Pervasive Signaling

B.D. Bernheim and A.L. Bodoh-Creed

August 5, 2019

Abstract

We ask how the increasing publicness of private decisions (due in part to social media and other aspects of the internet) affects the total costs people incur as a result of signaling distortions. While pervasive signaling may induce pervasive distortions, it may also permit people to signal successfully while distorting each choice to a smaller degree; hence, the overall impact on signaling costs is unclear. We show that, ironically, for a broad class of environments, a sufficient increase in the number of signaling opportunities allows senders to "live authentically" – i.e., to signal their types at arbitrarily low overall cost.

1 INTRODUCTION

Due in large part to the evolution of the internet, our personal and professional lives are becoming increasingly public. The most visible manifestation of this phenomenon is the explosive growth of online social networks (OSNs). People devote substantial time and effort to curating their postings on platforms such as Facebook, LinkedIn, Twitter, Snapchat, and Weibo, which have become virtually ubiquitous components of social interaction.[1] Less visibly (and some would say more nefariously), companies such as Facebook accumulate detailed knowledge of individuals' online activities through diligent monitoring, and draw inferences that inform advertisements, offers, and even pricing (Angwin et al. [3]).

[1] According to the 2016 General Social Survey, 88.6% of United States residents used at least one OSN; for 18 to 35 year olds, the figure is 95.6%. These statistics reflect aggregate rates of usage in the previous 3 months for respondents with ages between 16 and 74 years. In Q3 of 2018, Facebook alone reported 185 million daily active users in the United States and 242 million monthly active users. Nearly three quarters of those on Facebook use their accounts at least once a day, and more than half do so multiple times (Smith and Anderson [30]). Among OSN users between the ages of 18-49, this online activity consumes an average of roughly 6.5 hours per week (Neilsen [29]).

The increasing publicness of our actions has vastly expanded the scope for signaling. In the social sphere, people have almost unlimited opportunities to shape the way friends, family, and acquaintances perceive them. The same principle applies to the professional sphere.[2] Indeed, OSNs also convey information about pertinent personality characteristics to potential employers (e.g., Buffardi and Campbell [10], Kluemper and Rosen [24]). As a result, users perform a "balancing act between self-expression and self-promotion" (van Dijck [13]).

In this paper, we ask whether the increasing pervasiveness of signaling opportunities is socially helpful or harmful – that is, whether it increases or decreases the *aggregate* efficiency costs of signaling distortions. On the one hand, pervasive signaling leads to pervasive distortions. While the total cost of signaling cannot exceed the benefits of information transmission (otherwise people would not signal), it could in principle dissipate all of those benefits in the limit as signaling activities become more numerous. On the other hand, signaling through multiple activities might prove more efficient, in which case the total cost of signaling could decline in the limit. In principle, it could also remain unchanged.

Within the class of equilibria that treat multiple (identical) signaling activities symmetrically, increasing the number of such activities is equivalent to increasing the unit cost of a single signaling activity. We assume the unit costs are proportional to a parameter λ. Our investigation begins with an examination of the relationship between total signaling costs and the unit costs. In the settings we study, the sender's type is drawn from a continuum of possibilities. The marginal costs of the signaling activity are lower for higher types, and those perceived as higher types receive greater reputational benefits. We assume in addition that the higher types have higher bliss points, a property that is natural for a wide range of applications.[3] We also restrict attention to equilibria satisfying a dominance refinement. Our first main result shows that, ironically, the ratio of the total signaling cost to the total benefit of information transmission converges to zero as the unit cost of the signaling activity becomes arbitrarily large. We also characterize the rate of convergence, which is rapid ($O\left(\frac{1}{\sqrt{\lambda}}\right)$ or, under more stringent restrictions, $O\left(\frac{1}{\lambda}\right)$). Thus, in the limit,

[2]Sixty percent of employed Facebook users are "friends" with coworkers (Dourin et al. [17]).

[3]Signaling models with heterogeneous bliss points are widely employed in the literature. Examples include Spence [31] on signaling with productive educational investments; Mailath [26] on price signaling; Banks [4] on political competition; Miller and Rock [28] on dividend signaling; Bernheim [6] on conformity; Bagwell and Bernheim [5], Ireland [22], and Corneo and Jeanne [12] on conspicuous consumption; Bernheim and Severinov [8] on bequests; Bernheim and Andreoni [2] on fairness; or Bernheim and Bodoh-Creed [7] on decisive leadership.

senders reveal their information costlessly. Our analysis also establishes that this result does *not* hold for signaling environments with homogenous bliss points.

While our first main result may initially seem counterintuitive, consideration of a related class of models renders it less surprising. Suppose in particular that the sender's type is drawn from a *finite* set of possibilities, and that the environment is otherwise unchanged. In that case, when the unit costs of the signaling activity become sufficiently large, truth-telling is an equilibrium: finite reputational gains are insufficient to justify the costs of choosing a higher type's bliss point. However, this argument falls well short of providing a helpful intuitive account of our first result. One limitation is that it does not rule out the existence of other equilibria with different properties. Indeed, we show by way of example that the costs of signaling need not vanish in all equilibria. Another limitation is that the logic of the argument is not robust. In particular, suppose we allow the cardinality of the type space to increase so that it more closely approximates a continuum as the unit cost of the signaling activity grows. In that case, the question of whether truth-telling becomes an equilibrium for large unit costs hinges on the relative rates at which we increase unit costs and add types. If we add types sufficiently fast, then the simple argument breaks down and distortions never disappear; whether they grow or shrink becomes an open issue. Our analysis establishes that, for equilibria satisfying a plausible dominance requirement, total signaling costs vanish in the limit even when we take the type space to be a continuum from the outset.

The next step in our analysis is to study signal proliferation explicitly. We begin by assuming that the costs of signaling are additively separable and symmetric across N activities. We provide weak conditions under which the welfare-optimal separating equilibrium is one in which the sender chooses the same action for each activity, in effect spreading the cost of signaling distortions across the observable activities. Under those conditions, the optimal equilibrium is isomorphic to that of a model with a single signal where the weight attached to signaling costs is proportional to the number of actions. Applying our primary result to this setting, we conclude that the ratio of the total signaling costs to the benefits of information transmission converges to zero in the limit as N grows. Thus, the proliferation of observable activities enables a sender to "live authentically" – that is, to signal the truth at negligible overall cost. We also consider more general utility functions that allow for non-linear aggregation across the sender's activities, and we provide sufficient conditions under which our main conclusions continue to hold.[4]

[4]Such a generalization is necessary to study applications like conspicuous consumption, where the

We also explore the robustness of our conclusions concerning signal proliferation in settings where OSNs also amplify the benefits of signaling by enlarging the pertinent audience. We reach two main conclusions. First, as long as the relationship between the potential benefits of signaling and the size of the audience exhibits decreasing returns of sufficient magnitude, signal proliferation still drives the total cost of signaling to zero in the limit. Second, even in cases where equilibrium signaling costs do not vanish in the limit, they still dissipate a vanishing fraction of the OSN's potential informational benefits, except in cases where the effects of OSNs on the scale of informational benefits are proportionately larger than those involving signal proliferation.

The closest analogs to this paper appear in a branch of the theoretical biology literature concerning the "Handicap Principle," which holds that the credible transmission of information requires a costly signal (Zahavi [36], Grafen [19]). Follow-up papers focused attention on the fact that the credibility of a signal only requires out-of-equilibrium costs of mimicry, which in turn led to the realization that successful signaling in settings with discrete type spaces does not necessitate the use of costly signals in equilibrium (Hurd [21], Számado [33] and [34]). However, this literature also suggests that, except for knife-edge examples, costless signaling is typically impossible in settings where the type space is a continuum (Lachmann, Számado, and Bergstrom [25] and Bergstrom, Számado, and Lachmann [25]).[5]

With respect to the specific phenomenon of signaling through OSNs, there is also a small theoretical literature that examines the informational content of an individual's network connections. Donath [14] and Donath and Boyd [16] study social connections as a credible signal of identity in an otherwise anonymous virtual community, and Donath [15] draws out the implications of these ideas for OSN design. None of these studies focuses on the fact that OSNs contribute to the pervasiveness of signaling.

Section 2 introduces our baseline model, in which the sender signals with a single activity, the unit cost of which depends on an abstract parameter, λ. Section 3 builds intuition for our first main result by solving parametric examples of the baseline model analytically. Section 4 proves the result, and then links the parameter λ to the number of signaling activities in generalized versions of the model. Section 5 applies the results to signaling via OSNs. Section 6 concludes. All proofs appear in Appendix A. For additional

sender's utility is defined over bundles of goods that may be substitutes for one another.

[5]In an evolutionary context, preference parameters are endogenous, and evolutionary pressures may drive them to knife-edge values that would otherwise appear non-generic.

applications to job market signaling and politicians signaling decisiveness (as in Bernheim and Bodoh-Creed [7]), see Appendix B.

2 MODEL

The sender's private information pertains to her type $t \in [\underline{t}, \overline{t}] = T \subset \mathbb{R}$. For our benchmark model, we assume the sender chooses a single action $a \in \mathbb{R}_+$. The receiver observes a and draws inferences about the sender's type. The sender's direct utility from action a in the absence of any signaling incentive (i.e., with complete information) is $\lambda\pi(a, t)$, where $\lambda > 0$ parameterizes the costs of deviating from the sender's *bliss point*, defined as follows:

$$a_{BP}(t) = \underset{a}{arg\,max}\ \lambda\pi(a, t). \tag{1}$$

In words, $a_{BP}(t)$ is the action the sender would choose if her type were publicly observed. If the agent takes an action $a \neq a_{BP}(t)$ in equilibrium, then $\lambda [\pi(a, t) - \pi(a_{BP}(t), t)] < 0$ is the *total cost of signaling*.

Having observed a, the receiver uses Bayes's rule to form a posterior belief about the sender's type. We use $\delta(a) \in \Delta(T)$ to denote the belief of a receiver who observes action a, where $\Delta(T)$ is the set of Borel measures over T. We refer to $\delta(a)$ as the receiver's *perception* of the sender. In the case of fully separating equilibria, the receiver's equilibrium beliefs place probability 1 on the sender having the type $\widehat{t}(a)$, which is derived from the sender's strategy. When convenient, we suppress the arguments of δ and $\widehat{t}$, and refer to the sender as "choosing" the receiver's perception.

Given the receiver's perception, the sender receives benefits $B(t, \delta(a))$.[6] In many signaling models, the receiver responds to the sender's signal, and one can think of $B(t, \delta(a))$ as a reduced form representation of the utility the sender derives from this response. The sender's total utility is

$$U(a, t; \lambda) = B(t, \delta(a)) + \lambda\pi(a, t). \tag{2}$$

We make several assumptions that are easily verified in applications. We assume the existence of all referenced derivatives, and we use subscripts to denote partial derivatives

[6] It is straightforward to allow for the possibility that B depends on the action a. In that case, we would require that the partial derivatives $B_a(a, t, \delta)$, $B_{aa}(a, t, \delta)$, and $B_{aaa}(a, t, \delta)$ are defined for all (a, t, δ), and that $B_{aaa}(a, t, \delta)$ is uniformly bounded from above.

with respect to the subscripted variable. For ease of notation, we also assume that each agent has a unique optimal choice under complete information.

Assumption 1. *$a_{BP}(t)$ is unique.*

To analyze the costs of signaling when the equilibrium strategy approaches $a_{BP}(t)$, we use polynomial approximations of $\pi(a,t)$ around $(a_{BP}(t),t)$. Assumption 2 part (1) allows us to focus on second order effects and ignore 3^{rd} and higher order effects.[7] Part (2) of Assumption 2 generalizes part (1) to cases where (for example) $\pi_{aa}=0$ and $\pi_{aaaa}<0$ — in other words, it allows us to focus on the lowest relevant non-zero derivative and ignore higher order effects.[8]

Assumption 2. *There exists $C<\infty$ such that one of the following holds for all (a,t) in an open neighborhood of $\{(a,t): a=a_{BP}(t)\}$*

1. *For $a=a_{BP}(t)$ we have $\pi_{aa}(a,t)<0$ and $\pi_{aaa}(a,t)\leq C$.*
2. *For all pairs $(a_{BP}(t),t)$ we have $\frac{\partial^i \pi(a,t)}{\partial a^i}=0$ for $i\in\{2,..,k<\infty\}$, $\frac{\partial^{k+1}\pi(a,t)}{\partial a^{k+1}}\neq 0$, and $\frac{\partial^{k+2}\pi(a,t)}{\partial a^{k+2}}\leq C$.*

Our next assumption requires that first-order changes in type result in first-order changes in bliss points.

