shape (convexity).

Second, consider a canonical isoelastic class of costs and benefits: the cost for type θ of taking signaling action a is given by $C(a,\theta) = D(a)\theta^{-\sigma}$, while the benefit of being thought of as type $\hat{\theta}$ is $V(\hat{\theta}) = s\hat{\theta}^{\beta}$. Here, $\beta > 0$ is the elasticity of benefits (how steeply rewards rise with perceived type), $\sigma > 0$ is the elasticity of cost "strain" (how quickly higher types' comparative advantage grows), and $D(\cdot)$ and $s > 0$ are the difficulty and stakes respectively.

---

[1] Of course, signaling activities can also generate benefits: education builds human capital; and prosocial behavior brings positive externalities, which signaling can amplify (Bénabou and Tirole, 2006). Our paper focuses on the wasteful component of signaling.

Under such isoelasticity, we find that the waste ratio for all types is the constant

$$W = \frac{\beta}{\beta + \sigma}.$$

Waste thus depends only on β/σ, increasing from zero to one in that fraction. This constant waste ratio avoids issues of aggregation across types and delivers a simple answer to our motivating question.

We then show in Theorem 2 that (under multiplicative costs) waste is constant across types if and only if the costs and benefits satisfy a constant relative elasticity condition. Hence, up to the labeling of types, the isoelastic class is the unique setting for such uniform dissipation. This characterization provides a theoretical foundation for the isoelastic specification.

In Section 4, we apply the waste formula to signaling in tournaments. Under suitable assumptions, a winner-take-all market with N candidates competing for a prize has benefit elasticity $\beta = N - 1$ (from the prize structure) and strain elasticity $\sigma = 1$ (a cost normalization), yielding a waste ratio of $(N-1)/N$. Thus, the fraction of surplus burned increases in N, converging to full dissipation as $N \to \infty$. The $(N-1)/N$ formula is precisely the rent dissipation rate in symmetric Tullock lottery contests, revealing an unexpected connection between signaling and contest theory.

The conclusion, Section 5, discusses implications, interpretations, and limitations.

Related Literature. The costly signaling literature in economics, surveyed by Riley (2001) and Sobel (2009), emphasizes the conditions for separating equilibria and that information revelation entails surplus dissipation. However, we are aware of virtually no work that systematically analyzes this waste. One exception is Bernheim and Bodoh-Creed (2023), who show that dissipation can vanish when agents have heterogeneous bliss points and choose many actions (or, equivalently, costs are scaled up). Our paper instead quantifies waste in the canonical signaling setting with homogeneous bliss points and non-negligible distortions.

In the biological signaling literature, Nöldeke and Samuelson (1999) show that offspring's equilibrium cost is proportional to parents' fitness loss, with a constant depending only on genetic relatedness. Their analysis does not yield a constant waste ratio (cost relative to sender's benefit, which need not track parental loss), and their assumption of a linear cost precludes questions about signaling difficulty. But our paper shares with them a common theme that given some structure, dissipation can admit a simple formula based on primitive parameters, with certain invariance properties. By contrast, Bergstrom, Számadó and Lach-

mann (2002) point out that without structure, little can be said about the equilibrium level of signaling costs.

We discuss some other literature connections later in the paper.

2. Model

An agent has type $\theta \in \Theta := [0, \overline{\theta})$, where $0 < \overline{\theta} \leq \infty$.[2] The type is drawn from a continuous cumulative distribution F with support Θ. After privately learning his type, the agent chooses a publicly observable signal or action $a \in \mathbb{R}_{\geq 0}$. An observer sees the action and forms her belief $\hat{\theta} \in \Theta$ about the agent's type.[3] The agent's payoff is $V(\hat{\theta}) - C(a, \theta)$. We maintain throughout the following assumption (primes and subscripts on functions denote derivatives in the usual manner).

Assumption 1. *The benefit function* $V : \Theta \to \mathbb{R}_{\geq 0}$ *and cost function* $C : \mathbb{R}_{\geq 0} \times \Theta \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ *respectively satisfy:*

1. *V is twice differentiable, with $V(0) = 0$ and $V'(\theta) > 0$ for all $\theta > 0$.*

2. *On $\mathbb{R}_{\geq 0} \times (0, \overline{\theta})$, C is finite and continuous with $C(0, \theta) = 0$; on $\mathbb{R}_{>0} \times (0, \overline{\theta})$, C is differentiable with $C_a > 0$, and C_a is continuously differentiable with $C_{a\theta} < 0$. The lowest type has cost $C(a, 0) = \lim_{\theta \downarrow 0} C(a, \theta)$ for all a.*

Part 1 of Assumption 1 says that agents prefer to be perceived as higher types, with the benefit from the lowest perception normalized to zero. Part 2 says that higher actions are costlier, and higher types have lower marginal costs. While it would be natural for costs to be convex in the action, we don't need to assume that. Part 2 also normalizes $C(0, \theta) = 0$ for all θ, so that the payoff from taking the lowest action and receiving the lowest perception is zero. The technical conditions in the two parts are largely standard; note that we allow for type 0 to have infinite costs for actions $a > 0$ to encompass canonical isoelastic costs, detailed in Section 3.

Equilibrium. We study (fully-)*separating equilibria*. The equilibrium definition is standard and relegated to Appendix A, where Proposition 1 shows that any separating equilibrium can be described by a continuous, strictly increasing agent (pure) strategy $A : \Theta \to \mathbb{R}_{\geq 0}$ that is differentiable for $\theta > 0$ and satisfies $A(0) = 0$. Incentive compatibility requires that

---

[2] We use the notation $[0, x\rangle$ to mean $[0, x]$ if $x < \infty$ and $[0, \infty)$ if $x = \infty$.

[3] As we will focus on separating equilibria, we only need to consider degenerate beliefs on a single type.

each type θ optimally chooses $A(\theta)$ given that the observer correctly inverts the strategy on the equilibrium path, i.e., when beliefs satisfy $\hat{\theta}(a) = A^{-1}(a)$ for $a \in [0, A(\bar{\theta}))$. Off-path beliefs can simply be set to $\hat{\theta}(\cdot) = 0$.

Thus, in a separating equilibrium A, any type θ solves

$$\max_a \, [V(\hat{\theta}(a)) - C(a, \theta)],$$

where $\hat{\theta}(\cdot) = A^{-1}(\cdot)$. For $\theta > 0$, the first-order condition evaluated at the optimal action $A(\theta)$ is

$$C_a(A(\theta), \theta) = V'(\theta) \cdot \hat{\theta}'(A(\theta)) = \frac{V'(\theta)}{A'(\theta)}, \tag{1}$$

where the first equality uses $\hat{\theta}(A(\theta)) = \theta$ and the second uses $\hat{\theta}'(A(\theta)) = 1/A'(\theta)$.

