Information Selection and Media Bias: An Economic Foundation for Sentiment Analysis

Jinge Liu*

March 1, 2026

Textual sentiment analysis is widely used, yet the economic interpretation of sentiment remains unclear. I develop a theory of media in which outlets selectively report information under content capacity constraints based on endogenous newsworthiness for the receiver's decision problem. The model generates common media bias patterns, including audience appeal and sensationalism, where bias reflects the gap between reported content and the full underlying news. It yields an asymptotically explicit sentiment-scoring function linking textual sentiment to an underlying fundamental signal and economic parameters governing content production, providing a structural interpretation of sentiment and highlighting potential measurement issues in empirical work.

Keywords: media bias, sentiment, textual analysis
JEL Codes: C40, D83, D90

*City University of Hong Kong. Email: jingeliu@cityu.edu.hk. I am indebted to Arjada Bardhi, Simon Gervais, Felipe Varas and S. "Vish" Viswanathan for their invaluable advice and support. I also thank seminar and conference participants at the Asia Meetings of the Econometric Society at Shenzhen and Tokyo, City University of Hong Kong, Duke Finance Student Theory Workshop, Hong Kong Joint Finance Research Workshop, Stony Brook International Conference on Game Theory, and Paris Workshop on Games, Decisions, and Language. Financial support from the Research Grants Council of Hong Kong (Early Career Scheme, Grant No. 21508123) is gratefully acknowledged.

1 Introduction

The use of textual data has become increasingly prevalent across many fields of economic research.[1] A popular approach of converting text into measurable data is *sentiment scoring*: researchers statistically quantify how "positive" the language of a text appears, define the resulting value as a sentiment score, and empirically relate it to economic outcomes. This approach greatly expands the types of data and questions that researchers can investigate. Yet despite its rapidly growing popularity, a fundamental challenge remains: we still have no conceptual framework that clarifies exactly what economic information is embedded in such sentiment scores. We only have limited and suggestive reduced-form evidence: For instance, news text sentiment may sometimes proxy for nonfundamental hype (Tetlock 2007) and at other times reflect part of fundamental information (Tetlock, Saar-Tsechansky, and Macskassy 2008) of the news. Sentiment in message boards (Antweiler and Frank 2004) and corporate filings (Loughran and McDonald 2011) contains trade-relevant information, yet the underlying economics is unclear. More recently, data-mining papers increasingly treat text sentiment as a mere predictive factor, refraining from economic interpretation (e.g., Ke, Kelly, and Xiu 2019). This lack of clarity reflects the inherent complexity of text sentiment. As a product of communication, it conceivably links to various factors, such as the fundamental information, beliefs, utilities, and frictions even considering only full rationality. This conceptual ambiguity causes confusion and inconsistency in how researchers interpret and use sentiment as a measure. In fact, sentiment analysis and textual analysis more broadly remains largely statistical, rarely examined through economic theory. This issue is increasingly pressing in the age of large language models that generate black-box sentiment scores with unprecedented ease. Therefore, to explain empirical findings consistently and to address potential measurement problems, developing a clear economic foundation for textual sentiment is critical.

In this paper, I make progress on this question by building an economic model of sentiment analysis, the first to my knowledge. I study the sentiment of news reports. To microfound sentiment, the key is to identify a sentiment-scoring function

$$\text{sentiment} = s(\text{fundamental news}, \text{economic context}), \tag{1}$$

where fundamental news is the true signal underlying the report, sentiment is the positivity of the report's textual tone, and economic context consists of parameters capturing the economic environment of text production. Hence, this microfoundation question involves exploring how to link literal text to the underlying reality across different environments. Such theories often fall within the media bias literature (e.g., Suen 2004, Mullainathan and Shleifer 2005, Gentzkow and Shapiro 2006, Che and Mierendorff 2019, Perego and Yuksel 2022) and typically study specific persuasion

1. For surveys, see Gentzkow, Kelly, and Taddy (2019) and Korinek (2023).

or competition motives of the media, sometimes combined with communication capacity frictions. However, they are rarely written in a way that aligns with textual analysis. In this paper, I build a novel media model that matches sentiment analysis. Among the possible channels underlying content formation, I focus on the common media practice of information selection. I establish the relevance of this channel by showing the novel result that it independently generates stylized facts about news bias without persuasion or media competition and use it to derive a tractable sentiment-scoring function; the goal is to propose a conceptual framework for explaining what sentiment means and what it measures.

Information selection is central to the theory. To fix ideas, consider the example of a news outlet (sender) having access to $N = 5$ new stories about a binary state, each with a binary realization (positive or negative), but the newspaper must report $n = 3$ binary stories due to its size.[2] Had the receiver known all N stories, she would take a different action specific to that realization. However, she relies on the newspaper for information and even a benign newspaper cannot convey information to perfectly distinguish between all realizations due to the communication capacity constraint. Specifically, the positivity score (the share of positive stories) of the underlying news equals $0\%, 20\%, 40\%, 60\%, 80\%$, or 100%, while the report positivity takes values from a sparser set: $0\%, 33\%, 67\%$, or 100%. Under my assumptions, these scores are sufficient statistics for the information in the underlying stories and the report. This is because I restrict my attention to text where (1) the writer maintains credibility by presenting information without modification, and (2) the order of information is irrelevant. (Examples include newspaper pages as collections of bullish or bearish events, articles as collections of positive and negative evidence, etc.) Essentially, my theory posits that report text is a simplification of the complex reality. To establish a stable and easily interpretable reporting pattern, the news outlet designs an information structure and commits to it, assigning a distribution of report positivity scores to each fundamental positivity score. To isolate the information selection effect, I assume the news outlet serves the best interest of its readers; in the paper I discuss why this assumption may be reasonable. Even so, limited communication capacity makes it impossible for the news outlet to always recommend the perfectly optimal action through its report.

The two-step model roadmap is as follows. First, I model the discrete information selection problem, generalizing the notion in the example. I show that a benign sender uses only pure strategies. I show the optimal information structure is a cutoff strategy: in the example, the news outlet should use three cutoffs for the fundamental positivity scores to partition them into four nonempty sets, mapping the pooled scores within each set to the same reported score. The optimal information structure is thus characterized by the optimal cutoff locations, which I analyze in the

2. As *Seinfeld* put it, "It's amazing that the amount of news that happens in the world every day always just exactly fits the newspaper."

second step. Regardless of the specific cutoff locations, I focus on the case in which the news outlet assigns reported scores to partition sets in increasing order (this makes reports self-consistent since the sender reports higher positivity for higher fundamental positivity), which turns out to allow the reported stories to remain truthful.

Second, I take this model to its asymptotic limit by letting N and n tend to infinity while keeping n small compared to N by sequentially taking N to infinity first. I examine the limit of the optimal information structure, which is tractable. Report positivity remains a ratio; this is the *sentiment* we seek in the sentiment-scoring function of Equation (1). After scaling, the conditional Binomial fundamental positivity score converges to a conditional Gaussian signal, i.e., it becomes "(binary) truth plus (Gaussian) noise" that is common in economic models of learning; this is the *fundamental news* we seek in Equation (1). I tractably derive the distribution of now an infinite number of optimal cutoffs in the main theorem of this paper. Knowing where these cutoffs are located in the limit, I point out that the sentiment-scoring function of Equation (1) coincides exactly with the cumulative distribution of these cutoffs. That is, the asymptotically tractable answer to our microfoundation question is the formula *report positivity = cumulative distribution of the cutoffs*.

This paper makes two main contributions.

First, on the road to microfoundation, I identify information selection as a novel and independent channel generating media bias. In equilibrium, the selected content may exhibit two distinct types of bias: (i) audience appeal and (ii) sensationalism. The term "bias" refers to the disparity between the selected content and the underlying stories. It is captured in the discrete model as the difference between reported and underlying positivity scores[3] and in the asymptotic model as the nonlinearity of the sentiment-scoring function. Audience appeal means that the content favors the version of reality preferred or believed by the receiver. For example, if the receiver benefits when the state is positive or holds a prior belief favoring a positive state, the reported positivity score is higher. Sensationalism means that the content is more polarized than the underlying information: if the underlying positivity score is high (low), then the reported positivity score is higher (lower). Both biases coexist, with sensationalism prevailing under highly informative underlying information and audience appeal when it is less so. Here, audience appeal is a directional bias, while sensationalism is a nondirectional bias. I also discuss the possibility that under certain conditions, we have alarmism – the directional bias reverse of audience appeal.

We may illustrate bias using the earlier example (details are in Section 2.5). Figure I shows equilibria under two sets of contextual parameter values (prior belief, payoff relevance, news informativeness, and payoff shape). In the left panel, the receiver is inclined to take a positive-leaning action before reading the news; in the right panel, there is no obvious inclination. Audience appeal is evident in the left panel: when the underlying information is 20% (40%, 60%, 80%) positive,

3. This does not refer to the difference due to model discreteness which will disappear in the asymptotic analysis.

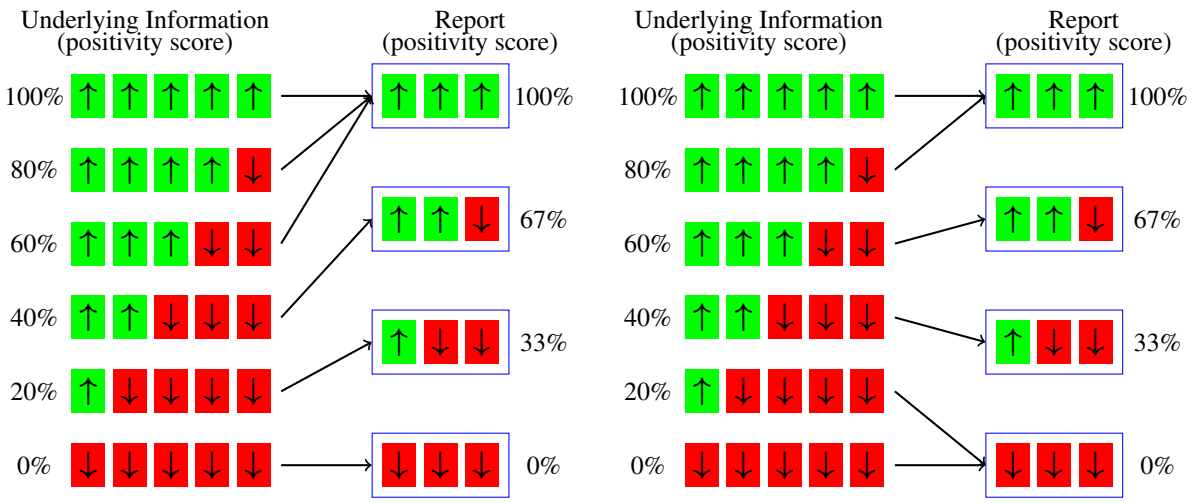


FIGURE I: Audience Appeal (left) and Sensationalism (right)

Notes: This illustration shows the equilibrium information structures of Example 2 in Section 2.5, with $N = 5$, $n = 3$, $\pi u_1/(1-\pi)u_0 = 9$ (left), $\pi u_1/(1-\pi)u_0 = 1.5$ (right), and $h(a) = \sin(\varpi a/2)$.

the report is 33% (67%, 100%, 100%). Sensationalism is evident in the right panel: when the underlying information is 20% (40%, 60%, 80%) positive, the report is 0% (33%, 67%, 100%).

This illustration shows that bias, or the apparent coverage imbalance, is the manifestation of different optimal pooling intensities across underlying information scenarios. In the left panel, underlying scenarios that confirm the inclined action are pooled aggressively. In the right panel, tail scenarios are pooled aggressively. Obviously, characterizing pooling intensity is equivalent to characterizing cutoff density. In the asymptotic model, I define the *newsworthiness* of underlying information as the cutoff density near the value of that underlying information. I derive an explicit formula for newsworthiness in the main theorem; I show that underlying information is more newsworthy if the conditional likelihoods and payoff sensitivities, as functions of underlying information, have higher values and elasticities. The model implies that media bias is simply how a news outlet and its readers communicate tacitly and efficiently, thereby maximizing welfare.

The second contribution is that the model offers a tractable economic foundation for sentiment analysis. Asymptotically, the sentiment-scoring function expressed as the cumulative distribution function of the cutoffs has a tractable form with explicit contextual parameters. This provides a clear answer to our microfoundation question. This answer has three appealing properties. First, it clarifies the concept of textual sentiment by showing its determinants and nonlinearity. Here, the main insight is that sentiment scoring is sensitive to the true signal when its realization is newsworthy, which varies across realizations and depends on the economic context. Second, it is simple: for common payoffs (e.g., quadratic, exponential, log, power, and cosine-difference), the cutoff density is the normalized product of a Gaussian density and another common density. Third,

it is based on information selection, which is not uncommon. This theory is best applicable when (1) text is selectively presented evidence without much modification; (2) the order of presentation hardly matters; and (3) persuasion and media competition are not the main focus – conditions not too restrictive for many text types.

Model tractability allows us to analyze interpretation and measurement issues transparently. I discuss how context and model nonlinearity affect bias and sentiment scoring. For example, I explain how prior beliefs and payoffs affect directional bias, and how increased informativeness of the underlying information leads to greater sensationalism. The model shows that text and its sentiment should be interpreted not literally but structurally. I derive the model-implied distribution for sentiment data and discuss the identifiability of information from such data. I further discuss the types of misspecification that can arise from not treating sentiment structurally: for example, I explain how confusing sentiment with the true signal may exaggerate the importance of extreme and contradictory shocks (thereby discussing the common empirical practice of treating positive- and negative-sentiment subsamples separately) and how neglecting context changes may lead to spurious significance in empirical results.

Beyond textual analysis and media bias, this paper contributes to several other literatures. It is related to work on communication with commitment (Kamenica and Gentzkow 2011, Bergemann and Morris 2019) under limited capacity (Gentzkow and Kamenica 2014, Bloedel and Segal 2021), but replaces information-theoretic capacity constraints (Sims 2003) with a practically motivated one and pins down report labels in addition to action recommendations. This paper also contributes to the literature on partial disclosure of evidence (Milgrom 1981, Milgrom and Roberts 1986, Dye 1985).

2 The Discrete Information Selection Problem

2.1 Setup

A sender (*she*) reports information about the binary state $\theta \in \{0,1\}$ to a receiver (*he*). The receiver's prior is $\Pr(\theta = 1) = \pi \in (0,1)$. After receiving the sender's report, he chooses a bet $a \in [0,1]$ with payoff

$$u(a;\theta) = u_\theta h(1-|a-\theta|) = \begin{cases} u_1 h(a) & \text{if } \theta = 1, \\ u_0 h(1-a) & \text{if } \theta = 0, \end{cases} \tag{2}$$

where $u_0, u_1 > 0$ measure the payoff relevance of each state and $h : [0,1] \to \mathbb{R}$ captures how payoffs vary with the distance between the bet and the realized state. I impose the following assumption on $h(\cdot)$.

Assumption 1. *(i) The function $h(\cdot)$ is twice continuously differentiable on $(0,1)$, with $h'(a) > 0$ and $h''(a) < 0$ for all $a \in (0,1)$.*
(ii) For any interior posterior belief π' that arises with positive probability under the maintained information environment, the unique maximizer

$$a^* := \arg\max_a (1-\pi')u_0 h(1-a) + \pi' u_1 h(a)$$

lies in the interior of the action space, i.e., $a^* \in (0,1)$.

