

Pervasive signaling

B. DOUGLAS BERNHEIM

Department of Economics, Stanford University and NBER

AARON L. BODOH-CREED

Amazon

How does the increasing publicness of decisions (due, e.g., to social media) affect the total costs of social signaling distortions? While pervasive signaling may induce pervasive distortions, it may also permit people to signal while distorting each choice to a smaller degree. Ironically, for a broad class of environments, a sufficient increase in the number of signaling opportunities allows senders to “live authentically,” that is, to signal their types at arbitrarily low overall cost. This result survives when social networking technologies expand signaling opportunities and audience size in tandem, provided the returns to the latter are not too great.

KEYWORDS. Multidimensional signaling, efficient information transmission, on-line social networks.

JEL CLASSIFICATION. D82.

1. INTRODUCTION

The last several decades have witnessed the emergence of information technologies that “may make modern life completely visible and permeable to observers” (Froomkin (2000)). This visibility is a consequence not only of the ways in which businesses and governments use technology to monitor our activities, but also of the tendency to reveal much of our lives online voluntarily. Because observation inevitably leads to inference, pervasive observation leads to pervasive signaling opportunities. In this paper, we ask whether the proliferation of these opportunities is socially helpful or harmful, that is, whether it increases or decreases the *aggregate* waste from signaling distortions. On the one hand, pervasive signaling leads to pervasive distortions. While the total waste from 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 waste from signaling could decline in the limit. In principle, it could also remain unchanged.

We examine settings with multiple signaling actions in which the sender’s single-dimensional type is drawn from a continuum of possibilities, and those perceived as

B. Douglas Bernheim: bernheim@stanford.edu

Aaron L. Bodoh-Creed: abodoh.creed@gmail.com

We are grateful to our anonymous referees for helpful comments.

higher types receive greater benefits. We assume in addition that the higher types have higher bliss points for each action, a property that is natural for a wide range of applications.¹ At the outset, we restrict attention to cases in which the sender's direct utility is additively separable and symmetric over the actions, but we also explore less restrictive settings. To study the effects of signal proliferation, we let the number of signals, N , grow without bound.

Ordinarily, we would not expect our model to yield a unique signaling equilibrium. Indeed, we present a simple motivating example in which signaling waste disappears in the limit for one type of separating equilibrium but remains undiminished for another. We address the multiplicity of equilibria through two complementary approaches. For the first approach, we propose a mild belief restriction that is weaker than, but similar in spirit to, the *Intuitive Criterion* of Cho and Kreps (1987). We then prove a convergence theorem that ensures the disappearance of aggregate signaling waste at the rate $1/\sqrt{N}$ for all belief-restricted separating equilibria. For our second approach, we avoid imposing belief restrictions, and instead provide an asymptotic bound on waste in symmetric separating equilibria, which are analytically tractable. It follows immediately that the same bound applies to the equilibria that reveal all private information with the least aggregate waste. Specifically, for symmetric separating equilibria, we demonstrate that aggregate waste disappears at the rate $1/N$. We also provide conditions under which the symmetric separating equilibrium is in fact waste-minimizing. We extend our main findings by examining more general utility functions that allow for nonlinear aggregation across the sender's activities, and we provide sufficient conditions under which our main conclusions continue to hold.² Thus, we conclude that, in the limit, senders fare as well as with complete information. In effect, the proliferation of observable activities enables each sender to "live authentically," that is, to signal the truth about her type at negligible overall cost.

Why do these convergence properties hold? To build intuition, focus only on the symmetric options. A sufficient proliferation of actions causes the utility sacrificed when choosing any given symmetric action other than the sender's bliss point to exceed any possible reputational gain. Therefore, choices in symmetric separating equilibria must converge to bliss points. However, because the number of actions also serves as a multiplier that magnifies the remaining signaling waste in each individual action, the convergence of equilibrium actions to bliss points does not by itself ensure that waste from signaling disappears in the limit. To establish that result, one must show that the action of each sender in the symmetric separating equilibrium converges to the sender's bliss point at a rate that is rapid enough to overcome the greater waste associated with the

¹Signaling models with heterogeneous bliss points are widely employed in the literature. Examples include Spence (2002) on signaling with productive educational investments; Mailath (1989) on price signaling; Banks (1990) on political competition; Miller and Rock (1985) on dividend signaling; Bernheim (1994) on conformity; Bagwell and Bernheim (1986), Ireland (1994), and Corneo and Jeanne (1997) on conspicuous consumption; Bernheim and Severinov (2016) on bequests; Bernheim and Andreoni (2009) on fairness; and Bernheim and Bodoh-Creed (2020) on decisive leadership.

²Such a generalization is necessary to study applications such as conspicuous consumption, where the sender's utility is defined over bundles of goods that may be substitutes or complements for one another.

rising number of actions. We show by way of example that, for models with homogeneous bliss points, the convergence of actions to bliss points can be too slow to ensure the disappearance of signaling waste in the limit. However, heterogeneous bliss points lead to more rapid convergence, and consequently allow us to establish our main result under reasonably general conditions.

In some cases, the very same technological developments that proliferate signaling opportunities also allow people to reach larger audiences. In particular, online social networks (OSNs) such as Facebook, Twitter, Snapchat, and Weibo have become virtually ubiquitous components of social interaction.³ These platforms make it possible for users to stay connected with family and friends while expanding their circle of acquaintances, and to present all of these online contacts with richly textured depictions of their lives. Any development that expands audiences also potentially increases the reputational benefits from signaling, which may prevent total signaling waste from declining. Modeling such phenomena therefore requires us to consider settings in which increases in the number of visible actions are accompanied by increases in the scale of signaling benefits.

We therefore consider an extension of our model that allows us to explore the interaction between signal proliferation and audience augmentation, and to characterize their joint effects on welfare. 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 informational stakes, except in cases where there are weakly increasing returns to audience size.

In Section 2, we place our work within the existing literature. Section 3 introduces our baseline model. Section 4 presents some examples that illustrate various aspects of our analysis. Section 5 focuses on separating equilibria satisfying a belief refinement, while Section 6 concerns the most efficient separating equilibria. Section 7 addresses the simultaneous effects of signal proliferation and audience augmentation. Section 8 concludes. In Appendix A, we extend the analysis to more general settings with nonadditive asymmetric aggregators. All proofs appear in Appendix B.

2. RELATED LITERATURE

Beginning with the seminal contribution of Spence (1973), the literature on signaling has emphasized that incentive-compatible information transmission generally entails inefficiencies (see, e.g., Spence (2002)). Exceptions to this principle have occasionally surfaced, primarily in the literature on *cheap talk*.

³According 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 (2018)). Among OSN users between the ages of 18–49, this online activity consumes an average of roughly 6.5 hours per week (Nielsen (2016)).

A defining feature of the cheap talk framework is that signals are costless (Crawford and Sobel (1982)). Consequently, inefficiencies arise when the transmitted information is incomplete, rather than when the sender fails to select the action she would most prefer in a setting with full information. A few papers in the cheap-talk literature identify conditions under which senders can reveal all their private information through costless statements. Battiglini (2002) shows that full revelation is possible in settings where multiple senders compete with each other to influence the receiver. Chakraborty and Harbaugh (2007, 2010) demonstrate that, if a sender seeks to influence multiple decisions by a receiver, she can often make credible *comparative* statements across those decisions (see also Lipnowski and Ravid (2020)). For example, although a professor may have difficulty convincing potential employers of her students' absolute abilities, she may be able to convey a credible ranking. Furthermore, if the number of students is sufficiently large, employers may be able to infer absolute ability from rank, in which case full revelation is achieved asymptotically.

While the aforementioned papers are concerned with the same ultimate possibility as our investigation—costless revelation of all private information—the mechanisms are only distantly related. We can think of signaling models as differentiated with respect to the magnitude of the loss the sender incurs when deviating from her first-best option. Cheap-talk models lie at one extreme end of that spectrum inasmuch as the aforementioned loss is zero. In our setting, increasing the number of signals while preserving the magnitude of the loss from any single deviation is much like increasing the cost of deviations from the first-best action in settings with a single signal. (To be clear, while these comparative statics are related, they are not equivalent, because the proliferation of actions allows for asymmetric choices.) Our focus is therefore, in effect, on the opposite end of the spectrum, where deviations from first-best choices lead to large losses. Asymptotic efficiency obtains in our framework not because signaling reveals more information as the number of signals grows, but rather because fully revealing actions converge sufficiently rapidly to first-best choices.

There is a closer formal connection between our work and that of Kartik (2009), who studies single-action signaling models in which agents report private information but incur costs if they lie (see also Kartik, Ottaviani, and Squintani (2007)). Applying the intuition from the preceding paragraph (and keeping in mind the same formal qualification concerning asymmetric choices), we see that increasing the costs of a single potential lie is much like increasing the number of potential lies while preserving the costs of each individual lie. Kartik (2009) shows that all statements converge to the truth as the costs of lying rise, which is the analog of actions converging to bliss points as the number of actions rise in our setting.⁴ However, he does not investigate the behavior of total lying costs in the limit, and it is not clear from his analysis whether those costs fall because people tell smaller lies, or rise because the cost of any given lie increases. Nor do these papers make a connection to models with multiple signaling actions.

Our paper is also related to a branch of the theoretical biology literature concerning the “Handicap Principle,” which holds that the credible transmission of information

⁴In his setting, an upper bound on the set of potential signals leads to pooling. He also shows that the pool shrinks, and consequently that the equilibrium becomes more revealing, as lying costs rise.

requires a costly signal (Zahavi (1979), Grafen (1990)). 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 (1995), Számado (1999, 2011)). 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 (2001) and Bergstrom, Számado, and Lachmann (2001)).⁵ In contrast, in our setting, asymptotically costless signaling is a robust phenomena.

With respect to the specific phenomenon of signaling through online social networks, there is also a small theoretical literature that examines the informational content of an individual's network connections. Donath (2002) and Donath and Boyd (2004) study social connections as a credible signal of identity in an otherwise anonymous virtual community, and Donath (2008) draws out the implications of these ideas for network design. None of these studies focuses on the fact that these technologies contribute to the pervasiveness of signaling.

3. THE BASIC MODEL

We consider a signaling model in which a sender takes N publicly observable actions simultaneously. Because our analysis concerns limiting outcomes as N grows without bound, we will write the action vector as $\mathbf{a} = (a_1, \dots, a_N) \in A^N \subset \mathbb{R}_+^N$. The sender possesses private information (her “type” $t \in [\underline{t}, \bar{t}] = T \subset \mathbb{R}$) pertaining to the costs and benefits of the action $\mathbf{a}$. A second party, the receiver, observes $\mathbf{a}$ and draws inferences about the sender's type. The sender cares about those inferences, and consequently takes them into account when choosing $\mathbf{a}$. We will take both T and A to be compact intervals.

The sender's direct utility from action $\mathbf{a}$ in the absence of any signaling incentive (i.e., with complete information) is given by a continuous function $\pi^N(\mathbf{a}, t)$. For most of our analysis, we will specialize to the case in which $\pi^N(\mathbf{a}, t)$ is additively separable and symmetric across actions:

$$\pi^N(\mathbf{a}, t) = \sum_{i=1}^N \pi(a_i, t)$$

In Appendix A, we develop an extension to cases with nonseparable and/or asymmetric utility.

We define the sender's *bliss point* (or *first-best* action) as follows:

$$a_{BP}(t) = \arg \max_a \pi(a, t). \quad (1)$$

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

In words, $\mathbf{a}_{BP}(t) = (a_{BP}(t), a_{BP}(t), \dots, a_{BP}(t)) \in \mathbb{R}_+^N$ is the action the sender would choose if her type were publicly observed. If the agent takes an action $\mathbf{a} \neq \mathbf{a}_{BP}(t)$ in equilibrium, then $\pi^N(\mathbf{a}, t) - \pi^N(\mathbf{a}_{BP}(t), t) < 0$ is the *total waste from signaling*.⁶

Having observed $\mathbf{a}$, the receiver uses Bayes's rule to form a posterior belief about the sender's type. We use $\delta(\mathbf{a}) \in \Delta(T)$ to denote the belief of a receiver who observes action $\mathbf{a}$, where $\Delta(T)$ is the set of Borel measures over T . We refer to $\delta \in \Delta(T)$ 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 $\hat{t}(\mathbf{a})$, which is derived from the sender's strategy. When convenient, we suppress the arguments of δ and $\hat{t}$, and refer to the sender as "choosing" the receiver's perception.

Given the receiver's perception (δ), the sender receives benefits $B(t, \delta)$.⁷ In many signaling models, the receiver responds to the sender's signal, and one can think of $B(t, \delta)$ as a reduced form representation of the utility the sender derives from this response. The sender's total utility is

$$U^N(\mathbf{a}, \delta, t) = B(t, \delta) + \pi^N(\mathbf{a}, t). \quad (2)$$

In cases where δ is a degenerate distribution that places all weight on some perception $\hat{t}$, with a slight abuse of notation we write the utility function as $U^N(\mathbf{a}, \hat{t}, t)$ and the benefit function as $B(t, \hat{t})$.