Assumption 3. *There exists $\beta>0$ such that for any $t>t'$, we have $a_{BP}(t)-a_{BP}(t')\geq \beta(t-t')$.*

Assumption 4 bounds the rate at which the sender's benefit can change with her perceived type in a fully separating equilibrium. Assumption 5 bounds the sender's benefit when the receiver is not confident about the sender's type, which may occur following a deviation from the equilibrium action by the sender. Together they imply that there exists $\underline{B}\leq\overline{B}$ such that $B(t,\delta(a))\in[\underline{B},\overline{B}]$.

Assumption 4. *There is $\gamma>0$ such that $0\leq B_{\hat{t}}(t,\hat{t})\leq\gamma$.*

Assumption 5. *If the support of $\delta(a)$ is $\mathcal{S}$, then $\max_{\widehat{t}\in S} B(t,\widehat{t})\geq B(t,\delta(a))\geq \min_{\widehat{t}\in S} B(t,\widehat{t})$.*

---

[7]Since we analyze actions in the neighborhood of the bliss point, first order effects are absent as $\pi_a(a_{BP}(t),t)=0$.

[8]$\pi_{aa}(a_{BP}(t),t)=0$ implies $\pi_{aaa}(a_{BP}(t),t)=0$, since otherwise $a_{BP}(t)$ would not be a local maximum of $\pi(a,t)$. This point is elaborated in Footnote 23.

Denote the strategy used in the separating equilibrium as $a_{SEP}(t;\lambda)$. We can write the equilibrium utility for an agent of type t who mimics the equilibrium action of an agent of type $\widehat{t}$ as

$$V(\widehat{t},t;\lambda) = B\left(t,\widehat{t}\right) + \lambda\pi\left(a_{SEP}\left(\widehat{t};\lambda\right),t\right).$$

By standard arguments, one separating equilibrium is the solution to the differential equation

$$\left.\frac{\partial a_{SEP}(\widehat{t};\lambda)}{\partial \widehat{t}}\right|_{\widehat{t}=t} = -\frac{1}{\lambda}\left.\frac{\partial B\left(t,\widehat{t}\right)}{\partial \widehat{t}}\right|_{(t,\widehat{t})=(t,t)}\left.\frac{1}{\pi_a(a,t)}\right|_{(a,t)=(a_{SEP}(t;\lambda),t)}. \tag{3}$$

Since the receiver makes the worst possible inference for type $\underline{t}$ senders, senders of this type can do no worse than choose their bliss-point action. This observation implies that Equation 3 has the initial condition $a_{SEP}(\underline{t};\lambda) = a_{BP}(\underline{t})$. One can prove that this solution is, in fact, a separating equilibrium under relatively mild conditions (see Mailath and von Thadden [27]). Instead of invoking sufficient conditions, we focus directly on cases in which a separating equilibrium exists.

Assumption 6. *A fully separating Bayes-Nash equilibrium exists.*

3 EXAMPLES

We illustrate some of our key insights through two examples. For the first, which has heterogeneous bliss points, we show that the difference between the bliss-point and the equilibrium action vanishes so rapidly that the total cost of signaling converges to zero.

Example 1. *Suppose $B(t,\widehat{t}) = \widehat{t}$, $\pi(a,t) = -(a-t)^2$, and $T = [0,1]$. The bliss points are (obviously) $a_{BP}(t) = t$. The ODE defining the fully separating equilibrium is*

$$\left.\frac{\partial a_{SEP}}{\partial \widehat{t}}\right|_{\widehat{t}=t} = \frac{1}{2\lambda(a_{SEP}(t)-t)}.$$

We use the change of variables $z(t) = a_{SEP}(t) - t$. Solving the inverse ODE yields

$$t = -\left[z + \frac{1}{2\lambda}\ln\left(\frac{1}{2\lambda} - z\right)\right] + C.$$

The initial condition $z(0) = 0$ *implies*

$$t = -\left[z + \frac{1}{2\lambda}\ln\left(1 - 2\lambda z\right)\right].$$

Reversing our change of variables and rearranging, we find

$$z(t) = \frac{1 - e^{2\lambda a_{SEP}(t)}}{2\lambda}.$$

The total cost of signaling is then

$$\lambda z(t)^2 \leq \lambda\left(\frac{1}{2\lambda}\right)^2 = \frac{1}{4\lambda},$$

which is $O\left(\lambda^{-1}\right)$ *as claimed.*

In the second example, the agents share the bliss point of $a_{BP} = 0$. As a result, $a_{SEP}(t)$ converges to $a_{BP}(t)$ at the rate $O\left(\lambda^{-0.5}\right)$. This slow convergence causes the total cost of signaling to be bounded away from 0. In fact, we show that while $a_{SEP}(t;\lambda) \to t$ as $\lambda \to \infty$, the total cost of signaling is invariant with respect to λ.

Example 2. *Suppose* $B(t,\hat{t}) = \hat{t}$, $\pi(a,t) = \frac{-a^2}{t+\gamma}$, $\lambda > 0$, *and* $T = [0,1]$. *The bliss points are all homogenous:* $a_{BP}(t) = 0$. *The ODE defining the fully separating equilibrium is*

$$\left.\frac{\partial a_{SEP}}{\partial \hat{t}}\right|_{\hat{t}=t} = \frac{t+\gamma}{2\lambda a_{SEP}(t)}.$$

We can write this equation in the following more convenient form:

$$2\lambda a_{SEP}(t)\left.\frac{\partial a_{SEP}}{\partial \hat{t}}\right|_{\hat{t}=t} = t + \gamma.$$

Integrating both sides and using our initial condition yields

$$\lambda a_{SEP}(t)^2 = \frac{1}{2}(t+\gamma)^2 - \frac{\gamma^2}{2}.$$

Thus, the total cost of signaling, $\lambda a_{SEP}(t)^2$, *is invariant with respect to* λ.

4 MAIN RESULTS

4.1 The case of a single costly activity

Our first result shows that the total cost of signaling with full separation of types vanishes in the limit as $\lambda \to \infty$. Because we hold the benefit of signaling fixed as λ grows, one can interpret this result as indicating that the ratio of total signaling costs to total benefits declines to 0 as λ increases. In the limit, the sender fares as well as with complete information.

Some of our results refer to equilibria that satisfy a dominance refinement. For each action a, we say that sender type t belongs to the *plausible set* if there is some pattern of conceivable receiver reactions (i.e., receiver beliefs about the sender) for which t would find a preferable to $\alpha_{BP}(t)$. The dominance refinement requires that, whenever possible, the receiver does not attribute any action a to any sender type t outside of the plausible set.[9]

Definition 1. *The* ***plausible set*** *for action a is*

$$\mathcal{P}(a) = \{t : \underline{B} + \lambda\pi(a_{BP}(t), t) \leq \overline{B} + \lambda\pi(a, t)\}.$$

The receiver's beliefs satisfy the ***dominance refinement*** *if, upon observing any a for which $P(a) \neq \emptyset$, the receiver is certain that $t \in \mathcal{P}(a)$.*

Let $a(t; \lambda)$ denote an equilibrium action function for the signaling game. Remark 1 highlights the fact that, with fixed bounds on the benefits of signaling, actions must converge to bliss points as λ (which parameterizes the unit costs of deviating from bliss points) grows.[10] In addition to serving as the first step in the proof of our main result, the remark is of independent interest since it applies to *all* equilibria of our signaling model, not just fully separating equilibria.[11]

Remark 1. *If $B\left(t, \widehat{t}\right)$ is bounded from above and below, then $a(t; \lambda_i) \to a_{BP}(t)$ for each t as $\lambda_i \to \infty$. If $a_{BP}(t)$ is continuous, then the convergence is uniform over t.*

Assuming the number of types is *finite* and bliss points are heterogeneous, it is straightforward to show that, as a consequence of Remark 1, the total cost of signaling vanishes completely (i.e., $\lambda\left[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t; \lambda), t)\right] = 0$) for $\lambda < \infty$ sufficiently large in *any*

[9] Our dominance refinement is similar to the one studied in Cho and Kreps [11].
[10] This result also implies that any pools become vanishingly small as λ_i increases.
[11] The interested reader can refer to Lemma 1 in Appendix A for a proof of the remark.

equilibrium that satisfies the dominance refinement. Example 3 exhibits the structure of such an equilibrium. For completeness, we also exhibit an equilibrium that violates the dominance refinement and for which the total costs of signaling do not vary with λ.

Example 3. *Suppose $B(t,\widehat{t}) = \widehat{t}$, $\pi(a,t) = -(a-t)^2$, and $T = \{0,1\}$. In all separating equilibria satisfying the dominance refinement, we have $a_{SEP}(0;\lambda) = 0$ and $a_{SEP}(1;\lambda) = max\{\lambda^{-0.5}, 1\}$. Various equilibrium beliefs are compatible with the refinement; e.g., the receiver could infer the sender is of type $t = 0$ if and only if $a < a_{SEP}(1;\lambda)$.*[12] *For $\lambda > 1$ we have $a_{SEP}(1;\lambda) = 1 = a_{BP}(1)$ since type $t = 0$ is unwilling to choose $a_{BP}(1)$ regardless of the receiver's reactions. Therefore the total cost of signaling vanishes for $\lambda \geq 1$.*

Next we exhibit an equilibrium for which the total cost of signaling does not shrink as λ grows. Specifically, suppose that $a_{SEP}(0;\lambda) = 0$, $a_{SEP}(1;\lambda) = 1 + \lambda^{-0.5}$, and the receiver infers the sender's type is $t = 0$ upon observing any action $a < a_{SEP}(1;\lambda)$. In this case, the type $t = 1$ sender's cost of signaling is $-\lambda(a_{SEP} - 1)^2 = 1$ for all values of λ. However, this equilibrium does not survive our dominance refinement.

In contrast, for the typical model with a continuum of types, there exist no truth-telling equilibria even for large λ. As a result, the existence of an equilibrium for which total signaling costs vanish as λ grows is far from obvious, despite the fact that the equilibrium actions converge to the respective bliss points. Consider Example 1, the continuum analog of Example 3. Since the cost function is quadratic, if we had $a_{SEP}(1;\lambda_i) - a_{BP}(t) = \Theta\left(\frac{1}{\sqrt{\lambda_i}}\right)$, we would then have $\lambda\left[\pi(a_{BP}(t),t) - \pi(a_{SEP}(t;\lambda_i),t)\right] = \Theta(1)$—in other words, the costs would not vanish.[13] Thus, to prove our main result, we need to show that $\|a_{SEP}(t;\lambda_i) - a_{BP}(t)\|$ converges at a sufficiently fast rate.

We provide two versions of our main result. We state both in terms of the rate of convergence of $\lambda\left[\pi(a_{SEP}(t;\lambda),t) - \pi(a_{BP}(t),t)\right]$ to 0. In the proof of each theorem we derive limits on the beliefs of the receiver following any sender action. In Theorem 1, the lower bound on the rate of change of $a_{BP}(t)$ (Assumption 3) allows us to use the dominance refinement to bound the inferences the receiver can make following a type t sender's choice to deviate to $a = a_{BP}(t)$. The bounds on the receiver's inferences combined with our bound on $B(\cdot,\cdot)$ and the rate of change of $B(t,\cdot)$ (Assumptions 4 and 5) yields an upper bound on the benefit of choosing $a_{SEP}(t;\lambda)$ relative to $a_{BP}(t)$ for senders of type t. The

[12] The non-uniqueness stems from the fact that the plausible set can be empty for very large a.

[13] The notation $f(x) = \Theta(g(x))$ means that there exists $k_1, k_2 > 0$ such that $k_1 g(x) \leq f(x) \leq k_2 g(x)$ as $x \to \infty$.

bound on the benefits of signaling one's type relative to choosing one's bliss point implicitly bounds the cost of signaling, and we show this bound implies that the total cost of signaling converges to 0 as $\lambda \to \infty$.

Theorem 1. *Under Assumptions 1 - 5, we have* $\lambda\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t;\lambda),t)\right]=O\left(\frac{1}{\sqrt{\lambda}}\right)$ *for equilibria in which the receiver's beliefs satisfy the dominance refinement.*

The second version of our main result focuses on cases in which both the bliss points and the fully separating equilibrium action function are continuous. If $a_{SEP}(t;\lambda)$ is continuous, then each type t's bliss point is the equilibrium action of some other type $t'<t$. Therefore, when we consider deviations from $a_{SEP}(t;\lambda)$ to $a_{BP}(t)$, we can ignore the issue of identifying the beliefs associated with off-path actions. Theorem 2 establishes a faster rate of convergence than Theorem 1 because we can use the continuity of $a_{SEP}(t;\lambda)$ to iterate our bounding argument and thereby obtain tighter bounds on $\lambda\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t;\lambda),t)\right]$.

Theorem 2. *Under Assumptions 1 - 5, if* $a_{BP}(t)$ *and* $a_{SEP}(t;\lambda)$ *are continuous in* t *for all* λ*, then* $\lambda\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t;\lambda),t)\right]=O\left(\frac{1}{\lambda}\right)$.