The *fundamental*, denoted by K, is a signal supported on $\{0, 1, ..., N\}$. The *report*, denoted by k, takes values in $\{0, 1, ..., n\}$, where $n < N$. The sender observes the realization of K and reports k to the receiver according to an ex ante designed information structure denoted by $\{\sigma_{Kk}\}$ where $\sigma_{Kk} = \Pr(k|K)$. I make the following assumption about the distribution of K.

Assumption 2. *The Bayes factor $\Lambda(K) := \frac{\Pr(K|\theta=1)}{\Pr(K|\theta=0)}$ is strictly increasing in K.*

The sender chooses an information structure to solve

$$\begin{aligned}
\max_{\{\sigma_{Kk}\}} \quad & U := \mathbb{E}[u(a^*(k);\theta) \mid \{\sigma_{Kk}\}] \\
\text{s.t.} \quad & \sigma_{Kk} \geq 0 \quad \forall\, K, k, \\
& \sum_k \sigma_{Kk} = 1 \quad \forall\, K, \\
& a^*(k) = \arg\max_{a \in [0,1]} \mathbb{E}[u(a;\theta) \mid k, \{\sigma_{Kk}\}] \quad \forall\, k.
\end{aligned}$$

2.2 Explanation of the setup

Information selection. The setup is a communication problem with a sparse symbol set. To connect it to an information selection problem, consider the following interaction. First, the sender receives N binary *signals* $\xi_1, \ldots, \xi_N \in \{0,1\}$ from nature. She then includes in her text n binary *elements* $r_1, \ldots, r_n \in \{0,1\}$. Finally, the receiver chooses an action a. Here, each signal ξ_i represents a piece of evidence—for example, a potential news story. I refer to a full profile of signal realizations $\xi = (\xi_1, \ldots, \xi_N) \in \{0,1\}^N$ as a *scenario*. I assume that $\xi_1, \ldots, \xi_N$ are conditionally independent given θ, and that $\Pr(\xi_i = \theta \mid \theta) = p > 1/2$ for every i. Each reported element r_i represents a story included in the text. For now, I do not require reported elements to be truthful. I refer to the collection of reported elements $r = (r_1, \ldots, r_n) \in \{0,1\}^n$ as the *content*, and assume that their order conveys no information.

The integer n represents the *content size*, which serves as the capacity constraint in this model. It should be understood as a requirement or convention governing content size: the sender must report exactly n elements and cannot report more or fewer. By imposing this equality constraint,

I am effectively assuming that the sender cannot convey information by leaving parts of the text blank. In practice, newspaper space, broadcast time slots, or report-length norms are capacities that are neither exceeded nor left partially empty. While these constraints may be endogenously determined ex ante due to various strategic considerations, they remain fixed once the communication stage begins. The choice of n is itself an interesting problem, but the paper's main results do not depend on it; rather, they depend on the most efficient way to communicate given n. Therefore, n is treated as exogenously fixed.

The information structure in this problem is $\{\sigma_{\mathbf{sr}}\}$ which makes the problem appear high dimensional, but it can be simplified substantially: we only need to consider the fundamental and the report, which are the sums of the signals and the reported elements and serve as their respective summary statistics. Under this formulation, the fundamental has a conditionally binomial distribution:

$$K \mid \theta \sim \mathrm{Bin}\big(N,\, p\theta + (1-p)(1-\theta)\big),$$

which satisfies Assumption 2.

Notably, the model generalizes this information selection problem by accommodating a broader range of fundamental distributions. Beyond textual data, it may apply to various data types, such as consumer ratings or academic grades. For instance, consider a sender who rates her experience K with a product on a five-star scale ($k = 4$) for a prospective shopper, who then chooses a purchase probability or quantity. The product type is either good or bad and affects K, which takes more than five values and satisfies Assumption 2. Similarly, a sender may design an exam with discrete grades to signal a student's performance K to a potential employer. Ultimately, the model's defining feature is that the fundamental realizations are richer than the available reports, creating capacity frictions that necessitate information compression.

Payoffs. The assumption $h'(a) > 0$ implies that the receiver is better off the closer his bet is to the state. The assumption $h''(a) < 0$ reflects diversification benefits: it allows for interior actions $a \in (0, 1)$ to be relevant.[4] Furthermore, part (ii) of Assumption 1 requires that the optimal action is *always* in the interior. This assumption is not essential: If some extreme K values induce an optimal action of 1 (0), then these K values should be treated as one equivalent K value, with its Bayes factor being still extreme. This reduces the effective number of fundamental values. The problem remains as long as this number still exceeds the number of reports.

The sender shares the receiver's preferences and prior belief; that is, she is benign. This assumption can be justified by simple or strategic loyalty. For the latter, suppose the sender can strategically position her reporting perspective described by $u_s(a;\theta)$ and π_s, and assume her profit

4. If $h''(a) \geq 0$, then the only possible optimal actions are $a = 0$ and $a = 1$, and so communication is trivial, since the sender has enough reports at her disposal to perfectly recommend one or the other.

from providing information services is increasing in the expected utility gain that her report delivers to the receiver. It is then optimal for the sender to align her perspective with the receiver's $u(a;\theta)$ and π and to aim to maximize his utility. This sender specification allows us to focus on information compression by muting persuasion.

2.3 Characterizing the solution: minimizing information loss

This part characterizes the optimal solution to narrow down the search for it. The next three propositions, proved in the appendix, give necessary conditions that the sender must meet to minimize information loss when assembling her report.

First, the sender must use pure strategies. Intuitively, since she is benign, she has no incentive to introduce unnecessary noise by mixing. Therefore, the optimal strategy is to partition the fundamental set and map each partition set to the same report.

Proposition 1 (pure strategy)**.** *Under Assumptions 1 and 2,* $\sigma^*_{Kk} \in \{0,1\}$ *for any* K*,* k*.*

Second, she must use all reports. The intuition here is that each additional report allows the sender to distinguish between more fundamentals by further refining the pooling structure (a map of partition sets to reports); thus, reports should not be wasted. In other words, given a pooling structure, if the sender has an extra report at her disposal, then she can split one of the partition sets into two parts, map one part to the original report, and map the second part to the new report; this will strictly increase expected utility. Therefore, an optimal strategy must use all $n+1$ reports.

Proposition 2 (surjection)**.** *Under Assumptions 1 and 2, for each* k *there exists* K *such that* $\sigma^*_{Kk} > 0$*.*

The third proposition helps describe how fundamental values are pooled.

Proposition 3 (cutoff structure)**.** *Under Assumptions 1 and 2, let* $\{B_0, ..., B_n\}$ *denote the partition of the fundamental set corresponding to an optimal information structure* $\{\sigma^*_{Kk}\}$*. Then there exist cutoffs* $\{K^*_1, ..., K^*_n\}$ *such that* $B_0 = \{0, ..., K^*_1 - 1\}$*,* $B_1 = \{K^*_1, ..., K^*_2 - 1\}$*, ... ,* $B_k = \{K^*_k, ..., K^*_{k+1} - 1\}$*, ... ,* $B_n = \{K^*_n, ..., N\}$*.*

Intuitively, to minimize information loss, the sender should pool "similar" fundamental values. The appropriate measure of similarity between two values is the closeness of their Bayes factors; by Assumption 2, this implies that an optimal solution involves an *ordered* partition, with adjacent fundamental values pooled. Thus, the partition is characterized by n cutoffs separating the fundamental set into $n+1$ partition sets. The following example visually illustrates this intuition.

Example 1. Let $u(a;\theta) = u_\theta \cos\left(\frac{\varpi}{2}|a-\theta|\right)$, where ϖ represents the mathematical constant pi.

(This corresponds to setting $h(a) = \sin\left(\frac{\varpi}{2}a\right)$.) We consider the information selection interpretation and have $\Lambda(K) = (p/(1-p))^{2K-N}$. The sender's objective is

$$U = \sum_{0 \leq i \leq n} \| \sum_{\mathbf{v}_K \in B_i} \mathbf{v}_K \|,$$

where $\mathbf{v}_K = \left((1-\pi)u_0\binom{N}{K}p^{N-K}(1-p)^K, \pi u_1\binom{N}{K}p^K(1-p)^{N-K}\right) \in \mathbb{R}^2$, $\|\cdot\|$ is the Euclidean norm, and B_i is an element of the partition $\{B_0, ..., B_n\}$.

Geometrically, each fundamental value K can be represented by a vector $\mathbf{v}_K$ (shown as an arrow in Fig. II). The sender calculates her utility as follows: First, she partitions the collection of all fundamental vectors into $n+1$ partition sets. Then, for each partition set, she calculates the vector sum of all of the fundamental vectors in that set. We call the resulting vector the *representative vector*; its Euclidean length equals the contribution of the fundamental vectors in that partition set to the sender's expected utility. Finally, she calculates the sum of the lengths of all of the representative vectors; this equals the expected utility.

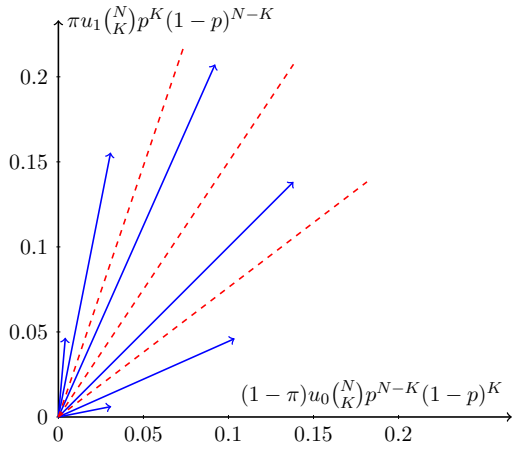


FIGURE II: Illustration of Example 1

Notes: This illustration shows the fundamentals (given by arrows, corresponding to $K = 0, ..., 5$ anticlockwise) and the partition (given by dashed lines) in Example 1. Parameter values are as follows: $N = 5$, $n = 3$, $\pi = p = 0.6$, $u_1 = u_0 = 1$.

To maximize the expected utility, the sender should choose a partition that minimizes the loss of length caused by summing the vectors in each partition set; thus, she should pool vectors with similar angles. The angle associated with $\mathbf{v_K}$ is $\tilde{a}(K) = \arctan\frac{\pi}{1-\pi}\frac{u_1}{u_0}\Lambda(K)$, which is strictly monotone in $\Lambda(K)$ and hence in K. (This angle stands for the optimal action if the receiver knows the fundamental value perfectly.) Therefore, the sender should pool adjacent K. Figure II depicts the optimal partition in an example with $N = 5$.

2.4 Modeling text content: connecting symbols with substance

Note that to maximize expected utility, it suffices to find an optimal partition; it does not matter which report is assigned to each partition set, as long as each set is associated with a distinct report. This means multiple equilibria always exist, as the sender can permute the reports arbitrarily without affecting utility. The assignment of reports is therefore merely a matter of labeling.

In this part, we refine equilibria by associating reports with meanings. We have assumed that the sender is a presenter of information. This is because in the real world, content data are not arbitrary symbols; readers of a newspaper, for example, will not accept its action recommendations unless they appear to be fact-based. Consequently, many senders insist on presenting evidence when other, more space-saving forms of communication are available, such as merely sending a short unsubstantiated summary or directly stating the recommended action. In this part, we further argue that the sender must not only frame her reports as a presentation of reality, but also ensure that they withstand common-sense scrutiny from readers. The goal is to give symbols substantive meaning.

I introduce two content criteria that reflect common audience expectations of trustworthiness and reliability. The first criterion is *(verifiable) honesty*, or the ability to survive fact-checks. This criterion requires that reported elements match actual signal realizations.

Definition 1 (the honesty criterion)**.** *An information structure is* honest *if, for each K, any k such that $\sigma_{Kk} > 0$ satisfies $K - (N - n) \leq k \leq K$.*

Importantly, this criterion does not require balanced coverage; it only requires that no constituent element of the content be fabricated. Suppose $N = 100$ and $n = 50$, and the realized scenario consists of 50 ones and 50 zeros. Then content consisting of 50 ones and no zeros would still be honest, though plainly biased. Hence the honesty criterion is not difficult to satisfy.

The second criterion is *(logical) self-consistency*, which requires that the sender not contradict herself. That is, the report must be increasing as a function of the fundamental.

Definition 2 (the self-consistency criterion)**.** *An information structure is* self-consistent *if for any K_1, K_2 such that $K_1 < K_2$, the conditions $\sigma_{K_1 k_1} > 0$ and $\sigma_{K_2 k_2} > 0$ imply $k_1 \leq k_2$.*

Under self-consistency, if one fundamental value favors $\theta = 1$ over $\theta = 0$ more than another, then the report associated with the former fundamental value should exhibit more favor toward $\theta = 1$ as well. To interpret this, note that there are two meanings associated with a report: a literal meaning, arising from the fact that the content appears as a collection of signals, and a true meaning regarding the fundamental, which is implied by the information structure. Self-consistency requires these two meanings to move in the same direction. For instance, suppose $\theta = 1$ and $\theta = 0$ represent good and bad states of the economy. Self-consistency then means that whenever the

sender receives better news about the economy, her report appears more optimistic. Information intermediaries that fail to maintain self-consistency may be perceived as untrustworthy.

Proposition 4 shows that we can refine the equilibria to obtain a desirable label structure.

Proposition 4. *Under Assumptions 1 and 2, there exists a solution that satisfies both honesty and self-consistency. Moreover, any solution that satisfies self-consistency also satisfies honesty.*

The proof is straightforward: For any solution, let the optimal partition sets $B_0, ..., B_n$ map to the reports $0, ..., n$, in that order. This information structure is the only self-consistent one under the optimal partition. It is also obviously honest: Because each B_i has i (and $n - i$) nonempty partition sets below (and above) it, the fundamental values within B_i must be no smaller than i and no higher than $N - (n - i)$; assigning report i to any fundamental value within this range does not violate the honesty criterion.

I call this self-consistent solution the *content-generating information structure*. It must exist,[5] although I have not yet specified where the optimal cutoffs are located. This task is deferred to Section 3, where I take the discrete model to an asymptotic limit. For now, the discrete problem can be solved by exhaustively computing utilities for all ordered partitions, a manageable task.

2.5 Illustration of content bias under the discrete model

In the discrete problem, *bias* refers to the difference between the two positivity scores k/n and K/N under the content-generating information structure.[6] Example 2 illustrates how bias may arise.

Example 2. The market state θ tomorrow is either a boom (1) or a bust (0). A typical investor has access to two assets: one pays off u_1 in a boom and zero in a bust, while the other pays zero in a boom and u_0 in a bust. He chooses a portfolio proxied by $a \in [0, 1]$, which represents his position in the former asset. The market belief about a boom is π. The newspaper editor has $N = 5$ news stories about θ, but must publish $n = 3$. The content-generating information structure depends on the contextual parameters, including π, u_1, u_0, p, and the shape of $h(\cdot)$. Here I focus on π, u_1, and u_0.

The left panels of Figures I and III present an example in which πu_1 greatly exceeds $(1-\pi)u_0$; that is, the investor prefers or expects the boom state. In this case, audience-appeal bias emerges. The content-generating information structure, shown in the left panel of Figure I, is given by the partition sets $\{0\}$, $\{1\}$, $\{2\}$, and $\{3, 4, 5\}$, which are mapped to the reports 0, 1, 2, and 3, in that

5. For its existence, see Appendix A. It is also generally speaking unique, in the sense that asymptotically it is unique by Theorem 1.