We make several assumptions that are easily verified in applications.⁸ Our first assumption, that each agent has a unique optimal choice under complete information, simplifies our notation and some of the arguments.

ASSUMPTION 1. $a_{BP}(t)$ exists, is unique, and is continuous in t .⁹

To analyze the costs of signaling when the equilibrium selection for type t approaches $a_{BP}(t)$, we use Taylor series expansions of $\pi(\mathbf{a}, t)$ around $(a_{BP}(t), t)$. Assumption 2 enables us to employ third-order expansions and ensures that second-order effects dominate any higher-order effects for large N .¹⁰

⁶One can think of this model as describing a setting with $M \geq N$ actions, N of which are public and $M - N$ of which are private. Because a type- t sender chooses $a_{BP}(t)$ for each private action, the total payoff from the private actions is fixed at $(M - N)\pi(a_{BP}(t), t)$, so they have no effect on the set of signaling equilibria. In taking the limit as $N \rightarrow \infty$, we must then either (i) take $M = +\infty$ from the outset, which requires us to normalize utility so that $\pi(t, t) = 0$ in order to avoid unbounded sums, or (ii) allow M to grow along with N . The latter alternative may at first seem somewhat unattractive because it brings new actions into being as the scope of observability expands. It is therefore important to bear in mind that, in characterizing limiting behavior, our object is to approximate behavior for large N and M . Because M ends up being inconsequential, the analytical fiction of increasing M is innocuous.

⁷It is straightforward to allow for the possibility that B depends on the action $\mathbf{a}$. In that case, we would require the existence of first, second, and third derivatives with respect to the elements of $\mathbf{a}$, as well as upper bounds on the third derivatives.

⁸By convention, throughout this paper, any assumption that references a derivative assumes the existence of that derivative. We use subscripts to denote partial derivatives with respect to the subscripted variable.

⁹Because π is continuous, uniqueness of $a_{BP}(t)$ implies continuity of $a_{BP}(t)$, so the continuity assumption is redundant. We nevertheless spell it out explicitly to avoid subsequent confusion.

¹⁰Since we analyze actions in a neighborhood of the bliss point, first-order effects are absent as $\pi_a(a_{BP}(t), t) = 0$.

ASSUMPTION 2. $\pi_{aa}(a, t)$ is continuous in t , and for all $t \in T$ we have $\pi_{aa}(a_{BP}(t), t) < 0$. Furthermore, there exists a finite number $C > 0$ such that $|\pi_{aaa}(a, t)| \leq C$.

Our next assumption requires that first-order changes in type produce first-order changes in bliss points.¹¹

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

Our next two assumptions pertain to the properties of B , the signaling benefit function. Focusing on cases in which the receiver's belief about the sender's type is degenerate, Assumption 4 ensures that the sender's benefit is increasing in her perceived type. It also imposes a uniform bound on the rate of change.¹² Assumption 5 bounds the sender's benefit when the receiver is not confident about the sender's type, which may occur even in a separating equilibrium following a deviation by the sender.

ASSUMPTION 4. There exists $\gamma > 0$ such that $0 \leq B_{\hat{t}}(t, \hat{t}) \leq \gamma$.

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

Together, Assumptions 4 and 5 imply that there exist finite values $\underline{B}$ and $\overline{B}$ such that $B(t, \delta(a)) \in [\underline{B}, \overline{B}]$.

We define a *separating equilibrium* in the usual way, restricting attention to Perfect Bayesian equilibria so that the support of beliefs, δ , always lies within T , even for unchosen actions. We say that the equilibrium is *symmetric* if it employs the same mapping from types to actions for every component choice.

The signaling structure of our model fits many applications. For instance, we can interpret it as capturing settings with costs of lying, as in Kartik (2009). In that context, each action a_i consists of a statement concerning the sender's type, t , and the bliss point involves telling the truth ($a_{BP}(t) = t$). Under that interpretation, our analysis shows that, if maintaining a falsehood requires a sufficiently large number of lies, the freedom to lie becomes inconsequential.¹³ As another example, we can interpret our model as describing conspicuous consumption, as in Ireland (1994), Bagwell and Bernheim (1986), and Corneo and Jeanne (1997). In that context, each action a_i involves some costly and visible purchase, which may serve to display the sender's wealth. The bliss-point function

¹¹When the first-order approach is applicable, implicit differentiation yields $a'_{BP}(t) = -\pi_{at}(a_{BP}(t), t)/\pi_{aa}(a_{BP}(t), t)$. In that case, Assumption 3 follows if π_{aa} is strictly negative and uniformly bounded while π_{at} is strictly positive and uniformly bounded away from zero.

¹²The assumption requires the benefit function B to be differentiable in the public perception, $\hat{t}$. In some applications, B may be discontinuous, for example, if the receiver chooses from a finite set of responses. For such settings, we can justify the differentiability assumption by positing a degree of uncertainty on the part of the sender concerning the receiver's objectives.

¹³To be clear, our analysis does not contemplate the possibility that telling many lies might desensitize the sender to the cost of lying. By way of analogy to our analysis of audience size in Section 7, we would expect the result to depend on the rate of desensitization.

$a_{BP}(t)$ captures the purchases the sender would make based solely on her intrinsic preferences. Bliss points are increasing in t because wealthier people tend to prefer more extravagant alternatives. Under that interpretation, our analysis shows that the pervasive observability of purchases renders signaling distortions unimportant not only for individual purchases, but also in total.

4. SOME EXAMPLES

This section exhibits two examples, the purpose of which is to show that convergence to costless revelation of private information obtains for some types of separating equilibria in some models, but not universally. The examples help build intuition concerning the prerequisites for the convergence property.

Our first example exploits the fact that the types of multidimensional signaling models described in Section 3 typically give rise to multiple separating equilibria. We exhibit two types of equilibria. In one, the sender chooses a symmetric action that converges to her bliss point, causing total signaling waste to disappear at the rate $1/N$ as N grows. In the other, the sender chooses an asymmetric action with a component that does not converge to her bliss point, leaving signaling waste undiminished as N grows. The contrast between these two outcomes focuses our subsequent attention on criteria for refining the set of separating equilibria.

EXAMPLE 1. Assume that $B(t, \hat{t}) = \hat{t}$ and $\pi(a, t) = -\lambda(a - t)^2$, and $T = [0, 1]$. The bliss points are (obviously) $a_{BP}(t) = t$.

In any symmetric separating equilibrium with actions $a_{SEP}(t, N)$, we can write the equilibrium utility of a type- t agent who mimics the equilibrium action of a type- $\hat{t}$ agent as

$$V(\hat{t}, t) = \hat{t} - N\lambda(a_{SEP}(\hat{t}, N) - t)^2$$

Because $a_{SEP}(t, N)$ is the optimal choice for type t , this expression is maximized at $\hat{t} = t$. The corresponding first-order condition for optimization over $\hat{t}$ yields an ODE that defines the symmetric separating equilibrium:

$$\left. \frac{\partial a_{SEP}(\hat{t}, N)}{\partial \hat{t}} \right|_{\hat{t}=t} = \frac{1}{2\lambda N(a_{SEP}(t, N) - t)}. \quad (3)$$

Treating N as fixed for the purpose of solving the ODE, we use the change of variables $z(t) = a_{SEP}(t, N) - t$. Noting that $z'(t) = a'_{SEP}(t, N) - 1$, we can write $dz/dt = (1 - 2\lambda Nz)/2\lambda Nz$, which in turn implies the inverse ODE, $dt/dz = 2\lambda Nz/(1 - 2\lambda Nz)$. Using the fact that $2\lambda Nz/(1 - 2\lambda Nz) = 1/(1 - 2\lambda Nz) - 1$, we can solve by integrating

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

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

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

Reversing our change of variables and rearranging, we find

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

The total signaling waste is then

$$N\lambda(a_{\text{SEP}}(t, N) - t)^2 \leq N\lambda\left(\frac{1}{2N\lambda}\right)^2 = \frac{1}{4\lambda N},$$

which plainly converges to zero at the rate $1/N$ as N grows without bound.

Continuous-type models do not give rise to inefficient *symmetric* separating equilibria under the sufficient conditions identified in Mailath (1987) and Mailath and von Thadden (2013), which this example satisfies. However, they can give rise to inefficient *asymmetric* separating equilibria. As an illustration, we will examine separating equilibria in which a type t sender selects $(a_{\text{BP}}(t), \dots, a_{\text{BP}}(t), a_{\text{ASEP}}(t, N))$. In other words, for her first $N - 1$ actions, she chooses her bliss point, but for her N th action, she chooses $a_{\text{ASEP}}(t, N)$. Using the fact that the derivative of her utility with respect to the first $N - 1$ actions is zero at her bliss point, we can convert her first-order condition for optimization over $\hat{t}$ into an ODE, exactly as before:

$$\left. \frac{\partial a_{\text{ASEP}}(t, N)}{\partial \hat{t}} \right|_{\hat{t}=t} = \frac{1}{2\lambda(a_{\text{ASEP}}(t, N) - t)}.$$

This ODE is identical to equation (3), except there is no N in the denominator. It follows that $a_{\text{ASEP}}(t, N)$ coincides with $a_{\text{SEP}}(t, 1)$ for all N .¹⁴ From this observation, we conclude that the total signaling waste is fixed for these asymmetric equilibria, and consequently does not vanish as $N \rightarrow \infty$. $\diamond$

From the contrast between the two equilibria described in Example 1, one might conjecture that signaling waste disappears in the limit as long every component action converges to the sender's bliss point. However, that conjecture is incorrect. Because N also serves as a multiplier that magnifies the remaining signaling waste in each individual action, waste does not vanish unless actions converge to bliss points rapidly enough to overcome the expanding scope of signaling. Our next example illustrates this point using a model involving a setting with *homogeneous* bliss points that otherwise satisfies our assumptions.¹⁵ Focusing on the (unique) symmetric separating equilibrium, we

¹⁴As long as λ is not too large, one obtains a Perfect Bayesian separating equilibrium by setting $\hat{t}(a^N) = \underline{t}$ for all unchosen a^N .

¹⁵To be clear, there are conditions under which signaling costs can converge to zero as $N \rightarrow \infty$ even with homogeneous bliss points. Suppose there are two types, $t \in \{1, 2\}$, who share the bliss point $a = 0$, and that the reputational benefits B depend only on the perceived type, $\hat{t}$. It is then straightforward to verify that the equilibrium costs of signaling converge to zero as $N \rightarrow \infty$ iff $\lim_{a \rightarrow 0} (\pi(a, 2) - \pi(0, 2)) / (\pi(a, 1) - \pi(0, 1)) = 0$ (henceforth, condition L). The idea is that the equilibrium separating action, a^* , is determined by the equation $N[\pi(0, 1) - \pi(a^*, 1)] = B(2)$. Thus we have $N[\pi(0, 2) - \pi(a^*, 2)] = [(\pi(0, 2) - \pi(a^*, 2)) / (\pi(0, 1) - \pi(a^*, 1))]B(2)$. Because $\lim_{N \rightarrow \infty} a^* = 0$, the conclusion is immediate. The argument extends directly to cases with finite type sets under the assumption that a condition analogous to L holds for all consecutive

show that actions converge to bliss points at the rate $1/\sqrt{N}$, which leaves the aggregate waste from signaling unchanged. Thus, in contrast to the first equilibrium exhibited in Example 1, the convergence of actions to bliss points is too slow to eliminate waste.

EXAMPLE 2. Suppose $B(t, \hat{t}) = \hat{t}$, $\pi(a, t) = -a^2/(t + \gamma)$ for some fixed $\gamma > 0$, and $T = [0, 1]$. Notice that the bliss points are homogenous ($a_{BP}(t) = 0$ for all t), and that the model satisfies the Spence–Mirrlees single-crossing property. The first-order condition for optimization over $\hat{t}$ yields an ODE that defines the symmetric separating equilibrium:

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

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

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

Integrating both sides and using the initial condition that $a_{SEP}(0, N) = 0$ yields

$$Na_{SEP}(t, N)^2 = \frac{1}{2}(t + \gamma)^2 - \frac{\gamma^2}{2}.$$

Thus, the total cost of signaling, $Na_{SEP}(t, N)^2/(t + \gamma)$, is invariant with respect to N , even though each action converges to the common bliss point at the rate $1/\sqrt{N}$. $\diamond$

To a limited extent, one can build intuition for our main results from the preceding examples by considering settings in which the number of sender types is *finite*. When bliss points are heterogeneous, it is easy to see that, with $|T| < +\infty$, $a_{BP}(t)$ (“truth telling”) is *always* a separating equilibrium for N sufficiently large: finite reputational gains eventually become insufficient to justify the costs of choosing a discretely higher type’s bliss point. It follows that, for efficient separating equilibria, signaling waste disappears entirely for large N .¹⁶ This reasoning obviously breaks down when bliss points are homogeneous.