4.2 Multiple Signals with Additive Aggregators

We now consider settings where the receiver observes multiple signals from the sender simultaneously. We consider a sequence of models indexed by $N\in\mathbb{N}$, where the N^{th} model features a sender who takes N simultaneous actions $a\in\mathbb{R}$ yielding an action vector $\boldsymbol{a}^N\in\mathbb{R}^N$ with the i^{th} component denoted a_i^N. The utility for the sender from these N actions is $\boldsymbol{\pi}^N(\boldsymbol{a}^N,t)$. Since we focus on separating equilibria, in equilibrium the receiver has a degenerate belief placing probability one on the sender's type being $\widehat{t}\left(\boldsymbol{a}^N\right)$, and the sender receives utility of $B\left(t,\widehat{t}\left(\boldsymbol{a}^N\right)\right)+\boldsymbol{\pi}^N(\boldsymbol{a}^N,t)$.

We start with the simplest case wherein $\boldsymbol{\pi}^N(\boldsymbol{a}^N,t)$ is additively separable and symmetric across actions, which allows us to characterize the sender-optimal equilibrium. In this model, the utility of the sender is

$$U_N(\boldsymbol{a}^N,t)=B\left(t,\widehat{t}\left(\boldsymbol{a}^N\right)\right)+\sum_{i=1}^{N}\pi\left(a_i^N,t\right).$$

Let $a_{BP}(t)=\underset{a}{arg\,max}\ \pi\left(a,t\right)$, so $\boldsymbol{a}_{BP}(t)=(a_{BP}(t),a_{BP}(t),...,a_{BP}(t))\in\mathbb{R}_+^N$ denotes the bliss point of the type t agent.

Consider a symmetric signaling equilibrium – i.e., one in which $a_1(t) = ... = a_N(t) = a_{SEP}(t;N)$. We can write the equilibrium utility of a type-t agent who mimics the action of a type-$\widehat{t}$ agent as

$$V_N(\widehat{t},t) = B\left(t,\widehat{t}\right) + N\pi\left(a_{SEP}\left(\widehat{t};N\right),t\right).$$

One symmetric separating equilibrium action function, $a_{SEP}(t;N)$, is defined by the following differential equation

$$\left.\frac{\partial a_{SEP}\left(\widehat{t};N\right)}{\partial \widehat{t}}\right|_{\widehat{t}=t} = -\frac{1}{N}\left.\frac{\partial B\left(t,\widehat{t}\right)}{\partial \widehat{t}}\right|_{(t,\widehat{t})=(t,t)}\left.\frac{1}{\pi_a(a,t)}\right|_{(a,t)=(a_{SEP}(t;N),t)}$$

with the initial condition $a_{SEP}(\underline{t};N) = a_{BP}(\underline{t})$.

We now provide conditions that imply the symmetric separating equilibrium is the welfare-optimal separating equilibrium for all types, which motivates our focus on it.

Theorem 3. *Assume that (i) $\pi(a,t)$ is supermodular in (a,t), (ii) for all t and $a > a_{BP}(t)$ we have $\pi_a(a,t) \leq 0$, and (iii) for all t, $\frac{\pi_{at}(a,t)}{\pi_a(a,t)}$ is weakly increasing in a. For any fixed N, $\boldsymbol{a}_{SEP}^N$ maximizes the payoff of each type of sender relative to any other separating equilibrium.*

The supermodularity requirement is standard. The second requirement simply states that utility is decreasing in a once a exceeds the bliss point for t. The third requirement implies that the marginal cost of a becomes more sensitive (in relative terms) to type as a increases. These requirements are satisfied for common specifications such as $\pi(a,t) = -(a-t)^2$.

When we restrict attention to the symmetric equilibrium action function $\boldsymbol{a}_{SEP}^N(t)$, the N-signal setting becomes almost identical to the model of Section 4, with N playing the role of λ. It then follows as a corollary to Theorem 1 that the ratio of the costs of signaling to the benefits of information transmission converges to zero as $N \to \infty$.

Corollary 1. *Consider the additive N-signal model. Under Assumptions 1 - 5,we have*

$$\sum_{i=1}^{N}\left[\pi\left(a_{BP}(t;R),t\right) - \pi\left(a_{SEP,i}^N(t),t\right)\right] = O\left(\frac{1}{\sqrt{N}}\right)$$

for equilibria in which the receiver's beliefs satisfy the dominance refinement.

Similarly, we also have a corollary to Theorem 2, which establishes a faster rate of convergence under more restrictive conditions:

Corollary 2. *Consider the additive N-signal model. Under Assumptions 1 - 5, if $a_{BP}(t)$ and $a_{SEP}(t;N)$ are continuous for all N, then $N\left[\pi(a_{BP}(t),t) - \pi(a_{SEP}(t;N),t)\right] = O\left(\frac{1}{N}\right)$*
.

4.3 Multiple Signals with General Aggregator Functions

In this section, we relax the additivity restriction and show that our conclusions hold within a broader class of environments. As before, we consider a sequence of models, the N^{th} of which allows the sender to choose N actions. The utility of the sender is

$$U_N(\boldsymbol{a}^N, t) = B\left(t, \widehat{t}\left(\boldsymbol{a}^N\right)\right) + \boldsymbol{\pi}^N\left(\boldsymbol{a}^N, t\right).$$

Throughout we assume that the second and third order partial derivatives of $\boldsymbol{\pi}^N$ exist and are bounded and continuous. We define $\boldsymbol{a}_{BP}(t)$ as

$$\boldsymbol{a}_{BP}^N(t) = \underset{\boldsymbol{a}^N}{arg\,max}\ \boldsymbol{\pi}^N(\boldsymbol{a}^N, t). \tag{4}$$

The analysis of these general settings is challenging because, in principle, the sender might choose to signal through all, some, or even just one of activity. There is no reason to think that signaling costs will vanish as the number of observable activities increases unless senders actually use a rising number of activities as signals.

Many of the assumptions we make in this section have straightforward analogs to those used in Sections 4.1 and 4.2. We have opted for assumptions that involve intuitive, easily understood restrictions, rather than ones that deliver the greatest technical generality. We indicate in footnotes where weaker (and less intuitive) assumptions would suffice.

Our first assumption, which requires that bliss points are monotone increasing in all dimensions, generalizes Assumption 3 of the additive model. Here we also impose the stronger requirement that the bliss point function is Lipschitz continuous.

Assumption 7. *$\boldsymbol{a}_{BP}^N(t)$ is the unique solution to (4). There exists scalars $\beta_H > \beta_L > 0$ such that for all N and any $t > t'$, we have*

$$\beta_H(t - t') \geq \boldsymbol{a}_{BP}^N(t) - \boldsymbol{a}_{BP}^N(t') \geq \beta_L(t - t')\mathbf{1}^N,$$

where $\mathbf{1}^N = (1, 1, ..., 1) \in \mathbb{R}^N$.

We now assume that as N grows, the cost of choosing any given action other than the agent's bliss-point grows as well. It is obvious that this assumption holds in the additive model.[14]

Assumption 8. *For any* $\boldsymbol{a}^N \neq \boldsymbol{a}^N_{BP}(t)$*, we have*

$$\underset{N\to\infty}{Lim}\, \boldsymbol{\pi}^N(\boldsymbol{a}^N_{BP}(t), t) - \boldsymbol{\pi}^N(\boldsymbol{a}, t) > \overline{B} - \underline{B}$$

Our next assumption requires that the effect of the second-order terms grows as N increases.[15]

Assumption 9. *There exists* $C < 0$ *such that for all* N

$$C^N \equiv \max_{\|s\|=1} \sum_{i=1}^{N} \sum_{j=1}^{N} \frac{\partial^2 \boldsymbol{\pi}^N(\boldsymbol{a}^N_{BP}(t'), t')}{\partial \boldsymbol{a}_i \partial \boldsymbol{a}_j} s_i s_j < C$$

Notice that $s = (0, ...0, 1, 0, ..., 0)$ singles out the second derivative with respect to a single action; thus, the assumption bounds each of these second derivatives away from zero. It follows immediately that the single-action and additive models both satisfy Assumption 9, and indeed one can think of it as a generalization of the requirement that $\pi_{aa}(a_{BP}(t), t) < 0$ in Assumption 2, Part (1).[16]

Our proof uses Taylor series approximations around the bliss points to identify which types fall within the plausible set for each action. The argument requires an assumption

[14] Here a weaker assumption will suffice: our proofs only require the costs to grow for all $\boldsymbol{a}^N$ such that $\boldsymbol{a}^N = \boldsymbol{a}^N_{BP}(t')$ for some $t' \neq t$.

[15] As in the case of Assumption 8 (see footnote 14), our results actually require a weaker (but less intuitive) condition involving comparisons to other types' bliss points:

Assumption. *For any* $t > t'$*, let* $\boldsymbol{\delta}^N = \boldsymbol{a}^N_{BP}(t) - \boldsymbol{a}^N_{BP}(t')$*. There exists a sequence* $\{\phi_N\}^\infty_{N=1}$ *such that* $\phi_N < 0$*,* $\phi_N \to -\infty$ *as* $N \to \infty$*, and* $\sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \boldsymbol{\pi}^N(\boldsymbol{a}^N_{BP}(t'), t')}{\partial \boldsymbol{a}_i \partial \boldsymbol{a}_j} \delta_i^N \delta_j^N < \phi_N (t - t')^2$ *for all* t*.*

[16] We could also provide an assumption along the lines of Assumption 2, Part (2), but it would be notationally intensive.

that allows us to ignore the effect of the third-order terms as N grows.[17,18]

Assumption 10. *Let* $D^N = \sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N} \max\left\{0, \frac{\partial^3 \boldsymbol{\pi}^N(\boldsymbol{a}^N_{BP}(t),t)}{\partial \boldsymbol{a}_i \partial \boldsymbol{a}_j \partial \boldsymbol{a}_k}\right\}$*. Then there exists* $D > 0$ *such that* $D^N < ND$ *for all* N*.*

In the additive model, we have $D^N = \sum_{i=1}^{N} \max\{0, \pi_{aaa}(a_{BP}(t),t)\}$; according to Part (1) of Assumption 2 (which requires $\pi_{aaa}(a_{BP}(t),t) < C$), this term is bounded above by NC. Thus, in limiting the aggregate importance of positive third derivatives, Assumption 10 generalizes the analysis of previous sections. There are some additional cases where the assumption obviously holds, such as when the third-order derivatives are all weakly negative. Assumption 10 is also satisfied when (a) all of the third derivative terms individually respect a common upon bound that is independent of N (i.e., , $\boldsymbol{\pi}^N_{ijk}(\boldsymbol{a}^N_{BP}(t),t) < K$ for some finite K), and (b) actions interact with each other (in the sense that cross-partial derivatives are non-zero) only within groups of size no greater than some fixed M.[19] The additive model can be viewed as an example of this class where $M = 1$.

We show that the total cost of signaling must vanish by spelling out the implications of the fact that the sender always has the option of deviating to her bliss-point. Since $\boldsymbol{a}^N_{BP}(t)$ is strictly increasing (by Assumption 7), the dominance refinement pins down the receiver's beliefs after observing $\boldsymbol{a}^N_{BP}(t)$.[20] The restricted beliefs of the receiver imply a bound on the benefit of signaling one's type in equilibrium, which we convert into a bound on the costs of signaling. Assumption 9 ensures that this bound tightens as $N \to \infty$.

[17] We approximate $\boldsymbol{\pi}^N$ around $\boldsymbol{a}^N_{BP}(t)$ for some t, and we evaluate the series in a neighborhood of $\boldsymbol{a}^N_{BP}(t)$. Moreover, this neighborhood shrinks as N goes to infinity. This would suggest the higher-order terms of the series can be neglected for large N without an additional assumption. The problem is that $\boldsymbol{\pi}^N$ changes with N, so we need a bound on the effect of higher order terms that holds across N, which no standard polynomial approximation result provides.

[18] As with the previous assumptions, we can provide a weaker (but less intuitive) condition that imposes the restriction in a more limited way, based on bliss points. The benefit of this approach is that it does not require an explicit bound on the rate at which the aggregate magnitude of third-order terms grows. The single-act and additive models also satisfy this assumption.

Assumption. *Fix* t *and suppose there exists a sequence* $\{\widehat{t}^N\}_{N=1}^{\infty}$ *such that* $\boldsymbol{\pi}^N\left(\boldsymbol{a}^N_{BP}(\widehat{t}^N), \widehat{t}^N\right) - \boldsymbol{\pi}^N\left(\boldsymbol{a}^N_{BP}(t), \widehat{t}^N\right) \leq \overline{B} - \underline{B}$*. If such a sequence exists, then let* $\delta^N = \boldsymbol{a}^N_{BP}(\widehat{t}^N) - \boldsymbol{a}^N_{BP}(t)$ *where* $\widehat{t}^N > t$*. Then we require that* $\underset{N\to\infty}{Lim}\left(\sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N} \frac{\partial^3 \boldsymbol{\pi}^N(\boldsymbol{a}^N_{BP}(t),t)}{\partial \boldsymbol{a}_i \partial \boldsymbol{a}_j \partial \boldsymbol{a}_k}\delta^N_i \delta^N_j \delta^N_k\right) / \left(\sum_{i=1}^{N}\sum_{j=1}^{N} \frac{\partial^2 \boldsymbol{\pi}^N(\boldsymbol{a}^N_{BP},t)}{\partial \boldsymbol{a}_i \partial \boldsymbol{a}_j}\delta^N_i \delta^N_j\right) \geq 0.$

[19] In that case, it is natural to treat N as indexing the number of groups, rather than the number of actions.