6. As previously noted, the discreteness of the model may contribute to this difference, but this contribution is unimportant and will disappear in the asymptotic model of Section 3.

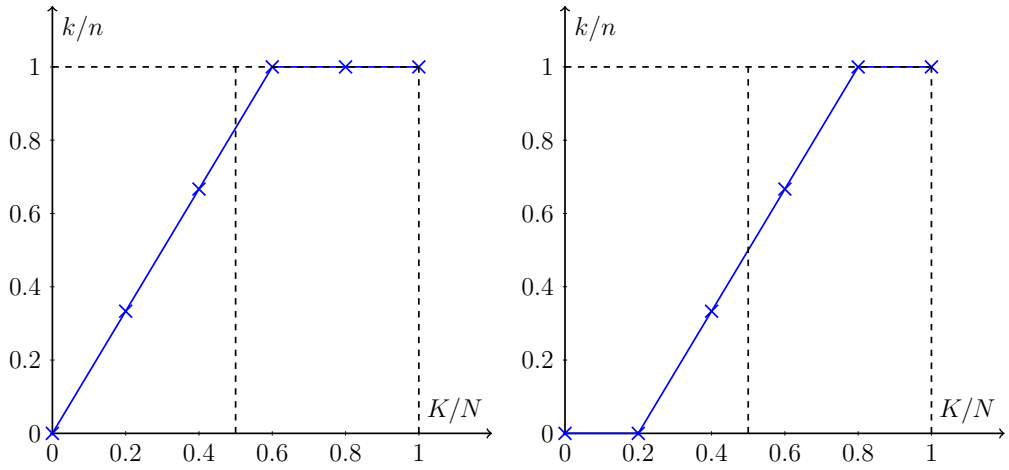


FIGURE III: Audience Appeal (left) and Sensationalism (right): Sentiment Curves

Notes: This illustration shows the (interpolated) sentiment curves of Example 2, with $N = 5$, $n = 3$, $\pi u_1/(1-\pi)u_0 = 9$ (left), $\pi u_1/(1-\pi)u_0 = 1.5$ (right), and $h(a) = \sin(\varpi a/2)$.

order. The left panel of Figure III shows the *sentiment curve*, a plot of k/n against K/N that serves as the graphical representation of the sentiment-scoring function (given here as a scatter plot). We see from the shape of this curve that the newspaper disproportionately omits negative stories, appearing to cater to the investor's belief or preference. For instance, when the scenario is 20% positive, with one positive story and four negative stories, the editor omits two of the latter, resulting in content that is 33% positive. When the scenario is 60% positive, with three positive and two negative stories, the editor omits both of the latter and publishes content that is 100% positive.

The right panels of Figures I and III illustrate a case where πu_1 and $(1-\pi)u_0$ are comparable, giving rise to sensationalism bias. The information structure consists of the partition sets $\{0,1\}$, $\{2\}$, $\{3\}$ and $\{4,5\}$ mapped to reports 0, 1, 2, and 3, respectively; that is, the newspaper provides exaggeratedly positive or negative content depending on the direction of the fundamental. For instance, when the scenario is 80% (20%) positive, the content is 100% (0%) positive. When the scenario is 60% (40%) positive, the content is 67% (33%) positive.

To analyze these biases, we must interpret the cutoff locations. As previously noted, I derive these locations tractably in the asymptotic model; here, I provide some intuition.

In the audience-appeal example, where πu_1 is high, the investor strongly tends to bet close to 1. As a result, he places little value on information confirming $\theta = 1$, because such information does not substantially alter his bet. Instead, he values precise information about scenarios far from $\theta = 1$, which may prompt him to reconsider his bet. To maximize the investor's utility, the newspaper editor pools positive-leaning fundamental values more aggressively, reserving higher precision for the negative news that triggers a bigger change in the investor's action.

In the sensationalism example, where audience-appeal effects are weak, the investor is not very interested in precise information about tail scenarios. There are two reasons for this: first, tail scenarios are statistically improbable and thus contribute minimally to expected utility; second, in such extreme scenarios, the investor will adopt a near-boundary bet regardless of the report's specific precision. Therefore, the editor pools tail scenarios more aggressively than moderate ones, concentrating informational precision in the middle of the distribution where the investor's action is most sensitive to new information.

Once the partition is established, the analysis of bias is straightforward. Two observations are helpful. First, the highest and lowest fundamental positivity scores are mapped to the highest and lowest reported scores, respectively, without bias. These polar scenarios serve as anchors for analyzing bias in the reporting of intermediate scenarios. Second, the assignment of reports is more sensitive to changes in the fundamental for more newsworthy fundamental values (e.g., negative-leaning values in the audience-appeal example, or moderate values in the sensationalism example). Consequently, the same incremental increase in K/N leads to a larger increase in k/n when K/N is more newsworthy. This variation in sensitivity stems from the non-uniform pooling intensity. Hence, for an interval of more newsworthy K/N, the associated range of k/n is larger.

In the audience-appeal case, fundamental positivities that contradict the receiver's preference or belief are more newsworthy, so the range of reported positivities is disproportionally wide, spanning reported values that might otherwise seem less contradictory or even confirmatory. Because the total range of reported scores is fixed between 0 and 1, the less newsworthy confirmatory fundamental scores are forced into a compressed range of even more extreme confirmatory reported scores. In the sensationalism case, moderate fundamental values have a reported-score range large enough to encompass values that appear more extreme. Within a fixed reported-score range between 0 and 1, near-polar fundamental values are thus associated with more extreme reported scores.

The sentiment curve, which depicts the relationship between fundamental and reported positivity scores, provides a graphical representation of the solution. In the discrete model, it takes the form of a scatter plot. An important observation is that the cumulative distribution of cutoffs passes through and extends the sentiment curve. Denoting by $\mathcal{K}$ the set of cutoffs (as is given in Proposition 3), we define

$$\tilde{s}(t) = \frac{1}{n} \sum_{K^* \in \mathcal{K}} \mathbf{1}_{K^*/N \leq t}$$

as the cumulative distribution of cutoffs, where t is the fundamental positivity score. This observation is striking: the positivity share of the content produced for a fundamental value coincides with the cumulative cutoff share up to that fundamental value. This observation is important for two reasons: (1) it links the sentiment curve to the cutoff distribution, which is easier to analyze

in the asymptotic model; and (2) it provides a novel perspective on capacity. We already know the capacity is n, corresponding to the content size or, equivalently, the number of cutoffs. Alternatively, we may view capacity through a cutoff random variable that takes values in $\mathcal{K}$ with a uniform probability of $1/n$ assigned to each realization. Under this interpretation, capacity corresponds to one unit of probability mass. The allocation of capacity across fundamental values is then characterized by the distribution of this random variable: the sender allocates more capacity to reporting a given fundamental value by placing a larger mass of cutoffs in its neighborhood.

2.6 Welfare

In this rational framework, bias is not a distortion of beliefs but a distortion of labels. The receiver rationally inverts the content to recover the underlying information. The model's key insight is to expose a systematic wedge between the surface meaning of a label and the information it truly conveys. What appears as bias at the level of language is the product of strategic information compression rather than a misperception of reality by either party of communication.

Contrary to the conventional wisdom that bias harms welfare, this model shows that biases may instead reflect efforts to maximize welfare. The view that bias reduces welfare is often based on the presumption that it arises from agency problems. This paper shows, however, that bias may emerge even in the absence of agency problems, simply as a means of increasing communication efficiency under a capacity constraint. In this model, bias enhances the welfare of both the receiver and society, since in both cases welfare is measured by the sender's ex ante maximized utility.

The parameters N and n both affect welfare. Maximized utility is strictly increasing in N, because richer fundamental information enables the sender to recommend better actions under the same communication constraints. Maximized utility is also strictly increasing in n, as a corollary of Proposition 2: given $n_1 < n_2$, the optimal information structure under n_1 is suboptimal under n_2, because it fails to exploit the additional capacity. Intuitively, welfare loss arises solely from information compression, and a bigger n implies a smaller loss.

The results also imply that bias-free communication policies are suboptimal for both the receiver and social efficiency. Such policies notably include two approaches that are widely regarded as ethical: (1) the sender produces a report that resembles the fundamental as closely as possible; and (2) the sender fully randomizes her reporting without any deliberate selection. In the latter case, the sender's ex ante expected utility is equivalent to that obtained by choosing n reported elements out of n signals, since she can simply report the first n of the N signals.

Should we therefore worry about bias in media coverage? The model suggests not. Of course, it assumes full rationality, which may not hold in practice, and it abstracts away from other potential channels for bias, as well as sociological, cultural and ethical factors. Nonetheless, it offers a

perspective on why the presence of news bias can be a socially reasonable outcome.

3 The Main Asymptotic Model

In this section I extend the discrete model to an asymptotic model. I begin by introducing the asymptotic setup and then justify its adoption by formalizing its connection to the discrete model.

3.1 Setup

The agents have the same payoffs as in the discrete model. I impose the following assumption, which is analogous to Assumption 1 in the discrete model.

Assumption 3. *The function $h(\cdot)$ is six times continuously differentiable, with $h'(a) > 0$ and $h''(a) < 0$ for all $a \in (0,1)$.*

The fundamental K in the asymptotic model is not necessarily positive or integer-valued, but takes values on the real line. I impose the following assumption on K.

Assumption 4. *The conditional distributions of the fundamental under $\theta = 1$ and $\theta = 0$, denoted by $F_1(K)$ and $F_0(K)$ respectively, satisfy the following properties:*
(i) They are absolutely continuous and six times continuously differentiable, with densities $f_1(K)$ and $f_0(K)$, respectively.
(ii) There exist $\underline{K}^{(0)},\ \underline{K}^{(1)},\ \bar{K}^{(0)}$ and $\bar{K}^{(1)}$ satisfying $-\infty \leq \underline{K}^{(0)} \leq \underline{K}^{(1)} < \bar{K}^{(0)} \leq \bar{K}^{(1)} \leq +\infty$, such that the conditional density satisfies $f_1(K) > 0$ if and only if $K \in (\underline{K}^{(1)}, \bar{K}^{(1)})$, and $f_0(K) > 0$ if and only if $K \in (\underline{K}^{(0)}, \bar{K}^{(0)})$.
(iii) The likelihood ratio $f_1(x)/f_0(x)$ is strictly increasing on $(\underline{K}^{(1)}, \bar{K}^{(0)})$.
(iv) The range of $(1-\pi)u_0 h'(1-a)/\pi u_1 h'(a)$ for $a \in (0,1)$ is a subset of the range of $f_1(K)/f_0(K)$ for $K \in (\underline{K}^{(1)}, \bar{K}^{(0)})$.

I consider the following optimal partition problem. Fix n and let $\mathcal{K}(n) = \{K_1, K_2, ..., K_n\}$ be a set of cutoffs that partitions the real line into $\{(-\infty, K_1), [K_1, K_2), ..., (K_n, +\infty)\}$. Denote these partition sets by $\{B_0, B_1, B_2, ..., B_n\}$, in that order. The sender chooses these cutoffs to maximize her ex ante expected payoff:

$$\sum_{i=0}^{n} \left\{ \pi u_1 \Pr(K \in B_i \mid \theta = 1)\, h(a_i^*) + (1-\pi) u_0 \Pr(K \in B_i \mid \theta = 0)\, h(1-a_i^*) \right\}. \tag{3}$$

For each i, the term a_i^* in Eq. (3) denotes the recommended action associated with B_i and therefore satisfies the first-order condition

$$\frac{\pi u_1 \Pr\left(K \in B_i | \theta = 1\right)}{(1-\pi) u_0 \Pr\left(K \in B_i | \theta = 0\right)} = \frac{h'(1-a_i^*)}{h'(a_i^*)}.$$

The reporting policy is as follows. Denote by $\mathcal{K}^*(n)$ the set of optimal cutoffs. For any realized signal $K \in B_i$, the sender issues a report i. This leads to a positivity score that coincides with the cumulative cutoff distribution

$$s_n(K) = \frac{1}{n} \sum_{K' \in \mathcal{K}^*(n)} \mathbf{1}_{K' \leq K}.$$

We investigate the following questions as $n \to \infty$: (1) whether $s_n(K)$ converges pointwise, and (2) if so, what its pointwise limit is. Denote this limit by $s(K)$ if it exists.

3.2 Explanation of the setup

Information selection. Let us consider the previously discussed discrete information selection problem. Now, let the binary signals ξ_i take values in $\{-\sigma/\sqrt{N}, \sigma/\sqrt{N}\}$ (where $\sigma > 0$ is a parameter) instead of $\{0, 1\}$. For a fixed N, this is simply an affine transformation of the signal space; the resulting information environment is equivalent to the original one. The fundamental $K = \sum_{i=1}^{N} \xi_i$ continues to serve as a sufficient statistic for the underlying scenario. Note that K is no longer necessarily a nonnegative integer; it is a real number and may be negative.

Let $\mu > 0$ be another parameter, and for each i, assume

$$\Pr\left(\xi_i = \frac{\sigma}{\sqrt{N}} \middle| \theta\right) = \frac{1}{2}\left(1 + \frac{\mu_\theta}{\sigma\sqrt{N}}\right),$$

where $\mu_1 = \mu$ and $\mu_0 = -\mu$. This assumption links the probability p in the discrete model to N. Note that the ratio μ/σ measures how informative each signal is about the state θ.

We require the discrete baseline model for the following insights. First, for fixed N and n, Propositions 1, 2, 3 and 4 apply, and the content-generating information structure is characterized by n cutoffs in the fundamental space. Second, the sentiment-scoring function can be extended to the cumulative cutoff distribution.

I am interested in the asymptotic limit of the solution to the information selection problem. There are two reasons. First, asymptotics brings tractability. Second, the asymptotic framework provides a reasonable approximation to the reality. When N and n are large, both reality and text consist of many small pieces of information ("details"), each of which is only weakly informative. Intuitively, both reality and text contain rich, detailed, and complex information. A large N captures real-world information abundance, while a large n captures real-world content complexity.

Despite the richness of information in text, I assume that reality is far more sophisticated, vastly exceeding the capacity of any text to capture it fully. I model this feature using the following two-step limiting procedure. First, I fix n and let N tend to infinity. For any finite N, the content-generating information structure is characterized by n cutoffs in the fundamental space; thus, in the limit as $N \to \infty$, the information structure is given by n cutoffs on the real line (denoted by $\mathbb{R}$).

Second, I let n tend to infinity. Then, the (limiting) content-generating information structure is then characterized by an infinite collection of cutoffs, normalized to a unit measure and continuously distributed on $\mathbb{R}$, thereby forming a cutoff density.

For the first step, with n fixed and $N \to \infty$, the analysis proceeds as follows. By standard arguments for infill asymptotics, the fundamental under N, denoted by $K^{(N)}$, satisfies

$$K^{(N)}|\theta \Rightarrow N(\mu_\theta, \sigma^2),$$

where $\Rightarrow$ denotes convergence in distribution. The limiting fundamental K therefore follows a Gaussian mixture distribution and is supported on $\mathbb{R}$. That is, in the limit, we have

$$K|\theta \sim N(\mu_\theta, \sigma^2),$$

satisfying Assumption 4. The sender's problem is to choose n cutoffs for the Gaussian-mixture fundamental on $\mathbb{R}$. The content-generating information structure specifies reporting $s_n(K)$ for each realization of K.