Equation 1 has a simple interpretation. Its left-hand side is the marginal cost of increasing the action; the right-hand side is the marginal benefit of inducing a higher belief scaled by the marginal increase in action required for that higher belief. Together with $A(0) = 0$, Equation 1 defines a boundary-value differential equation in A. There is a unique solution by standard arguments.[4] That solution, which we continue to refer to as just A subsequently, characterizes the unique separating equilibrium (uniqueness is up to the specification of off-path beliefs); sufficiency is verified by Proposition 2 in Appendix A.

The Waste Ratio. To measure signaling inefficiency we define three quantities. The opt-out payoff $U^O(\theta) := V(0) - C(0, \theta) = 0$ is what a type would get if it chose the least-cost action and was perceived as the lowest type. The complete-information payoff $U^{CI}(\theta) := V(\theta) - C(0, \theta) = V(\theta)$ is what a type would get if it revealed itself costlessly. Lastly, $U(\theta) := V(\theta) - C(A(\theta), \theta)$ is a type's separating equilibrium payoff.

Definition 1. The *waste ratio* for type $\theta > 0$ is the fraction of its payoff from costless separation that is dissipated through costly signaling:

$$W(\theta) := \frac{U^{CI}(\theta) - U(\theta)}{U^{CI}(\theta) - U^O(\theta)} = \frac{C(A(\theta), \theta)}{V(\theta)}. \tag{2}$$

We refer to the denominator $V(\theta)$ as *surplus*: it is the payoff that type θ would hypothetically get by verifying her type at zero cost. The numerator $C(A(\theta), \theta)$ is the deadweight

---

[4] More precisely, standard existence and uniqueness results for ordinary differential equations can be applied on $(0, \bar{\theta})$ and extended to the lower boundary by continuity; see the arguments in, for example, Mailath (1987) or Kartik (2009).

loss from signaling. The ratio $W(\theta)$ thus measures the effective "tax" that the separating equilibrium imposes on the agent to secure his surplus.

Note that our definition of waste compares the agent's cost of information revelation relative to a frictionless benchmark in which information is revealed at no private cost. This benchmark is, of course, unachievable. Relatedly, we are not defining waste relative to a pooling equilibrium or any other equilibrium. Waste is also only defined in terms of the agent's private surplus, not necessarily social surplus from information.

Our goal is to understand how the waste ratio (2) depends on the parameters of the signaling environment.[5]

3. The Constant of Dissipation

We hereafter focus on multiplicatively separable costs that are commonplace in signaling models. Formally, we assume that

$$C(a, \theta) = D(a) \cdot S(\theta), \tag{3}$$

where $D : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ and $S : \Theta \to \mathbb{R}_{>0} \cup \{\infty\}$. Here $D(a)$ represents the *difficulty* of action a (relative to other actions) and $S(\theta)$ represents the *strain* experienced by type θ (relative to other types). Assumption 1 part 2 implies (i) $D(0) = 0$ and $D'(a) > 0$ for $a > 0$; and (ii) for $\theta > 0$, we have $S(\theta)$ finite and $S'(\theta) < 0$, while $S(0) = \lim_{\theta \downarrow 0} S(\theta)$. Note that $S(0) = \infty$ corresponds to type 0 facing prohibitive signaling costs for any $a > 0$.[6]

It is also useful to write, without loss,

$$V(\theta) = s \cdot B(\theta),$$

where $s > 0$ represents the agent's *stakes* in signaling and $B : \Theta \to \mathbb{R}_{\geq 0}$.

Definition 2. An *isoelastic environment* is defined by

$$B(\theta) = \theta^{\beta} \quad \text{and} \quad S(\theta) = \theta^{-\sigma},$$

---

for some *benefit elasticity* $\beta > 0$ and *strain elasticity* $\sigma > 0$.

Note that the definition stipulates isoelasticity in B and S, but not D. Our first result says that multiplicative costs ensure the waste ratio is independent of stakes and difficulty, and isoelasticity further implies a constant waste ratio across types.

Theorem 1. *Under multiplicative costs:*

1. *The waste ratio $W(\theta)$ is invariant to stakes (s) and difficulty ($D(\cdot)$).*

2. *In an isoelastic environment, the waste ratio is constant: for any $\theta > 0$, it is*

$$W(\theta) = \frac{\beta}{\beta + \sigma}. \tag{4}$$

Equation 4 is a remarkably simple formula for how much surplus is wasted by signaling. The textbook example (e.g. Fudenberg and Tirole, 1991, p. 329) with $B(\theta) = \theta$ and $C(a, \theta) = a/\theta$ corresponds to $\beta = \sigma = 1$, and so it is precisely 50% of the surplus that is dissipated there. More generally, only the ratio β/σ matters; waste is monotonically increasing in β/σ, ranging all the way from 0 to 1.

These directional effects are intuitive. Higher β means a greater incentive to separate from lower types; the rat race for higher beliefs becomes fiercer and more of the surplus is burned. Conversely, higher σ confers a stronger relative cost advantage to higher types, so separation requires less waste. What is perhaps surprising in Theorem 1 is the irrelevance of stakes s and difficulty $D(\cdot)$. To explain that irrelevance and how isoelasticity delivers a constant waste, we present the theorem's straightforward proof.

Proof of Theorem 1. Substituting $C_a(a, \theta) = D'(a)S(\theta)$ and $V(\theta) = sB(\theta)$ into Equation 1, the separating strategy A satisfies (for $\theta > 0$) the differential equation

$$D'(A(\theta))A'(\theta) = \frac{sB'(\theta)}{S(\theta)}.$$

As the left-hand side is $\frac{d}{d\theta}D(A(\theta))$, integrate from 0 to θ to obtain

$$D(A(\theta)) = s \int_0^\theta \frac{B'(t)}{S(t)} \, dt,$$

using $D(A(0)) = D(0) = 0$. Thus, equilibrium costs are

$$C(A(\theta), \theta) = D(A(\theta))S(\theta) = sS(\theta) \int_0^\theta \frac{B'(t)}{S(t)} \, dt, \tag{5}$$

6

and the waste ratio is

$$W(\theta) = \frac{C(A(\theta), \theta)}{V(\theta)} = \frac{S(\theta)}{B(\theta)} \int_0^\theta \frac{B'(t)}{S(t)} \, dt. \tag{6}$$

Both the stakes s and the difficulty $D(\cdot)$ have canceled, establishing part 1. For part 2, substituting the isoelastic form $B(\theta) = \theta^\beta$ and $S(\theta) = \theta^{-\sigma}$ into (6) and simplifying yields

$$W(\theta) = \theta^{-\sigma-\beta} \int_0^\theta \beta t^{\beta-1+\sigma} \, dt = \theta^{-\beta-\sigma} \frac{\beta}{\beta+\sigma} \theta^{\beta+\sigma} = \frac{\beta}{\beta+\sigma}. \qquad \square$$

While Theorem 1 shows invariance to both stakes and difficulty, there is a sense in which the latter runs deeper. As seen in Equation 5, changes in difficulty do not affect equilibrium costs: agents adjust their actions to any change in D—whether scale or shape—exactly enough to keep equilibrium costs unchanged.[7] Stakes operate differently: higher stakes do increase equilibrium costs, but benefits rise proportionally, leaving the waste ratio unchanged. The following example illustrates.