While instructive, these observations fall short of providing a helpful intuitive account of our main results. In particular, the finite-type case provides no insight as to why, with heterogeneous bliss points and a continuum of types, the reduction in signaling waste per action resulting from the convergence of actions to bliss points outpaces the scale effect of increasing N . Even if we take the view that the number of types, $|T|$,

pairs of types. We suspect, but have not verified, that the argument also extends to cases with continuous type sets. The function $\pi(a, t) = -a^t$ satisfies condition L , as well as the analogous condition for finite type sets (provided $t > 0$). Because this function violates the single-crossing property, existence of a separating equilibrium is guaranteed only for sufficiently large N . While it is interesting that signaling costs can in principle vanish as $N \rightarrow \infty$ even with homogeneous bliss points, this outcome requires assumptions that are not typical of the literature. In contrast, with heterogeneous bliss points, the result obtains under standard assumptions.

¹⁶In finite-type models, there are also typically inefficient symmetric separating equilibria for which total signaling costs need not disappear as N grows without bound.

is *actually* finite (but large), studying cases with fixed $|T|$ and $N \rightarrow \infty$ could be misleading. With large N and large $|T|$, truth telling is only an equilibrium in settings with heterogeneous bliss points if N is large *relative* to $|T|$. If, on the contrary, $|T|$ is large relative to N , the limiting ($N \rightarrow \infty$) case of the continuous-type model likely does a better job of capturing the signaling distortions that remain. In effect, our analysis establishes that “truth telling” holds as an approximation for large $|T|$ and N even when $|T|$ is large relative to N .

Together, Examples 1 and 2 show that the asymptotic magnitude of signaling waste may depend on the class of equilibria considered. As noted in Section 1, we handle this issue through two distinct approaches. For the first approach (Section 5), we prove a convergence theorem that ensures the disappearance of aggregate signaling waste at the rate $1/\sqrt{N}$ for *all* separating equilibria satisfying a mild belief restriction. For our second approach, we avoid imposing belief restrictions, and prove a stronger convergence result for symmetric separating equilibria, which in turn has implications for waste-minimizing separating equilibria.

5. SIGNALING WASTE IN EQUILIBRIA SATISFYING A BELIEF RESTRICTION

Our first approach to issues arising from the potential multiplicity of signaling equilibria involves the imposition of a minimal “plausibility” restriction on equilibrium beliefs. We characterize the signaling waste associated with *all* separating equilibria satisfying this belief restriction.

We formulate the restriction as follows. For each action a , we say that sender type t belongs to the *plausible set* $\mathcal{P}(a)$ if there is some pattern of conceivable receiver reactions (i.e., receiver beliefs about the sender) for which t would find a preferable to t ’s bliss point, $a_{BP}(t)$. Formally,

$$\mathcal{P}(a) \equiv \{t : \underline{B} + \pi^N(a_{BP}(t), t) \leq \bar{B} + \pi^N(a, t)\}.$$

Imagine the receiver observes a sender choosing an out-of-equilibrium action a , and that $T \supset \mathcal{P}(a) \neq \emptyset$. The receiver may try to rationalize the existence of this deviation by attributing it to a misunderstanding on the part of the sender concerning the receiver’s reactions (i.e., the sender may be wrong about $\hat{t}(\tilde{a})$ for certain $\tilde{a} \in \mathbb{R}_+^N$). By construction, the sender’s misunderstanding could in principle justify the choice of a if $t \in \mathcal{P}(a)$, but not if $t \notin \mathcal{P}(a)$. Therefore, the receiver cannot reasonably attribute the choice of a to any type outside of $\mathcal{P}(a)$.¹⁷ Through this reasoning, we arrive at the following belief restriction.

DEFINITION 1. The receiver’s beliefs satisfy the *dominance refinement* if, upon observing any a for which $\mathcal{P}(a) \neq \emptyset$, the receiver is certain that $t \in \mathcal{P}(a)$.

¹⁷Our dominance refinement is similar to the concept of “elimination of type-message pairs by dominance” (ETMPD) proposed in Cho and Kreps (1987), although the notation we use to describe our games is different (e.g., we do not explicitly model a reaction choice by the receiver). Since ETMPD is a weaker refinement than the *Intuitive Criterion* defined by Cho and Kreps (1987), our refinement is also weaker than the Intuitive Criterion. For similar reasons, our refinement is also weaker than the *Divinity* refinement proposed by Banks and Sobel (1987).

Our interest lies in the properties of equilibria when they exist, rather than in the conditions that ensure existence (a question addressed in other research). Accordingly, we assume existence directly. We do not, however, insist on the existence of a separating equilibrium satisfying regularity properties such as continuity or symmetry.

ASSUMPTION 6. *A separating equilibrium satisfying the dominance refinement exists for all N .*

The main result for this section tells us that type t 's total waste from signaling in *all* equilibria satisfying the dominance refinement converges to zero as the number of activities grows without bound, and that this convergence is uniform for $t \in T$. Uniformity is important because it implies that aggregate waste exhibits convergence at the same rate.¹⁸

THEOREM 1. *Under Assumptions 1–6, there exists a finite constant $K > 0$ such that, for sufficiently large N , the total waste from signaling in all separating equilibria satisfying the dominance refinement, $\pi^N(a_{BP}(t), t) - \pi^N(a(t; N), t)$ (where $a(t; N)$ is the equilibrium signaling action function) is bounded above by $K/\sqrt{N}$ for all types $t \in T$.*

To build intuition for this result, start by thinking about type t 's first-best action, $a_{BP}(t)$. As N grows, there are fewer and fewer types t' who prefer $a_{BP}(t)$ along with the best possible inference to $a_{BP}(t')$ along with the worst possible inference. The plausible set for $a_{BP}(t)$ therefore lies within a shrinking neighborhood of t , which means the associated inference, $\hat{t}(a_{BP}(t))$ converges to t . It follows that $B(\hat{t}(a_{BP}(t)), t)$ converges to $B(t, t)$. Alternatively, type t 's equilibrium action, $a(t, N)$, leads to the inference $\hat{t}(a(t, N)) = t$, and hence to the signaling benefit $B(t, t)$. The total reputational gain from choosing $a(t, N)$ rather than $a_{BP}(t)$, that is, the difference $B(t, t) - B(\hat{t}(a_{BP}(t)), t)$ therefore converges to zero. But incentive compatibility implies that this reputational gain is an upper bound on the signaling waste, $N[\pi(a_{BP}(t), t) - \pi(a(t, N), t)]$. The total signaling waste must therefore converge to zero as N grows without bound.

It is instructive to examine the implications of the dominance refinement in Example 1 from Section 4. Consider the inefficient asymmetric separating equilibrium. As we have just explained, the plausible set for $a_{BP}(t)$ lies within a vanishing neighborhood of t as N grows. Once this neighborhood becomes sufficiently small, a type t sender would prefer to choose her bliss point, $a_{BP}(t)$, and accept an inference about her type that may be slightly less than t , rather than choose $(a_{BP}(t), \dots, a_{BP}(t), a_{ASEP}(t))$ and have others recognize her type correctly. Consequently, these equilibria do not satisfy the dominance refinement for sufficiently large N . The symmetric separating equilibria does not suffer from this problem because the choice between, on the one hand, $a_{BP}(t)$ along with an inference about the sender's type that is slightly below t , and on the other hand, $a_{SEP}(t, N)$ along with an inference that her type is t , does not tip toward the former as N grows (because $a_{SEP}(t, N)$ converges to $a_{BP}(t)$).

¹⁸The proof of this result makes use of the upper bound on π_{aaa} in Assumption 2, but not the lower bound. We use both bounds in the proof of Theorem 2.

Although reasonably general, the convergence property identified in Theorem 1 is not universal. The following example shows that the total costs of signaling need not shrink when the model violates the technical assumptions listed in Section 3.

EXAMPLE 3. Let $T = [0, 1]$. Suppose

$$B(t, \hat{t}) = \begin{cases} \sqrt{\hat{t} - t} & \text{for } \hat{t} \geq t \\ -\infty & \text{otherwise.} \end{cases}$$

Also assume that there is some constant $q > 0$ such that

$$\pi(a, t) = \begin{cases} -(a - t)^2 & \text{for } |a - t| \leq q \\ -\infty & \text{otherwise.} \end{cases}$$

Notice that this model entails heterogeneous bliss points. However, it violates our assumptions about differentiability and boundedness.

We claim that, for all N , there exists a separating equilibrium such that $\mathbf{a}(t, N) = (t + q, t, \dots, t)$, in which receivers attribute any off-path action $\mathbf{a}^o$ to type $t^o(\mathbf{a}^o) = \max_i a_i^o - q$. A deviation to any $\mathbf{a}'$ (whether on-path or off-path) yields infinitely negative direct utility for type t if $a'_i > t + q$ for some i , and infinitely negative signaling benefits for type t if all $a'_i < t + q$. The only remaining options entail $a'_i \leq t + q$ for all i , with equality for some i . But receivers would attribute any such action to type t , and consequently the sender's payoff would be no higher than in equilibrium. It is also easily verified that this equilibrium satisfies the dominance refinement.¹⁹ Notice that the signaling waste associated with these equilibria, q^2 , does not shrink with N . $\diamond$

A prominent feature of our basic model is that utility is additive over actions. This functional restriction may be reasonable in some settings but not in others. In Appendix A, we identify conditions under which Theorem 1 generalizes to settings with nonadditive asymmetric utility aggregators.

6. SIGNALING WASTE IN WASTE-MINIMIZING SEPARATING EQUILIBRIA

Our objective in this section is to evaluate the possibility of near-costless transmission of all private information by bounding the signaling waste in the most efficient (i.e., waste-minimizing) separating equilibria. As explained in Section 1, our strategy is to provide an asymptotic bound on waste in symmetric separating equilibria, which are analytically tractable, and then to note that the same bound must apply to waste-minimizing separating equilibria. We also provide conditions under which the symmetric separating

¹⁹If we follow the convention that $t \in \mathcal{P}(\mathbf{a})$ whenever $\underline{B} + \pi^N(\mathbf{a}_{BP}(t), t)$ is $-\infty$, even if $\bar{B} + N\pi^N(\mathbf{a}, t)$ is also $-\infty$, then $\mathcal{P}(\mathbf{a}^o) = T$ for all off-path choices $\mathbf{a}^o$, from which it follows that all equilibria satisfy the dominance refinement. If instead we follow the convention that $t \notin \mathcal{P}(\mathbf{a})$ whenever $\bar{B} + N\pi^N(\mathbf{a}, t)$ is $-\infty$, then for any off-path action $\mathbf{a}^o$, if $\max_i a_i^o - \min_i a_i^o \leq 2q$, we have $t^o(\mathbf{a}^o) \in \mathcal{P}(\mathbf{a}^o)$, and if $\max_i a_i^o - \min_i a_i^o > 2q$, we have $\mathcal{P}(\mathbf{a}^o) = \emptyset$. In either case, the prescribed beliefs satisfy the dominance refinement.

equilibrium is in fact waste-minimizing. It is worth noting that several leading refinements direct our attention to the most efficient separating equilibria in broad classes of signaling models (see, e.g., Cho and Kreps (1987), and Banks and Sobel (1987)).

If we imposed symmetry of actions as a restriction, the sender's utility function would become

$$\hat{U}^N(a, \delta, t) = B(t, \delta) + N\pi(a, t).$$

Interpreted as a setting with a single signal, the symmetry-restricted formulation falls within the general frameworks studied by Mailath (1987) and Mailath and von Thadden (2013), and our assumptions are compatible with (but do not imply) the sufficient conditions provided in those papers for existence and uniqueness of a separating function $a_{\text{SEP}}(t, N)$. Furthermore, when the symmetry restriction is removed, we can construct a separating equilibrium based on the action function $\mathbf{a}_{\text{SEP}}(t, N) \equiv (a_{\text{SEP}}(t, N), \dots, a_{\text{SEP}}(t, N))$, where we extend beliefs to all action vectors $\mathbf{a}^N$ so that any such equilibrium remains an equilibrium; moreover, no new symmetric equilibria can appear.²⁰ It follows that we could in principle guarantee the existence and uniqueness of a symmetric separating equilibrium by invoking sufficient conditions from the literature. Because that approach would potentially distract from the assumptions that most directly drive our results and would also sacrifice some generality, we instead simply assume existence, uniqueness, monotonicity, and continuity.