For the second step, as $n \to \infty$, the sender seeks to find the curve $s(K) := \lim_{n\to\infty} s_n(K)$, which is both the sentiment curve and the cumulative distribution of cutoffs in the limit.

Payoffs. Assumption 3 imposes more smoothness than Assumption 1, since the asymptotic analysis involves higher-order derivatives. Moreover, there is no requirement analogous to Assumption 1(ii), so we allow some fundamental values to induce the same extreme action.

To discuss this situation, I introduce the *perfect-information optimal action* $\tilde{a}(K)$, defined as the receiver's hypothetical optimal action if he knew K. Let

$$R(K) = \frac{\pi u_1 f_1(K)}{(1-\pi) u_0 f_0(K)}. \tag{4}$$

When $h'(1)/h'(0) < R(K) < h'(0)/h'(1)$, we have $\tilde{a}(K) \in (0,1)$ determined by the equation $R(K) = h'(1-\tilde{a})/h'(\tilde{a})$. Otherwise, if $R(K) \leq h'(1)/h'(0)$, then $\tilde{a}(K) = 0$; if $R(K) \geq h'(0)/h'(1)$, then $\tilde{a}(K) = 1$. Under Assumption 4(iv), the range of $\tilde{a}(K)$ includes the interval $(0,1)$. If some fundamental values induce the extreme action of 1 (or 0), then they can be pooled without loss and mapped to the extreme report 1 (or 0). Hence it suffices to solve for $s(K)$ for values of K such that $\tilde{a}(K) \in (0,1)$, an interval that I denote by $(\underline{K}, \bar{K})$.

3.3 Main Result

The limiting sentiment curve $s(K)$, if it exists, is non-decreasing, with limits 0 and 1 as $K \to -\infty$ and ∞, respectively. If it is absolutely continuous, it is a cumulative distribution function with an associated density. This probabilistic perspective on the sentiment curve has two implications.

First, the sender's communication capacity is captured by a unit measure. As in the discrete

model, we can define a cutoff random variable. Asymptotically it is defined on the probability space $(\mathbb{R}, \mathcal{B}(\mathbb{R}), \mathbf{P}')$, where, for any $S \in \mathcal{B}(\mathbb{R})$, $\mathbf{P}'(S)$ equals the Lebesgue measure of $s(S)$. Communication capacity is then given by one unit of probability, corresponding to the total mass of cutoffs. If s admits a density $s'(K)$, I refer to $s'(K)$ as the *capacity allocation curve*, by analogy with the attention allocation curve in the literature on limited attention. It describes how the sender allocates the scarce resource of cutoffs across the fundamental space. In fact, although I view the capacity as a constraint binding the sender, it can also be interpreted as an information-receiving capacity constraint binding the receiver. Under this interpretation, $s'(K)$ can also be called the attention allocation curve. This novel perspective on capacity or attention is microfounded and complements the information-theoretic perspective prevalent in the literature.

I also refer to $s'(K)$ as the *newsworthiness curve*, because it captures the importance of each fundamental value: If $s'(K)$ is large at a given K, the sender places more cutoffs in the neighborhood of K, reflecting her assessment that this value warrants more elaborate coverage. This observation echoes and makes rigorous the discussion in Section 2.5 on the value of each scenario.

Second, instead of solving for $s(K)$, the sender may equivalently solve for its derivative $s'(K)$. This leads to Theorem 1, the main result of this paper. Define

$$H_1(K) := h(\tilde{a}(K)), \qquad H_0(K) := h(1 - \tilde{a}(K)).$$

Define the relative curvature components

$$\kappa_H(K) := H_1''(K)H_0'(K) - H_1'(K)H_0''(K),$$
$$\kappa_F(K) := F_1''(K)F_0'(K) - F_1'(K)F_0''(K),$$

and the effective relative curvature

$$\kappa(K) := \sqrt{\kappa_H(K)\kappa_F(K)}.$$

The product $\kappa_H(K)\kappa_F(K)$ is guaranteed to be nonnegative under Assumptions 3 and 4.

Theorem 1 (Asymptotic Capacity Allocation)**.** *Suppose Assumptions 3 and 4 hold. Then on* $(\underline{K}, \bar{K})$*,*

$$s'(K) \propto \kappa(K)^{1/3}, \tag{5}$$

provided the right-hand side is integrable.

Theorem 1 decomposes newsworthiness into two components: a payoff-sensitivity component $\kappa_H(K)$ which is related to the payoff function under perfect information, and an informational-sensitivity component $\kappa_F(K)$ which is related to the fundamental distribution. In logarithmic form, the log cutoff density is linear in the logs of $\kappa_H(K)$ and $\kappa_F(K)$, each with a coefficient of $1/6$. The sentiment curve is an antiderivative of $s'(K)$, normalized to satisfy the properties of a

cumulative distribution function. The proof of Theorem 1 is provided in Appendix B. I discuss the proof intuition in Section 3.4.

By the theorem, for the information selection problem where $K \sim N(\mu_\theta, \sigma^2)$, we immediately obtain

$$\kappa_F(K) \propto \exp\left(-\frac{K^2}{\sigma^2}\right).$$

Theorem 1 makes it easy to compute $s'(K)$ for many common payoff functions. Table I lists several examples. In some cases, the sentiment curve takes a particularly nice form. For instance, with exponential utility, the sentiment curve is the cumulative distribution function of a truncated $N(0, 3\sigma^2)$ distribution with asymmetric tail cutoffs. For quadratic utility, the term $\kappa_H(K)^{\frac{1}{6}}$ is proportional to a hyperbolic secant density, and for cosine-difference, log and power utilities it is proportional to certain powers of a hyperbolic secant density. A useful computational identity is: $\kappa_H(K) = -(h''(\tilde{a}(K))h'(1-\tilde{a}(K)) + h'(\tilde{a}(K))h''(1-\tilde{a}(K)))\tilde{a}'(K)^3$ which significantly simplify these derivations. Figure IV illustrates the capacity allocation and sentiment curves for an example with a cosine-difference payoff.

TABLE I: Equilibrium for Some Common Payoff Functions with $K|\theta \sim N(\mu_\theta, \sigma^2)$

	$h(a)$	$s'(K) \propto \exp\left(-\frac{K^2}{6\sigma^2}\right) \times ...$
Cosine difference	$\sin\left(\frac{\varpi}{2}a\right)$	$\left((\pi u_1)^2 \exp\left(\frac{2\mu}{\sigma^2}K\right) + ((1-\pi)u_0)^2 \exp\left(-\frac{2\mu}{\sigma^2}K\right)\right)^{-\frac{1}{2}}$
Quadratic	$A(1-a)^2 + B$	$\left(\pi u_1 \exp\left(\frac{\mu}{\sigma^2}K\right) + (1-\pi)u_0 \exp\left(-\frac{\mu}{\sigma^2}K\right)\right)^{-1}$
	$(A < 0)$	
Log	$\ln(a)$	$\left(\pi u_1 \exp\left(\frac{\mu}{\sigma^2}K\right) + (1-\pi)u_0 \exp\left(-\frac{\mu}{\sigma^2}K\right)\right)^{-\frac{1}{3}}$
Power	$\frac{1}{1-\gamma}a^{1-\gamma}$	$\left((\pi u_1)^{\frac{1}{\gamma}} \exp\left(\frac{\mu}{\gamma\sigma^2}K\right) + ((1-\pi)u_0)^{\frac{1}{\gamma}} \exp\left(-\frac{\mu}{\gamma\sigma^2}K\right)\right)^{\frac{\gamma-2}{3}}$
	$(\gamma > 0,\ \gamma \neq 1)$	
Exponential	$C_0 - C_1 \exp(-Aa)$	$\begin{cases} 1 & \text{for } -\frac{\sigma^2}{2\mu}\ln\left(\frac{\pi u_1}{(1-\pi)u_0}\right) - \frac{A\sigma^2}{2\mu} < K \ldots \\ & \quad \ldots < -\frac{\sigma^2}{2\mu}\ln\left(\frac{\pi u_1}{(1-\pi)u_0}\right) + \frac{A\sigma^2}{2\mu}; \\ 0 & \text{otherwise} \end{cases}$
	$(A,\ C_1 > 0)$	

As a final observation, let s^{-1} denote the inverse of s on $(\underline{K}, \bar{K})$, and let Φ be the standard Gaussian cumulative distribution function. Under $K|\theta \sim N(\mu_\theta, \sigma^2)$, the receiver's posterior belief upon observing a report $\rho \in (0,1)$ is determined by $K = s^{-1}(\rho)$, i.e.,

$$\Pr(\theta = 1|\rho) = \frac{\pi \exp\left(\frac{2\mu}{\sigma^2}s^{-1}(\rho)\right)}{\pi \exp\left(\frac{2\mu}{\sigma^2}s^{-1}(\rho)\right) + 1 - \pi}. \tag{6}$$

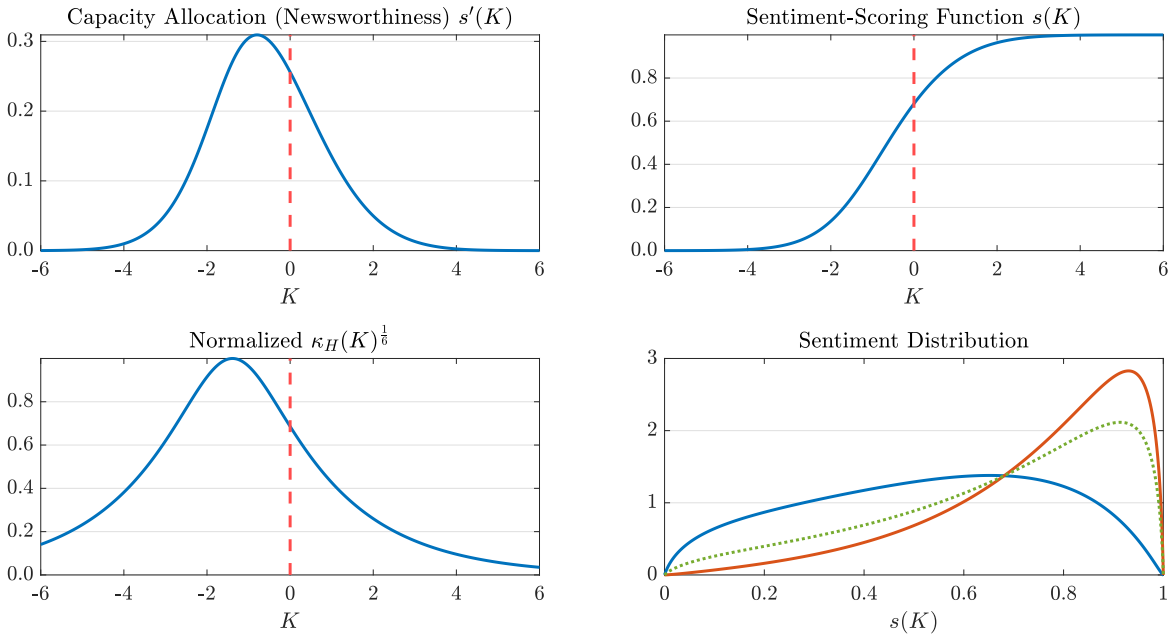


FIGURE IV: Illustration of Theorem 1

Notes: This is an illustration of Theorem 1 with cosine-difference utility. Parameter values are as follows: $\mu = 0.5$, $\sigma = 1$, $u_1/u_0 = 2$, $\pi = 2/3$. The term $\kappa_H(K)^{\frac{1}{6}}$ is normalized to equal 1 at $-\sigma^2 \ln(\pi u_1/(1-\pi)u_0)/2\mu$. In the plot of the sentiment distributions, the blue (less negatively skewed), orange (more negatively skewed), and dotted curves respectively show the conditional density on $\theta = 0$, the conditional density on $\theta = 1$, and the unconditional density.

For $\rho = 0$ and $\rho = 1$, the corresponding posteriors are respectively

$$\frac{\pi\Phi(\frac{\underline{K}-\mu}{\sigma})}{\pi\Phi(\frac{\underline{K}-\mu}{\sigma}) + (1-\pi)\Phi(\frac{\underline{K}+\mu}{\sigma})} \quad \text{and} \quad \frac{\pi(1-\Phi(\frac{\bar{K}-\mu}{\sigma}))}{\pi(1-\Phi(\frac{\bar{K}-\mu}{\sigma})) + (1-\pi)(1-\Phi(\frac{\bar{K}+\mu}{\sigma}))}.$$

3.4 Interpretation

Intuition of newsworthiness. Eq. (5) expresses capacity allocation or newsworthiness using derivatives of $H_1(K)$, $H_0(K)$, $F_1(K)$, and $F_0(K)$. They are intrinsic properties of K. The two components, $\kappa_H(K)$ and $\kappa_F(K)$, can be alternatively expressed as

$$\kappa_H(K) = H_0'(K)H_1'(K)\left(\frac{H_1''(K)}{H_1'(K)} - \frac{H_0''(K)}{H_0'(K)}\right) = H_0'(K)H_1'(K)\left(\ln\frac{h'(\tilde{a}(K))}{h'(1-\tilde{a}(K))}\right)',$$

$$\kappa_F(K) = f_0(K)f_1(K)\left(\frac{f_1'(K)}{f_1(K)} - \frac{f_0'(K)}{f_0(K)}\right) = f_0(K)f_1(K)\left(\ln\frac{f_1(K)}{f_0(K)}\right)'.$$

That is, $\kappa_F(K)$ is increasing in the two conditional likelihoods and the difference in the percentage rate of change of two likelihoods (or equivalently, the slope of the log likelihood ratio), while

$\kappa_H(K)$ is increasing in the two conditional payoff sensitivities to the fundamental and difference in the percentage rate of change of the two sensitivities (or equivalently, the slope of the log payoff sensitivity ratio).

This expression provides a clear dual-source intuition for newsworthiness. A fundamental value K warrants high communication capacity for two primary reasons:

(1) *High baseline importance*: This is reflected in high conditional likelihoods (f_1, f_0) or high marginal perfect-information payoffs (H_1', H_0'). Intuitively, a fundamental value K is significant if it is likely to occur or if it has a strong impact on payoffs regardless of which state θ is realized.

(2) *High action sensitivity*: As established in Eq. (4), the optimal action $\tilde{a}(K)$ is determined by the alignment of the marginal payoff ratio against the likelihood ratio. The log-slope components in the expressions for $\kappa_H(K)$ and $\kappa_F(K)$ represent the "velocities" of these two ratios. A value K is thus important if it strongly induces the receiver to reconsider his action; the strength of this induction stems from two distinct sources: how rapidly the evidence shifts its support from one state to the other, and how rapidly the payoff incentives shift between those states.

Proof intuition. The intuition for the proof of Theorem 1 is as follows. To find the limiting cutoff distribution, it is sufficient to determine the ratio of the cutoff density at any two fundamental values K_1 and K_2. This density ratio is inversely proportional to the ratio of the distances between adjacent cutoffs in the neighborhoods of K_1 and K_2. Hence, the proof focuses on the local spacing of cutoffs, governed by two key conditions: (1) the optimality of the receiver's action a_i^* given the observation of B_i, and (2) the optimality of the cutoff K_i^*.