Example 1. An isoelastic environment with difficulty $D(a) = a^\gamma$ for $\gamma > 0$ yields the following separating equilibrium quantities:

$$A(\theta) = \left(\frac{s\beta}{\beta+\sigma}\right)^{1/\gamma} \theta^{(\beta+\sigma)/\gamma} \quad \text{and} \quad C(A(\theta), \theta) = \frac{s\beta}{\beta+\sigma} \theta^\beta.$$

The difficulty convexity parameter γ affects equilibrium actions, but not equilibrium cost, and hence not waste. Stakes s affect actions and cost, but not waste. ◇

Multiplicative separability of costs is important for Theorem 1 part 1; Appendix C confirms that more generally the waste ratio can either decrease or increase in stakes.[8] Similarly, the isoelastic environment is important for part 2 of the theorem. In fact, up to a normalization of types, the constant-waste property characterizes isoelasticity under multiplicative costs. That is the content of our next result, whose proof is in Appendix B.

---

[7] Equilibrium costs act similarly to payments in mechanism design, and the derivation of Equation 5 is akin to that of the payment identity there (Myerson, 1981). Indeed, based on that link, Appendix D shows that the constant-waste formula (4) under isoleasticity can be recovered by mapping our signaling game to an all-pay auction and exploiting revenue equivalence and order statistics.

[8] Multiplicative costs are immaterial for another invariance: the waste ratio $W(\theta)$ does not depend on the type distribution F. This invariance owes to the well-known property that the separating equilibrium strategy only depends on the support of F. The strategy discontinuity at complete information carries over to waste; in particular, under isoelasticity, waste equals $\beta/(\beta+\sigma)$ for any full-support F, even though it would be zero under complete information.

Theorem 2. *Under multiplicative costs, the waste ratio $W(\theta)$ is constant for all $\theta > 0$ if and only if the benefit function V and the strain function S satisfy*

$$-\frac{d \ln S(\theta)}{d \ln V(\theta)} = \rho \tag{7}$$

for some constant $\rho > 0$. The waste ratio then is $W(\theta) = 1/(1+\rho)$.

Equation 7 is a constant relative elasticity condition: the (log) rate at which strain declines must be proportional to the (log) rate at which benefits increase. When this proportionality holds uniformly across types, the tension between signaling incentives and costs resolves identically for all types.

To see why Theorem 2 characterizes the isoelastic environment up to relabeling types, note that integrating Equation 7 yields $S = \kappa V^{-\rho}$ for some $\kappa > 0$. Since V is strictly increasing and $V(0) = 0$, it can be reparametrized as $V(\tilde{\theta}) = s\tilde{\theta}^{\beta}$ via the change of variables $\tilde{\theta} = (V(\theta)/s)^{1/\beta}$, which then gives $S(\tilde{\theta}) = s^{-\rho}\tilde{\theta}^{-\sigma}$ with $\sigma = \rho\beta$. The constant $s^{-\rho}$ can be absorbed into $D(\cdot)$, yielding the isoelastic form.

We note that the assumption of multiplicative costs cannot be dropped from Theorem 2. Example 2 in Appendix C shows that a non-multiplicative cost can have constant waste even though $-\partial \ln C(a, \theta)/\partial \ln V(\theta)$ varies with θ.

4. A Signaling Tournament

We now apply the constant-waste formula to signaling in a tournament, such as workers competing for a job. Richer models of tournament-like signaling and matching have been studied by Hoppe, Moldovanu and Sela (2009) and Hopkins (2012); they do not focus on quantifying waste.

Specifically, consider a market modeled as a tournament among $N \geq 2$ candidates for a single prize of value or size $s > 0$ (e.g., a promotion or a job offer). Candidates' types are their private information, drawn independently from a common distribution F on $[0, 1]$. Candidates simultaneously choose their observable signaling actions. The prize is awarded to the candidate with the highest perceived type.

If candidate i is perceived as type $\hat{\theta}_i$, her probability of winning is $\mathbb{P}(\hat{\theta}_i > \max_{j \neq i} \hat{\theta}_j)$. In a symmetric separating equilibrium, each candidate's expected benefit is

$$V(\theta) = s\left(F(\theta)\right)^{N-1}.$$

Under the uniform distribution the expected benefit simplifies to $V(\theta) = s\theta^{N-1}$; this is an isoelastic benefit with $\beta = N - 1$. Assuming isoelastic signaling costs with unit strain elasticity $\sigma = 1$ (i.e., $C(a, \theta) = D(a)/\theta$ for some D), Theorem 1 immediately implies a constant waste ratio:

$$W_N(\theta) := \frac{N-1}{N}. \tag{8}$$

This expression is $1/2$ when $N = 2$; it increases in N; and as $N \to \infty$, the waste ratio approaches 1. In other words, greater competition exacerbates waste, with the entire surplus dissipated in the limit.[9]

The Tullock Connection. The formula (8) is exactly the rent dissipation rate in a symmetric Tullock lottery contest with N players (Tullock, 1980). There, each player i chooses effort $x_i \geq 0$ to win a prize of size $s > 0$, with a linear cost and winning probability $x_i / \sum_j x_j$.[10] In the symmetric equilibrium, each player exerts effort $x^* = s(N-1)/N^2$, and total effort is $Nx^* = s(N-1)/N$, hence $(N-1)/N$ of the prize is dissipated.

For an arbitrary strain elasticity $\sigma > 0$ and power distribution $F(\theta) = \theta^k$ with $k > 0$ (which yields benefit elasticity $\beta = k(N-1)$), Theorem 1 implies that the waste ratio (8) generalizes to $k(N-1)/(k(N-1) + \sigma)$. Tullock (1980) showed that in a contest with winning probability $(x_i)^r / \sum_j (x_j)^r$ for $r \in (0, 1]$, the rent dissipation rate is $r(N-1)/N$. We see that when $r \neq 1$, there is an important difference in large markets. As $N \to \infty$, the rent dissipation rate in the Tullock contest asymptotes to r, whereas in the signaling tournament it goes to 1 regardless of σ and k. Intuitively, the noise in a non-lottery Tullock contest can preserve some surplus even with extreme competition; but separating from a dense field of competitors forces full dissipation under signaling.