ASSUMPTION 7. *There exists a unique symmetric separating equilibrium for the symmetry-constrained signaling game, $\mathbf{a}_{\text{SEP}}(t, N)$, that is increasing and continuous in t .*

Our analysis of symmetric separating equilibria begins with the simple observation that actions must converge to bliss points as N grows without bound. In the Appendix, we formalize this observation as Lemma 1, which also provides a bound on the rate of convergence (which we subsequently tighten in the proof of Theorem 2), and establishes that convergence is uniform. The intuition for this result is straightforward: because N multiplies the direct utility function π , a sufficient increase in N must cause the utility sacrificed when choosing any given symmetric action other than the sender's bliss point to exceed any possible reputational gain.

To establish our main convergence result, we show that the action of each sender in the symmetric separating equilibrium converges to the sender's bliss point at a rate that is rapid enough to overcome the greater waste associated with the rising number of actions. Thus, the total waste from signaling converges to zero in the limit as the number

²⁰The simplest way to extend beliefs is to let $\hat{t}(\mathbf{a}) = \underline{t}$ for all actions not chosen in equilibrium, that is, $\mathbf{a} \notin A_e^N \equiv \{\mathbf{a}' \mid \mathbf{a}' = \mathbf{a}_{\text{SEP}}(t, N) \text{ for some } t \in [\underline{t}, \hat{t}]\}$. Combining the fact that $U^N(\mathbf{a}, \hat{t}(\mathbf{a}), t) = U^N(\mathbf{a}, \underline{t}, t) \leq U^N(\mathbf{a}_{\text{BP}}(t), \hat{t}(\mathbf{a}_{\text{BP}}(t)), t)$ for $\mathbf{a} \notin A_e^N$ with the incentive-compatibility condition for the symmetry-restricted equilibrium, $U^N(\mathbf{a}_{\text{BP}}(t), \hat{t}(\mathbf{a}_{\text{BP}}(t)), t) \leq U^N(\mathbf{a}_{\text{SEP}}(t, N), t, t)$, we have $U^N(\mathbf{a}, \hat{t}(\mathbf{a}), t) \leq U^N(\mathbf{a}_{\text{SEP}}(t, N), t, t)$, which ensures incentive compatibility over the entire action space. Under the additional assumptions listed in Theorem 3, one can establish the existence of another set of beliefs defined over the entire action space that ensures incentive compatibility while satisfying the dominance refinement. To see that no new symmetric equilibria can appear, notice that any symmetric equilibrium of the unrestricted game remains an equilibrium when asymmetric choices are disallowed.

of actions grows, so that the equilibrium approximates the first-best outcome for large N . Moreover, the convergence is uniform over $t \in T$.

THEOREM 2. *Under Assumptions 1–5 and 7, there exists a finite constant $K > 0$ such that, for sufficiently large N , the total waste from signaling in the unique symmetric separating equilibrium, $N[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t; N), t)]$, is bounded above by K/N for all types $t \in T$.*

As with Theorem 1, uniformity is important because it implies that aggregate waste converges to zero at least as rapidly as the uniform bound implies. It follows that per capita waste in the waste-minimizing separating equilibrium is also bounded above by K/N .

Although reasonably general, the convergence property identified in Theorem 2 is not universal. As we now show by way of example, it may fail when the model violates the technical assumptions listed in Section 3.

EXAMPLE 4 (continued). Consider once again the model described in Example 3. We claim that, for all N , there exists a symmetric separating equilibrium such that $a(t, N) = t + q$, in which receivers attribute any off-path action to type $\underline{t}$. A deviation by type t to $a(t', n)$ for some $t' > t$ yields infinitely negative direct utility, and any other deviation yields infinitely negative signaling benefits. Notice that the signaling waste associated with these equilibria, Nq^2 , grows without bound as $N \rightarrow \infty$. Notice also that, regardless of N , every symmetric separating function must assign $a(t, N) = t + q$ at any point t for which a has a bounded derivative. Otherwise, type t 's marginal benefit from mimicking some $\hat{t}$ slightly greater than t , $B_{\hat{t}}(t, t) = +\infty$, would exceed the associated marginal cost, $2\lambda N(a(t, N) - t)$. $\diamond$

We conclude this section by providing sufficient conditions under which the symmetric separating equilibrium is in fact waste-minimizing. Specifically, we have the following.

CONDITION (i). $\pi(a, t)$ is supermodular in (a, t) .

CONDITION (ii). For all t and $a > a_{BP}(t)$, we have $\pi_a(a, t) < 0$.

CONDITION (iii). For all t , $\pi_{at}(a, t)/\pi_a(a, t)$ is increasing in a on $\{a \mid a > a_{BP}(t)\}$.

CONDITION (iv). There exists some finite $K > 0$ such that for all $t, \hat{t} \in T$, and $a \in A^N$, we have $|U_t^N(t, \hat{t}, a)| < K$.

CONDITION (v). There exists $a \in A$ such that $a > a_{BP}(\bar{t})$ and, for all $t \in T$, $\pi(a_{BP}(t), t) - \pi(a, t) > \bar{B} - \underline{B}$.

Condition (i) is a form of single-crossing. Condition (ii) simply states that utility is decreasing in a once a exceeds the bliss point for t . Condition (iii) implies that the

marginal cost of a becomes more sensitive (in relative terms) to type as a increases. Condition (iv) bounds the derivative of utility with respect to type. Condition (v) ensures that $\mathcal{A}$ includes large actions that no sender would be willing to take irrespective of the resulting inference. The first four requirements are satisfied for common specifications such as $\pi(a, t) = -(a - t)^2$.²¹ For the same specification, the last condition only requires the upper bound on $\mathcal{A}$ to be sufficiently large.

The following theorem establishes the efficiency of symmetric separating equilibrium under the aforementioned conditions.

THEOREM 3. *Under Assumption 7 and Conditions (i)–(v), for any fixed N , the symmetric separating equilibrium maximizes the payoff for each type of sender relative to any other separating equilibrium. Moreover, if $\mathbf{a}^o(t, N)$ is the action function for a separating equilibrium such that $\mathbf{a}^o(t, N) \neq \mathbf{a}_{\text{SEP}}(t, N)$ for all $t > \underline{t}$, then the equilibrium associated with $\mathbf{a}_{\text{SEP}}(t, N)$ yields strictly higher payoffs than the one associated with $\mathbf{a}^o(t, N)$ on an open-dense subset of T .²²*

7. AUDIENCE AUGMENTATION VERSUS SIGNAL PROLIFERATION

At the outset of Section 1, we observed that technological developments have rendered modern life pervasively observable, thereby creating pervasive signaling opportunities. For instance, the typical user of a social networking platform such as Facebook or Instagram has almost unlimited opportunities to shape the way friends, family, and acquaintances perceive them by proliferating postings to provide detailed accounts of their daily experiences. The same principle applies to the professional sphere.²³ As a general matter, people respond to these opportunities as signaling theory predicts, in that they perform a “balancing act between self-expression and self-promotion” (van Dijck (2013)).

The proliferation of signals is, however, only part of the story. The same technological advances allow people to present richly textured depictions of their lives to substantially larger audiences. For example, in 2014, the typical Facebook user had 338 “friends” on the platform (Smith (2014)). Any development that expands audiences also potentially increases the reputational benefits from signaling, which may prevent total signaling waste from declining. Modeling such phenomena therefore requires us to consider settings in which increases in N are accompanied by increases in the scale of B .

In this section, we explore the interaction between signal proliferation and audience augmentation, and we characterize their joint effects on welfare. For this purpose, we employ an extension of the additive model from Section 3, which we interpret as follows. We think of the sender as creating M independent experiences, in each instance

²¹For condition (iv), recall that $T \times T \times \mathcal{A}$ is compact.

²²We conjecture that $\mathbf{a}_{\text{SEP}}(t, N)$ yields strictly higher payoffs than the one associated with $\mathbf{a}^o(t, N)$ on all $T \setminus \{t\}$. The challenge is to rule out the possibility that a discontinuity in $\mathbf{a}^o(t, N)$ appears at some t^* such that the difference in the utility type t^* derives from the two equilibria is zero. We leave this small detail unresolved.

²³Indeed, OSNs also convey information about pertinent personality characteristics to potential employers (e.g., Buffardi and Campbell (2008), Kluemper and Rosen (2009)). Notably, 60% of employed Facebook users are “friends” with coworkers (Drouin, O’Connor, Schmidt, and Miller (2015)).

by taking an action a . The resulting direct utility is given by $\pi(a, t)$. If actions were private, someone of type t would select $a_{BP}(t)$ in each case. People consider higher values of t more impressive, and those with larger t prefer higher actions when social reputation is not at stake. For example, wealthier people like to take more expensive vacations and dine at fancier restaurants, more adventurous people like to engage in more extreme leisure activities, and more popular people like to frequent more fashionable parties. Because each day presents more opportunities for additional activities, M is a large number.²⁴

The observability of actions depends on the prevailing technology, τ . (For example, we think of OSNs, collectively, as a single technology.) A technology τ provides opportunities for each person to make $N_\tau \leq M$ of their actions visible to F_τ social contacts.²⁵ We assume that the size of the audience potentially impacts the magnitude of signaling benefits. In particular, the type- t sender's utility function is

$$U_F^N(\mathbf{a}, \delta, t) = B(t, \delta, F) + \sum_{i=1}^M \pi(a_i, t).$$

We also make the following assumption concerning the signaling benefit function B .

ASSUMPTION 8. *There is a function $\gamma(F) > 0$ such that $0 \leq B_\tau(t, \hat{t}, F) \leq \gamma(F)$.*

As an example, suppose $B(t, \hat{t}, F) = \zeta(F)B(t, \hat{t})$. Then, for the bound $\gamma(F)$, we can use $\zeta(F) \max_{t, \hat{t} \in T} B_\tau(t, \hat{t})$. To the extent a larger audience provides greater incentives to signal, we would expect γ to be increasing in F . Within the social sphere, $\gamma(F)$ may flatten out rather quickly because people presumably derive most of the signaling benefits from interactions with immediate friends and family. Similarly, within the professional sphere, someone seeking a fixed number of discrete assignments (such as a job or a portfolio of projects) may experience decreasing returns to audience size.²⁶

We are interested in characterizing signaling outcomes for settings in which both N_τ and F_τ are large due to the prevailing technology τ . We can think of any such setting

²⁴For simplicity, we proceed as if all postings are simultaneous. In practice, the fact that OSN users post their experiences sequentially may complicate the signaling problem. The proper approach to modeling these dynamic considerations is, however, far from clear. The sequentiality of experience may be less analytically consequential for a user who curates a portfolio of postings to impress new viewers, or who cares primarily about the “long-run” impression created by her cumulative postings. We leave questions about signaling dynamics for future work.

²⁵Technically, our model assumes that all N_τ actions are *automatically* visible. In applications involving online social networks, posting is voluntary. Arguably, it might therefore be appropriate to modify our model to allow costless shrouding of each action. However, a close reading of Example 1 reveals that our characterization of equilibria would be unaffected by this modification. If someone shrouds an action, they presumably select their bliss point, which receivers can infer from other actions (in a separating equilibrium). Thus, for every equilibrium with voluntarily shrouded actions in the modified model, there is an equivalent equilibrium of the original model wherein all actions are visible and the signaling distortion is confined to the other actions.

²⁶For 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.

as a member of a sequence of signaling problems, the N th of which involves N visible actions and $\phi_\tau N$ audience members, where $\phi_\tau \equiv F_\tau/N_\tau$. We approximate behavior for large- N settings in which the ratio of the audience size to the number of visible actions equals the fixed number ϕ_τ .²⁷ In that way we obtain an approximation for the N_τ -th element of the series, which corresponds to technology τ .

For this modified model, a slight adjustment to the proof of Theorem 1 yields the following conclusion: under Assumptions 1–6, there exists a finite constant $K > 0$ such that, for sufficiently large N , the total waste from signaling in equilibria satisfying the dominance refinement, $\pi^N(a_{BP}(t), t) - \pi^N(a(t; N), t)$ (where $a(t; N)$ is the equilibrium signaling action function) is bounded above by $K\sqrt{\gamma(\phi_\tau N)^3/N}$ for all types $t \in T$.²⁸ For the purpose of illustration, assume that $\gamma(F) = \xi F^\sigma$. Then as long as $\sigma < \frac{1}{3}$, the aggregate costs of signaling vanish in the limit as $N \rightarrow \infty$. However, even when signaling waste does not become vanishingly small, it may still represent a vanishing fraction of the information stakes, which are proportional to $\gamma(F)$. In particular, the modified convergence result implies that $N[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t; N), t)]/\gamma(\phi_\tau N)$ is bounded above by $K\sqrt{\gamma(\phi_\tau N)/N}$ for large N . For the same functional form, this bound converges to zero for all $\sigma < 1$ – in other words, any degree of decreasing returns to audience size guarantees that signaling waste becomes negligible relative to the information stakes for large N .