The optimality condition for a_i^* yields $\kappa_F(K)$. Specifically,

$$\frac{F_1(K_{i+1}^*) - F_1(K_i^*)}{F_0(K_{i+1}^*) - F_0(K_i^*)} = \frac{(1-\pi)u_0 h'(1-a_i^*)}{\pi u_1 h'(a_i^*)} \equiv \frac{f_1(\tilde{K}(a_i^*))}{f_0(\tilde{K}(a_i^*))}. \tag{7}$$

Here, $\tilde{K}(a)$ is the inverse function of $\tilde{a}(K)$ on $(\underline{K}, \bar{K})$ and denotes the fundamental value that induces action a under perfect information. The second equality in Eq. (7) follows from the definition of $\tilde{a}(K)$. Using a high-order Taylor expansion, I show that this condition implies that the relative location of $\tilde{K}(a_i^*)$ within $[K_i^*, K_{i+1}^*]$ is[7]

$$\frac{\tilde{K}(a_i^*) - K_i^*}{K_{i+1}^* - K_i^*} \approx \frac{1}{2} + \frac{1}{24}\left(\frac{d\ln\kappa_F(K)}{dK}\bigg|_{K\in[K_i^*, K_{i+1}^*]}\right)(K_{i+1}^* - K_i^*). \tag{8}$$

The optimality condition for K_i^* yields $\kappa_H(K)$. Specifically,

$$\frac{h(a_i^*) - h(a_{i-1}^*)}{h(1-a_{i-1}^*) - h(1-a_i^*)} = \frac{(1-\pi)u_0 f_0(K_i^*)}{\pi u_1 f_1(K_i^*)} \equiv \frac{h'(\tilde{a}(K_i^*))}{h'(1-\tilde{a}(K_i^*))}.$$

7. Eq. (8) and Eq. (9) are variants of Eq. (B.3) and Eq. (B.4), respectively, in Appendix B.

The position of K_i^* within $[\tilde{K}(a_{i-1}^*), \tilde{K}(a_i^*)]$, is

$$\frac{K_i^* - \tilde{K}(a_{i-1}^*)}{\tilde{K}(a_i^*) - \tilde{K}(a_{i-1}^*)} \approx \frac{1}{2} + \frac{1}{24}\left(\frac{d\ln\kappa_H(K)}{dK}\bigg|_{K\in[\tilde{K}(a_{i-1}^*),\tilde{K}(a_i^*)]}\right)(\tilde{K}(a_i^*) - \tilde{K}(a_{i-1}^*)). \tag{9}$$

With these conditions linking the distances between $K_i^*, K_{i+1}^*, \tilde{K}(a_{i-1}^*)$ and $\tilde{K}(a_i^*)$ for any i, we can relate $K_{i+1}^* - K_i^*$ to $K_i^* - K_{i-1}^*$, that is, two adjacent-cutoff distances. Iterating this relationship allows us to relate adjacent-cutoff distances near any two fundamental values K_1 and K_2, yielding their distance ratio. This iteration involves integration of small terms along the interval between K_1 and K_2. Eventually, the limiting adjacent-cutoff distance ratio pins down the limiting cutoff distribution.

Information revelation. Importantly, the theorem is not about information revelation, even though it implies that the limiting equilibrium reporting policy entails no compression loss. In fact, now that the fundamental and the report both take values from a continuum, there exist many one-to-one maps from the fundamental to the report, any of which yields a loss-free reporting policy. The key point of Theorem 1 is that it identifies an *economically meaningful* one-to-one map—namely, the map that emerges as the limiting solution to the discrete information selection problem.

4 Media Bias and Sentiment

The theorem provides a tractable sentiment-scoring function, laying the foundation for analyzing media bias and its effect on sentiment analysis. This section examines the shape of this function. Throughout this section, we consider the information selection problem with $K|\theta \sim N(\mu_\theta, \sigma^2)$.

4.1 Two types of bias reflected in the sentiment-scoring function

Audience appeal. To identify audience-appeal bias in the reporting policy, I compare the location of $s'(K)$ with the "midpoint" of the fundamental space. It is appropriate to take this midpoint to be zero, since it is zero for any finite N. Relative to zero, the more $s'(K)$ leans against the direction given by the sign of $\frac{\pi u_1}{(1-\pi)u_0} - 1$, the more the reporting policy should be viewed as biased in that direction. Let $K_{1/2}$ denote the quantity $-\frac{\sigma^2}{2\mu}\ln\frac{\pi u_1}{(1-\pi)u_0}$, which solves $\tilde{a}(K) = 1/2$ under the information selection problem. Its sign is opposite to that of $\frac{\pi u_1}{(1-\pi)u_0} - 1$.

Definition 3. *The reporting policy given by the sentiment curve $s(K)$ is* strongly appealing *if $s'(K) \geq s'(-K)$, and* strongly alarmist *if $s'(K) \leq s'(-K)$, for any K such that $KK_{1/2} > 0$.*

This is a strong definition: If a strongly appealing $s(K)$ is the distribution of a random variable, then its mean (if well-defined), its median, and the average of its upper and lower αth percentiles for any α all have the same sign as $K_{1/2}$.

We know that $\kappa_F(K)^{\frac{1}{6}}$ is proportional to a $N(0, 3\sigma^2)$ density and does not introduce directional bias. Hence, any directional bias must arise from $\kappa_H(K)^{\frac{1}{6}}$, which is symmetric about $K_{1/2}$.[8]

Proposition 5. *Under Assumption 3 and $K|\theta \sim N(\mu_\theta, \sigma^2)$, $\kappa_H(K)$ is symmetric about $K_{1/2}$.*

This implies that $s'(K)$ is given by the product of two symmetric curves: one symmetric about zero and the other symmetric about $K_{1/2}$. This observation leads to Proposition 6, which provides sufficient conditions for the strongly appealing and strongly alarmist properties.

Proposition 6. *Assume that the right-hand side of Eq. (5) is integrable. Under Assumptions 3 and $K|\theta \sim N(\mu_\theta, \sigma^2)$, the following hold:*
(i) $s(K)$ is strongly appealing if $\kappa_H(K)^{\frac{1}{6}}$ is bell-shaped, i.e., increasing on $(\underline{K}, K_{1/2})$ and decreasing on $(K_{1/2}, \bar{K})$;
(i) $s(K)$ is strongly alarmist if $\kappa_H(K)^{\frac{1}{6}}$ is U-shaped, i.e., decreasing on $(\underline{K}, K_{1/2})$ and increasing on $(K_{1/2}, \bar{K})$;*
(ii) $\kappa_H(K)^{\frac{1}{6}}$ is bell-shaped (U-shaped) if and only if, for $a < \frac{1}{2}$,

$$\frac{d}{da}\ln h'(a) + \frac{d}{da}\ln h'(1-a) + \left(-2\frac{d}{da}\ln\left(\frac{d}{da}\ln\frac{h'(1-a)}{h'(a)}\right)\right) \geq (\leq)\, 0. \tag{10}$$

(iii) $\kappa_H(K)^{\frac{1}{6}}$ is bell-shaped if both $\frac{h''(a)}{h'(a)}$ and $\frac{h'''(a)}{h'(a)}$ are decreasing in a.

The conditions in Proposition 6 involve only h and are simple to verify. Many common payoffs, including cosine-difference, quadratic, log, power (with $\gamma \leq 2$), and exponential utilities, satisfy Proposition 6(i), and all of these except for the log and power utilities satisfy Proposition 6(iii).

Intuitively, Proposition 6 states that the reporting policy is strongly appealing whenever customizing extreme action recommendations is not overly important for overall utility from the perspective of $\kappa_H(K)^{\frac{1}{6}}$: specifically, $\kappa_H(K)^{\frac{1}{6}}$ must not explode near extreme scenarios. If it explodes, the model instead predicts that we observe an interesting alarmist bias: the sender reports more information that contradicts the receiver's prior or preference.

Eq. (10) provides a condition that characterizes the relative importance of customizing extreme actions versus middle actions. The left-hand side consists of three terms. The first two terms are related to the slopes of log marginal payoff levels. The sign of their sum depends on the absolute risk-aversion: For exponential payoff, it is zero; if the absolute risk-aversion is increasing (quadratic, cosine-difference), the sum is positive; if it is decreasing (log, power), the sum is negative, reflecting carefulness in choosing a near zero driven by concerns about losses from extreme bets. The third term captures a form of "ratio sensitivity" measuring the relative stakes of small deviations in action across alternative states. It is closely related to action sensitivity and, for the payoffs in Table I, is nonnegative, reflecting the importance of diversification and tailoring actions

8. The proofs of Propositions 5, 6 and 7 are in Appendix C.

near $\frac{1}{2}$ when the receiver is uncertain about which state to favor. For many utilities, the sum of all three terms is indisputably nonnegative, leading to the strongly appealing property. For log and power utilities, the sign depends on the relative risk-aversion γ: with $\gamma \leq 2$, the third term dominates the first two, leading to the strongly appealing property. By contrast, with $\gamma > 2$, the opposite occurs, resulting in the strongly alarmist property.

Example 3. Consider the constant relative risk-aversion (CRRA) utility functions $h(a) = \frac{a^{1-\gamma}}{1-\gamma}$ ($\gamma > 0$ and $\gamma \neq 1$) and $h(a) = \ln(a)$ ($\gamma = 1$). The parameter γ captures payoff curvature.

Figure V shows $\kappa_H(K)^{\frac{1}{6}}$ and $s'(K)$ for various values of γ. When $\gamma \leq 2$, $\kappa_H(K)^{\frac{1}{6}}$ satisfies Proposition 6(i), implying the strongly appealing property. When $\gamma > 2$, by contrast, $\kappa_H(K)^{\frac{1}{6}}$ is U-shaped, fitting Proposition 6(i*). We know that $s'(K)$ is given by the product of $\kappa_F(K)^{\frac{1}{6}}$ and $\kappa_H(K)^{\frac{1}{6}}$. Because $\kappa_H(K)^{\frac{1}{6}}$ is symmetric about $K_{1/2}$ and increases with the distance from $K_{1/2}$, it amplifies $\kappa_F(K)$ more strongly for values of K with the opposite sign to $K_{1/2}$, thereby making the capacity allocation curve larger at confirmatory values of K. This generates an interesting alarmist bias for large values of γ.

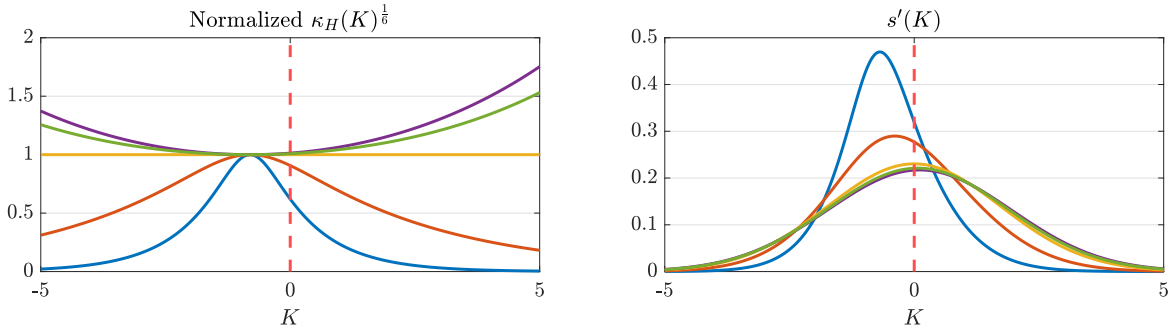


FIGURE V: Illustration of Example 3

Notes: This is an illustration of Example 3 with audience-appeal and alarmist biases. Parameter values are as follows: $\mu = 1$, $\sigma = 1$, $\pi u_1/(1-\pi)u_0 = 5$, $\pi = 2/3$. The term $\kappa_H(K)^{\frac{1}{6}}$ is normalized to equal 1 at $-\sigma^2 \ln(\pi u_1/(1-\pi)u_0)/2\mu$. The curves in blue, orange, yellow, green, and purple (from low to high for the left panel; from high to low in terms of values at $K = 0$ for the right panel) are for $\gamma = 0.5, 1, 2, 5, 10$, respectively.

Sensationalism. Under the integrability condition in Theorem 1, sensationalism is inevitable. Sensationalism can be described as the sender compressing two large regions of tail scenarios into two small bins of extreme reports; more precisely, it occurs when the slope of the sentiment curve becomes very small as K tends to $\pm\infty$. Essentially, as long as the sentiment-scoring function converges to a well-behaved, fixed-range function of $K \in \mathbb{R}$, sensationalism arises. This outcome is guaranteed by the structural assumption that N greatly exceeds n as both quantities tend to infinity. Intuitively, sensationalism is inevitable whenever the complexity of the underlying scenarios

far exceeds what an information intermediary can represent, however nuanced her reports may be—a feature common to many real-world situations.

Of the two components of newsworthiness, $\kappa_F(K)^{\frac{1}{6}}$ promotes sensationalism through its well-behaved tails, whereas $\kappa_H(K)^{\frac{1}{6}}$ may either reinforce or counteract it, depending on its tail behavior.

4.2 Comparative statics: the role of economic context

In this part, I discuss how several contextual parameters affect the sentiment curve under $K|\theta \sim N(\mu_\theta, \sigma^2)$. To begin with directional bias, Definition 4 compares the degrees of audience appeal and alarmism across sentiment curves.

Definition 4. *The sentiment curve $s^{(1)}(K)$* leans more positive *than another sentiment curve $s^{(2)}(K)$ if $s^{(2)}(K)$ first-order stochastically dominates $s^{(1)}(K)$, denoted by $s^{(2)} \succeq_{FOSD} s^{(1)}$.*

The sentiment curve depends on three effective parameters: (i) the directional-bias parameter $\psi := u_1\pi/u_0(1-\pi)$, (ii) the informativeness parameter $\varphi := \mu/\sigma$, and (iii) any parameter governing the shape of $h(\cdot)$. We have discussed the shape of $h(\cdot)$ in the previous part. This part focuses on (i) and (ii) for common payoffs.

Directional bias. First, consider parameter ψ. Its comparative statics depends on the shape of $\kappa_H(K)^{\frac{1}{6}}$. Let $\psi^{(1)} > \psi^{(2)}$ with corresponding sentiment curves $s^{(1)}(K)$ and $s^{(2)}(K)$. For cosine-difference, quadratic, log, and power utilities with $\gamma < 2$, the density ratio $s^{(1)\prime}(K)/s^{(2)\prime}(K)$ is strictly decreasing. Consequently, $s^{(2)}(K) \succeq_{FOSD} s^{(1)}(K)$ and thus $s^{(1)}(K)$ leans more positive. For an exponential utility, it is immediate that $s^{(1)}(K)$ leans more positive. In these cases, the reporting policy tilts toward the state with a higher level of ψ. By contrast, for power utilities with $\gamma > 2$, the ratio $s^{(1)\prime}(K)/s^{(2)\prime}(K)$ is strictly increasing, implying that $s^{(1)}(K)$ leans less positive.