[20] The formal statement of the dominance refinement for the multidimensional case is the same as for the single dimensional case, except that one interprets the action a as a vector.

Theorem 4. *Under Assumptions 4, 5, and 7—10, if the receiver's beliefs satisfy the dominance refinement, then* $\boldsymbol{\pi}^N\left(\boldsymbol{a}_{BP}^N(t), t\right) - \boldsymbol{\pi}^N\left(a_{SEP,i}^N(t), t\right) = O\left(\frac{1}{\sqrt{N}}\right)$.

As we have noted, the additive model studied in Section 4.2 satisfies the assumptions of Theorem 4. Therefore, the following corollary holds for any fully separating equilibrium of the additive model, symmetric or otherwise (where $a_{SEP,i}^N(t)$ denotes the i^{th} component of an arbitrary fully separating equilibrium):

Corollary 3. *Consider the (potentially asymmetric) additive model. Assume* $\pi(a, t)$ *satisfies assumptions 1 - 5, and that the receiver's beliefs satisfy the dominance refinement. Then*

$$\sum_{i=1}^{N}\left[\pi\left(a_{BP}(t), t\right) - \pi\left(a_{SEP,i}^N(t), t\right)\right] = O\left(\frac{1}{\sqrt{N}}\right).$$

5 Audience Augmentation versus Signal Proliferation on Social Networks

So far, our analysis has focused on one effect of OSNs: signal proliferation. OSNs also enable people to reach much wider audiences, and hence magnify the payoffs from signaling activities. In this section, we explore the robustness of our conclusions concerning signal proliferation in settings with audience augmentation effects.

We employ the additive model of Section 4.2. The sender has the opportunity to signal N experiences. Senders have these types of experiences in the absence of OSNs, but OSNs allow them to showcase more experiences to more receivers. The action a associated with a particular experience represents the "impressiveness" of the experience, while the sender's type t governs the ease with which she can create impressive experiences. The direct utility she receives from an experience with an impressiveness index of a is $\pi(a, t)$. This benefit flows either from her intrinsic enjoyment or from instrumental motivations aside from signaling (e.g., accumulating experiences that are valuable for career development). We assume the user's bliss point $a_{BP}(t)$ is increasing in t: those who can enhance the impressiveness of experiences at lower incremental cost will naturally choose more impressive experiences even in the absence of signaling opportunities.

What counts as an impressive experience depends on the application of interest. For example, if the goal is to signal professional success, then impressive experiences might include job promotions, publicity in news releases, or prestigious awards. In contrast, if the goal is to signal wealth, then impressive experiences might include expensive vacations or costly dinners.

We depart from our previous formulation of the additive model by assuming that the expansion of OSNs not only contributes to the pervasiveness of signaling, but also enlarges the sender's audience, and hence the benefits from favorable information transmission. In particular, we assume that, when the sender has the opportunity to showcase N experiences, she has an audience of F_N social contacts. We are interested in applications for which an increase in the size of the audience magnifies the impact of information. In particular, we allow for the possibility that an increase in F_N raises the bound on the marginal benefit of signaling:

Assumption 11. *There is a function $\gamma(F_N) > 0$ such that $0 \leq B_{\widehat{t}}(\widehat{t}, F_N) \leq \gamma(F_N)$.*

Our first result characterizes limiting behavior as $N \to \infty$ for symmetric fully separating equilibrium in which the receivers' beliefs satisfy the dominance refinement.

Theorem 5. *Under Assumptions 1 - 3 and 11, we have*

$$\sum_{i=1}^{N} \left[\pi\left(a_{BP}(t), t\right) - \pi\left(a^{N}_{SEP,i}(t), t\right)\right] = O\left(\sqrt{\frac{\gamma(F_N)^3}{N}}\right)$$

for fully separating equilibria of the additive model in which receivers' beliefs satisfy the dominance refinement.

According to Theorem 5, the total cost of signaling vanishes as long as $\gamma(F_N)^3/N \to 0$. For the purpose of illustration, assume that $\gamma(F_N) = \xi N^{\sigma}$. Then, according to the theorem, if $\sigma < \frac{1}{3}$, the aggregate costs of signaling vanish in the limit as $N \to \infty$. However, even when aggregate signaling costs do not become vanishingly small, they may still represent a vanishing fraction of the information transmission benefits. As an immediate corollary of Theorem 5, we know that

$$\sum_{i=1}^{N} \left[\frac{\pi\left(a_{BP}(t), t\right) - \pi\left(a^{N}_{SEP,i}(t), t\right)}{\gamma(F_N)}\right] = O\left(\sqrt{\frac{\gamma(F_N)}{N}}\right) \tag{5}$$

It follows that signaling costs become arbitrarily small relative to potential informational benefits as long as $\gamma(F_N)/N \to 0$. For the class of functions $\gamma(F_N) = \xi N^{\sigma}$, the corresponding condition is $\sigma < 1$. For $\sigma \in (\frac{1}{3},1)$, the spread of OSNs can indeed lead people to incur progressively greater signaling costs, but those costs are dwarfed by the incremental benefits.

As before, our second result establishes more rapid convergence under more restrictive conditions.

Theorem 6. *Under Assumptions 1 - 3 and 11, if $a_{BP}(t)$ and $a_{SEP}(t)$ are continuous for all N then*[21]

$$N\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t)\right]=O\left(\frac{\gamma(F_N)^2}{N}\right).$$

According to Theorem 6, the total cost of signaling vanishes as long as $\gamma(F_N)^2/N \to 0$. For the case of $\gamma(F_N)=\xi N^{\sigma}$, this condition becomes $\sigma<\frac{1}{2}$. Turning to the question of whether the ratio of costs to potential benefits vanishes as $N\to\infty$, we obtain as a corollary a characterization analogous to equation 5:

$$\sum_{i=1}^{N}\left[\frac{\pi\left(a_{BP}(t),t\right)-\pi\left(a^{N}_{SEP,i}(t),t\right)}{\gamma(F_N)}\right]=O\left(\frac{\gamma(F_N)}{N}\right)$$

It follows once again that signaling costs become arbitrarily small relative to potential informational benefits as long as $\gamma(F_N)/N\to 0$. Thus, for the class of functions $\gamma(F_N)=\xi N^{\sigma}$, the spread of OSNs can indeed lead people to incur progressively greater signaling costs, but those costs are dwarfed by the incremental benefits as long as $\sigma\in(\frac{1}{3},1)$.

Several core messages emerge from this analysis. First, as long as the relationship between the potential benefits of signaling and the number of available signals exhibits decreasing returns of sufficient magnitude, signal proliferation still drives the total cost of signaling to zero in the limit. Second, even in cases where equilibrium signaling costs do not vanish in the limit, they still dissipate a vanishing fraction of the OSN's potential informational benefits, except in cases where the effects of OSNs on the scale of informational benefits are proportionately larger than those involving signal proliferation. While OSNs have arguably had an enormous effect on the sizes of audiences (e.g., roughly 60 percent of Facebook users have more than 200 network friends, Statista [32]), it is unclear whether their impact on the benefits of signaling are comparable. Within the social sphere, people presumably derive most of those benefits from interaction with immediate friends and family; consequently, $\gamma(F_N)$ may flatten out rather quickly. Within the professional

[21] The additional factor of $\gamma(F)$ in Theorem 5 is the result of a $\sqrt{\overline{B}-\underline{B}}$ term that vanishes over the course of the iterating argument used in the proof of Theorem 6.

sphere, the issue is more complex. On the one hand, the value of distant contacts may be comparable to that of close ones, e.g., for an academic striving to build a national or international reputation. On the other hand, someone seeking a fixed number of discrete assignments (such as a job or a portfolio of projects) may experience rapidly decreasing returns to audience size.[22] It is therefore entirely possible that the effects of OSNs on signaling costs differ by context.

6 CONCLUSION

The growing prevalence of OSNs has led many people to live their lives more publicly. As a result, opportunities to signal have proliferated, and audiences for signals have grown. Because signaling leads to costly distortions, it is natural to conjecture that pervasive signaling would be highly inefficient. On the contrary, our analysis demonstrates that the ability to signal through many activities simultaneously can actually mitigate welfare losses. Instead of increasing the cost of signaling, a public life may allow one to "live authentically" – that is, to credibly reveal private information at a negligible cost.

Our findings may also shed light on specific signaling phenomena. First, signaling is a natural explanation for conspicuous consumption (see Ireland [22] and Bagwell and Bernheim [5]). The avenues for signaling affluence have expanded immensely with the growth of OSNs. In the past, people could advertise their wealth through specific durable goods such as expensive cars, jewelry, and clothing. Now they can display wealth through OSN posts describing a wide variety of experiences, such as high-end vacations, expensive dinners, and premium seating at concerts. Thoumrungroje [35] finds that increased social media use is indeed correlated with an intensification of conspicuousness as a driver of consumption. According to other surveys, consumers are spending an increasing fraction of their resources on live events, which are the subject of frequent OSN postings, particularly among younger users (Eventbrite and Harris [18]). While the expanding scope of conspicuous consumption would seem to suggest greater wastefulness, our results suggest that this effect may be swamped by the reduction in waste per conspicuous good. In other words, although we observe more forms of conspicuous consumption, welfare is actually greater because signaling distorts the consumption of each good to a much smaller degree.

[22] For example, suppose each job offer includes a wage drawn from some fixed distribution, and that the worker accepts the offer with the highest wage. In that case, the expected benefit of an additional offer declines quickly with the number of offers.

A second application involves signaling by politicians in lower office who hope to win either reelection or higher office. Greater transparency in government and closer monitoring by news media (which now include 24 hour news networks, specialized Twitter feeds, political blogs, and the like) has the effect of increasing politicians' opportunities to signal (see, for example, our analysis of "decisiveness" in Bernheim and Bodoh-Creed [7]). Our findings suggest that signaling motives may distort politicians' choices to a smaller degree as a result of greater transparency. A third application involves job market signaling. A direct application of our results implies that employers can reduce the cost of signaling by evaluating potential employees holistically (according to many criteria), rather than on the basis of a few criteria (such as college grades). We develop these additional applications in Appendix B.

Finally, our results also have potential implications for applications that highlight the importance of pooling equilibria. For example, Bernheim [6] models social conformity as a partial pooling equilibrium wherein agents who value social esteem converge on the preferred action of the most esteemed type. Banks [4] studies a model of political conformity wherein politicians' political platforms converge on the median voter's preferred policy in order to signal more moderate outlooks. Remark 1 suggests that the set of agents who join a central pool will shrink as the number of observable actions increases. This observation suggests that the proliferation of signals may erode conformity by status seekers, leaving us with groups of conformists who are "true believers" in the norms they practice.

References

[1] "2016 Nielsen Social Media Report," retrieved from *https://www.nielsen.com/content/dam/corporate/us/en/reports-downloads/2017-reports/2016-nielsen-social-media-report.pdf* on 5 July 2018.

[2] Andreoni, J. and B.D. Bernheim (2009) "Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects," *Econometrica*, 77 (5), pp. 1607 - 1636.

[3] Angwin, J.; T. Parris Jr.; and S. Mattu (2016) "Breaking the Black Box: What Facebook knows about you," retrieved from *https://www.propublica.org/article/breaking-the-black-box-what-facebook-knows-about-you* on 25 July 2019.

[4] Banks, J. (1990) "A Model of Electoral Competition with Incomplete Information," *Journal of Economic Theory,* 50, pp. 309 - 325.

[5] Bagwell, L. and B.D. Bernheim (1986) "Veblen Effects in a Theory of Conspicuous Consumption," *The American Economic Review*, 86 (3), pp. 349 - 373.

[6] B.D. Bernheim (1994) "A Theory of Conformity," *The Journal of Political Economy*, 102 (4), pp. 841-877.

[7] Bernheim, B.D. and A. Bodoh-Creed (2017) "A Theory of Decisive Leadership," *mimeo*.

[8] Bernheim, B.D. and S. Severinov (2016) "Bequests as Signals: An Explanation for the Equal Division Puzzle," *Journal of Political Economy*, 111 (4), pp. 733-764.

[9] boyd, d. (2004) "Friendster and Publicly Articulated Social Networks," *Conference on Human Factors and Computing Systems*, Vienna: ACM.