The two models also differ in their sensitivity to the cost structure. Consider a Tullock lottery contest (so $r = 1$) with an isoelastic cost of effort, $D(x) = x^\gamma$ with $\gamma \geq 1$. As this contest is isomorphic to one with a linear cost but winning probability parameter $r = 1/\gamma$, the rent dissipation rate is now $(1/\gamma)(N-1)/N$. Greater cost convexity (higher γ) reduces rent dissipation by steepening marginal costs, which discourages effort. By contrast, Theorem 1 implies that in the signaling tournament, the difficulty $D(\cdot)$ is irrelevant; separation forces agents to scale their efforts with $D(\cdot)$ exactly enough to leave waste unchanged. Unlike in

---

[9] The formula (8) is reminiscent of auction theory. Indeed, a signaling tournament (with its separating equilibrium) is equivalent to an all-pay auction with N bidders (with its usual symmetric equilibrium). See also Appendix D.

[10] The "lottery" descriptor for the Tullock contest refers to effort entering the winning probability linearly in both the numerator and denominator. Nitzan (1994) and Corchón (2007) survey generalizations of this and many other aspects of contests.

a contest, then, signaling tournaments' waste cannot be reduced by simply making better performance costlier (increasing γ); one must reduce the number of competitors or flatten the prize gradient (the benefit elasticity β).

5. Conclusion

We have proposed the waste ratio—the fraction of a type's surplus dissipated through signaling (in the separating equilibrium)—as a natural measure of signaling inefficiency. Our main result, Theorem 1, has two parts. First, under multiplicative costs, the waste ratio is invariant to signaling stakes (a scale parameter of benefits) and signaling difficulty (the type-independent scale or shape of the cost function). Second, in isoelastic environments, the waste ratio has a simple formula that is constant across types: $\beta/(\beta + \sigma)$, where β is the elasticity of benefits and σ measures signaling strain, or higher types' comparative advantage.

The invariance to difficulty undermines some common intuitions about the inefficiency of signaling. Consider the debate about the difficulty of standardized tests for college admissions. Recent trends favor shorter, less complex tests (such as the digital SAT) to reduce student stress. Conversely, some critics call for harder exams to restore selectivity. Our results suggest that, when viewed through the canonical signaling lens, neither approach may address the underlying waste. Adjusting the difficulty of the test uniformly for all students—whether making it easier or harder—need not change the total resource dissipation; it could merely rescale equilibrium effort while leaving the waste ratio constant.[11] As long as admissions at selective colleges resemble winner-take-all signaling tournaments with many competitors, the process is likely to dissipate significant surplus, regardless of how the testing technology is calibrated.

Of course, college admission itself is not the final prize. There are concerns about the "winner-take-all" nature of the broader society (e.g., Frank and Cook, 1996). For a fixed distribution of underlying types, our isoelastic specification captures the inequality of socio-economic outcomes via the benefit elasticity β. In particular, with value function $V(\theta) = s\theta^\beta$, a higher β corresponds to more inequality via a more convex mapping from types to benefits. Our results show that a higher β—corresponding, perhaps, to the US versus lower inequality countries like Canada or Sweden—goes hand in hand with more waste from signaling.

---

[11] This invariance does rely on a single-dimensional framework; it could break if students also differ in test-taking aptitude separate from underlying ability, which would lead to "muddled information" (Frankel and Kartik, 2019).

The strain elasticity σ also matters for waste. Another approach to reducing waste would be to make signaling instruments more discriminating in the sense of increasing σ. Returning to exams, a redesigned test that amplifies high-ability students' comparative advantage—rather than scaling difficulty uniformly—would correspond to increasing σ and would indeed reduce waste. Interestingly, this echoes a discussion in the biological signaling literature: animals can often reliably convey information while incurring minimal waste. The mechanism stems from sharply different marginal costs across types—originally proposed by Zahavi (1977) to refine his earlier "handicap" hypothesis—rather than difficulty. Scholars have argued that because Darwinian selection favors efficiency, it leads to biological signals that are cheap for high-quality types but prohibitive for low-quality types (Penn and Számadó, 2020), corresponding to a high strain elasticity σ.

We close by commenting on some limitations of our analysis. First, we only study separating equilibria. That is consistent with much of the literature's emphasis, often justified by stability-based arguments (e.g., Cho and Sobel, 1990). But there are also critiques of the exclusive focus on separating equilibria (e.g., Mailath, Okuno-Fujiwara and Postlewaite, 1993). Equilibria with some pooling, where certain types choose identical actions, can reduce signaling costs and waste.

Second, our waste ratio only factors in the cost of signaling relative to the *private* benefits from being thought of as a higher type. The social value of learning agents' types may be different from agents' private value—higher, lower, or even zero.

Third, we focus on the dissipative cost of signaling activities. Such activities can, of course, sometimes be productive; see fn. 1. We expect that even in a broader welfare calculus, our waste ratio is a useful input: the net social value must weigh signaling's intrinsic benefit against its waste.

Lastly, Theorem 2 indicates that the waste ratio is less tractable outside the isoelastic class. It will then generally vary with type—so aggregation will depend on the type distribution—and on the full functional forms of costs and benefits. We do not suggest that isoelasticity should be taken literally. Rather, we view it as focal by analogy to how CRRA utility is canonical not because preferences literally exhibit constant relative risk aversion, but because it affords tractable analyses and scale-free results. We hope the waste formula $\beta/(\beta + \sigma)$ is a similarly useful benchmark for signaling's welfare cost. A broader lesson is the invariance of waste to difficulty and stakes under multiplicative costs.

A. Separating Equilibria

A *(mixed) strategy* for the agent is a measurable mapping $\alpha : \Theta \to \Delta(\mathbb{R}_{\geq 0})$, where $\Delta(\mathbb{R}_{\geq 0})$ denotes the set of probability distributions over actions. A *pure strategy* is a strategy α such that $\alpha(\theta)$ has singleton support for all θ; we denote a pure strategy more simply by $A : \Theta \to \mathbb{R}_{\geq 0}$. Since a belief concentrated on type 0 is the "worst belief" (by monotonicity of the benefit function V), and hence is the most conducive off-path belief to support an equilibrium, we say that strategy α defines a *separating equilibrium* if:

1. (Separation.) For any $\theta \neq \theta'$, the distributions are mutually singular: $\alpha(\theta) \perp \alpha(\theta')$.[12]

2. (Incentive compatibility.) For each $\theta \in \Theta$, $\alpha(\theta)$-a.e. a, and all $a' \in \mathbb{R}_{\geq 0}$:

$$V(\theta) - C(a, \theta) \geq V(\hat{\theta}(a')) - C(a', \theta),$$

where $\hat{\theta}(a') = \theta'$ if $a' \in \text{supp}(\alpha(\theta'))$ for a unique θ', and $\hat{\theta}(a') = 0$ otherwise.[13]

Proposition 1. *Any separating equilibrium is a pure-strategy equilibrium. Moreover, its strategy $A : \Theta \to \mathbb{R}_{\geq 0}$ is continuous and strictly increasing on $[0, \overline{\theta}\rangle$, differentiable on $(0, \overline{\theta}\rangle$, and satisfies*

$$A'(\theta) = \frac{V'(\theta)}{C_a(A(\theta), \theta)} \quad \text{for } \theta > 0, \tag{9}$$

with boundary condition $A(0) = 0$.