Informativeness and sensationalism. Second, parameter φ contributes to sensationalism in the limit. To obtain this limiting behavior, let h satisfy Proposition 6(i), and fix μ. Let σ tend to zero. Then the bell-shaped curves $\kappa_H(K)^{\frac{1}{6}}$ and $\kappa_F(K)^{\frac{1}{6}}$ concentrate their mass around $K = 0$, and the sentiment curve $s(K)$ becomes extremely sensationalist:

Proposition 7. *Assume that the right-hand side of Eq. (5) is integrable. Under Assumptions 3 and 4, if $\kappa_H(K)^{\frac{1}{6}}$ is bell-shaped, then as $\sigma \to 0$,*

$$s(K) \to 0 \text{ for } K < 0 \quad \text{and} \quad s(K) \to 1 \text{ for } K \geq 0.$$

This happens because, as informativeness increases, the state becomes increasingly clear across wide ranges of the fundamental on each side. As a result, the sender allocates most capacity to a small intermediate region of the fundamental space, generating extreme sensitivity of sentiment to the fundamental.

5 Discussion: Using Sentiment As a Measure

The model provides a framework for microfounding sentiment analysis and makes progress in addressing the long-standing interpretational challenges inherent in using textual sentiment as an economic variable. Adopting a microfounded approach to textual analysis can improve our understanding of empirical results and enhance the design of empirical methods. Within the theoretical framework of this paper, this section discusses three benefits of incorporating microfoundations. First, microfoundations clarify the economic information embedded in text. Second, they mitigate the risk of spurious empirical findings arising from model misspecification. Third, they help guide the choice of appropriate text data and discipline the processing approach. Throughout this section, I assume $K|\theta \sim N(\mu_\theta, \sigma^2)$.

5.1 Interpreting sentiment and identifying economic information

Conceptual ambiguity has long been a challenge for sentiment analysis. There is no consistent or clear interpretation of what economic information textual sentiment represents, which in turn complicates its use as a measure in reduced-form analysis. As this model demonstrates, this difficulty is inherent to the problem: textual sentiment is a nonlinear function not only of fundamental information but also of various contextual parameters. As a result, ambiguity in interpretation is unavoidable without further structure. A microfoundation thus enables researchers to interpret their data with clarity and to structurally retrieve the specific information needed for analysis.

To exemplify the advantage of a microfounded approach, consider the case of news text analysis. Suppose a researcher believes that news sentiment contains information about the underlying fundamental K. Even if this interpretation is correct, the learning process involves three distinct economic objects: the prior belief before exposure to the news, the information shock (the news itself), and the posterior belief after learning. Without a microfounded model, it is often difficult to determine which of these three objects sentiment is conceptually measuring. The framework in this paper, for instance, provides a clear distinction: If the researcher's text data satisfy the model assumptions, then the news sentiment score ρ equals $s(K)$. The three economic variables are then: the prior belief π, the information shock K recovered from the sentiment score via the inverse mapping $s^{-1}(\rho)$, and the posterior belief obtained from a Bayesian update using this recovered shock, as given in Eq. (6). This structure conceptually separate these quantities.

Identifiable information in sentiment data. A natural question is what information can be retrieved from text sentiment. While text sentiment is often observed alongside other data, I consider text data in isolation here for simplicity. The model predicts that the conditional distribution of the

sentiment score ρ is

$$\begin{cases} \Pr(\rho = 0|\theta) = \Phi\left(\frac{\underline{K}-\mu_\theta}{\sigma}\right) & \text{if } \rho = 0, \\ F_{\rho|\theta}(\rho) = \Phi\left(\frac{s^{-1}(\rho)-\mu_\theta}{\sigma}\right) & \text{if } \rho \in (0,1), \\ \Pr(\rho = 1|\theta) = 1 - \Phi\left(\frac{\bar{K}-\mu_\theta}{\sigma}\right) & \text{if } \rho = 1, \end{cases} \tag{11}$$

and its unconditional distribution is a mixture with mixing probability π. Here, Φ is the cumulative distribution function of a standard Gaussian distribution.

Researchers are often interested in the relationship between sentiment scores and the fundamental information, which in the model is the "truth-plus-noise" Gaussian signal K. We begin by discussing the identifiability of K and its informativeness μ/σ from sentiment data. The fundamental K and parameters μ and σ are not separately identifiable from the data, although two of the three ratios μ/σ, K/μ and K/σ are. To see this, consider Setting 1 with parameters $\mu^{(1)}$ and $\sigma^{(1)}$ and Setting 2 with parameters $\mu^{(2)} = C\mu^{(1)}$ and $\sigma^{(2)} = C\sigma^{(1)}$, where C is a constant. Then a fundamental $K^{(1)} = K$ in Setting 1 is equivalent to the fundamental $K^{(2)} = CK$ in Setting 2, for any K. That is, $K^{(1)}$ and $K^{(2)}$ are the same quantiles in their respective conditional distributions and also have the same sentiment score. Therefore, the conditional and unconditional sentiment distributions in both settings are the same. Essentially, rescaling the fundamental set does not change the problem; what matters is the standardized fundamental K/μ. For the same reason, to identify the fundamental $K = s^{-1}(\rho)$ corresponding to some observed $\rho \in (0,1)$, we must first fix either μ or σ, which can be done without loss of generality.

In the limit described in Proposition 7, the conditional sentiment distributions tend to a mass of one at 0 and 1, respectively, and the unconditional distribution is a highly dispersed mixture.

Regarding parameters related to the directional bias, the prior belief π and the relative payoff relevance u_1/u_0 are separately identifiable from the data, conditional on other parameters, for common utility functions listed in Table I. It is sufficient to show the separate identifiability of π and the directional-bias parameter ψ. We already know that ψ shifts the sentiment-scoring function, thereby changing the sentiment distribution. To see this, fix π and consider two settings that differ only in ψ (equivalently, in u_1/u_0), with respective values $\psi^{(1)}$ and $\psi^{(2)}$ and sentiment-scoring functions $s^{(1)}(K)$ and $s^{(2)}(K)$. Without loss of generality, suppose that $s^{(1)}(K)$ leans more positive. Then we have $s^{(2)} \succeq_{FOSD} s^{(1)}$, which implies $s^{(1)}(K) \geq s^{(2)}(K)$. Consequently, we have $s^{(1)-1}(\rho) \leq s^{(2)-1}(\rho)$ and $\underline{K}^{(1)} \leq \underline{K}^{(2)}$ in Eq. (11), implying that both the conditional and unconditional sentiment distributions under $\psi^{(1)}$ first-order stochastically dominate their counterparts under $\psi^{(2)}$.

To establish the separate identifiability of π from ψ, note that π affects the sentiment distribution in two distinct ways: in addition to its role in determining ψ, which shapes the sentiment-scoring function, π also enters directly as the mixing probability. Holding ψ fixed, different values

of π therefore generate different unconditional sentiment distributions. This distinct second role implies that there is room for π and ψ to be separately identifiable, and therefore so are π and u_1/u_0. In fact, for the common utility functions listed in Table I, variation in ψ changes the shape of the conditional sentiment distributions in a way that induces changes in the unconditional sentiment distribution that cannot be replicated by holding ψ fixed and reweighting π. Hence, ψ and π are separately identifiable.

5.2 Mitigating the risks of misspecification error

Failing to adopt a microfounded approach can lead to model misspecification, which in turn poses risks to the validity of empirical results. This paper argues that researchers should be especially vigilant about such errors when structural analysis is not employed.

Directional bias and sentiment asymmetry. Since Tetlock (2007), it has been a very common observation that positive and negative sentiment articles exhibit different empirical properties. A standard empirical practice is therefore to examine asymmetries across subsamples with different sentiment directions. This model proposes a hypothesis that can explain some of these asymmetries through the lens of directional bias.

According to the model, audience-appeal (or alarmist) bias implies that the effects of confirmatory (or contradictory) shocks may not be fully captured in a misspecified empirical model, potentially generating a spurious asymmetry between positive and negative sentiment environments. To see this, consider the following example. Suppose a researcher classifies texts into discrete levels of positivity, for instance by labeling each text as "positive" or "negative" using an indicator variable $x := \mathbf{1}_{\rho \geq 1/2}$. Such practices are common in empirical work. Assume without loss of generality that positive news is confirmatory. Under audience appeal, a text with sentiment above 50% may correspond to either positive or moderately negative fundamental information. If this indicator variable is used to explain an outcome y whose true data-generating process depends on the fundamental K, the impact of the fundamental shock on y may be understated conditional on $x = 1$. By contrast, the impact may be overstated conditional on $x = 0$, because negative-sentiment texts are associated with more extreme negative fundamental shocks.

To illustrate this point concretely, consider an outcome y with a pre-shock value of 0, which evolves according to the true model

$$y = \begin{cases} (2 \times \mathbf{1}_{K \geq 0} - 1) + \varepsilon & \text{if } K \notin (-d, d), \\ \varepsilon & \text{if } K \in (-d, d), \end{cases}$$

where ε is an independent mean-zero noise. In this setting, the true impact of positive and negative fundamental shocks is symmetric. Consider an audience-appealing sentiment-scoring function that

rates any fundamental value above $-d$ as above $\frac{1}{2}$. Suppose a researcher examines the post-shock impact of y across the positive- and negative-sentiment subsamples. He would find that

$$\mathbb{E}[\,y\,|\,x=1] < \mathbb{E}\,[\,y\,|\,K \geq 0] < 1, \qquad \mathbb{E}[\,y\,|\,x=0] = -1 < \mathbb{E}\,[\,y\,|\,K \leq 0]\,,$$

The researcher therefore observes an apparent asymmetry: negative sentiment appears to be associated with larger and statistically more significant effects than positive sentiment. This asymmetry, however, is purely mechanical; it arises from failing to correct for directional bias as described in this paper. If the researcher goes further and treats x as a direct proxy for the fundamental K without adjusting for audience appeal, he risks drawing spurious conclusions about asymmetric effects of positive and negative fundamental shocks.

This mechanism provides a model-based explanation for observed asymmetries between positive and negative sentiment environments. Without proper corrections for directional bias, researchers may systematically understate the explanatory power of text-based measures of fundamentals in confirmatory sentiment settings. An analogous argument applies in the presence of alarmist bias.

Nondirectional bias and selection of signal strength. On the other hand, sensationalism, if not properly accounted for, may lead to a spurious overestimation of the impact of large shocks and an underestimation of the impact of moderate shocks. Consider a setting in which a researcher uses text-based information to explain an outcome y. Suppose the true data-generating process is $y = f(K) + \varepsilon$, where ε is an independent mean-zero noise. If the researcher mistakenly treats sentiment ρ as a proxy for K, then sensationalism implies that the effect of extreme fundamental shocks on y is overstated relative to that of moderate shocks. This follows because

$$\frac{d\,\mathbb{E}[y \mid \rho]}{d\rho} = \frac{d\,\mathbb{E}[y \mid K]}{dK}\,\frac{dK}{d\rho},$$

where $dK/d\rho$ is the slope of $s^{-1}(\rho)$ which is large for extreme values of ρ and small for moderate ones.

Variation in context. A final source of misspecification arises when changes in the economic context are not properly accounted for, potentially leading to spurious conclusions. The following example provides a concrete illustration. Consider a dataset in which some observations are generated under a high value of u_1/u_0 and others under a low value. Even if the data-generating process for K is identical across the entire sample, the resulting sentiment scores will have different conditional means across the two subsamples generated under different contexts. If a researcher mistakenly treats sentiment as a proxy for K and uses it to explain an outcome variable y that is related to payoff relevance u_1/u_0 but orthogonal to K, she obtains a statistically significant relationship between y and sentiment. Such significance is spurious and may lead the researcher to

draw incorrect conclusions about the effect of fundamental information.

5.3 Using model applicability to discipline text selection and processing

The choice of text data and their processing methods in empirical work involves substantial flexibility and is often guided primarily from statistical considerations. A microfoundation complements this approach by providing an economic perspective.

Microfoundations require researchers to restrict attention to text data that are consistent with the model's assumptions, thereby disciplining the choice of data. In practice, media outlets face many considerations when forming content, and any given model can capture only a subset of them underlying content formation. These considerations may include not only information selection, but also persuasion, media competition, and other mechanisms. Adopting the model in this paper entails focusing on text for which information selection rather than other channels is the primary consideration, and for which the model's remaining assumptions are reasonably well satisfied.

Microfoundations may also guide data processing from an economic perspective. One example is how this model may discipline the tokenization of text in an interpretable way. Gentzkow, Kelly, and Taddy (2019) summarize textual analysis as proceeding in three steps: (i) quantifying text into tokens, (ii) using the tokenized text to construct a measure of an economic outcome, and (iii) using this measure for further analysis. Tokenization refers to dividing text into smaller components, such as words, phrases, sentences, or stories and is a critical step in sentiment analysis. In the model, each token corresponds to a reported element. Applying the model requires that each reported element be interpretable as positive or negative evidence about two alternative states, with ordering largely uninformative. Therefore, tokenizing text into words is not ideal for the model because word order conveys information. By contrast, the model is more naturally compatible with tokenizations that decompose a newspaper page into economic news stories, an intelligence briefing into discrete events, or an article into individual stories or evidence units.

Miscellanies on model applicability. This paper does not attempt to provide a complete theory of textual sentiment. Instead, it develops a structural framework that clarifies the economic information embedded in sentiment measures through the lens of information selection and illustrates how microfounded models can discipline empirical practice.

The information selection channel highlighted in this paper operates independently of discrepancies in agents' preferences or beliefs. When such discrepancies are present, additional persuasion motives may arise, potentially leading to mixed reporting strategies and trade-offs between honesty, self-consistency, and optimality. Nevertheless, as long as information compression and selection remain relevant, the forces identified in this paper continue to shape observed sentiment and reporting behavior.

The framework extends beyond the media setting. First, although the analysis is framed around news reporting, the underlying mechanism applies more broadly to information intermediaries, such as government agencies, firms, or advisors, that selectively transmit information to facilitate receivers' decisions under communication constraints. Second, the notion of a sentiment score admits a more general interpretation. In the model, the score is a statistic that aggregates reported evidence to discriminate between competing hypotheses, e.g., boom versus bust, certainty versus uncertainty, or political left versus right. In this sense, sentiment is not simply "positivity," but a score used for differentiating states for making a decision. Third, abstracting from the information-selection interpretation, the model can be viewed as a general communication problem with a sparse symbol set. Whenever rich underlying states must be compressed into a limited number of labels, such as text scores, star ratings, or categorical evaluations, the same logic applies and bias may emerge.

6 Concluding Remarks

This paper provides, to the best of my knowledge, the first microfoundation of textual sentiment analysis. I develop a model of content formation in which a sender selectively reports information to facilitate a receiver's decision-making. The model generates familiar media bias patterns as an implication of endogenous information selection, independently of standard persuasion or media competition mechanisms. Within this framework, I derive an asymptotically tractable sentiment-scoring function that links observed textual sentiment to an underlying fundamental signal and to contextual economic parameters governing content production. The analysis clarifies how text sentiment should be dissected and highlights implications for empirical specification.

The core mechanism is information compression under communication constraints. When complex information must be pooled into a sparse symbol set for reporting, optimal pooling depends on the perceived newsworthiness of realizations for the receiver's decision problem. In this setting, the sentiment-scoring function emerges as an antiderivative of newsworthiness. Information that is more newsworthy (because it has higher baseline importance or greater action sensitivity) induces greater local sensitivity in sentiment labeling. Variation in this sensitivity generates labeling patterns that appear as bias. The same mechanism applies beyond text, including ratings and exam scores.