As one would expect, a slight adjustment to the proof of Theorem 2 yields the following conclusion: under Assumptions 1–5 and 7, there exists a finite constant $K > 0$ such that, for sufficiently large N , the total waste from signaling in the unique symmetric separating equilibrium, $N[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t; N), t)]$, is bounded above by $K\gamma(\phi_\tau N)^2/N$ for all types $t \in T$.²⁹ For the case of $\gamma(F) = \xi F^\sigma$, the aggregate costs of signaling vanish in the limit as $N \rightarrow \infty$ if and only if $\sigma < 1/2$. However, observe that $N[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t; N), t)]/\gamma(\phi_\tau N)$ is bounded above by $\gamma(\phi_\tau N)/N$ for large N . It follows once again that, despite the slower rate of convergence, signaling waste becomes a vanishing small fraction of the information stakes as long as $\sigma < 1$ (i.e., for any degree of decreasing returns to audience size).

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 informational stakes, except in cases where there are weakly increasing returns to audience size.

²⁷To understand how we impose the constraint $N \leq M$ as N grows, see footnote 6.

²⁸Consider equation (7) in the proof of Theorem 1. The claim follows once we replace γ with $\gamma(F)$, and note that $K_1 = (2/\beta)\sqrt{(\underline{B} - \bar{B})/\bar{\pi}_{aa}} \leq (2/\beta)\sqrt{\gamma(F)(\underline{t} - \bar{t})/\bar{\pi}_{aa}}$ by Assumption 8.

²⁹Consider equation (14). Replacing γ with $\gamma(F)$ and $\bar{B} - \underline{B}$ with $\gamma(F)(\bar{t} - \underline{t})$, we obtain an expression proportional to $(\gamma(F)/\lambda N)^{0.75}$. A third iteration of the argument yields an expression with a form similar to equation (14) that is proportional to $(\gamma(F)/\lambda N)^{0.875}$, while a fourth iteration yields an expression proportional to $(\gamma(F)/N)^{0.9375}$, and so forth. Combining these observations with the Taylor approximation of $\pi(a, t)$ in equation (12) yields the conclusion.

8. CONCLUSION

This paper analyzes signaling games wherein a large collection of actions potentially serve as a signal of an underlying characteristic. It is easy to see that as the number of opportunities to send signals grows, the distortion of any single choice in a symmetric separating equilibrium shrinks, vanishing in the limit. However, it is not clear whether the total cost of the distortions aggregated over the growing number of actions shrinks, remains constant, or grows, or whether significant waste persists in asymmetric equilibria. We show that, for a broad class of signaling games, access to pervasive signaling opportunities leads to costless information revelation in the limit. The primary economic assumption we require for this result is that different types of senders possess different bliss points. We show by way of example that, with homogenous bliss points, the total cost of signaling need not shrink.

Our results may be broadly relevant in light of technological developments, such as the emergence of OSNs, that allow people to live their lives more publicly. A countervailing consideration is that, as opportunities to signal have proliferated, audiences for signals have also grown. While the former trend reduces signaling distortions, the latter presumably increases them. We have examined the interaction between these trends, and have identified conditions under which signaling distortions vanish, either in absolute terms or relative to informational stakes. When these conditions hold, a publicly visible life may allow one to “live authentically,” that is, to credibly reveal private information at a negligible cost rather than suffer from the cumulative burden of pervasive distortions.

Our findings may also shed light on specific signaling phenomena. First, signaling is a natural explanation for conspicuous consumption (see Ireland (1994) and Bagwell and Bernheim (1986)). 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. Thourmrungrroje (2014) 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 (2017)). While the expanding scope of conspicuous consumption would seem to indicate greater wastefulness, our results suggest that this effect may be swamped by the reduction in waste per conspicuous action. In other words, although we observe more forms of conspicuous consumption, welfare may actually be greater because signaling distorts the consumption of each good to a much smaller degree. 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, e.g., our analysis of “decisiveness” in Bernheim and Bodoh-Creed (2020)). 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).

Finally, our results also have potential implications for applications that highlight the importance of pooling equilibria. For example, Bernheim (1994) 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 (1990) 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. We conjecture that, due to the types of considerations that arise in the current paper, the set of agents who join a central pool will shrink as the number of observable actions increases. This observation raises the possibility that the proliferation of signals may erode conformity by status seekers, leaving groups of conformists who are "true believers" in the norms they practice.

APPENDIX A: SIGNALING WASTE IN MODELS WITH GENERAL AGGREGATOR FUNCTIONS

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

$$U^N(\mathbf{a}, \delta, t) = B(t, \delta) + \pi^N(\mathbf{a}, t).$$

Throughout, we assume that the second- and third-order partial derivatives of π^N exist and are bounded and continuous. We define $\mathbf{a}_{\text{BP}}^N(t)$ as follows:

$$\mathbf{a}_{\text{BP}}^N(t) \equiv \arg \max_{\mathbf{a}} \pi^N(\mathbf{a}, t). \quad (4)$$

The analysis of these general settings is challenging because, in principle, the sender might choose to signal through all, some, or even just one 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 Appendix have straightforward analogs to those used in Section 3. We have opted for assumptions that involve intuitive, easily understood restrictions, rather than ones that deliver the greatest technical generality.

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 9. $\mathbf{a}_{\text{BP}}^N(t)$ 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

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

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

Next, we assume that as N grows, the cost of choosing any bliss point other than the agent's own bliss point grows as well. This assumption obviously holds in the additive model.

ASSUMPTION 10. *There exists a function $g : \mathbb{Z} \rightarrow \mathbb{R}$ with $\lim_{N \rightarrow \infty} g(N) = +\infty$ and a strictly increasing function $h : \mathbb{R} \rightarrow \mathbb{R}$ with $h(0) = 0$ such that for all N and $t, t' \in T$,*

$$\pi^N(a_{\text{BP}}^N(t), t) - \pi^N(a_{\text{BP}}^N(t'), t) > g(N)h(|t - t'|).$$

Our next assumption bounds weighted averages of the second-order terms.

ASSUMPTION 11. *There exists $\Pi < 0$ such that for all N , $t \in T$, and $(s_1, \dots, s_N)$ such that $s_i \geq 0$ and $\sum_{i=1}^N s_i = 1$,*

$$\sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \pi^N(a_{\text{BP}}^N(t), t)}{\partial a_i \partial a_j} s_i s_j < \Pi.$$

Notice that $s = (0, \dots, 0, 1, 0, \dots, 0)$ 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 additive model satisfies Assumption 11, and indeed one can think of it as a generalization of the requirement that $\pi_{aa}(a_{\text{BP}}^N(t), t) < 0$ in Assumption 2.

Our proof uses Taylor series expansions around the bliss points to identify which types fall within the plausible set for each action. The assumption enables us to employ third-order expansions and ensures that second-order effects dominate any higher-order effects for large N .

ASSUMPTION 12. *Let*

$$D^N(a) = \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N \max \left\{ 0, \frac{\partial^3 \pi^N(a, t)}{\partial a_i \partial a_j \partial a_k} \right\}.$$

Then there exists a finite number $D > 0$ such that $D^N(a) < ND$ for all N and $a \in A^N$.

In the additive model, we have $D^N = \sum_{i=1}^N \max\{0, \pi_{aaa}(a, t)\}$; according to Assumption 2 (which requires $\pi_{aaa}(a, t) < C$), this term is bounded above by NC . Thus, in limiting the aggregate importance of positive third derivatives, Assumption 12 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 12 is also satisfied when (a) all of the third derivative terms individually respect a common upper bound that is independent of N (i.e., $\pi_{ijk}^N(a, t) < K$ for some finite K), and (b) actions interact with each other (in the sense that cross-partial derivatives are nonzero) only within groups of size no greater than some fixed M .³⁰ One can think of the additive model as an example of this class where $M = 1$.

³⁰In that case, it is natural to treat N as indexing the number of groups, rather than the number of actions.

In relaxing symmetry and additive separability, we sacrifice the ability to identify a tractable class of reliably efficient separating equilibria. For that reason, we use the approach employed in Section 5, where we proposed and applied the dominance refinement. As before, 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 $a_{BP}(t)$ is strictly increasing (by Assumption 9), the dominance refinement ensures that, for large N , the receiver's beliefs after observing $a_{BP}(t)$ are close to t . 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 waste from signaling. Assumption 10 ensures that this bound tightens as $N \rightarrow \infty$.

THEOREM 4. *Under Assumptions 4, 5, and 9–12, there exists a finite constant $K > 0$ such that, for sufficiently large N , the total waste from signaling in all separating equilibria satisfying the dominance refinement, $\pi^N(a_{BP}^N(t), t) - \pi^N(a(t; N), t)$ (where $a(t; N)$ is the equilibrium separating action function) is bounded above by $K/\sqrt{N}$ for all types $t \in T$.*

One way to prove Theorem 1, the corresponding convergence result for the additively separable and symmetric formulation, is to demonstrate that the restricted model satisfies Assumptions 9–12. This analytic strategy repositions the earlier result as a corollary of Theorem 4. We have already noted that Assumptions 9, 11, and 12 have direct counterparts in Section 3, which they generalize. The remaining step is to verify that the restricted model satisfies Assumption 10, which has no direct counterpart in Section 3. Appendix B instead includes a direct proof of Theorem 1 because its greater simplicity makes the logic of the argument easier to follow.

APPENDIX B: PROOFS

PROOF OF THEOREM 1. The proof proceeds in two steps.

Step 1: There exists $K_1 > 0$ such that, for all $t \in T$ and $t' \in \mathcal{P}(a_{BP}(t))$, we have $t - t' \leq K_1/\sqrt{N}$.

For t' to lie in $\mathcal{P}(a_{BP}(t))$, we must have

$$N\pi(a_{BP}(t'), t') - N\pi(a_{BP}(t), t') \leq \bar{B} - \underline{B}. \quad (5)$$

Our object is to establish a lower bound on the set of types satisfying (5).

First, we claim that, for all $\varepsilon > 0$, there exists N_ε^* such that if $N > N_\varepsilon^*$, then $|a_{BP}(t) - a_{BP}(t')| < \varepsilon$ for all $t \in T$ and $t' \in \mathcal{P}(a_{BP}(t))$. Define $\Phi_\varepsilon \equiv \{(t, t') \in T^2 \mid |a_{BP}(t) - a_{BP}(t')| \geq \varepsilon\}$. Because because π and a_{BP} are continuous, and because Φ_ε is (therefore) compact, we can define

$$d(\varepsilon) \equiv \min_{(t, t') \in \Phi_\varepsilon} [\pi(a_{BP}(t'), t') - \pi(a_{BP}(t), t')].$$

In light of Assumptions 1 and 3 and the fact that $t \neq t'$ in Φ_ε , we have $d(\varepsilon) > 0$ for all $\varepsilon > 0$. Now let $N_\varepsilon^* \equiv (\bar{B} - \underline{B})/d(\varepsilon)$. Supposing $N > N_\varepsilon^*$, we have, by construction,

$$N\pi(a_{BP}(t'), t') - N\pi(a_{BP}(t), t') > \bar{B} - \underline{B},$$

for all $(t, t') \in \Phi_\varepsilon$, which means $t' \notin \mathcal{P}(a_{BP}(t))$ by (5). But then $|a_{BP}(t) - a_{BP}(t')| < \varepsilon$ for all $t \in T$ and $t' \in \mathcal{P}(a_{BP}(t))$, which establishes the claim.

Using a Taylor expansion, the left side of equation (5) is equal to

$$-\frac{1}{2}N\pi_{aa}(a_{BP}(t'), t')(a_{BP}(t) - a_{BP}(t'))^2 - \frac{1}{6}N\pi_{aaa}(\xi, t')(a_{BP}(t) - a_{BP}(t'))^3, \quad (6)$$

where ξ lies between $a_{BP}(t')$ and $a_{BP}(t)$. Assumption 3 implies $|a_{BP}(t) - a_{BP}(t')| \geq \beta|t - t'|$. Noting that the left-hand side of equation (5) is strictly positive, we can then write

$$\frac{\beta^2}{2}N(t - t')^2 \left[-\pi_{aa}(a_{BP}(t'), t') - \frac{1}{3}\pi_{aaa}(\xi, t')(a_{BP}(t) - a_{BP}(t')) \right] \leq \bar{B} - \underline{B}.$$

Because α_{BP} and π_{aa} are continuous (Assumptions 1 and 2), $\pi_{aa}(a_{BP}(t'), t')$ achieves a maximum, $\bar{\pi}_{aa} < 0$, for $t' \in T$. It follows that $-\pi_{aa}(a_{BP}(t'), t') \geq -\bar{\pi}_{aa}$. Using the fact that $\pi_{aaa}(\xi, t') \leq C$ for some finite $C > 0$ (Assumption 2), and focusing on $t' < t$ (which implies $a_{BP}(t) > a_{BP}(t')$), we then have

$$\frac{\beta^2}{2}N(t - t')^2 \left[-\bar{\pi}_{aa} - \frac{1}{3}C(a_{BP}(t) - a_{BP}(t')) \right] \leq \bar{B} - \underline{B}.$$

Using our initial claim, we see that there exists N^* such that, for $N > N^*$, the bracketed term exceeds $-\bar{\pi}_{aa}/2$ for $t' < t$ such that $t' \in \mathcal{P}(a_{BP}(t))$. It follows that, for all $t \in T$ and $t' \in \mathcal{P}(a_{BP}(t))$, we have $t - t' \leq K_1/\sqrt{N}$ for $K_1 = (2/\beta)\sqrt{(\bar{B} - \underline{B})/\bar{\pi}_{aa}}$, as desired.