Although we are not aware of an existing result that directly implies Proposition 1, the proof follows familiar lines (cf. Mailath, 1987) and is provided in the Supplementary Appendix. It is worth noting that because Proposition 1 only establishes necessary conditions for a separating equilibrium, it does not require all of Assumption 1; in particular, it is enough that V is differentiable (rather than twice differentiable), that C_a is continuous (rather than differentiable), and that $C(\cdot, \theta)$ is strictly increasing—where finite for type 0—for all θ (it is not necessary that $C_{a\theta} < 0$). The additional properties are instead used to verify sufficiency in the next result, whose proof—which largely follows Mailath (1987)—is also in the Supplementary Appendix.

---

[12] Recall that two distributions are mutually singular if each assigns probability one to a set that the other assigns probability zero. In other words, the two types choose distinct actions with probability one.

[13] Note that the belief is stipulated as zero for any action in the support of multiple types' mixtures. This is for convenience; it is justified because Bayes' rule (formally, $\hat{\theta}(\cdot)$ being the expectation under a regular conditional distribution) only pins down beliefs μ-a.e., where μ is the marginal distribution over actions induced by the prior F and strategy α. Mutual singularity from point 1 of the definition ensures that for each θ and $\alpha(\theta)$-a.e. action a, the belief $\hat{\theta}(a) = \theta$.

Proposition 2. *A continuous function $A : \Theta \to \mathbb{R}_{\geq 0}$ that is differentiable on $(0, \bar{\theta})$ and satisfies* (9) *and $A(0) = 0$ constitutes a separating equilibrium.*

B. Proof of Theorem 2

Necessity. In a separating equilibrium with strategy A, type θ's payoff is $U(\theta) := V(\theta) - C(A(\theta), \theta)$. Observe that

$$U'(\theta) = -C_\theta(A(\theta), \theta) = -D(A(\theta))S'(\theta), \tag{10}$$

where the first equality is by the Envelope Theorem, and the second equality is by (3). Note that $U'(\theta) > 0$ for $\theta > 0$ because on this domain $S'(\theta) < 0$ and $A(\theta) > 0$ and hence $D(A(\theta)) > 0$. Consequently, for any $\theta > 0$, it holds that $U(\theta) > 0$ (as $U(0) = 0$) and also $U(\theta) < V(\theta)$ (as $C(A(\theta), \theta) > 0$ because $A(\theta) > 0$).

Suppose $W(\theta) = k$ is constant. Then the definition of W and Equation 3 combine to imply

$$D(A(\theta))S(\theta) = kV(\theta) \tag{11}$$

and

$$U(\theta) = (1 - k)V(\theta). \tag{12}$$

Note that $k \in (0, 1)$ because $U(\theta) \in (0, V(\theta))$ for $\theta > 0$.

Differentiating (12) yields

$$U'(\theta) = (1 - k)V'(\theta). \tag{13}$$

Equating the expressions in (10) and (13) gives

$$-D(A(\theta))S'(\theta) = (1 - k)V'(\theta),$$

into which we substitute for $D(A(\theta))$ from (11) to obtain

$$-\frac{kV(\theta)}{S(\theta)}S'(\theta) = (1 - k)V'(\theta),$$

or after rearranging,

$$-\frac{S'(\theta)}{S(\theta)} \cdot \frac{V(\theta)}{V'(\theta)} = \frac{1 - k}{k}.$$

The left-hand side is $-d\ln S/d\ln V$. The result follows by letting $\rho = (1 - k)/k > 0$.

Sufficiency. Suppose (7) holds for some $\rho > 0$. Integrating yields $S(\theta) = \kappa \cdot V(\theta)^{-\rho}$ for some $\kappa > 0$. Following the proof of Theorem 1, the waste ratio is

$$W(\theta) = \frac{S(\theta)}{V(\theta)} \int_0^\theta \frac{V'(t)}{S(t)} \, dt = \frac{\kappa \, (V(\theta))^{-\rho}}{V(\theta)} \int_0^\theta \frac{V'(t)}{\kappa V(t)^{-\rho}} \, dt = (V(\theta))^{-(1+\rho)} \int_0^\theta V'(t) V(t)^\rho \, dt.$$

Using $V(0) = 0$, we evaluate the integral

$$\int_0^\theta V'(t) V(t)^\rho \, dt = \frac{(V(\theta))^{1+\rho}}{1 + \rho}.$$

Hence $W(\theta) = 1/(1 + \rho)$. $\qquad\square$

C. Non-Multiplicative Cost Examples

C.1. Waste Decreasing in Stakes

The following example shows that when costs are not multiplicative, there can be a constant waste ratio that decreases with stakes.

Example 2. Let $V(\theta) = s\theta$ for $s > 0$, and $C(a, \theta) = a^2/\theta + a^3/\theta^2$. These functions satisfy Assumption 1 but the cost is not multiplicatively separable. Moreover, by contrast to Equation 7, we have

$$-\frac{\partial \ln C(a, \theta)}{\partial \ln V(\theta)} = \frac{\theta + 2a}{\theta + a},$$

which varies with θ. We show that the waste ratio depends on stakes s, but is nevertheless constant across types.

Let $c(s) > 0$ be the unique solution to

$$2c^2 + 3c^3 = s, \tag{14}$$

and consider the strategy $A(\theta) = c(s)\theta$. This linear (hence continuous and differentiable) separating strategy satisfies $A(0) = 0$ and the differential equation (9), because

$$C_a(A(\theta), \theta) A'(\theta) = \left(\frac{2c(s)\theta}{\theta} + \frac{3c(s)^2\theta^2}{\theta^2} \right) c(s) = 2c(s)^2 + 3c(s)^3 = s = V'(\theta),$$

where the penultimate equality is by (14). Hence, by Proposition 2, A constitutes a separating equilibrium.