[10] Buffardi, L. and W.K. Campbell (2008) "Narcissism and Social Networking Web Sites," *Personality and Social Psychology Bulletin*, 34 (10), pp. 1303-1314.

[11] Cho, I-K. and D. Kreps (1987) "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 102 (2, pp. 179-222.

[12] Corneo, G. and O. Jeanne (1997) "Conspicuous Consumption, Snobbism, and Conformism," *Journal of Public Economics*, 66 (1), pp. 55 - 71.

[13] van Dijck, Jose (2013) "You have one identity: performing the self on Facebook and LinkedIn," *Media, Culture & Society*, 35 (2), pp. 199-215.

[14] Donath, J. (2002) "Identity and deception in the virtual community." *Communities in Cyberspace*. Eds. P. Kollock and M. Smith. London: Taylor & Francis Group, pp. 27-58.

[15] Donath, J. (2008) "Signals in Social Supernets," *Journal of Computer-Mediated Communication*, 13, pp. 231-351.

[16] Donath, J. and d boyd (2004) "Public displays of connection," *BT Technology Journal*, 22 (4), pp. 71-82.

[17] Drouin, M.; K. O'Connor; G. Schmidt; and D. Miller (2015) "Facebook fired: Legal perspectives and young adults' opinions on the use of social media in hiring and firing decisions," *Computers in Human Behavior*, 46, pp. 123-128.

[18] Eventbrite (2017) "Millenials: Fueling the Experience Economy," downloaded on 16 August 2018 from *https://www.eventbrite.com/blog/academy/millennials-fueling-experience-economy/*.

[19] Grafen, A. (1990) "Biological Signals as Handicaps," *Journal of Theoretical Biology*, 144, pp. 517-546.

[20] Granovetter, M. (1973) "The Strength of Weak Ties," *American Journal of Sociology*, 78 (6) , pp. 1360-1380.

[21] Hurd, P. (1995) "Communication in Discrete Action-Response Games," *Journal of Theoretical Biology*, 174, pp. 217-222.

[22] Ireland, N. (1994) "On limiting the market for status signals," *Journal of Public Economics*, 53, pp. 91-110.

[23] Jobvite (2013), *2013 Social Recruiting Survey*, downloaded from https://www.jobvite.com/jobvite-news-and-reports/2013jobvitesocialrecruitingsurvey/.

[24] Kluemper, D. and P. Rosen (2009) "Future employment selection methods: evaluating social networking web sites," *Journal of Managerial Psychology*, 24 (6), pp. 567-580.

[25] Lachmann, M.; S. Számado; and C. Bergstrom (2001) "Cost and conflict in animal signals and human language," *Proceedings of the National Academy of Sciences of the United States of America*, 98, pp. 13189-13194.

[26] Mailath, G. (1989) "Simultaneous Signaling in an Oligopoly Model," *The Quarterly Journal of Economics*, 104 (2), pp. 417 - 427.

[27] Mailath, G. and E. von Thadden (2013) "Incentive compatibility and differentiability: New results and classic applications," *Journal of Economic Theory*, 148, pp. 1841-1861.

[28] Miller, M. and K. Rock (1985) "Dividend Policy Under Asymmetric Information," *Journal of Finance*, 40 (4), pp. 1031- 1051.

[29] Neilsen (2016), *Neilsen 2016 Social Media Report*, downloaded from http://www.nielsen.com/us/en/insights/reports/2017/2016-nielsen-social-media-report.html.

[30] Smith, A. and M. Anderson (2018) "Social Media Use in 2018," (Pew Research Center, Washington, D.C.).

[31] Spence, M. (2002) "Signaling in Retrospect and the Informational Structure of Markets," *The American Economic Review*, 92 (3), pp. 434 - 459.

[32] Statista Research Department (2016) "Average number of Facebook friends of users in the United States in 2016," retrieved from https://www.statista.com/statistics/398532/us-facebook-user-network-size/ on 30 July 2019.

[33] Számado, S. (1999) "The Validity of the Handicap Principle in Discrete Action-Response Games," *Journal of Theoretical Biology*, 198, pp. 593-602.

[34] Számado, S. (2011) "The cost of honesty and the fallacy of the handicap principle," *Animal Behavior*, 81, pp. 3-10.

[35] Thoumrungroje, A. (2014) "The Influence of Social Media Intensity and EWOM on Conspicuous Consumption," *Procedia - Social and Behavioral Sciences*, 148, pp. 7-15.

[36] Zahavi, A. (1979) "Mate Selection—A Selection for a Handicap," *Journal of Theoretical Biology*, 53, pp. 205-214.

A PROOF APPENDIX

Lemma 1. *Let Assumptions 1, 2 and 5 hold. Then* $\|a(t;\lambda_i) - a_{BP}(t)\| = O\left(\frac{1}{\sqrt{\lambda_i}}\right)$ *as* $\lambda_i \to \infty$*. If* $a_{BP}(t)$ *is continuous, then the convergence is uniform over* t*.*

Proof. The bound on $a(t;\lambda_i) - a_{BP}(t)$ will be derived from the following inequality, which must hold in equilibrium:

$$\lambda_i \left[\pi(a_{BP}(t), t) - \pi(a(t;\lambda_i), t)\right] \leq \overline{B} - \underline{B}. \tag{6}$$

We first prove our result for the case where $\pi_{aa}(a_{BP}(t),t) < 0$, and then consider what occurs when $\pi_{aa}(a_{BP}(t),t) = 0$.

The uniqueness of the bliss point for type t (Assumption 1) implies that $\pi(a_{BP}(t),t) - \pi(a(t;\lambda_i),t) > 0$, and the continuity of $\pi(\circ,t)$ implies $a(t;\lambda_i) - a_{BP}(t) \to 0$ as $\lambda \to \infty$. The Taylor expansion of $\pi(a,t)$ around $(a_{BP}(t),t)$ is:

$$\pi(a_{BP}(t),t) - \pi(a(t;\lambda_i),t) = \frac{-1}{2}\pi_{aa}(a_{BP}(t),t)\left(a(t;\lambda_i) - a_{BP}(t)\right)^2 - \frac{\pi_{aaa}(\xi)}{6}\left(a(t;\lambda_i) - a_{BP}(t)\right)^3, \tag{7}$$

where $\xi \in [a_{BP}(t), a(t;\lambda_i)]$. Suppose $\pi_{aaa}(\xi)\left(a(t;\lambda_i) - a_{BP}(t)\right) < 0$. Then

$$\pi(a_{BP}(t),t) - \pi(a(t;\lambda_i),t) \geq \frac{-1}{2}\pi_{aa}(a_{BP}(t),t)\left(a(t;\lambda_i) - a_{BP}(t)\right)^2.$$

Suppose $\pi_{aaa}(\xi)\left(a(t;\lambda_i) - a_{BP}(t)\right) > 0$. Then since $\|\pi_{aaa}(a,t)\| \leq C$ (Assumption 2) and $a(t;\lambda_i) - a_{BP}(t) \to 0$ as $\lambda_i \to \infty$, we can choose λ_i^* such that for all $\lambda_i > \lambda_i^*$

$$\left\|\frac{\pi_{aaa}(\xi)}{6}\left(a(t;\lambda_i) - a_{BP}(t)\right)\right\| \leq \frac{-1}{4}\pi_{aa}(a_{BP}(t),t),$$

which means we have

$$\pi(a_{BP}(t),t) - \pi(a(t;\lambda_i),t) \geq \frac{-1}{4}\pi_{aa}(a_{BP}(t),t)\left(a(t;\lambda_i) - a_{BP}(t)\right)^2. \tag{8}$$

In either case, using equation 6 we can write

$$\left(a(t;\lambda_i) - a_{BP}(t)\right)^2\left(\frac{-1}{4}\pi_{aa}(a_{BP}(t),t)\right) \leq \frac{\overline{B} - \underline{B}}{\lambda},$$

which in turn yields

$$\|a(t;\lambda_i) - a_{BP}(t)\| \leq \frac{1}{\sqrt{\lambda}}\sqrt{\frac{-4(\overline{B} - \underline{B})}{\pi_{aa}(a_{BP}(t),t)}}. \tag{9}$$

In the case where $\pi_{aa}(a_{BP}(t),t) = 0$, the first higher-order partial derivative that is nonzero must be of an even-numbered order.[23] If the order of the first nonzero derivative with respect to a is k, we can use Assumption 2, Part (2) to make an argument analogous to that provided above to show convergence at the rate

$$\|a(t;\lambda_i) - a_{BP}(t)\| = O\left(\frac{1}{\lambda_i^{k/4}}\right)$$

Because the type space is compact, Dini's theorem implies uniform convergence if the limit, $a_{BP}(t)$, is continuous. □

Before proving our main theorem, we present a lemma showing that each agent takes an action above her bliss point.

Lemma 2. *Let Assumptions 1, 3, and 4 hold. Then in any fully separating equilibrium we have for all t that $a_{SEP}(t;\lambda) \geq a_{BP}(t)$ for λ sufficiently large.*

Proof. First note that the claim holds for $\underline{t}$ since $a_{SEP}(\underline{t}) = a_{BP}(t)$ in any fully separating equilibrium. Now consider an arbitrary $t > \underline{t}$ and suppose our claim fails to hold. Then for any $\lambda^* > 0$ there exists $\lambda > \lambda^*$ such that $a_{SEP}(t) < a_{BP}(t)$. Equation 9 implies that for such a choice of λ sufficiently large that $a_{SEP}(\underline{t}) \leq a_{SEP}(t) < a_{BP}(t)$. But from the continuity of $a_{BP}(t)$ (Assumption 1), there exists $t' \in [\underline{t}, t)$ such that $a_{BP}(t') = a_{SEP}(t)$. Since the equilibrium is fully separating, it must be that $a_{SEP}(t') \neq a_{BP}(t')$ or types t and t' would pool. Since $B(t, \widehat{t})$ is increasing in $\widehat{t}$ (Assumption 4), t' can profitably deviate to $a_{BP}(t')$ and have the receiver infer her type to be $\widehat{t} = t > t'$. This contradiction proves our claim. □

Theorem 1. *Under Assumptions 1 - 5, we have $\lambda\left[\pi(a_{BP}(t),t) - \pi(a_{SEP}(t;\lambda),t)\right] = O\left(\frac{1}{\sqrt{\lambda}}\right)$ for equilibria in which the receiver's beliefs satisfy the dominance refinement.*

Proof. For the duration of the proof we assume that λ is sufficiently large that $a_{SEP}(t) \geq a_{BP}(t)$ as per Lemma 2. The goal of this proof is to tighten the bound provided by Lemma

[23] To see this more formally, if $\pi_{aa}(a_{BP}(t),t) = 0$, then Equation 7 has the form

$$\pi(a_{BP}(t),t) - \pi(a(t;\lambda_i),t) = \frac{-1}{6}\pi_{aaa}(a_{BP}(t),t)\left(a_{BP}(t) - a(t;\lambda_i)\right)^3 - \frac{\pi_{aaaa}(\xi)}{24}\left(a_{BP}(t) - a(t;\lambda_i)\right)^4.$$

The fourth order term is negligible relative to the third order term for λ sufficiently large (i.e., $a_{BP}(t)$ and $a(t;\lambda_i)$ sufficiently close), so if $\pi_{aaa}(a_{BP}(t),t) \neq 0$ utility would be increased by either a slight increase or decrease in a from $a_{BP}(t)$. But then $a_{BP}(t)$ cannot be optimal, and from this contradiction we conclude that $\pi_{aa}(a_{BP}(t),t) = 0$ entails $\pi_{aaa}(a_{BP}(t),t) = 0$. We conclude that if $\pi_{aa}(a_{BP}(t),t) = 0$, then the first higher-order partial derivative that is nonzero must be of an even-numbered order.