Step 2: Proof of the theorem.

Suppose a sender with type t deviates from $a(t, N)$ to $a_{BP}(t)$. If a type t' agent chooses $a_{BP}(t)$ as part of an equilibrium, then plainly we must have $t' \in \mathcal{P}(a_{BP}(t))$. If no type chooses $a_{BP}(t)$ as part of an equilibrium, then under the dominance refinement receivers are certain that any deviation to $a_{BP}(t)$ is attributable to $t' \in \mathcal{P}(a_{BP}(t))$. In either case, from Step 1, we know that any inference the receiver makes following an observation of $a^N = a_{BP}(t)$ is at least as favorable (from the sender's perspective) as $t - K_1/\sqrt{N}$. Combined with Assumptions 4 and 5, this observation implies that the cost of choosing $a(t, N)$ for a sender with type t satisfies

$$\pi^N(a_{BP}(t), t) - \pi^N(a(t, N), t) \leq B(t, t) - B\left(t, t - \frac{K_1}{\sqrt{N}}\right) \leq \frac{\gamma K_1}{\sqrt{N}}. \quad (7)$$

Taking $K = \gamma K_1$ completes the proof of the theorem. $\square$

Before proving Theorem 2, we first establish in the following lemma that actions converge to bliss points.

LEMMA 1. *Under Assumptions 1, 2, 5, and 7, there exists a finite constant $K > 0$ such that, for all $t \in T$, we have $|a_{SEP}(t, N) - a_{BP}(t)| \leq K/\sqrt{N}$ as $N \rightarrow \infty$.*

PROOF. Let $a_{SEP}(t, N)$ be the unique symmetric separating equilibrium, the existence of which Assumption 7 guarantees.

First, we claim that actions converge to bliss points uniformly: for all $\varepsilon > 0$, there exists N_ε^* such that if $N > N_\varepsilon^*$, then $|a_{\text{SEP}}(t) - a_{\text{BP}}(t)| < \varepsilon$ for all $t \in T$. Define $\Psi_\varepsilon \equiv \{(a, t) \in \mathcal{A} \times T \mid |a - a_{\text{BP}}(t)| \geq \varepsilon\}$. Because π and a_{BP} are continuous, and because Ψ_ε is (therefore) compact, we can define

$$\delta(\varepsilon) \equiv \min_{(a,t) \in \Psi_\varepsilon} [\pi(a_{\text{BP}}(t), t) - \pi(a, t)].$$

In light of Assumption 1, we have $\delta(\varepsilon) > 0$ for all $\varepsilon > 0$. Now let $N_\varepsilon^* \equiv (\bar{B} - \underline{B})/\delta(\varepsilon)$. Supposing $N > N_\varepsilon^*$, we have, by construction,

$$N\pi(a_{\text{BP}}(t), t) - N\pi(a, t) > \bar{B} - \underline{B}.$$

for all $(t, a) \in \Psi_\varepsilon$. But then, for all t , we must have $(t, a_{\text{SEP}}(t, N)) \notin \Psi_\varepsilon$, else $a_{\text{SEP}}(t, N)$ would not be an equilibrium choice for type t . It follows immediately that, if $N > N_\varepsilon^*$, then for all t we have $|a_{\text{SEP}}(t, N) - a_{\text{BP}}(t)| < \varepsilon$, as desired.

Before establishing the speed of convergence, we make two simple observations about $a_{\text{SEP}}(t, N)$.

OBSERVATION 1. $a_{\text{SEP}}(\underline{t}, N) = a_{\text{BP}}(\underline{t})$.

Suppose on the contrary that $a_{\text{SEP}}(\underline{t}, N) \neq a_{\text{BP}}(\underline{t})$. In a Perfect Bayesian equilibrium, $\hat{t}(a_{\text{BP}}(\underline{t})) \geq \underline{t}$. Thus, a deviation by type $\underline{t}$ to $a_{\text{BP}}(\underline{t})$ would increase direct utility (π) without reducing reputational utility (B), a contradiction.

OBSERVATION 2. $a_{\text{SEP}}(t, N) \geq a_{\text{BP}}(t)$ for all $t \in T$.

Suppose on the contrary that $a_{\text{SEP}}(t, N) < a_{\text{BP}}(t)$. Recalling that $a_{\text{SEP}}(s, N)$ is increasing in s (Assumption 7), we know that $a_{\text{SEP}}(t, N) > a_{\text{BP}}(\underline{t})$. Recalling that $a_{\text{BP}}(s)$ is continuous in s (Assumption 1), and applying Observation 1, we have $a_{\text{SEP}}(t, N) = a_{\text{BP}}(t')$ for some $t' \in (0, t)$. But then, in deviating from $a_{\text{SEP}}(t', N)$ to $a_{\text{BP}}(t')$, type t' would increase direct utility (π) without decreasing reputational utility (given Assumption 4 inasmuch as $t > t'$), a contradiction.

To establish the speed of convergence, we derive a bound on $a_{\text{SEP}}(t, N) - a_{\text{BP}}(t)$ from the following inequality, which must hold in equilibrium:

$$N[\pi(a_{\text{BP}}(t), t) - \pi(a_{\text{SEP}}(t, N), t)] \leq \bar{B} - \underline{B}. \quad (8)$$

The Taylor expansion of $\pi(a, t)$ around $(a_{\text{BP}}(t), t)$ yields

$$\begin{aligned} \pi(a_{\text{BP}}(t), t) - \pi(a_{\text{SEP}}(t, N), t) &= -\frac{1}{2}\pi_{aa}(a_{\text{BP}}(t), t)(a_{\text{SEP}}(t, N) - a_{\text{BP}}(t))^2 \\ &\quad - \frac{\pi_{aaa}(\xi, t)}{6}(a_{\text{SEP}}(t, N) - a_{\text{BP}}(t))^3, \end{aligned} \quad (9)$$

where ξ lies between $a_{\text{BP}}(t)$ and $a_{\text{SEP}}(t, N)$. Combining equations (8) and (9), we have

$$N(a_{\text{SEP}}(t, N) - a_{\text{BP}}(t))^2 \left[-\frac{1}{2}\pi_{aa}(a_{\text{BP}}(t), t) - \frac{\pi_{aaa}(\xi, t)}{6}(a_{\text{SEP}}(t, N) - a_{\text{BP}}(t)) \right] \leq \bar{B} - \underline{B}.$$

As in the proof of Theorem 1, under Assumptions 2 and 7, the following quantity is well-defined: $\bar{\pi}_{aa} \equiv \max_{t \in T} \pi_{aa}(a_{BP}(t), t) < 0$. Plainly, we have $-\pi_{aa}(a_{BP}(t'), t') \geq -\bar{\pi}_{aa}$. Using the facts that $\pi_{aaa}(a, t) \leq C$ for some finite $C > 0$ (Assumption 2) and $a_{SEP}(t, N) - a_{BP}(t) \geq 0$ (Observation 2), we then have

$$\frac{1}{2}N(a_{SEP}(t, N) - a_{BP}(t))^2 \left[-\bar{\pi}_{aa} - \frac{1}{3}C(a_{SEP}(t, N) - a_{BP}(t')) \right] \leq \bar{B} - \underline{B}.$$

Our first claim implies that we can choose N^* such that for all $t \in T$ and $N > N^*$,

$$\frac{1}{3}C(a_{SEP}(t, N) - a_{BP}(t')) \leq -\frac{1}{2}\bar{\pi}_{aa}.$$

It then follows that, for $N > N^*$, we have, for all $t \in T$,

$$-\frac{1}{4}\bar{\pi}_{aa}(a_{SEP}(t, N) - a_{BP}(t))^2 \leq \frac{\bar{B} - \underline{B}}{N}, \quad (10)$$

which in turn yields

$$a_{SEP}(t, N) - a_{BP}(t) \leq \frac{1}{\sqrt{N}} \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}} \quad (11)$$

for all $t \in T$ (provided $N > N^*$). Taking

$$K \equiv \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}},$$

and remembering that $a_{SEP}(t, N) \geq a_{BP}(t)$ completes the proof. $\square$

PROOF OF THEOREM 2. The arguments here reference Observations 1 and 2 from the proof of Lemma 1.

Suppose $a_{SEP}(t, N) > a_{BP}(t)$. Consider a deviation by agent t to from $a_{SEP}(t, N)$ to $a_{BP}(t)$. Because $a_{BP}(t) > a_{BP}(\underline{t}) = a_{SEP}(\underline{t}, N)$ (where the inequality follows from Assumption 7 and the equality follows from Observation 1), and because $a_{SEP}(s, N)$ is continuous in s (Assumption 7), there exists a type $t' < t$ such that $a_{SEP}(t', N) = a_{BP}(t)$. Thus, in equilibrium, we must have $\hat{t}(a_{BP}(t)) = t'$. Applying equation (11), we can then infer that

$$a_{BP}(t) = a_{SEP}(t', N) \leq a_{BP}(t') + \frac{1}{\sqrt{N}} \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}}.$$

From Assumption 3, we then have

$$\beta(t - t') \leq a_{BP}(t) - a_{BP}(t') \leq \frac{1}{\sqrt{N}} \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}},$$

so

$$t' \geq t - \frac{1}{\beta\sqrt{N}} \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}}.$$

From the preceding expression, we draw the following inference:

$$\begin{aligned}\pi(a_{BP}(t), t) - \pi(a_{SEP}(t), t) &\leq \frac{1}{N} \left[B(t, t) - B\left(t, t - \frac{1}{\beta\sqrt{N}} \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}}\right) \right] \\ &\leq \frac{1}{N^{1.5}} \frac{\gamma}{\beta} \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}},\end{aligned}$$

where the first inequality follows from the fact that type t chooses $a_{SEP}(t, N)$ over $a_{BP}(t)$ in equilibrium, while the second follows from Assumption 4.

Using the same argument employed in the derivation of equation (10) (based on Assumption 2) and then combining the result with the previous inequality yields:

$$-\frac{1}{4}\bar{\pi}_{aa}[a_{SEP}(t, N) - a_{BP}(t)]^2 \leq \pi(a_{BP}(t), t) - \pi(a_{SEP}(t, N), t) \quad (12)$$

$$\leq \frac{1}{N^{1.5}} \frac{\gamma}{\beta} \sqrt{\frac{-4(\bar{B} - \underline{B})}{\bar{\pi}_{aa}}}. \quad (13)$$

Simplifying yields

$$a_{SEP}(t) - a_{BP}(t) \leq \frac{1}{N^{3/4}} \sqrt{\frac{\gamma}{\beta}} \left(-\frac{4}{\bar{\pi}_{aa}}\right)^{3/4} (\bar{B} - \underline{B})^{1/4}. \quad (14)$$

Iterating this argument L times yields

$$a_{SEP}(t, N) - a_{BP}(t) \leq \frac{C_L}{N^{1-0.5L}},$$

where

$$C_L = \left(\frac{\gamma}{\beta}\right)^{0.5^{L-1}} \left(-\frac{4}{\bar{\pi}_{aa}}\right)^{1-0.5^L} (\bar{B} - \underline{B})^{0.5^L}.$$

Because the inequality must hold for all L , we have, for any given N ,

$$a_{SEP}(t, N) - a_{BP}(t) \leq \lim_{L \rightarrow \infty} \frac{C_L}{N^{1-0.5^L}} \equiv \frac{C_\infty}{N},$$

where

$$C_\infty = -\frac{4}{\bar{\pi}_{aa}}.$$

Using this expression in our Taylor expansion yields

$$\begin{aligned}\pi(a_{BP}(t), t) - \pi(a_{SEP}(t, N), t) &= (a_{SEP}(t) - a_{BP}(t))^2 \left(-\frac{1}{2}\bar{\pi}_{aa}(a_{BP}(t), t) \right. \\ &\quad \left. - \frac{\pi_{aaa}(\xi, t)}{6}(a_{SEP}(t) - a_{BP}(t))\right) \\ &\leq \left(\frac{C_\infty}{N}\right)^2 \left(-\frac{1}{2}\bar{\pi}_{aa}(a_{BP}(t), t) + \frac{C}{6}\left(\frac{C_\infty}{N}\right)\right),\end{aligned}$$

where C is the bound appearing in Assumption 2. Under Assumptions 2 and 7, we can define $\underline{\pi}_{aa} \equiv \min_{t \in T} \pi_{aa}(a_{BP}(t), t) < 0$. We then have:

$$\pi(a_{BP}(t), t) - \pi(a_{SEP}(t, N), t) \leq \left(\frac{4}{N\overline{\pi}_{aa}} \right)^2 \left(-\frac{1}{2}\underline{\pi}_{aa} - \frac{1}{N} \left(\frac{4C}{6\overline{\pi}_{aa}} \right) \right).$$

Let $N^* = 4C/3\overline{\pi}_{aa}\underline{\pi}_{aa}$. Then for $N > N^*$, we have

$$\pi(a_{BP}(t), t) - \pi(a_{SEP}(t, N), t) \leq -\underline{\pi}_{aa} \left(\frac{4}{N\overline{\pi}_{aa}} \right)^2.$$

It follows immediately that, for $N > N^*$, we have

$$N[\pi(a_{BP}(t), t) - \pi(a_{SEP}(t, N), t)] \leq \frac{K}{N}$$

where the scaling factor, $K \equiv -\underline{\pi}_{aa}(4/\overline{\pi}_{aa})^2$, is independent of t . $\square$

PROOF OF THEOREM 4. The proof parallels that of Theorem 1. It involves the same two steps, and the proof of the second step is unchanged. Here, we provide a proof of Step 1 for the general case.