The waste ratio is

$$W(\theta) = \frac{C(A(\theta), \theta)}{V(\theta)} = \frac{c(s)^2\theta + c(s)^3\theta}{(2c(s)^2 + 3c(s)^3)\theta} = \frac{1 + c(s)}{2 + 3c(s)},$$

which is constant across types but depends on s. In particular, since c is increasing in s (Equation 14) and W is decreasing in c, we see that W is decreasing in s with range $(1/3, 1/2)$. ◇

Example 2 can be generalized to a class of examples in which there is constant waste that is decreasing in stakes. Building on Example 1, consider benefits $V(\theta) = s\theta^\beta$ with $s, \beta > 0$, and costs that are a combination of isoelastic terms:

$$C(a, \theta) = \sum_{i=1}^{n} w_i \cdot a^{\gamma_i} \cdot \theta^{-\sigma_i},$$

with $w_i, \gamma_i, \beta_i > 0$. Assume there is a constant $\alpha > 0$ such that the exponents satisfy

$$\frac{\beta + \sigma_i}{\gamma_i} = \alpha \quad \text{for all } i, \tag{15}$$

and γ_i is not constant across i. Then the separating strategy is $A(\theta) = c(s)\theta^\alpha$, where the coefficient $c(s)$ satisfies

$$\alpha \sum_i w_i \gamma_i c(s)^{\gamma_i} = s\beta.$$

Hence, $c(s)$ is increasing in s. Using the above equation, the waste ratio can be computed as

$$\frac{\beta}{\alpha} \cdot \frac{\sum_i w_i c(s)^{\gamma_i}}{\sum_i w_i \gamma_i c(s)^{\gamma_i}},$$

which is decreasing in c, and hence decreasing in s. Intuitively, higher stakes lead to higher actions, putting more weight on cost terms a^{γ_i} with higher γ_i, which correspondingly have higher σ_i (from (15)); this magnifies higher types' comparative advantage and lowers waste, akin to the waste ratio decreasing in σ in Theorem 1.

C.2. Waste Increasing in Stakes

Now we provide an example in which there is a constant waste that increases in stakes.

Example 3. Let $V(\theta) = s\theta$ for $s > 0$, and $C(a, \theta) = a^2/(\theta + a)$. The idea here is that for small a (and hence small stakes) cost is of the order a^2/θ (giving higher types an advantage), whereas for large a (and hence large stakes) cost is of the order a (which is type independent).

15

Conjecture a linear separating strategy $A(\theta) = c\theta$ for some constant $c > 0$. Since

$$C_a(c\theta, \theta) = \frac{c(2+c)}{(1+c)^2},$$

the differential equation (9) reduces to

$$s = \frac{c^2(2+c)}{(1+c)^2}, \tag{16}$$

which has a unique solution $c > 0$ that is increasing in s.

Waste is

$$W(\theta) = \frac{C(A(\theta), \theta)}{V(\theta)} = \frac{\frac{c^2}{1+c}\theta}{s\theta} = \frac{c^2}{s(1+c)} = \frac{1+c}{2+c},$$

using (16) in the last equality. Thus W is constant across types but increasing in stakes s (since W increases in c, and c increases in s), with range $(1/2, 1)$.   ◇

D. The All-pay Auction Equivalence

Consider the isoelastic environment (with stakes $s = 1$, without loss):

$$V(\hat{\theta}) = \hat{\theta}^\beta \quad \text{and} \quad C(a, \theta) = D(a)\theta^{-\sigma}.$$

For simplicity, assume $\bar{\theta} = 1$; the argument below extends more generally by using a quantile transformation. Define $b(\theta) := D(A(\theta))$ and multiply costs and benefits by θ^σ (which is strategically equivalent) to write the payoff for type θ choosing to mimic type $\hat{\theta}$ as

$$\theta^\sigma \hat{\theta}^\beta - b(\hat{\theta}). \tag{17}$$

This payoff corresponds to that in an all-pay auction. Specifically, consider a symmetric N-bidder ($N \geq 2$) independent private value all-pay auction with bidder value $v := \theta^\sigma$ and type distribution $G(v) := v^\alpha$ on $[0, 1]$, where $\alpha := \beta/(\sigma(N-1))$. Consider a symmetric equilibrium in which each bidder uses a differentiable strictly increasing bidding strategy $\tilde{b}(v)$ with $\tilde{b}(0) = 0$. A bidder with true value v who bids to mimic value $\hat{v}$ wins with probability $G(\hat{v})^{N-1}$ and pays $\tilde{b}(\hat{v})$, and hence has payoff

$$vG(\hat{v})^{N-1} - \tilde{b}(\hat{v}). \tag{18}$$

When $\tilde{b}(v) = b(v^{1/\sigma})$, we have $vG(\hat{v})^{N-1} = \theta^\sigma(\hat{\theta}^\sigma)^{\alpha(N-1)} = \theta^\sigma \hat{\theta}^\beta$ and $\tilde{b}(\hat{v}) = b(\hat{\theta})$, and so

the payoffs (17) and (18) match. Consequently, the two settings become strategically equivalent, and the usual all-pay equilibrium bidding strategy matches the signaling separating equilibrium strategy.

This means that we can also use auction results to derive the waste ratio (4). In the all-pay auction, bids are dissipated. By revenue equivalence, a type-v bidder bids his expected payment in the corresponding second-price auction: the win probability $G(v)^{N-1}$ times the expected second-highest value conditional on winning, $\mathbb{E}\left[\max_{j \neq i} v_j \mid \max_{j \neq i} v_j < v\right]$. Standard order statistics for the power distribution $G(v) = v^\alpha$ yield

$$\mathbb{E}\left[\max_{j \neq i} v_j \mid \max_{j \neq i} v_j < v\right] = v \frac{\alpha(N-1)}{\alpha(N-1)+1} = v \frac{\beta}{\beta + \sigma}.$$

Hence, the equilibrium bidding function is

$$\tilde{b}(v) = G(v)^{N-1} v \frac{\alpha(N-1)}{\alpha(N-1)+1} = v G(v)^{N-1} \frac{\beta}{\beta + \sigma}.$$

A type-v bidder's expected gross value is $vG(v)^{N-1}$, so the dissipation rate is

$$\frac{vG(v)^{N-1} \frac{\beta}{\beta + \sigma}}{vG(v)^{N-1}} = \frac{\beta}{\beta + \sigma}.$$

References

Bénabou, Roland and Jean Tirole, "Incentives and Prosocial Behavior," *American Economic Review*, 2006, *96* (5), 1652–1678.

Bergstrom, Carl T., Szabolcs Számadó, and Michael Lachmann, "Separating Equilibria in Continuous Signalling Games," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 2002, *357* (1427), 1595–1606.

Bernheim, B. Douglas and Aaron L. Bodoh-Creed, "Pervasive Signaling," *Theoretical Economics*, 2023, *18* (1), 163–196.

Cho, In-Koo and Joel Sobel, "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory*, 1990, *50* (2), 381–418.

Corchón, Luis C, "The Theory of Contests: a Survey," *Review of Economic Design*, 2007, *11* (2), 69–100.

Frank, Robert H. and Philip J. Cook, *The Winner-Take-All Society: Why the Few at the Top Get So Much More Than the Rest of Us*, New York: Penguin Books, 1996.