1. To that end, suppose agent t deviates from $a_{SEP}(t)$ to $a_{BP}(t)$. The proof of Lemma 1 showed that even if a deviation from her bliss point could shift $B(t, \delta(a))$ from $\underline{B}$ to $\overline{B}$, the largest she would be willing to deviate is of $O(\lambda^{-0.5})$. If there exists t' such that $a_{SEP}(t') = a_{BP}(t)$, then the receiver infers (incorrectly) after observing $a = a_{BP}(t)$ that the sender is of type t'. Using Equation 9 we can write:

$$a_{BP}(t) = a_{SEP}(t') \leq a_{BP}(t') + \frac{1}{\sqrt{\lambda}} \sqrt{\frac{-4\left(\overline{B} - \underline{B}\right)}{\pi_{aa}(a_{BP}(t'), t')}}.$$

From Assumption 3 we have

$$\begin{aligned} \beta(t - t') &\leq a_{BP}(t) - a_{BP}(t') \leq \frac{1}{\sqrt{\lambda}} \sqrt{\frac{-4\left(\overline{B} - \underline{B}\right)}{\pi_{aa}(a_{BP}(t'), t')}} \\ t' &\geq t - \frac{1}{\beta\sqrt{\lambda}} \sqrt{\frac{-4\left(\overline{B} - \underline{B}\right)}{\pi_{aa}(a_{BP}(t'), t')}}. \end{aligned} \tag{10}$$

If there is no t' such that $a_{SEP}(t') = a_{BP}(t)$, then Lemma 1 combined with the dominance refinement implies that the receiver must believe that the sender has some type t' that satisfies $\|a_{BP}(t') - a_{BP}(t)\| = O(\lambda^{-0.5})$. Using Equation 9 we can write this formally as:

$$a_{BP}(t) \leq a_{BP}(t') + \frac{1}{\sqrt{\lambda}} \sqrt{\frac{-4\left(\overline{B} - \underline{B}\right)}{\pi_{aa}(a_{BP}(t'), t')}}.$$

As before, Assumption 3 implies

$$t' \geq t - \frac{1}{\beta\sqrt{\lambda}} \sqrt{\frac{-4\left(\overline{B} - \underline{B}\right)}{\pi_{aa}(a_{BP}(t'), t')}}. \tag{11}$$

We now use Equation 10 or 11 as appropriate to bound the effect on the signaling incentive $B(t, \widehat{t})$ more tightly. The core idea is that the the cost of signaling can be no larger than the benefit received by having the receiver infer that the sender has type t instead of type

t'. Using Assumption 4 we have:

$$\begin{aligned}\pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t) &\leq \frac{1}{\lambda}\left[B(t,t)-B\left(t,t-\frac{1}{\beta\sqrt{\lambda}}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}\right)\right] \\ &\leq \frac{\gamma}{\lambda^{1.5}}\left[\frac{1}{\beta}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}\right].\end{aligned}$$

This then yields:

$$\begin{aligned}\lambda\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t)\right] &\leq \frac{\gamma}{\sqrt{\lambda}}\left[\frac{1}{\beta}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}\right] \quad (12)\\ &= O\left(\frac{1}{\sqrt{\lambda}}\right). \quad (13)\end{aligned}$$

□

Theorem 2. *Under Assumptions 1 - 5, if $a_{BP}(t)$ and $a_{SEP}(t;\lambda)$ are continuous in t for all λ, then $\lambda\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t;\lambda),t)\right]=O\left(\frac{1}{\lambda}\right)$.*

Proof. The goal of this proof is also to tighten the bound provided by Lemma 1. The key difference with the proof of Theorem 1 is that due to the continuity of $a_{BP}(t)$ and $a_{SEP}(t)$, if an agent deviates from $a_{SEP}(t)$ to $a_{BP}(t)$, then there exists a t' such that the receiver believes that sender has type $t'<t$ — namely t' such that $a_{SEP}(t')=a_{BP}(t)$. This allows us to avoid the issue of off-path beliefs.

To that end, suppose agent t deviates from $a_{SEP}(t)$ to $a_{BP}(t)$. The type t' such that $a_{SEP}(t')=a_{BP}(t)$ defines the inference made by the receiver following the deviation by type t. Equation 9 yields:

$$a_{SEP}(t')=a_{BP}(t)\leq a_{BP}(t')+\frac{1}{\sqrt{\lambda}}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}.$$

From Assumption 3 we have

$$\begin{aligned}
\beta(t-t') &\leq a_{BP}(t)-a_{BP}(t') \leq \frac{1}{\sqrt{\lambda}}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t),t)}} \\
t' &\geq t-\frac{1}{\beta\sqrt{\lambda}}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}.
\end{aligned}$$

Using Assumption 4 we have:

$$\begin{aligned}
\pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t) &\leq \frac{1}{\lambda}\left[B(t,t)-B\left(t,t-\frac{1}{\beta\sqrt{\lambda}}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}\right)\right] \\
&\leq \frac{\gamma}{\lambda^{1.5}}\left[\frac{1}{\beta}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}\right].
\end{aligned}$$

If Assumption 2, Part (1) holds, we can then write

$$\begin{aligned}
\frac{-1}{4}\pi_{aa}(a_{BP}(t),t)\left[a_{SEP}(t)-a_{BP}(t)\right]^2 &\leq \pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t) && (14) \\
&\leq \frac{1}{\lambda^{1.5}}\frac{\gamma}{\beta}\sqrt{\frac{-4\left(\overline{B}-\underline{B}\right)}{\pi_{aa}(a_{BP}(t'),t')}}. && (15)
\end{aligned}$$

where the first inequality can be derived using an argument essentially identical to that used to derive equation 8. Simplifying yields

$$a_{SEP}(t)-a_{BP}(t) \leq \frac{1}{\lambda^{3/4}}\sqrt{\frac{\gamma}{\beta}}\left(\frac{-4}{\pi_{aa}(a_{BP}(t'),t')}\right)^{3/4}\left(\overline{B}-\underline{B}\right)^{1/4}. \qquad (16)$$

If Assumption 2, Part (2) applies, we can make an analogous argument that yields an even tighter bound on $a_{SEP}(t)-a_{BP}(t)$ as in the proof for Lemma 1.

Iterating this process K times yields

$$a_{SEP}(t)-a_{BP}(t) \leq \frac{C_K}{\lambda^{1-0.5^K}}.$$

When we use this in our Taylor expansion we get

$$\begin{aligned}\pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t) &= (a_{SEP}(t)-a_{BP}(t))^2\left(\frac{-1}{2}\pi_{aa}(a_{BP}(t),t)-\right.\\ &\qquad\left.\frac{\pi_{aaa}(\xi)}{6}\left(a_{SEP}(t)-a_{BP}(t)\right)\right)\\ &= \frac{C_K^2}{\lambda^{2-0.5^{K-1}}}\left(\frac{-1}{2}\pi_{aa}(a_{BP}(t),t)-\frac{\pi_{aaa}(\xi)}{6}\left(\frac{C_K}{\lambda^{1-0.5^K}}\right)\right).\end{aligned}$$

Using the negligibility of the third order terms, we find

$$\begin{aligned}\lambda\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t)\right] &\le \frac{C_K^2}{\lambda^{1-0.5^{K-1}}}\left(\frac{-1}{2}\pi_{aa}(a_{BP}(t),t)-\frac{\pi_{aaa}(\xi)}{6}\left(\frac{C_K}{\lambda^{1-0.5^K}}\right)\right)\\ &= O\left(\frac{1}{\lambda^{1-0.5^{K-1}}}\right).\end{aligned}$$

as desired. Taking the limit as $K\to\infty$ yields our result. □

Theorem 3. *Assume that (i) $\pi(a,t)$ is supermodular in (a,t), (ii) for all t and $a>a_{BP}(t)$ we have $\pi_a(a,t)\le 0$, and (iii) for all t $\frac{\pi_{at}(a,t)}{\pi_a(a,t)}$ is weakly increasing in a. For any fixed N, $\boldsymbol{a}_{SEP}^N$ maximizes the payoff of each type of sender relative to any other separating equilibrium.*

Proof. Suppose we have a separating equilibrium with action functions $\mathbf{a}(t)=(a_1(t),...,a_N(t))$. Defining

$$\Gamma(\mathbf{a},\mathbf{t})\equiv\sum_{i=1}^{N}\pi\left(a_i,t\right), \tag{17}$$

we can write the first-order condition for type t's optimal choice as

$$\left.\frac{\partial B(t,\widehat{t})}{\partial\widehat{t}}\right|_{\widehat{t}=t}+\sum_{i=1}^{N}\frac{\partial\pi\left(a_i,t\right)}{\partial a_i}\frac{da_i(t)}{dt}=0. \tag{18}$$

We are interested in determining type t's total payoff in equilibrium. If we let $V\left(t,\widehat{t}\right)$ denote the payoff of a type t sender having chosen the action of type $\widehat{t}$, we have by definition

$$V\left(t,\widehat{t}\right)=B\left(t,\widehat{t}\right)+\Gamma(\mathbf{a}(\widehat{t}),t), \tag{19}$$

and the Envelope Theorem yields

$$\frac{dV\left(t,t\right)}{dt}=\frac{dB\left(t,t\right)}{dt}+\frac{\partial\Gamma(\mathbf{a}(t),t)}{\partial t}.$$

Notice that only the final term depends on the particular separating equilibrium. Let $\mathbf{a}^0$ denote the symmetric separating equilibrium with payoffs V^0, and $\mathbf{a}^A$ denote an asymmetric separating equilibrium with payoffs V^A. To demonstrate that payoffs in the symmetric separating equilibrium are strictly higher than in the asymmetric separating equilibrium, we will establish the following Property (capitalized for clarity of subsequent references): if it were the case for some t that either (i) $V^0(t,t)=V^A(t,t)$ and $\mathbf{a}^0(t)\neq\mathbf{a}^A(t)$, or (ii) $V^0(t,t)<V^A(t,t)$, then we would have $\frac{dV^0(t,t)}{dt}>\frac{dV^A(t,t)}{dt}$.[24]

To understand why this Property delivers the desired conclusion, note that $V^A(t',t')-V^0(t',t')$ would shrink as t' rises over $[\underline{t},t]$ if the property holds. But then we would have a violation of the boundary condition $V^0(\underline{t},\underline{t})=V^A(\underline{t},\underline{t})=B(\underline{t},\underline{t})-\Gamma(\boldsymbol{a}_{BP}(\underline{t}),\underline{t})$ where $\boldsymbol{a}_{BP}\left(t\right)=\left(a_{BP}(t),...,a_{BP}(t)\right)\in\mathbb{R}^N$. In light of Equation 19, we can rewrite the Property as follows: if it were the case for some t that either (i)$'$ $\Gamma(\mathbf{a}^0(t),t)=\Gamma(\mathbf{a}^A(t),t)$ and $\mathbf{a}^0(t)\neq\mathbf{a}^A(t)$, or (ii)$'$ $\Gamma(\mathbf{a}^0(t),t)>\Gamma(\mathbf{a}^A(t),t)$, then we would have $\frac{\partial\Gamma(\mathbf{a}^0(t),t)}{\partial t}>\frac{\partial\Gamma(\mathbf{a}^A(t),t)}{\partial t}$.

We now establish the Property. Supposing condition (i)$'$ were satisfied for some $t>\underline{t}$, we would begin by defining[25]

$$\overline{a}_m=\begin{cases}a_m^A(t)\text{ if }a_m^A(t)\geq a_{BP}(t)\\ a\geq a_{BP}(t)\text{ s.t. }\pi\left(a,t\right)=\pi(a_m^A(t),t)\text{ otherwise}\end{cases},$$

where this can be done in an arbitrary order over the dimensions of $\boldsymbol{a}^A$. Let $Q\equiv\left\{m\mid a_m^A(t)<a_{BP}(t)\right\}$. Then from supermodularity we have

$$\frac{\partial\Gamma(\overline{\boldsymbol{a}},t)}{\partial t}-\frac{\partial\Gamma(\boldsymbol{a}^A(t),t)}{\partial t}=\sum_{m\in Q}\pi_t(\overline{a}_m,t)-\pi_t(a_m^A,t)\geq 0,$$

with strict inequality if Q is non-empty.

If $\boldsymbol{a}^0(t)=\overline{\boldsymbol{a}}$, we are done. If not, then since $\Gamma(\boldsymbol{a}^A(t),t)=\Gamma(\overline{\boldsymbol{a}},t)$ by construction, there must exist i and j such that $\overline{a}_i>a^0(t)>\overline{a}_j$. Define the function $\widetilde{\boldsymbol{a}}(a_i)$ as follows:

[24] Suppose our claim is true. Then if either condition (i) or (ii) holds for t, then condition (ii) must hold for all $t'\in(\underline{t},t)$.

[25] This step sets computes a cost-equivalent signal to $\boldsymbol{a}^A$ that has the intuitive property that $a_m^A\geq a_{BP}(t)$.