For t' to lie in $\mathcal{P}(a_{BP}(t))$, we must have

$$\pi^N(a_{BP}(t'), t') - \pi^N(a_{BP}(t), t') \leq \overline{B} - \underline{B} \quad (15)$$

Our object is to establish bounds on the set of types satisfying (15).

First, we claim that, for all $\varepsilon > 0$, there exists N_ε^* such that if $N > N_\varepsilon^*$, then $|t - t'| < \varepsilon$ for all $t \in T$ and $t' \in \mathcal{P}(a_{BP}(t))$. Define $\Phi_\varepsilon \equiv \{(t, t') \in T^2 \mid |t - t'| \geq \varepsilon\}$. Let g and h be the functions referenced in Assumption 10. Because h is strictly increasing and $h(0) = 0$, we have $h(\varepsilon) > 0$. Moreover, because $\lim_{N \rightarrow \infty} g(N) = +\infty$, we can choose N_ε^* such that, for all $N > N_\varepsilon^*$, we have $g(N)h(\varepsilon) > \overline{B} - \underline{B}$. But then, for all $N > N_\varepsilon^*$ and $(t, t') \in \Phi_\varepsilon$, we have

$$\pi^N(a_{BP}(t'), t') - \pi^N(a_{BP}(t), t') > g(N)h(|t - t'|) \geq g(N)h(\varepsilon) > \overline{B} - \underline{B},$$

where the first inequality follows from Assumption 10, the second follows from the fact that h is strictly increasing and $(t, t') \in \Phi_\varepsilon$, while the third follows from the fact that $N > N_\varepsilon^*$. In light of (15), we then have $t' \notin \mathcal{P}(a_{BP}(t))$. But then $|t - t'| < \varepsilon$ for all $t \in T$ and $t' \in \mathcal{P}(a_{BP}(t))$, which establishes the claim.

Consider some t and $t' < t$. Define $\delta_i = a_{BP,i}(t) - a_{BP,i}(t')$, and note that $\delta_i^N > 0$ by Assumption 9. Using a Taylor expansion, the left side of equation (15) is equal to

$$-\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \pi^N(a_{BP}(t'), t')}{\partial a_i \partial a_j} \delta_i \delta_j - \frac{1}{6} \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N \frac{\partial^3 \pi^N(\xi, t')}{\partial a_i \partial a_j \partial a_k} \delta_i \delta_j \delta_k,$$

for some ξ satisfying $a_{BP,i}(t') \leq \xi_i \leq a_{BP,i}(t)$ for $i = 1, \dots, N$. Using Assumptions 9 and 11, we can write

$$-\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \pi^N(a_{BP}(t'), t')}{\partial a_i \partial a_j} \delta_i \delta_j = -\frac{\|\delta\|^2}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \pi^N(a_{BP}(t'), t')}{\partial a_i \partial a_j} \left(\frac{\delta_i}{\|\delta\|} \right) \left(\frac{\delta_j}{\|\delta\|} \right)$$

$$> -\frac{\|\delta\|^2}{2}\Pi \geq -\frac{N\beta_L^2(t-t')^2\Pi}{2}.$$

Assumptions 9 (which guarantees $\delta_i > 0$) and 12 together imply:

$$-\frac{1}{6} \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N \frac{\partial^3 \pi^N(\xi, t)}{\partial a_i \partial a_j \partial a_k} \delta_i \delta_j \delta_k > -\frac{ND\beta_H(t-t')^3}{6}.$$

Therefore, we have

$$\begin{aligned} & -\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2 \pi^N(\mathbf{a}_{BP}(t'), t')}{\partial a_i \partial a_j} \delta_i \delta_j - \frac{1}{6} \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N \frac{\partial^3 \pi^N(\xi, t')}{\partial a_i \partial a_j \partial a_k} \delta_i \delta_j \delta_k \\ & \geq N(t-t')^2 \left[-\frac{\beta_L^2 \Pi}{2} - \frac{D\beta_H(t-t')}{6} \right]. \end{aligned}$$

From our opening claim (and the fact that $D, \beta_H > 0$ and $\Pi < 0$), we know there exists N^* such that, if $N > N^*$, then for all t, t' with $t > t'$ and $t' \in \mathcal{P}(\mathbf{a}_{BP}(t))$, we have

$$\frac{D\beta_H(t-t')}{6} < -\frac{\beta_L^2 \Pi}{4}.$$

From equation (5), it then follows that, for such N ,

$$-N \frac{\beta_L^2(t-t')^2 \Pi}{4} \leq \pi^N(\mathbf{a}_{BP}^N(t'), t') - \pi^N(\mathbf{a}_{BP}^N(t), t') \leq \bar{B} - \underline{B},$$

which in turn implies

$$t - t' \leq \sqrt{\frac{-4(\bar{B} - \underline{B})}{N\beta_L^2 \Pi}}.$$

It follows that, for all t and $t' \in \mathcal{P}(\mathbf{a}_{BP}(t))$, we have $t - t' \leq \frac{K_1}{\sqrt{N}}$ for $K_1 = (2/\beta_L) \times \sqrt{(\underline{B} - \bar{B})/\Pi}$, as desired. $\square$

PROOF OF THEOREM 3. Suppose we have a separating equilibrium with action functions $\mathbf{a}(t, N) = (a_1(t, N), \dots, a_N(t, N))$. We are interested in determining type t 's total payoff in equilibrium and comparing it to the payoff type t receives in the symmetric separating equilibrium.

The first part of the proof establishes the following Claim (capitalized for clarity of subsequent references): if it were the case for some t that either (a) $\pi^N(\mathbf{a}(t, N), t) = \pi^N(\mathbf{a}_{SEP}(t, N), t)$ and $\mathbf{a}(t, N) \neq \mathbf{a}_{SEP}(t, N)$, or (b) $\pi^N(\mathbf{a}(t, N), t) > \pi^N(\mathbf{a}_{SEP}(t, N), t)$, then we would have $\pi_t^N(\mathbf{a}_{SEP}(t, N), t) > \pi_t^N(\mathbf{a}(t, N), t)$. (Ultimately, we will show that case (b) cannot arise, but for the purpose of establishing the Claim we treat it as a possibility.)

We now prove the Claim. Suppose case (a) arises for some $t > \underline{t}$. Define

$$\bar{a}_m = \begin{cases} a_m(t, N) & \text{if } a_m(t, N) \geq a_{BP}(t, N) \\ a \geq a_{BP}(t) \text{ such that } \pi(a, t) = \pi(a_m(t, N), t) & \text{otherwise.} \end{cases}$$

Existence of $\bar{a}_m$ follows from condition (v) and the continuity of π in a . In effect, this step replaces $a_m(t, N)$ with an action $\bar{a}_m \geq a_{BP}(t)$ that is equally costly for type t . Let $Q \equiv \{m \mid a_m^A(t) < a_{BP}(t)\}$. Then from supermodularity (condition (i)) we have

$$\pi_t^N(\bar{a}, t) - \pi_t^N(a(t, N), t) = \sum_{m \in Q} [\pi_t(\bar{a}_m, t) - \pi_t(a_m(t, N), t)] \geq 0, \quad (16)$$

with strict inequality if Q is nonempty.

If $a_{SEP}(t) = \bar{a}$, we are done. If not, then since $\pi^N(a_{SEP}(t, N), t) = \pi^N(\bar{a}, t)$ by construction, by condition (ii) there must exist i and j such that $\bar{a}_i > a_{SEP}(t) > \bar{a}_j$. Define the function $\tilde{a}(a_i)$ as follows: $\tilde{a}_i(a_i) = a_i$, $\tilde{a}_k(a_i) = \bar{a}_k$ for $k \neq i, j$, and $\pi^N(\tilde{a}(a_i), t) = \pi^N(\bar{a}, t)$. In other words, $\tilde{a}_j(a_i)$ indicates how a_j must vary in response to changes in a_i to keep the value of π^N constant at its equilibrium value for type t . Implicit differentiation reveals that, as long as a_i and $\tilde{a}_j(a_i)$ exceed $a_{BP}(t)$, we have

$$\frac{d\tilde{a}_j}{da_i} = -\frac{\pi_a(a_i, t)}{\pi_a(\tilde{a}_j(a_i), t)} < 0.$$

Plainly, there exists a unique value $a_i^e > a_{BP}(t)$ such that $\tilde{a}_j(a_i^e) = a_i^e$. For $a_i \in (a_i^e, \bar{a}_i)$, we have $\bar{a}_i > a_i > a_i^e > \tilde{a}_j(a_i) > \bar{a}_j \geq a_{BP}(t)$, as well as

$$\begin{aligned} \frac{d}{da_i} \pi_t^N(\tilde{a}(a_i), t) &= \frac{d}{da_i} \left(\sum_{i=1}^N \pi_t(\tilde{a}_i(a_i), t) \right) \\ &= \pi_{at}(a_i, t) + \pi_{at}(\tilde{a}_j(a_i), t) \frac{d\tilde{a}_j}{da_i} \\ &= \pi_{at}(a_i, t) - \pi_{at}(\tilde{a}_j(a_i), t) \frac{\pi_a(a_i, t)}{\pi_a(\tilde{a}_j(a_i), t)}, \\ &= \pi_a(a_i, t) \left[\frac{\pi_{at}(a_i, t)}{\pi_a(a_i, t)} - \frac{\pi_{at}(\tilde{a}_j(a_i), t)}{\pi_a(\tilde{a}_j(a_i), t)} \right] \\ &< 0. \end{aligned}$$

To understand the final inequality, first recall that we have $a_i > a_{BP}(t)$, which implies $\pi_a(a_i, t) < 0$ by condition (ii). Next, recall that $a_i > \tilde{a}_j(a_i) > a_{BP}(t)$, which implies that the term in square brackets is positive by condition (iii).

Now imagine reducing a_i from $\bar{a}_i$ while increasing a_j according to $\tilde{a}_j(a_i)$ until either $a_i = a_{SEP}(t, n)$ or $\tilde{a}_j(a_i) = a_{SEP}(t, n)$, which must occur for some $a_i \in [a_i^e, \bar{a}_i]$. Call the resulting vector of actions $\mathbf{a}^{(1)}$. Because $d\pi_t^N(\tilde{a}(a_i), t)/da_i < 0$ over this range, we can infer that $\pi_t^N(\mathbf{a}^{(1)}, t) > \pi_t^N(\bar{\mathbf{a}}, t)$. If $\mathbf{a}^{(1)} = \mathbf{a}_{SEP}(t, N)$, we are done. Otherwise, because $\pi^N(\mathbf{a}^{(1)}, t) = \pi^N(\bar{\mathbf{a}}, t) = \pi^N(\mathbf{a}_{SEP}(t, N), t)$, we know there must exist k

and l such that $a_k^{(1)} > a_{\text{SEP}}(t, N) > a_l^{(1)}$. Iterating this step, we produce a sequence of vectors, $\mathbf{a}^{(n)}$, such that $\mathbf{a}^{(n)}$ has at least n elements in common with $\mathbf{a}_{\text{SEP}}(t, N)$, and $\pi_t^N(\mathbf{a}^{(n)}, t) > \dots > \pi_t^N(\mathbf{a}^{(1)}, t) > \pi_t^N(\bar{\mathbf{a}}, t)$. Consequently, there exists $n^* \leq N$ for which $\mathbf{a}^{(n^*)} = \mathbf{a}_{\text{SEP}}(t, N)$. It follows that $\pi_t^N(\mathbf{a}_{\text{SEP}}(t, N), t) > \pi_t^N(\bar{\mathbf{a}}, t)$. Recalling expression (16), we then have $\pi_t^N(\mathbf{a}_{\text{SEP}}(t, N), t) > \pi_t^N(\mathbf{a}(t, N), t)$, as desired.