References

Antweiler, Werner, and Murray Z Frank. 2004. "Is all that talk just noise? The information content of internet stock message boards." *The Journal of Finance* 59 (3): 1259–1294.

Bergemann, Dirk, and Stephen Morris. 2019. "Information design: A unified perspective." *Journal of Economic Literature* 57 (1): 44–95.

Bloedel, Alex, and Ilya Segal. 2021. "Persuading a rationally inattentive agent." *Working Paper.*

Che, Yeon-Koo, and Konrad Mierendorff. 2019. "Optimal dynamic allocation of attention." *American Economic Review* 109 (8): 2993–3029.

Dye, Ronald A. 1985. "Disclosure of nonproprietary information." *Journal of Accounting Research,* 123–145.

Gentzkow, Matthew, and Emir Kamenica. 2014. "Costly persuasion." *American Economic Review* 104 (5): 457–62.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as data." *Journal of Economic Literature* 57 (3): 535–74.

Gentzkow, Matthew, and Jesse M Shapiro. 2006. "Media bias and reputation." *Journal of Political Economy* 114 (2): 280–316.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101 (6): 2590–2615.

Ke, Zheng Tracy, Bryan T Kelly, and Dacheng Xiu. 2019. "Predicting returns with text data." *Working Paper.*

Korinek, Anton. 2023. "Generative AI for economic research: Use cases and implications for economists." *Journal of Economic Literature* 61 (4): 1281–1317.

Loughran, Tim, and Bill McDonald. 2011. "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks." *The Journal of Finance* 66 (1): 35–65.

Milgrom, Paul, and John Roberts. 1986. "Relying on the information of interested parties." *The RAND Journal of Economics,* 18–32.

Milgrom, Paul R. 1981. "Good news and bad news: Representation theorems and applications." *The Bell Journal of Economics,* 380–391.

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The market for news." *American Economic Review* 95 (4): 1031–1053.

Perego, Jacopo, and Sevgi Yuksel. 2022. "Media competition and social disagreement." *Econometrica* 90 (1): 223–265.

Scheffé, Henry. 1947. "A useful convergence theorem for probability distributions." *The Annals of Mathematical Statistics* 18 (3): 434–438.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics* 50 (3): 665–690.

Suen, Wing. 2004. "The self-perpetuation of biased beliefs." *The Economic Journal* 114 (495): 377–396.

Tetlock, Paul C. 2007. “Giving content to investor sentiment: The role of media in the stock market.” *The Journal of Finance* 62 (3): 1139–1168.

Tetlock, Paul C, Maytal Saar-Tsechansky, and Sofus Macskassy. 2008. “More than words: Quantifying language to measure firms’ fundamentals.” *The Journal of Finance* 63 (3): 1437–1467.

Supplementary Appendix for "Information Selection and Media Bias: An Economic Foundation for Sentiment Analysis"

Jinge Liu

A Propositions 1, 2, and 3

Let $p_K := \Pr(K|\theta = 1)$, $q_K := \Pr(K|\theta = 0)$. The posterior beliefs are $\pi'_k := \Pr(\theta = 1|k) = \pi \Pr(k|\theta = 1)/\Pr(k)$. Hence the ex ante utility is

$$\begin{aligned} U &= \sum_{k=0}^{n} \Pr(k) \left\{\pi'_k u_1 h(a^*(k)) + (1 - \pi'_k) u_0 h(1 - a^*(k))\right\} \\ &= \sum_{k=0}^{n} \left\{\pi \Pr(k|\theta = 1) u_1 h(a^*(k)) + (1 - \pi) \Pr(k|\theta = 0) u_0 h(1 - a^*(k))\right\} := \sum_{k=0}^{n} U_k, \end{aligned}$$

where $a^*(k)$ is $\arg\max_a \pi'_k u_1 h(a) + (1 - \pi'_k) u_0 h(1 - a)$ if $\sum_K \sigma_{Kk} > 0$, and may be any value otherwise. The domain of U is $[0, 1]^{(N+1)\times(n+1)}$ for $\{\sigma_{Kk}\}_{K=0,\ldots,N;k=0,\ldots,n}$. The utility is continuous on a compact set, so there always exists a solution to maximizing this utility.

Proofs of Propositions 1 and 2. We will show that Proposition 1′, below, implies Proposition 2, which in turn implies Proposition 1.

Proposition 1′. *Under Assumption 1, there exists a pure-strategy equilibrium.*

Proof of Proposition 1′. For σ_{Kk},

$$\begin{aligned} \frac{\partial U}{\partial \sigma_{Kk}} = \frac{\partial U_k}{\partial \sigma_{Kk}} &= \frac{\partial U_k(\{\sigma_{Kk}\}_{K=0,\ldots,N}, a^*(k))}{\partial \sigma_{Kk}} + \frac{\partial U_k(\{\sigma_{Kk}\}_{K=0,\ldots,N}, a^*(k))}{\partial a^*(k)} \frac{\partial a^*(k)}{\partial \sigma_{Kk}} \\ &= \frac{\partial U_k(\{\sigma_{Kk}\}_{K=0,\ldots,N}, a^*(k))}{\partial \sigma_{Kk}} \qquad \text{(by optimality of } a^*(k)\text{)} \\ &= \pi u_1 p_K h(a^*(k)) + (1 - \pi) u_0 q_K h(1 - a^*(k)). \end{aligned} \tag{A.1}$$

The optimality of $a^*(k)$ for k such that $\sum_K \sigma_{Kk} > 0$ is characterized by the first-order condition $0 = \pi'_k u_1 h'(a) - (1-\pi'_k) u_0 h'(1-a)$. By the implicit function theorem,

$$\frac{\partial a^*(k)}{\partial \sigma_{Kk}} = -\frac{\pi u_1 p_K h'(a^*(k)) - (1-\pi) u_0 q_K h'(1-a^*(k))}{\pi u_1 (\sum_{K'=1}^N p_{K'} \sigma_{K'k}) h''(a^*(k)) + (1-\pi) u_0 (\sum_{K'=1}^N q_{K'} \sigma_{K'k}) h''(1-a^*(k))}.$$

Hence,

$$\begin{aligned} \frac{\partial^2 U}{\partial \sigma_{Kk}^2} &= (\pi u_1 p_K h'(a^*(k)) - (1-\pi) u_0 q_K h'(1-a^*(k))) \frac{\partial a^*(k)}{\partial \sigma_{Kk}} \qquad \text{(A.2)} \\ &= -\frac{(\pi u_1 p_K h'(a^*(k)) - (1-\pi) u_0 q_K h'(1-a^*(k)))^2}{\pi u_1 (\sum_{K'=1}^N p_{K'} \sigma_{K'k}) h''(a^*(k)) + (1-\pi) u_0 (\sum_{K'=1}^N q_{K'} \sigma_{K'k}) h''(1-a^*(k))}, \end{aligned}$$

which is ≥ 0 since $h''(\cdot) < 0$. Also,

$$\frac{\partial^2 U}{\partial \sigma_{Kk_1} \partial \sigma_{Kk_2}} = 0, \qquad \text{(A.3)}$$

implying that $\partial U / \partial \sigma_{Kk}$ does not depend on $\sigma_{Kk'}$ for $k' \neq k$.

Suppose $\sigma^*_{Kk_1} \in (0,1)$ is in an optimal information structure. Then there must exist $\sigma^*_{Kk_2} \in (0,1)$. Both $\sum_K \sigma^*_{Kk_1} > 0$ and $\sum_K \sigma^*_{Kk_2} > 0$ hold, so Eqs. (A.2) and (A.3) hold. Writing U for $U(\sigma_{Kk_1}, \sigma_{Kk_2})$, we consider the following cases.

Case (i): $\frac{\partial U}{\partial \sigma_{Kk_1}}|_{\sigma^*_{Kk_1}} \neq \frac{\partial U}{\partial \sigma_{Kk_2}}|_{\sigma^*_{Kk_2}}$ (without loss of generality, we may assume $\frac{\partial U}{\partial \sigma_{Kk_1}}|_{\sigma^*_{Kk_1}} > \frac{\partial U}{\partial \sigma_{Kk_2}}|_{\sigma^*_{Kk_2}}$). Then for a small ε, $\sigma^{**}_{Kk_1} = \sigma^*_{Kk_1} + \varepsilon$ and $\sigma^{**}_{Kk_2} = \sigma^*_{Kk_2} - \varepsilon$ will improve U. This is because $\frac{\partial U}{\partial \sigma_{Kk_1}}$ is continuous in σ_{Kk_1}, as is $\frac{\partial U}{\partial \sigma_{Kk_2}}$ in σ_{Kk_2}; thus there exists $\varepsilon > 0$ such that $\frac{\partial U}{\partial \sigma_{Kk_1}}|_{\sigma^*_{Kk_1}+\varepsilon'} > \frac{\partial U}{\partial \sigma_{Kk_2}}|_{\sigma^*_{Kk_2}-\varepsilon'}$ for all $\varepsilon' \in (0, \varepsilon]$. The utility increase from changing to $\sigma^{**}_{Kk_1}$ and $\sigma^{**}_{Kk_2}$ is $U(\sigma^{**}_{Kk_1}, \sigma^{**}_{Kk_2}) - U(\sigma^*_{Kk_1}, \sigma^*_{Kk_2}) = U(\sigma^{**}_{Kk_1}, \sigma^{**}_{Kk_2}) - U(\sigma^{**}_{Kk_1}, \sigma^*_{Kk_2}) + U(\sigma^{**}_{Kk_1}, \sigma^*_{Kk_2}) - U(\sigma^*_{Kk_1}, \sigma^*_{Kk_2})$, which equals $\int_0^\varepsilon \frac{\partial U}{\partial \sigma_{Kk_1}}|_{\sigma^*_{Kk_1}+\varepsilon'} d\varepsilon' - \int_0^\varepsilon \frac{\partial U}{\partial \sigma_{Kk_2}}|_{\sigma^*_{Kk_2}-\varepsilon''} d\varepsilon'' > 0$ by Eq. (A.3). This contradicts optimality.

Case (ii): $\frac{\partial U}{\partial \sigma_{Kk_1}}|_{\sigma^*_{Kk_1}} = \frac{\partial U}{\partial \sigma_{Kk_2}}|_{\sigma^*_{Kk_2}}$, and there is no $\delta > 0$ such that both $\frac{\partial^2 U}{\partial \sigma^2_{Kk_1}} = 0$ on $(\sigma^*_{Kk_1} - \delta, \sigma^*_{Kk_1} + \delta)$ and $\frac{\partial^2 U}{\partial \sigma^2_{Kk_2}} = 0$ on $(\sigma^*_{Kk_2} - \delta, \sigma^*_{Kk_2} + \delta)$; assume without loss that $\frac{\partial^2 U}{\partial \sigma^2_{Kk_1}} > 0$ on $(\sigma^*_{Kk_1}, \sigma^*_{Kk_1} + \delta_0]$. Then $\sigma^{**}_{Kk_1} = \sigma^*_{Kk_1} + \delta_0$ and $\sigma^{**}_{Kk_2} = \sigma^*_{Kk_2} - \delta_0$ will improve U. Here is why: The utility increase is $U(\sigma^{**}_{Kk_1}, \sigma^{**}_{Kk_2}) - U(\sigma^*_{Kk_1}, \sigma^*_{Kk_2}) = U(\sigma^{**}_{Kk_1}, \sigma^{**}_{Kk_2}) - U(\sigma^{**}_{Kk_1}, \sigma^*_{Kk_2}) +$

$U(\sigma^{**}_{Kk_1}, \sigma^{*}_{Kk_2}) - U(\sigma^{*}_{Kk_1}, \sigma^{*}_{Kk_2})$, which equals $U(\sigma^{*}_{Kk_1}, \sigma^{**}_{Kk_2}) - U(\sigma^{*}_{Kk_1}, \sigma^{*}_{Kk_2}) + U(\sigma^{**}_{Kk_1}, \sigma^{*}_{Kk_2}) - U(\sigma^{*}_{Kk_1}, \sigma^{*}_{Kk_2})$ by Eq. (A.3). Define a continuous function f by $f(x) := U(\sigma^{*}_{Kk_1}, \sigma^{*}_{Kk_2} + x)$ if $x \in [-\delta_0, 0]$ and $f(x) := U(\sigma^{*}_{Kk_1} + x, \sigma^{*}_{Kk_2})$ if $x \in (0, \delta_0]$; then f is convex on $[-\delta_0, \delta_0]$ and strictly convex on $(0, \delta_0]$. The utility increase is $f(\delta_0) + f(-\delta_0) - 2f(0)$, which is positive by Jensen's inequality, contradicting optimality.

Case (iii): $\frac{\partial U}{\partial \sigma_{Kk_1}}|_{\sigma^{*}_{Kk_1}} = \frac{\partial U}{\partial \sigma_{Kk_2}}|_{\sigma^{*}_{Kk_2}}$, and there exists $\delta > 0$ such that both $\frac{\partial^2 U}{\partial \sigma^2_{Kk_1}} = 0$ on $(\sigma^{*}_{Kk_1} - \delta, \sigma^{*}_{Kk_1} + \delta)$ and $\frac{\partial^2 U}{\partial \sigma^2_{Kk_2}} = 0$ on $(\sigma^{*}_{Kk_2} - \delta, \sigma^{*}_{Kk_2} + \delta)$. Then one of the following two cases will occur.

Case (iii-a): $\frac{\partial^2 U}{\partial \sigma^2_{Kk_1}} > 0$ or $\frac{\partial^2 U}{\partial \sigma^2_{Kk_2}} > 0$ somewhere on $(0, \sigma^{*}_{Kk_1} + \sigma^{*}_{Kk_2})$; assume without loss that for some constant $\sigma^0 \in (\sigma^{*}_{Kk_1}, \sigma^{*}_{Kk_1} + \sigma^{*}_{Kk_2})$, $\frac{\partial^2 U}{\partial \sigma^2_{Kk_1}} > 0$ when $\sigma_{Kk_1} \in (\sigma^0, \sigma^{*}_{Kk_1} + \sigma^{*}_{Kk_2})$, $\frac{\partial^2 U}{\partial \sigma^2_{Kk_1}} = 0$ when $\sigma_{Kk_1} \in (\sigma^{*}_{Kk_1}, \sigma^0]$, and $\frac{\partial^2 U}{\partial \sigma^2_{Kk_2}} = 0$ when $\sigma_{Kk_2} \in (\sigma^{*}_{Kk_2} - (\sigma^0 - \sigma^{*}_{Kk_1}), \sigma^{*}_{Kk_2})$. Letting $\sigma^0_{Kk_1} = \sigma^0$ and $\sigma^0_{Kk_2} = \sigma^{*}_{Kk_2} - (\sigma^0 - \sigma^{*}_{Kk_1})$, we get a new information structure falling under case (ii), with the same utility. Applying the reasoning of case (ii) by taking $\sigma^0_{Kk_1}$ and $\sigma^0_{Kk_2}$ as the new $\sigma^{*}_{Kk_1}$ and $\sigma^{*}_{Kk_2}$, we get a contradiction to optimality.