$\widetilde{a}_i(a_i) = a_i$, $\widetilde{a}_k(a_i) = \overline{a}_k$ for $k \neq i, j$, and $\Gamma(\widetilde{a}(a_i), t) = \Gamma(\overline{\boldsymbol{a}}, t)$. In other words, $\widetilde{a}_j(a_i)$ indicates how a_j must vary in response to changes in a_i to keep the value of Γ constant at its equilibrium value. Implicit differentiation reveals that

$$\left.\frac{d\widetilde{a}_j}{da_i}\right|_{a_i=\overline{a}_i} = -\frac{\pi_a\left(\overline{a}_i, t\right)}{\pi_a\left(\widetilde{a}_j(\overline{a}_i), t\right)} < 0.$$

Plainly, there exists a unique value $a_i^e > a_{BP}(t)$ such that $\widetilde{a}_j(a_i^e) = a_i^e$. For $a_i \in [a_i^e, \overline{a}_i(t))$ we have

$$\begin{aligned}
\frac{d}{da_i}\left(\frac{\partial\Gamma(\widetilde{\mathbf{a}}(a_i), t)}{\partial t}\right) &= \frac{d}{da_i}\left(\sum_{i=1}^{N}\pi_t\left(\widetilde{a}_i(a_i), t\right)\right)_{a_i=\overline{a}_i} \\
&= \pi_{at}\left(a_i, t\right) + \pi_{at}\left(\widetilde{a}_j(a_i), t\right)\left.\frac{d\widetilde{a}_j}{da_i}\right|_{a_i=a_i} \\
&= \pi_{at}\left(a_i, t\right) - \pi_{at}\left(\widetilde{a}_j(a_i), t\right)\frac{\pi_a\left(a_i, t\right)}{\pi_a\left(\widetilde{a}_j(a_i), t\right)} < 0,
\end{aligned}$$

where we have used the fact that since $\overline{a}_i \geq a_i \geq a_i^e \geq \widetilde{a}_j(a_i) > a_{BP}(t)$ (which implies $\pi_a\left(\overline{a}_i, t\right) < 0$) and our assumption that for $a > a_{BP}(t)$ we have $\pi_a(a, t) \leq 0$ and

$$\frac{\pi_{at}\left(a_i, t\right)}{\pi_a\left(a_i, t\right)} > \frac{\pi_{at}\left(\widetilde{a}_j(\overline{a}_i), t\right)}{\pi_a\left(\widetilde{a}_j(\overline{a}_i), t\right)}.$$

It follows that $\frac{\partial\Gamma(\widetilde{a}(a_i^e), t)}{\partial t} > \frac{\partial\Gamma(\overline{a}, t)}{\partial t}$ since $\overline{a}_i > a^0(t)$ is being reduced in this equalization step. Through repeated application of this equalization argument, we conclude that $\frac{\partial\Gamma(a^0(c), t)}{\partial t} > \frac{\partial\Gamma(\overline{a}, t)}{\partial t} \geq \frac{\partial\Gamma(a^A(t), t)}{\partial t}$, as desired.

Next, supposing condition (ii)$'$ were satisfied for some $t > \underline{t}$, we would begin by defining a' s.t. $a_1' = a_2' = ... = a_N' > a_{BP}(t)$ and $\Gamma(a', t) = \Gamma(a^A(t), t)$. By the same argument as for condition (i)$'$, we infer $\frac{\partial\Gamma(a', t)}{\partial t} \geq \frac{\partial\Gamma(a^A(t), t)}{\partial t}$.[26] Because $\Gamma(a^0(t), t) > \Gamma(a^A(t), t) = \Gamma(a', t)$ by assumption, we have $a_m^0(t) > a_m'$. From our assumption of supermodularity we conclude:

$$\frac{\partial\Gamma(a^0(c), t)}{\partial t} - \frac{\partial\Gamma(a', t)}{\partial t} = \sum_{m=1}^{N}\pi_t(a_m^0, t) - \pi_t(a_m', t) \geq 0.$$

It follows that $\frac{\partial\Gamma(a^0(c), t)}{\partial t} > \frac{\partial\Gamma(a^A(c), t)}{\partial t}$, as desired.

[26] The inequality is weak because we include the possibility that $a' = a^A(t)$.

Having established that the Property holds, the Proposition follows for the reasons given above. $\square$

Theorem 4. *Under Assumptions 4, 5, and 7—10, if the receiver's beliefs satisfy the dominance refinement, then* $\boldsymbol{\pi}^N\left(\boldsymbol{a}_{BP}^N(t),t\right)-\boldsymbol{\pi}^N\left(a_{SEP,i}^N(t),t\right)=O\left(\frac{1}{\sqrt{N}}\right)$.

Proof. Suppose a sender with type t deviates from $\boldsymbol{a}_{SEP}^N(t)$ to $\boldsymbol{a}_{BP}^N(t)$. For an agent with type t' to choose $\boldsymbol{a}^N=\boldsymbol{a}_{BP}^N(t)$ as part of a separating equilibrium or for $t'\in\mathcal{P}(\boldsymbol{a}_{BP}^N(t))$, we must have

$$\boldsymbol{\pi}^N\left(\boldsymbol{a}_{BP}^N(t'),t'\right)-\boldsymbol{\pi}^N\left(\boldsymbol{a}_{BP}^N(t),t'\right)\leq\overline{B}-\underline{B} \tag{20}$$

Assumption 7 implies $\boldsymbol{a}_{BP}^N(t)-\boldsymbol{a}_{BP}^N(t')\geq\beta_L\delta\mathbf{1}^N$. Using a Taylor expansion and letting $\delta^N=\boldsymbol{a}_{BP}^N(t)-\boldsymbol{a}_{BP}^N(t')$, the left side of equation 20 is equal to

$$\frac{-1}{2}\sum_{i=1}^N\sum_{j=1}^N\frac{\partial^2\boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t')t')}{\partial a_i\partial a_j}\delta_i^N\delta_j^N-\frac{1}{6}\sum_{i=1}^N\sum_{j=1}^N\sum_{k=1}^N\frac{\partial^3\boldsymbol{\pi}^N(\boldsymbol{\xi},t')}{\partial\boldsymbol{a}_i\partial\boldsymbol{a}_j\partial\boldsymbol{a}_k}\delta_i^N\delta_j^N\delta_k^N \tag{21}$$

where $\boldsymbol{\xi}$ lies on the line connecting $\boldsymbol{a}_{BP}^N(t')$ and $\boldsymbol{a}_{BP}^N(t)$. From Assumptions 8 and the continuity of the partial derivatives, as $N\to\infty$ we have

$$\begin{aligned}&\frac{-1}{2}\sum_{i=1}^N\sum_{j=1}^N\frac{\partial^2\boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t')t')}{\partial a_i\partial a_j}\delta_i^N\delta_j^N-\frac{1}{6}\sum_{i=1}^N\sum_{j=1}^N\sum_{k=1}^N\frac{\partial^3\boldsymbol{\pi}^N(\boldsymbol{\xi},t')}{\partial\boldsymbol{a}_i\partial\boldsymbol{a}_j\partial\boldsymbol{a}_k}\delta_i^N\delta_j^N\delta_k^N\\&\qquad\to\frac{-1}{2}\sum_{i=1}^N\sum_{j=1}^N\frac{\partial^2\boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t')t')}{\partial a_i\partial a_j}\delta_i^N\delta_j^N-\frac{1}{6}\sum_{i=1}^N\sum_{j=1}^N\sum_{k=1}^N\frac{\partial^3\boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t'),t')}{\partial\boldsymbol{a}_i\partial\boldsymbol{a}_j\partial\boldsymbol{a}_k}\delta_i^N\delta_j^N\delta_k^N\end{aligned}$$

Using Assumption 9, we can write

$$\begin{aligned}-\frac{1}{2}\sum_{i=1}^N\sum_{j=1}^N\frac{\partial^2\boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t'),t')}{\partial\boldsymbol{a}_i\partial\boldsymbol{a}_j}\delta_i^N\delta_j^N&=\frac{-\left\|\delta^N\right\|^2}{2}\sum_{i=1}^N\sum_{j=1}^N\frac{\partial^2\boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t'),t')}{\partial\boldsymbol{a}_i\partial\boldsymbol{a}_j}\left(\frac{\delta_i^N}{\left\|\delta^N\right\|}\right)\left(\frac{\delta_j^N}{\left\|\delta^N\right\|}\right)\\&>\frac{-\left\|\delta^N\right\|^2}{2}C\geq\frac{-N\beta_L^2(t-t')^2C}{2}\end{aligned}$$

Noting that $\delta_i^N>0$ from Assumption 7, Assumptions 7 and 10 together imply

$$\frac{-1}{6}\sum_{i=1}^N\sum_{j=1}^N\sum_{k=1}^N\frac{\partial^3\boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t),t)}{\partial\boldsymbol{a}_i\partial\boldsymbol{a}_j\partial\boldsymbol{a}_k}\delta_i^N\delta_j^N\delta_k^N>\frac{-ND\beta_H(t-t')^3}{6}$$

Therefore, we have

$$\frac{-1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}\frac{\partial^2 \boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t')t')}{\partial a_i \partial a_j}\delta_i^N\delta_j^N - \frac{1}{6}\sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N}\frac{\partial^3 \boldsymbol{\pi}^N(\boldsymbol{a}_{BP}^N(t'),t')}{\partial \boldsymbol{a}_i \partial \boldsymbol{a}_j \partial \boldsymbol{a}_k}\delta_i^N\delta_j^N\delta_k^N$$
$$\geq -N\left[\frac{\beta_L^2(t-t')^2 C}{2} + \frac{D\beta_H(t-t')^3}{6}\right]$$

Since the third order terms are $o\left((t-t')^2\right)$, equation 20 implies for any $\gamma > 0$ and large enough N

$$-N\frac{\beta_L^2(t-t')^2 C}{2} - \gamma \leq \boldsymbol{\pi}^N\left(\boldsymbol{a}_{BP}^N(t'),t'\right) - \boldsymbol{\pi}^N\left(\boldsymbol{a}_{BP}^N(t),t'\right) \leq \overline{B} - \underline{B}$$

which implies

$$t - t' \leq \sqrt{\frac{-2(\overline{B}-\underline{B})}{N\beta_L^2 C}}$$

Therefore, $t' \in \mathcal{P}(\boldsymbol{a}_{BP}(t))$ implies $t' \geq t - \sqrt{\frac{-2(\overline{B}-\underline{B})}{N\beta_L^2 C}}$. Using the dominance refinement, this means the worst (for the sender) inference the receiver can make following an observation of $\boldsymbol{a}^N = \boldsymbol{a}_{BP}(t)$ is that the sender has type $t - \sqrt{\frac{-2(\overline{B}-\underline{B})}{N\beta_L^2 C}}$. Combined with Assumption 4, this means that the cost of choosing $\boldsymbol{a}_{SEP}(t)$ for a sender with type t satisfies

$$\boldsymbol{\pi}^N\left(\boldsymbol{a}_{BP}^N(t),t\right) - \boldsymbol{\pi}^N\left(\boldsymbol{a}_{SEP}^N(t),t\right) \leq B(t,t) - B\left(t, t - \sqrt{\frac{-2(\overline{B}-\underline{B})}{N\beta_L^2 C}}\right) \leq \gamma\sqrt{\frac{-2(\overline{B}-\underline{B})}{N\beta_L^2 C}}$$

which proves our claim. □

Theorem 5. *Under Assumptions 1 - 3 and 11, we have*

$$\sum_{i=1}^{N}\left[\pi\left(a_{BP}(t),t\right) - \pi\left(a_{SEP,i}^N(t),t\right)\right] = O\left(\sqrt{\frac{\gamma(F)^3}{N}}\right).$$

Proof. Consider equations 12 and 13 in the proof of Theorem 1. The claim follows once we replace γ with $\gamma(F)$ and note that $\overline{B} - \underline{B} \leq \gamma(F)\bar{t}$ by Assumption 11. □

Theorem 6. *Under Assumptions 1 - 3 and 11, if $a_{BP}(t)$ and $a_{SEP}(t)$ are continuous for* *all N then*[27]

$$N\left[\pi(a_{BP}(t),t)-\pi(a_{SEP}(t),t)\right]=O\left(\frac{\gamma(F)^2}{N}\right)$$

..

Proof. Consider equation 16. Replacing γ with $\gamma(F)$ and $\overline{B}-\underline{B}$ with $\gamma(F)\bar{t}$ as per Assumption 11, we find an equation proportional to $(\gamma(F)/\lambda)^{0.75}$. A third iteration of the argument yields an equation with a similar form to equation 16 that is proportional to $(\gamma(F)/\lambda)^{0.875}$, while a fourth iteration yields an equation proportional to $(\gamma(F)/\lambda)^{0.9375}$. Combining this with the Taylor approximation of $\pi(a,t)$ in equation 14 yields our result. □

[27]The additional factor of $\gamma(F)$ in Theorem 5 is the result of a $\sqrt{\overline{B}-\underline{B}}$ term that vanishes over the course of the iterating argument used in the proof of Theorem 6.

B Online Appendix: Additional Applications

B.1 Political Decisiveness

Bernheim and Bodoh-Creed [7] presents a model of politician decisiveness. We provide a reduced form version of the model here for conciseness. The type of politician, denoted t, reflects the politician's innate aversion to delay. This can be interpreted as either the politician's perception of the opportunity costs of delaying a policy decision or the strength of a politician's policy preferences (i.e., weak policy preferences imply a high t). As intuition would suggest, politicians with a high aversion to delay prefer to make policy decisions more quickly. The politician's ideal policy is $p^* = \theta + x_p \in \mathbb{R}$, where θ represents a policy preference that is common across the population and x_p is a policy preference idiosyncratic to the politician. Both θ and x_p are uncertain. A politician in office gets to tune the policy based on what she learns during the period of delay about both her idiosyncratic preferences and the common component of preferences.