Frankel, Alex and Navin Kartik, "Muddled Information," *Journal of Political Economy*, 2019, *129* (4), 1739–1776.

Fudenberg, Drew and Jean Tirole, *Game Theory*, Cambridge, MA: MIT Press, 1991.

Hopkins, Ed, "Job Market Signaling of Relative Position, or Becker Married to Spence," *Journal of the European Economic Association*, 2012, *10* (2), 290–322.

Hoppe, Heidrun C., Benny Moldovanu, and Aner Sela, "The Theory of Assortative Matching Based on Costly Signals," *Review of Economic Studies*, 2009, *76* (1), 253–281.

Kartik, Navin, "Strategic Communication with Lying Costs," *Review of Economic Studies*, 2009, *76* (4), 1359–1395.

Koutsoupias, Elias and Christos Papadimitriou, "Worst-case equilibria," in "Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)" Springer 1999, pp. 404–413.

Mailath, George, "Incentive Compatibility in Signaling Games with a Continuum of Types," *Econometrica*, 1987, *55* (6), 1349–1365.

_ , **Masahiro Okuno-Fujiwara, and Andrew Postlewaite**, "Belief-Based Refinements in Signalling Games," *Journal of Economic Theory*, 1993, *60* (2), 241–276.

Myerson, Roger B., "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

Nitzan, Shmuel, "Modelling Rent-seeking Contests," *European Journal of Political Economy*, 1994, *10*, 41–60.

Nöldeke, Georg and Larry Samuelson, "How Costly is the Honest Signaling of Need?," *Journal of Theoretical Biology*, 1999, *197* (4), 527–539.

Penn, Dustin J. and Szabolcs Számadó, "The Handicap Principle: How an Erroneous Hypothesis became a Scientific Principle," *Biological Reviews*, 2020, *95* (1), 267–290.

Riley, John G., "Silver Signals: Twenty-Five Years of Screening and Signaling," *Journal of Economic Literature*, 2001, *39* (2), 432–478.

Roughgarden, Tim, *Selfish Routing and the Price of Anarchy*, MIT Press, 2005.

_ , **Vasilis Syrgkanis, and Eva Tardos**, "The Price of Anarchy in Auctions," *Journal of Artificial Intelligence Research*, 2017, *59*, 59–101.

Sobel, Joel, "Signaling Games," in Robert A. Meyers, ed., *Encyclopedia of Complexity and Systems Science*, Springer, 2009, pp. 8116–8143.

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

Tullock, Gordon, "Efficient Rent Seeking," in James M. Buchanan, Robert D. Tollison, and Gordon Tullock, eds., *Toward a Theory of the Rent-Seeking Society*, College Station: Texas A&M University Press, 1980, pp. 97–112.

Zahavi, Amotz, "The Cost of Honesty (Further Remarks on the Handicap Principle)," *Journal of Theoretical Biology*, 1977, *67* (3), 603–605.

Supplementary Appendix

E. Omitted Proofs

The following lemma is used in the proof of Proposition 1.

Lemma 1. $C(\cdot, 0)$ *is strictly increasing on* $\{a \geq 0 : C(a, 0) < \infty\}$.

Proof. Since $C(\cdot, \theta)$ is strictly increasing for $\theta > 0$, the limit definition of $C(\cdot, 0)$ implies it is weakly increasing. To establish strict monotonicity, consider $a'' > a' \geq 0$ with $C(a'', 0) < \infty$. For any $\theta > 0$, we have

$$C(a'', \theta) - C(a', \theta) = \int_{a'}^{a''} C_a(x, \theta)\, dx.$$

For any $x > 0$, the function $C_a(x, \cdot)$ is strictly decreasing on $(0, \bar{\theta}]$ because $C_{a\theta} < 0$ on this domain, and so $L(x) := \lim_{\theta \downarrow 0} C_a(x, \theta)$ exists in $\mathbb{R}_{>0} \cup \{\infty\}$. By monotone convergence,

$$C(a'', 0) - C(a', 0) = \int_{a'}^{a''} L(x)\, dx > 0. \qquad \square$$

Proof of Proposition 1. We proceed in three steps.

Step 1: Any separating equilibrium has a pure strategy.

Consider any separating equilibrium strategy α and type θ. Incentive compatibility requires $C(a, \theta) = C(a', \theta)$ for $\alpha(\theta)$-a.e. a and a', since any such actions induce the same belief θ. By the strict monotonicity of $C(\cdot, \theta)$ for $\theta > 0$ and using Lemma 1 for $\theta = 0$ (noting that this type will never choose an action with infinite cost), we have $a = a'$ for $\alpha(\theta)$-almost every a and a', which implies $\alpha(\theta)$ has singleton support. Hence α is a pure strategy.

Step 2: The boundary condition $A(0) = 0$.

Let A be a separating equilibrium (pure) strategy. If $A(0) > 0$, then $V(0) - C(A(0), 0) < V(0) - C(0, 0)$, contradicting incentive compatibility (IC) for type 0.

Step 3: Continuity, differentiability, and monotonicity.

Let A be a separating equilibrium (pure) strategy. We first show A is continuous. Consider any $\theta^* \in (0, \bar{\theta})$ and let $a^- := \lim_{\theta \uparrow \theta^*} A(\theta)$ and $a^+ := \lim_{\theta \downarrow \theta^*} A(\theta)$ (passing to subsequences if necessary, noting that A is bounded on compact subsets by IC). IC for type θ^*

implies that for all $\theta < \theta^*$, we have

$$V(\theta^*) - C(A(\theta^*), \theta^*) \geq V(\theta) - C(A(\theta), \theta^*).$$

Taking $\theta \uparrow \theta^*$ and using continuity of V and C yields

$$V(\theta^*) - C(A(\theta^*), \theta^*) \geq V(\theta^*) - C(a^-, \theta^*). \tag{19}$$

Applying IC in the reverse direction (for a type $\theta < \theta^*$ to not mimic θ^*) and taking the same limit yields the opposite inequality to (19). Hence $C(A(\theta^*), \theta^*) = C(a^-, \theta^*)$. An analogous argument using types above θ^* gives $C(A(\theta^*), \theta^*) = C(a^+, \theta^*)$. Since $C(\cdot, \theta^*)$ is strictly increasing, this implies $a^- = A(\theta^*) = a^+$, so A is continuous at θ^*. The same argument, but using only a^- or a^+ as applicable, also establishes continuity when $\theta^* \in \{0, \bar{\theta}\}$.