Next, suppose case (b) arises for some $t > \underline{t}$. Let $\mathbf{a}' = (a', \dots, a')$ be a symmetric vector with $a' \geq a_{\text{BP}}(t)$ and $\pi^N(\mathbf{a}', t) = \pi^N(\mathbf{a}(t, N), t)$. Existence of $\mathbf{a}'$ follows from condition (v) and the continuity of π^N . By the same argument as for condition (a), we infer $\pi_t^N(\mathbf{a}', t) \geq \pi_t^N(\mathbf{a}(t, N), t)$.³¹ Because $\pi^N(\mathbf{a}(t, N), t) > \pi^N(\mathbf{a}_{\text{SEP}}(t, N), t)$ by assumption, we have $a_{\text{SEP}}(t, N) > a'_m$ (by condition (ii)). From our assumption of supermodularity (condition (i)), we have

$$\pi_t^N(\mathbf{a}_{\text{SEP}}(t, N), t) - \pi_t^N(\mathbf{a}', t) = \sum_{m=1}^N \pi_t(a_{\text{SEP}}(t, N), t) - \pi_t(a'_m, t) \geq 0.$$

It follows that $\pi_t^N(\mathbf{a}_{\text{SEP}}(t, N), t) > \pi_t^N(\mathbf{a}(t, N), t)$, as desired.

This concludes the proof of the Claim. We now use the Claim to establish the theorem.

Let $V^N(t, \hat{t}; \mathbf{a})$ denote the payoff a type t sender receives when choosing the action assigned to type $\hat{t}$:

$$V^N(t, \hat{t}; \mathbf{a}) \equiv B(t, \hat{t}) + \pi^N(\mathbf{a}(\hat{t}, N), t). \quad (17)$$

In equilibrium, type t chooses $\hat{t} = t$ and receives a payoff of

$$v^N(t; \mathbf{a}) \equiv V^N(t, t; \mathbf{a}).$$

Condition (iv) implies that $V_t^N(t, \hat{t}; \mathbf{a})$ exists and $|V_t^N(t, \hat{t}; \mathbf{a})| < K$ for some finite $K > 0$. We can therefore invoke the integral version of the envelope theorem (Milgrom and Segal (2002, Theorem 2)), which implies

$$v^N(t; \mathbf{a}) - v^N(\underline{t}; \mathbf{a}) = \int_{\underline{t}}^t V_t^N(s, s; \mathbf{a}) ds.$$

Now we compare the payoffs received with the symmetric separating action function $\mathbf{a}_{\text{SEP}}$ and any alternative separating action function $\mathbf{a}$. For this purpose, define $\Delta(t) \equiv v^N(t; \mathbf{a}_{\text{SEP}}) - v^N(t; \mathbf{a})$. Consider any $t' \in [\underline{t}, t)$ such that $\Delta(t') = 0$. (Note, e.g., that $\Delta(\underline{t}) = 0$.) Then

$$\begin{aligned} \Delta(t) &= [v^N(t; \mathbf{a}_{\text{SEP}}) - v^N(t'; \mathbf{a}_{\text{SEP}})] - [v^N(t; \mathbf{a}) - v^N(t'; \mathbf{a})] \\ &= \int_{t'}^t [V_t^N(s, s; \mathbf{a}_{\text{SEP}}) - V_t^N(s, s; \mathbf{a})] ds \\ &= \int_{t'}^t [\pi_t^N(\mathbf{a}_{\text{SEP}}(s, N), s) - \pi_t^N(\mathbf{a}(s, N), s)] ds. \end{aligned} \quad (18)$$

³¹The inequality is weak because we include the possibility that $\mathbf{a}' = \mathbf{a}(t, N)$.

We will prove the theorem by showing that there exists no $t \in T$ with $\Delta(t) < 0$. Suppose on the contrary that such a t exists. Consider the set $T^*(t) \equiv \{t' < t \mid \Delta(t') \geq 0\}$. We clearly have $\underline{t} \in T^*(t)$, so the set is nonempty. Because Δ is continuous, $T^*(t)$ is compact. Therefore, it contains a maximal element, $t^* < t$; furthermore, the continuity of Δ implies $\Delta(t^*) = 0$. From equation (18), we then have

$$\Delta(t) = \int_{t^*}^t [\pi_t^N(\mathbf{a}_{\text{SEP}}(s, N), s) - \pi_t^N(\mathbf{a}(s, N), s)] ds.$$

But, by the Claim, we have $\pi_t^N(\mathbf{a}_{\text{SEP}}(s, N), s) - \pi_t^N(\mathbf{a}(s, N), s) > 0$ for all $s \in [t^*, t]$. It follows that $\Delta(t) > 0$, which contradicts our hypothesis that $\Delta(t) < 0$.

Now we turn to the theorem's final claim. Let $T^+ \equiv \{t \in T \mid \Delta(t) > 0\}$. Because Δ is continuous, T^+ is open. Assume T^+ is not dense in T . Then there exists a nondegenerate interval $[t_l, t_h] \subset T \setminus T^+$. The last part of the proof establishes that $\Delta(t) = 0$ for $t \in [t_l, t_h]$. From equation (18), we have

$$\Delta(t_h) = \int_{t_l}^{t_h} [\pi_t^N(\mathbf{a}_{\text{SEP}}(s, N), s) - \pi_t^N(\mathbf{a}(s, N), s)] ds.$$

But, by the Claim, since we have assumed $\mathbf{a}_{\text{SEP}}(s, N) \neq \mathbf{a}(s, N)$ for $t > \underline{t}$, we have $\pi_t^N(\mathbf{a}_{\text{SEP}}(s, N), s) - \pi_t^N(\mathbf{a}(s, N), s) > 0$ for all $s \in [t_l, t_h]$ (except for $s = t_l$ when $t_l = \underline{t}$). It follows that $\Delta(t_h) > 0$, which contradicts our hypothesis. Thus, T^+ is open-dense in T . $\square$

REFERENCES

- Andreoni, James and B. Douglas Bernheim (2009), "Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects." *Econometrica*, 77, 1607–1636. [164]
- Bagwell, Laurie S. and B. Douglas Bernheim (1986), "Veblen effects in a theory of conspicuous consumption." *The American Economic Review*, 86, 349–373. [164, 169, 181]
- Banks, Jeffrey (1990), "A model of electoral competition with incomplete information." *Journal of Economic Theory*, 50, 309–325. [164, 182]
- Banks, Jeffrey and Joel Sobel (1987), "Equilibrium selection in signaling games." *Econometrica*, 55, 647–661. [173, 176]
- Battaglini, Marco (2002), "Multiple referrals and multidimensional cheap talk." *Econometrica*, 70, 1379–1401. [166]
- Bernheim, B. Douglas (1994), "A theory of conformity." *The Journal of Political Economy*, 102, 841–877. [164, 182]
- Bernheim, B. Douglas and Aaron L. Bodoh-Creed (2020), "A theory of decisive leadership." *Games and Economic Behavior*, 121, 146–168. [164, 181]
- Bernheim, B. Douglas and Sergei Severinov (2016), "Bequests as signals: An explanation for the equal division puzzle." *Journal of Political Economy*, 111, 733–764. [164]

- Buffardi, Laura E. and W. Keith Campbell (2008), "Narcissism and social networking web sites." *Personality and Social Psychology Bulletin*, 34, 1303–1314. [178]
- Chakraborty, Archishman and Rick Harbaugh (2007), "Comparative cheap talk." *Journal of Economic Theory*, 132, 70–94. [166]
- Chakraborty, Archishman and Rick Harbaugh (2010), "Persuasion by cheap talk." *American Economic Review*, 100, 2361–2382. [166]
- Cho, In-Koo and David Kreps (1987), "Signaling games and stable equilibria." *The Quarterly Journal of Economics*, 102, 179–222. [164, 173, 176]
- Corneo, Giacomo and Oliver Jeanne (1997), "Conspicuous consumption, snobbism, and conformism." *Journal of Public Economics*, 66, 55–71. [164, 169]
- Crawford, Vincent P. and Joel Sobel (1982), "Strategic information transmission." *Econometrica*, 50, 1431–1451. [166]
- Donath, Judith S. (2002), "Identity and deception in the virtual community." In *Communities in Cyberspace* (Peter Kollock and Marc Smith, eds.), 27–58, Taylor & Francis Group, London. [167]
- Donath, Judith S. (2008), "Signals in social supernets." *Journal of Computer-Mediated Communication*, 13, 231–351. [167]
- Donath, Judith S. and Danah M. Boyd (2004), "Public displays of connection." *BT Technology Journal*, 22, 71–82. [167]
- Drouin, Michelle, Kimberly W. O'Connor, Gordon B. Schmidt, and Daniel A. 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, 123–128. [178]
- Eventbrite Blog (2017), "Millennials: Fueling the experience economy." Downloaded on 16 August 2018 from <https://www.eventbrite.com/blog/academy/millennials-fueling-experience-economy/>. [181]
- Froomkin, A. Michael (2000), "The death of privacy?" *Stanford Law Review*, 52, 1461–1543. [163]
- Grafen, Alan (1990), "Biological signals as handicaps." *Journal of Theoretical Biology*, 144, 517–546. [167]
- Hurd, Peter L. (1995), "Communication in discrete action-response games." *Journal of Theoretical Biology*, 174, 217–222. [167]
- Ireland, Norman J. (1994), "On limiting the market for status signals." *Journal of Public Economics*, 53, 91–110. [164, 169, 181]
- Kartik, Navin (2009), "Strategic communication with lying costs." *Review of Economic Studies*, 76, 1359–1395. [166, 169]

- Kartik, Navin, Marco Ottaviani, and Francesco Squintani (2007), “Credulity lies and costly talk.” *Journal of Economic Theory*, 134, 93–116. [166]
- Kluemper, Donald H. and Peter A. Rosen (2009), “Future employment selection methods: Evaluating social networking web sites.” *Journal of Managerial Psychology*, 24, 567–580. [178]
- Lachmann, Michael, Szabolcs Számado, and Carl T. 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, 13189–13194. [167]
- Lipnowski, Elliot and Ravid Doron (2020), “Cheap talk with transparent motives.” *Econometrica*, 88, 1631–1660. [166]
- Mailath, George J. (1987), “Incentive compatibility in signaling games with a continuum of types.” *Econometrica*, 55, 1349–1365. [171, 176]
- Mailath, George J. (1989), “Simultaneous signaling in an oligopoly model.” *The Quarterly Journal of Economics*, 104, 417–427. [164]
- Mailath, George J. and Ernst-Ludwig von Thadden (2013), “Incentive compatibility and differentiability: New results and classic applications.” *Journal of Economic Theory*, 148, 1841–1861. [171, 176]
- Milgrom, Paul and Ilya Segal (2002), “Envelope theorems for arbitrary choice sets.” *Econometrica*, 70, 583–601. [192]
- Miller, Merton H. and Kevin Rock (1985), “Dividend policy under asymmetric information.” *Journal of Finance*, 40, 1031–1051. [164]
- Nielsen (2016), “Nielsen 2016 social media report.” Downloaded from <http://www.nielsen.com/us/en/insights/reports/2017/2016-nielsen-social-media-report.html> on 5 July 2018. [165]
- Smith, Aaron (2014), “What people like and dislike about Facebook.” Pew Research Center, Washington, D.C. [178]
- Smith, Aaron and Monica Anderson (2018), “Social media use in 2018.” Pew Research Center, Washington, D.C. [165]
- Spence, A. Michael (1973), “Job market signaling.” *Quarterly Journal of Economics*, 87, 355–374. [165]
- Spence, A. Michael (2002), “Signaling in retrospect and the informational structure of markets.” *The American Economic Review*, 92, 434–459. [164, 165]
- Számado, Szabolcs (1999), “The validity of the handicap principle in discrete action-response games.” *Journal of Theoretical Biology*, 198, 593–602. [167]
- Számado, Szabolcs (2011), “The cost of honesty and the fallacy of the handicap principle.” *Animal Behavior*, 81, 3–10. [167]

Thoumrungroje, Amonrat (2014), “The influence of social media intensity and EWOM on conspicuous consumption.” *Procedia—Social and Behavioral Sciences*, 148, 7–15. [181]

van Dijck, Jose (2013), “You have one identity: Performing the self on Facebook and LinkedIn.” *Media, Culture & Society*, 35, 199–215. [178]

Zahavi, Amotz (1979), “Mate selection—a selection for a handicap.” *Journal of Theoretical Biology*, 53, 205–214. [167]

Co-editor Marina Halac handled this manuscript.

Manuscript received 27 November, 2020; final version accepted 17 January, 2022; available online 13 April, 2022.