The agency conflict stems from two features. First, the politician delays her decision to learn about both the common and idiosyncratic components of her preferences. While voters would like the politician to learn about the common component, the voters do not gain any benefit from what the politician learns about her idiosyncratic preferences. The net effect is that the politician delays longer than voters would prefer. Second, the politician tunes her policy choice to her own idiosyncratic preferences and not those of the voters, so the politician's idiosyncratic preferences act as a source of risk for voters. Because of the agency conflict, politicians have an incentive to signal their aversion to delay.

If only one decision can be observed, the signaling incentives can cause the politician to make the decision more quickly than either the politician or the voters would prefer. However, the rise of cable news and online media has increased the volume of political coverage, which has in turn increased the opportunities for politicians to make their electorate aware of their actions. In addition, as the current U.S. president has demonstrated, OSNs such as Twitter provide ample opportunities to advertise purported accomplishments. Pew Research reports that 93% of American report getting news online, and users often finds news on social networks such as Facebook or Twitter.[28] The online webpages for the top 50 newspapers reported 11.7 million unique viewers in 2016, and natively online news orga-

[28] Downloaded on 3 August 2018 from http://assets.pewresearch.org/wp-content/uploads/sites/13/2018/07/11183646/State-of-the-News-Media2017-Archive.pdf

nizations were able to attract 22.8 million unique users per news organization per month.[29] And none of these statistics count individuals who are made aware of headlines through social networks. Our analysis suggests that the increasing attention paid to politicians has reduced the costs they must incur to signal their decisiveness (or many other traits they might wish to signal).[30]

In period 1, the voters observe a politician make a decision in lower office. In particular, the voters observe how quickly the politician makes the decision, and will later use this as a signal of the politician's decisiveness. In period 2, the politician competes against a randomly drawn opponent for higher office. Politicians are motivated to signal t because voters prefer more decisive politicians that are inclined to make decisions more quickly. In other words, a candidate of type t that signals her type is $\widehat{t}$ will win the election if voters believe her opponent has a lower value of t (i.e., a lower aversion to delay). In period 3, the winning politician holds office and reaps the rewards from her decisiveness. The losing politician gets the same utility as the voters.

Let a denote the amount of time taken to make a decision in period 1. Note that in this model, the politician wants to signal a high aversion to delay by choosing a *lower* (i.e., quicker) action. This is, of course, a minor reconfiguration of the model, but the definition of many of the assumptions must be slightly adjusted to account for the fact that $a_{SEP}(t) < a_{BP}(t)$. The direct benefit of action a is $\frac{-1}{a+\phi}$, which captures the fact that delaying the decision (increasing a) reduces the expected difference between the policy chosen and the ideal policy. The direct cost of delay is at. Therefore we have

$$\pi(a,t) = \frac{-1}{a+\phi} - at$$

The first order condition yields $a_{BP}(t) = \sqrt{\frac{1}{a}} - \phi$. Therefore, the model satisfies Assumptions 1-3, and $\frac{\pi_{at}(a,t)}{\pi_a(a,t)}$ is weakly decreasing in a (which is the desired condition since $a_{SEP}(t) < a_{BP}(t)$).

To simplify our exposition, we let the payoff of winning the election given a type t be denoted $B_{win}(t)$. Given that a politician wins, his payoff does not depend on who his

[29] Natively digital means the primary outlet is online. Websites that received at least 10 million unique visitors per month were analyzed.

[30] On the other hand, the opening of these alternate channels for news reporting has been at the expense of local television and local and national print journalism, however. This means that our argument applies more readily to national level politicians, and, in fact, signaling costs may get higher as local news outlets close.

opponent was in period 2, so $B_{win}(t)$ depends only on the winner's type. The payoff in expectation for a politician of type t that is believed by voters to have type $\widehat{t}$ is $B_{loss}(\widehat{t}, t)$. $B_{loss}(\widehat{t}, t)$ depends on $\widehat{t}$ to account for the expected type of the winner, and it depends on t to account for the fact that the amount of time the winner takes to make decisions has different welfare effects on different types of losing politicians. If the distribution of t amongst the politicians is $F(t)$, then we can write

$$B(t, \widehat{t}) = F(\widehat{t})B_{win}(\widehat{t}) + (1 - F(\widehat{t}))B_{loss}(t, \widehat{t})$$

where $F(\widehat{t})$ is the probability that sender wins the election in a separating equilibrium (i.e., that the opponent's type is less than $\widehat{t}$). We assume that $B(t, \widehat{t})$ satisfies Assumptions 4 and 5, but one can prove this from the microfoundations provided in Bernheim and Bodoh-Creed [7].

In reality, politicians make many decisions while in office, and we classify a political institution as more or less transparent according to whether voters can observe a greater or lesser fraction of these decisions. Greater transparency leads politicians to "spread" their signals across many decisions. Formally, suppose the voters observe N of these choices. If we assume that $\pi(a, t)$ is the same for each decision, Theorem 3 implies that the symmetric equilibrium where $a_{SEP}(t)$ is the choice for each of the N observed decisions is welfare optimal for the politicians. Bodoh-Creed and Bernheim [7] show that that $a_{BP}(t)$ and $a_{SEP}(t)$ are continuous for all N. Corollary 2 then implies

Corollary 4. $N\left[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t), t)\right] = O\left(\frac{1}{N}\right).$

This corollary shows that as transparency increases (N grows), the welfare losses from signaling fade. This implies that increasing transparency is good for politicians. However, the case is ambiguous for voters, as discussed at length in Bodoh-Creed and Bernheim [7]. Too little transparency causes politicians to behave hastily and make decision more quickly than voters would prefer. However, too much transparency reduces the signaling distortion, which allows politicians to signal their type while acting less decisively than voters would prefer.

B.2 Job Market Signaling

We now consider how our results play out in a model of job market signaling. The model of Spence [31] assumed that students vary in terms of an underlying attribute t that is the private information of the students. The students signal this attribute through the choice of years of eduction, a, which we interpret as a measure of human capital. The individual's productivity for an employer is $S(a,t)$, so the choice of a is productive. In a competitive separating equilibrium $a_{SEP}(t)$, the wages are $S(a_{SEP}(\widehat{t}),\widehat{t})$, and we assume that students obtain no intrinsic benefit from human capital except as it influences their wages. We also assume that $S(a,\widehat{t}) = B(\widehat{t}) + P(a)$, and accruing human capital level a has a cost equal to $C(a,t)$.[31]

Now we encode this application into our model. The benefit of signaling is captured by $B(\widehat{t})$. The direct utility from choosing a given t is $\pi(a,t) = P(a) - C(a,t)$. The bliss point is

$$a_{BP}(t) = \underset{a}{arg\,max}\ P(a) - C(a,t)$$

We assume that $a_{BP}(t)$ is unique, $a_{BP}(\underline{t}) > 0$, and $C(a,t)$ is strictly supermodular in $(-a,t)$ so that $a_{BP}(t)$ is strictly increasing. To insure that the productivity effect of effort is nontrivial, we also assume there exists $\beta > 0$ such that for $t > t'$ we have $a_{BP}(t) - a_{BP}(t') \geq \beta(t-t')$.

Letting $\widehat{t}(a)$ denote the inferences of the firms after observing a, utility maximization in a fully separating equilibrium requires the following first-order condition be satisfied

$$B_{\widehat{t}}\frac{d\widehat{t}(a)}{da} + P_a(a) - C_a(a,t) = 0$$

where the first term captures the marginal signaling incentive, the second term reflects the marginal productivity of effort, and the third time is the marginal cost of effort. In a complete information model, the first order condition that defines the agents' bliss points is

$$S_a(a_{BP}(t),t) + \pi_a(a_{BP}(t),t) = 0$$

Since the signaling incentive provides an additional incentive to accrue capital relative to the complete-information model, the students accrue too much human capital relative to the first best, even though the human capital is productive.

The model laid out here assumes that students can only differentiate themselves to employers through a single action, which may initially appear to be a useful and innocuous

[31] Assuming separability of $(a,\widehat{t})$ is done so that we can encode this application into our additive model. If $(a,\widehat{t})$ are not separable, then we need to let $B(\boldsymbol{a}^N,\widehat{t}) = S(\boldsymbol{a}^N,\widehat{t})$ in our multi-action setting, impose conditions along the lines of Footnote 6, and use Theorem 4 on general aggregator functions.

simplification. However, our results reveal that this simplifying assumption actually has substantive economic implications. A more realistic model would allow the students to distinguish themselves along many different dimensions of human capital accumulation. For example, subject-specific human capital can be measured with grades. Leadership and communication skill is reflected in leadership positions in student organizations. The ability to work in teams can be assessed via internships or references from previous employers.

Assume there are N human capital metrics and, for simplicity, assume that the direct utility and cost functions are the same across all of these metrics. Letting a_i^N be the i^{th} component of $\boldsymbol{a}^N$, the sender's utility is

$$B(\widehat{t}(\boldsymbol{a}^N)) + \sum_{i=1}^{N} \left[P(a_i^N) - C(a_i^N, t)\right]$$

where for simplicity we have assumed the receiver's beliefs, $\widehat{t}(\boldsymbol{a}^N)$, are degenerate. We assume that

$$\frac{\pi_{at}(a,t)}{\pi_a(a,t)} = \frac{-C_{at}(a,t)}{P_a(a,t) - C_a(a,t)}$$

is weakly increasing in a, and from Theorem 3 we conclude that an equilibrium where the same action is taken across all N metrics is welfare optimal. Letting $a_{SEP}(t)$ be the symmetric action taken across the N metrics, Corollary 2 implies that the total cost of signaling vanishes as the number of human capital metrics increases.

Corollary 5. *Suppose that $a_{BP}(t)$ and $a_{SEP}(t)$ are continuous for all N. Then*

$$N\left[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t), t)\right] = O\left(\frac{1}{N}\right).$$

C Online Appendix: Nonadditive Aggregators

Consider the following utility function:

$$U_N(\boldsymbol{a}^N, t) = B\left(t, \widehat{t}\left(\boldsymbol{a}^N\right)\right) + \boldsymbol{\pi}_{Agg}\left(a_1, .., a_N, t\right)$$

As an example, suppose the cost function has the form:

$$\boldsymbol{\pi}_{Agg}\left(a_1, .., a_N, t\right) = \sqrt{\sum_{m=1}^{N} \pi(a_m, t)}$$

where $\sqrt{\pi(a,t)}$ satisfies Assumptions 2 and 3.[32] Then for a separating equilibrium with symmetric actions we can write:

$$U_N(\boldsymbol{a}^N, t) = B\left(t, \widehat{t}\left(\boldsymbol{a}^N\right)\right) + \sqrt{N}\sqrt{\pi(a_{SEP}(t), t)}$$

Theorem 2 implies $U_N(a_{BP}(t), t; \lambda) - U_N(a_{SEP}(t; \lambda), t; \lambda) = O\left(\frac{1}{N}\right)$.

More generally, suppose one can write the equilibrium utility for a symmetric separating equilibrium, $\boldsymbol{a}_{SEP}(t) = (a_{SEP}(t), a_{SEP}(t), ..., a_{SEP}(t))$, as:

$$U_N(\boldsymbol{a}_{SEP}(t), t) = B\left(t, t\right) + \lambda(N) g\left(a_{SEP}(t), t\right)$$

Assume the functions g and λ satisfy the following two assumptions:

Assumption 12. $g_{aa}(a,t) < 0$ *and there exists* $C < \infty$ *such that* $g_{aaa} \leq C$*.*

Assumption 13. $\lambda(N) \to \infty$ *as* $N \to \infty$*.*

Assumption 12 is analogous to Assumption 2 and allows us to focus our analysis on the second order expansion of the function $g(a,t)$. Assumption 13 requires that $\lambda(N)$, which is analogous to the decision weight of Section 4, diverges to infinity as $N \to \infty$. The additively separable model of Subsection 4.2 and the example leading this subsection satisfy these requirements.

[32] Theorem 3 applies in any situation where $\boldsymbol{\pi}^N$ is a function of $\sum_{m=1}^{N} \pi(a_m, t)$, which would justify our focus on symmetric separating equilibria.

Under these assumptions, an argument analogous to the proof of Theorem 2 shows that aggregate signaling costs vanish as N grows.[33] For completeness, we state this result as the following corollary, where $a_{BP}(t) = \underset{a}{arg\,max}\ g\,(a,t)$.[34]

Corollary 6. *Let Assumptions 1, 4 - 3, 12, and 13 hold and assume that $a_{SEP}(t)$ and $a_{BP}(t)$ are continuous. Then*

$$\lambda(N)\,[g(a_{BP}(t),t) - g(a_{SEP}(t),t)] = O\left(\frac{1}{\lambda(N)}\right).$$

[33]To see this point, if one replaces λ with $\lambda(N)$ throughout the proof of Theorem 2, the same technical argument applies.

[34]Theorem 1 admits an analogous corollary, which we omit for the sake of brevity.