We now establish differentiability. Fix $\theta > 0$ and $\theta' \neq \theta$. Adding the IC inequalities for θ to not mimic θ' and vice-versa, and rearranging, yields

$$C(A(\theta'), \theta) - C(A(\theta), \theta) \geq V(\theta') - V(\theta) \geq C(A(\theta'), \theta') - C(A(\theta), \theta'). \tag{20}$$

Since C is differentiable in its first argument, the mean value theorem implies that the left-hand side of (20) equals $C_a(\tilde{a}, \theta) \cdot (A(\theta') - A(\theta))$ for some $\tilde{a}$ between $A(\theta)$ and $A(\theta')$, and analogously for the right-hand side with $C_a(\tilde{a}', \theta')$ for some $\tilde{a}'$ between $A(\theta)$ and $A(\theta')$. Substituting into (20) and dividing by $\theta' - \theta$ gives

$$C_a(\tilde{a}, \theta) \left(\frac{A(\theta') - A(\theta)}{\theta' - \theta} \right) \geq \frac{V(\theta') - V(\theta)}{\theta' - \theta} \geq C_a(\tilde{a}', \theta') \left(\frac{A(\theta') - A(\theta)}{\theta' - \theta} \right), \tag{21}$$

where the inequalities are written for $\theta' > \theta$ and would flip if $\theta' < \theta$. Either way, take $\theta' \to \theta$: since $\tilde{a}, \tilde{a}' \to A(\theta)$ by continuity of A, and C_a is continuous, both $C_a(\tilde{a}, \theta)$ and $C_a(\tilde{a}', \theta')$ converge to $C_a(A(\theta), \theta) > 0$. Since the middle term of (21) converges to $V'(\theta)$ and the two outer terms share the common factor $\frac{A(\theta') - A(\theta)}{\theta' - \theta}$ with coefficients converging to the same positive limit, we conclude that $A'(\theta)$ exists with

$$A'(\theta) = \frac{V'(\theta)}{C_a(A(\theta), \theta)}.$$

Note that if V' is continuous on $(0, \bar{\theta})$, then A' is continuous on that domain because C_a and A are all continuous on that domain and $C_a(A(\theta), \theta) > 0$ for $\theta > 0$ (noting that $A(\theta) > 0$ for $\theta > 0$ by separation).

Finally, since we have established $A'(\theta) > 0$ for $\theta > 0$ and that A is continuous, it follows that A is strictly increasing on $[0, \overline{\theta}\rangle$. $\qquad\square$

Proof of Proposition 2. Let A have the stated properties. Note that for $\theta > 0$, if $A(\theta) = 0$ then $C_a(0, \theta) > 0$: the derivative must exist and be nonzero to satisfy (9), as $V'(\theta) > 0$, and the derivative cannot be negative because $C(\cdot, \theta)$ is continuous and $C_a(a, \theta) > 0$ for $a > 0$. Using Assumption 1, it follows that for $\theta > 0$, no matter the value of $A(\theta)$, both $V'(\theta) > 0$ and $C_a(A(\theta), \theta) > 0$, and hence (9) implies $A'(\theta) > 0$. By continuity of A, it is strictly increasing on $[0, \overline{\theta}\rangle$. Moreover, A is twice differentiable on $(0, \overline{\theta}\rangle$ since the right-hand side of (9) is differentiable in θ (because Assumption 1 entails V twice differentiable and C_a differentiable on the relevant domain).

We now verify that A defines a separating equilibrium. Since any off-path action is met with belief $\hat{\theta} = 0$, it is strictly worse than action 0. So any type θ can be viewed as only choosing which type $\tilde{\theta}$ to mimic, with payoff

$$\Pi(\theta, \tilde{\theta}) := V(\tilde{\theta}) - C(A(\tilde{\theta}), \theta).$$

For any $\theta > 0$, the first order-condition obviously holds, since A solves (9). Let us verify the second-order condition. Using numeric subscripts for partial derivatives of Π in the usual way, and restricting attention to the domain $\theta > 0$, we have

$$\Pi_{22}(\theta, \tilde{\theta}) = V''(\tilde{\theta}) - C_{aa}(A(\tilde{\theta}), \theta) \cdot (A'(\tilde{\theta}))^2 - C_a(A(\tilde{\theta}), \theta) \cdot A''(\tilde{\theta}). \qquad (22)$$

Totally differentiating (9) with respect to θ yields

$$V''(\theta) = C_{aa}(A(\theta), \theta) \cdot (A'(\theta))^2 + C_{a\theta}(A(\theta), \theta) \cdot A'(\theta) + C_a(A(\theta), \theta) \cdot A''(\theta). \qquad (23)$$

Substituting from (23) into (22) and then evaluating at $\tilde{\theta} = \theta$ yields

$$\Pi_{22}(\theta, \theta) = C_{a\theta}(A(\theta), \theta) \cdot A'(\theta) < 0,$$

where the inequality follows from $A'(\theta) > 0$ and, by Assumption 1, $C_{a\theta} < 0$ on $\mathbb{R}_{>0} \times (0, \overline{\theta}\rangle$. This verifies the second-order condition and hence local optimality for each type $\theta > 0$. It follows from the single-crossing assumption $C_{a\theta} < 0$ that (global) incentive compatibility holds on the domain $(0, \overline{\theta}\rangle$, i.e., for any pair of types in this domain, neither wants to mimic the other.

Finally, we address type 0. Consider an arbitrary other type $\theta > 0$. We must show

$\Pi(\theta, \theta) \geq \Pi(\theta, 0)$ and $\Pi(0,0) \geq \Pi(0,\theta)$. Taking each in turn:

1. Incentive compatibility on $(0, \overline{\theta}\rangle$ implies $\Pi(\theta, \theta) \geq \Pi(\theta, \tilde{\theta})$ for all $\tilde{\theta} > 0$, while continuity of V, A and $C(\cdot, \theta)$ imply $\Pi(\theta, \tilde{\theta}) \to \Pi(\theta, 0)$ as $\tilde{\theta} \to 0$. Hence, $\Pi(\theta, \theta) \geq \Pi(\theta, 0)$.

2. For any $\tilde{\theta} > 0$, the previous point implies $\Pi(\tilde{\theta}, \tilde{\theta}) \geq \Pi(\tilde{\theta}, 0) = 0$, and hence $0 \leq C(A(\tilde{\theta}), \tilde{\theta}) \leq V(\tilde{\theta})$. Since $V(\tilde{\theta}) \to 0$ as $\tilde{\theta} \to 0$, it follows that $\Pi(\tilde{\theta}, \tilde{\theta}) \to 0 = \Pi(0,0)$. Moreover, for any $\tilde{\theta} > 0$ we have $\Pi(\tilde{\theta}, \tilde{\theta}) \geq \Pi(\tilde{\theta}, \theta)$ by incentive compatibility on $(0, \overline{\theta}\rangle$, and $\Pi(\tilde{\theta}, \theta) \to \Pi(0, \theta)$ by the limit property of $C(\cdot, 0)$ in Assumption 1. Hence $\Pi(0,0) \geq \Pi(0, \theta)$. $\qquad\square$