Case (iii-b): $\frac{\partial^2 U}{\partial \sigma^2_{Kk_1}} = 0$ and $\frac{\partial^2 U}{\partial \sigma^2_{Kk_2}} = 0$ for $\sigma_{Kk_1}, \sigma_{Kk_2} \in (0, \sigma^{*}_{Kk_1} + \sigma^{*}_{Kk_2})$. Let $\sigma^{**}_{Kk_1} = \sigma^{*}_{Kk_1} + \sigma^{*}_{Kk_2}$ and $\sigma^{**}_{Kk_2} = 0$.

Case (iii-b-1): $\sigma^{*}_{Kk_1} + \sigma^{*}_{Kk_2} = 1$. In this case, the new strategy has the same utility with no mixing for K. If there is no mixing for any other fundamental K', then we have found a pure strategy that delivers the same utility as the optimal mixed strategy. If some other K' has mixing, then let K' be the new K and iterate the discussion of cases (i)–(iii).

Case (iii-b-2): $\sigma_{Kk^{*}_1} + \sigma_{Kk^{*}_2} < 1$. Then there exists k_3 such that $\sigma^{*}_{Kk_3} \in (0, 1)$. Let k_3 be the new k_2 and iterate the discussion of cases (i)–(iii).

Hence the only case in which the existence of mixing is plausible is (iii-b-1). An equilibrium mixed strategy exists only when it delivers the same utility as a pure strategy. Therefore, a pure-strategy equilibrium always exists. □

Lemma 1. *Under Assumptions 1 and 2, the equilibrium involves either a pure strategy, or a mixed strategy that yields the same utility as a pure strategy in which* $\sum_K \sigma_{Kk} = 0$ *for some* k.

Proof of Lemma 1. From the proof of Proposition 1′, mixing can occur only in case (iii-b-1), when $\sigma^*_{Kk_1} + \sigma^*_{Kk_2} = 1$. We restrict our attention to the case where mixing only happens to K and not to any other $K' \neq K$, because we can always do the iteration described in (iii-b-1), which can be completed in finite steps, and focus on the last step. In that case, $\partial U / \partial \sigma_{Kk_1}$ is constant, so $a^*(k_1)$ is constant with respect to σ_{Kk_1} by Eq. (A.1). Hence, by the first-order condition, σ_{Kk_1} does not affect π'_{k_1}. Analogously, σ_{Kk_2} does not affect π'_{k_2}. This implies $\pi'_{k_1} = \pi'_{k_2}$. Therefore, this information structure yields the same utility as an information structure given by $\sigma^{**}_{Kk_1} = 1$, $\sigma^{**}_{K'k_1} = \sigma^*_{K'k_1} + \sigma^*_{K'k_2}$, and $\sigma^{**}_{K'k_2} = 0$ for every $K' \neq K$. □

Proof of Proposition 2. By Proposition 1′ and Lemma 1, it suffices to consider pure-strategy equilibria. The proof is by contradiction. Suppose, in a pure-strategy equilibrium, $\sum_K \sigma_{Kk_1} = 0$. Then $N+1$ fundamentals are mapped to at most $(n+1)-1 = n \leq N$ reports, so some report is associated with $m \geq 2$ fundamentals. Let k_2 be such a report, representing the set of fundamentals $\{K^{(1)}, ..., K^{(m)}\}$. The contribution of k_2 to the utility is U_{k_2}, and the optimal action for k_2 is $a^*(k_2)$. Consider an alternative pure strategy in which $K^{(1)}$ is mapped to k_1, $\{K^{(2)}, ..., K^{(m)}\}$ is mapped to k_2, and the rest of the strategy is unchanged. Let U'_{k_1} and U'_{k_2} denote the contributions of k_1 and k_2 to the utility, and $a'^*(k_1)$ and $a'^*(k_2)$ the optimal actions. Then

$$
\begin{aligned}
U_{k_2} &= \Pr\Big(K \in \{K^{(1)}, ..., K^{(m)}\}\Big) E[u(a^*(k_2); \theta) | K \in \{K^{(1)}, ..., K^{(m)}\}] \\
&= \Pr\Big(K \in \{K^{(1)}, ..., K^{(m)}\}\Big) E[E[u(a^*(k_2); \theta) | \tilde{1}_{K=K^{(1)}}] | K \in \{K^{(1)}, ..., K^{(m)}\}] \\
&< \Pr\Big(K \in \{K^{(1)}, ..., K^{(m)}\}\Big) E[\max_a E[u(a; \theta) | \tilde{1}_{K=K^{(1)}}] | K \in \{K^{(1)}, ..., K^{(m)}\}] \\
&= \Pr\Big(K = K^{(1)}\Big) E[u(a'^*(k_1); \theta) | K = K^{(1)}; K \in \{K^{(1)}, ..., K^{(m)}\}] + \\
&... + \Pr\Big(K \in \{K^{(2)}, ..., K^{(m)}\}\Big) E[u(a'^*(k_2); \theta) | K \neq K^{(1)}, K \in \{K^{(1)}, ..., K^{(m)}\}] \\
&= U'_{k_1} + U'_{k_2},
\end{aligned}
$$

where $\tilde{1}_{K=K^{(1)}}$ is a random variable that equals 1 when $K = K^{(1)}$ and 0 otherwise. The inequality is strict because $a^*(k_2)$, $a'^*(k_1)$ and $a'^*(k_2)$ cannot coincide, since p_K / q_K is strictly monotone. This

contradicts optimality. □

Proof of Proposition 1. Proposition 2 rules out the mixed-strategy case in Lemma 1. □

Proof of Proposition 3. Let $x_K = \pi u_1 p_K$ and $y_K = (1-\pi)u_0 q_K$. Each K is fully characterized by (x_K, y_K). Let $\eta_K := \frac{x_K}{y_K} = \frac{\pi}{1-\pi}\frac{u_1}{u_0}\Lambda(K)$.

Lemma 2. *Let Assumption 1 hold, let the fundamental space be $\{K_0, ..., K_N\}$ with all η_{K_i} different, denote the optimal partition by $\{B_k\}_{k=0}^n$, and suppose B_n consists of $m \geq 2$ fundamentals (by Proposition 2, $m \leq N - n + 1$). Define the fundamental v by $x_v = \sum_{K\in B_n} x_K$ and $y_v = \sum_{K\in B_n} y_K$. Then, for the alternative problem with fundamental space $\{K_0, ..., K_N\} \cup \{v\}\backslash B_n$ and $n+1$ reports, the optimal partition is $\{B_k\}_{k=0}^{n-1} \cup \{\{v\}\}$.*

Proof of Lemma 2. Under the partition $\{B_k\}_{k=0}^n$, the utility is

$$U = \sum_{k=0}^{n} \left\{ \left(\sum_{K\in B_k} x_K \right) h(a^*(k)) + \left(\sum_{K\in B_k} y_K \right) h(1 - a^*(k)) \right\},$$

where $a^*(k)$ is determined by the conditional probabilities of fundamentals through $\sum_{K\in B_k} x_K / \sum_{K\in B_k} y_K$. Therefore, any utility delivered by a feasible partition of $\{K_0, ..., K_N\} \cup \{v\}\backslash B_n$ can be delivered by a feasible partition of $\{K_0, ..., K_N\}$, specifically by replacing v, in the partition set of $\{K_0, ..., K_N\} \cup \{v\}\backslash B_n$ containing v, by all of the elements of B_n. That is, the utilities obtained by partitioning $\{K_0, ..., K_N\} \cup \{v\}\backslash B_n$ form a subset of those obtained by partitioning $\{K_0, ..., K_N\}$. Therefore, if $\{B_k\}_{k=0}^n$ is optimal, then $\{B_k\}_{k=0}^{n-1} \cup \{\{v\}\}$, as a partition of $\{K_0, ..., K_N\} \cup \{v\}\backslash B_n$ that delivers the same utility, must be optimal. □

Proof of Proposition 3. The proof is by induction. Let $P(N, n)$ be the problem of choosing n reported elements from N signals. We want to show that the solution of $P(N, n)$ involves an ordered partition of the fundamental space. We first induct from the base cases $P(2, 1)$ and $P(3, 1)$ to establish the assertion for $P(N, 1)$, then deduce it for $P(N, n)$.

Step 1. We prove that the solution to $P(N,1)$ involves an ordered partition.

Step 1.1. Here we prove that the solution to $P(2,1)$ involves an ordered partition.

The proof is by contradiction. Consider the following strategy: (1) The fundamentals $0,1,2$ map to k_0 with probabilities $1-s, 0, 1-t$ respectively; the optimal action is a. (2) The fundamentals $0,1,2$ map to k_1 with probabilities $s,1,t$ respectively; the optimal action is b. Thus,

$$U = ((1-s)x_0 + (1-t)x_2)h(a) + ((1-s)y_0 + (1-t)y_2)h(1-a) + ... \\ + (sx_0 + x_1 + tx_2)h(b) + (sy_0 + y_1 + ty_2)h(1-b).$$

We need to show that $s = t = 0$ is a suboptimal strategy. By the envelope theorem,

$$\frac{\partial U}{\partial s} = x_0(h(b) - h(a)) + y_0(h(1-b) - h(1-a)), \tag{A.4}$$

$$\frac{\partial U}{\partial t} = x_2(h(b) - h(a)) + y_2(h(1-b) - h(1-a)). \tag{A.5}$$

We consider two cases.

Case (i): $a \neq b$ at $s = t = 0$. Then as long as one of the two partial derivatives is positive at $s = 0$ or $t = 0$—say, $\frac{\partial U}{\partial s}|_{s=0} > 0$—the pure strategy at $s = t = 0$ is strictly worse than a strategy with a small positive s. But the latter is a mixed strategy and thus suboptimal by Proposition 1. Therefore we only need to show that Eq. (A.4) or Eq. (A.5) is positive. We do so by contradiction. Suppose both Eq. (A.4) and Eq. (A.5) are nonpositive at $s = t = 0$. Then $\eta_1 = \frac{h'(1-b)}{h'(b)}$ by the optimality of b. On the other hand, if $a > b$, then

$$\eta_0 \geq \frac{h(1-b) - h(1-a)}{h(a) - h(b)} = \frac{(h(1-b) - h(1-a))/((1-b) - (1-a))}{(h(a) - h(b))/(a-b)}$$

by Eq. (A.4) being nonpositive, so $\eta_1 > \eta_0 \geq \frac{(h(1-b)-h(1-a))/((1-b)-(1-a))}{(h(a)-h(b))/(a-b)}$. Since $h(\cdot)$ is strictly increasing and concave, we have $\frac{h(a)-h(b)}{a-b} < h'(b)$ and $\frac{h(1-b)-h(1-a)}{(1-b)-(1-a)} > h'(1-b)$; hence $\eta_1 > \frac{h'(1-b)}{h'(b)}$,

a contradiction. And if $a < b$, then

$$\eta_2 \leq \frac{h(1-b)-h(1-a)}{h(a)-h(b)} = \frac{(h(1-b)-h(1-a))/((1-b)-(1-a))}{(h(a)-h(b))/(a-b)}$$

by Eq. (A.5) being nonpositive, and hence, by analogous arguments, $\eta_1 < \frac{h'(1-b)}{h'(b)}$, a contradiction.

Case (ii): $a = b$ at $s = t = 0$. Then $\frac{\partial U}{\partial s}|_{s=0} = \frac{\partial U}{\partial t}|_{t=0} = 0$. The utility is $U|_{s=t=0} = (x_0 + x_1 + x_2)h(a) + (y_0 + y_1 + y_2)h(1-a)$. Consider the strategy at $s = t = 1$. By the optimality of a and b, $\frac{y_0+y_2}{x_0+x_2} = \frac{h'(a)}{h'(1-a)} = \frac{h'(b)}{h'(1-b)} = \frac{y_1}{x_1}$, so $\frac{h'(a)}{h'(1-a)} = \frac{h'(b)}{h'(1-b)} = \frac{y_0+y_1+y_2}{x_0+x_1+x_2}$. Therefore, for $s = t = 1$, the optimal action is also a. The utility satisfies $U|_{s=t=1} = U|_{s=t=0}$. However, the strategy at $s = t = 1$ does not use both reports and thus is suboptimal by Proposition 2. Hence the strategy at $s = t = 0$ is also suboptimal. This completes Step 1.1.

Step 1.2. Prove that the solution to $P(3, 1)$ involves an ordered partition.

The proof is by contradiction. The possible non-ordered partitions $\{B_k\}_{k=0}^1$ are $\{\{1\}, \{0, 2, 3\}\}$, $\{\{2\}, \{0, 1, 3\}\}$, $\{\{0, 3\}, \{1, 2\}\}$, and $\{\{0, 2\}, \{1, 3\}\}$. Below we examine each.

Case (i): $\{\{1\}, \{0, 2, 3\}\}$ or $\{\{2\}, \{0, 1, 3\}\}$ is optimal. Suppose without loss of generality that $\{\{1\}, \{0, 2, 3\}\}$ is optimal. Consider another problem with fundamentals $\{0, 1, v\}$, where v is defined by $x_v = x_2 + x_3$ and $y_v = y_2 + y_3$. By Lemma 2, the optimal partition must be $\{\{1\}, \{0, v\}\}$. However, the problem is $P(2, 1)$, and since $\eta_0 < \eta_1 < \eta_v$, by Step 1.1, $\{\{1\}, \{0, v\}\}$ is suboptimal because it is not ordered, a contradiction.

Case (ii): $\{\{0, 3\}, \{1, 2\}\}$ is optimal. Consider another problem with fundamentals $\{0, v, 3\}$ where v is defined by $x_v = x_1 + x_2$ and $y_v = y_1 + y_2$. By Lemma 2, the optimal partition must be $\{\{v\}, \{0, 3\}\}$. However, the problem is $P(2, 1)$, and since $\eta_0 < \eta_v < \eta_3$, by Step 1.1, $\{\{v\}, \{0, 3\}\}$ is suboptimal because it is not ordered, a contradiction.

Case (iii): $\{\{0, 2\}, \{1, 3\}\}$ is optimal. We consider three cases.

Case (iii-a): $a > b$. Consider the following strategy: (1) The fundamentals $0, 1, 2, 3$ map to k_0 with probabilities $1 - t, 0, 1, s$ respectively; the optimal action is a. (2) The fundamentals $0, 1, 2, 3$

map to k_1 with probabilities $t, 1, 0, 1-s$ respectively; the optimal action is b. Then

$$U = ((1-t)x_0 + x_2 + sx_3)h(a) + ((1-t)y_0 + y_2 + sy_3)h(1-a) + ...$$
$$+ (tx_0 + x_1 + (1-s)x_3)h(b) + (ty_0 + y_1 + (1-s)y_3)h(1-b).$$

We need to show that $s = t = 0$ is a suboptimal strategy. By the envelope theorem,

$$\frac{\partial U}{\partial s} = x_3(h(a) - h(b)) + y_3(h(1-a) - h(1-b)),$$
$$\frac{\partial U}{\partial t} = -x_0(h(a) - h(b)) - y_0(h(1-a) - h(1-b)).$$

As in Step 1.1, it suffices to show that one of the two partial derivatives is positive at $s = 0$ or $t = 0$. Suppose both are nonpositive; then $\eta_3 \leq \frac{h(1-b)-h(1-a)}{h(a)-h(b)} \leq \eta_0$, contradicting the assumption that $\eta_3 > \eta_0$.

Case (iii-b): $a < b$. Consider the following strategy: (1) The fundamentals $0, 1, 2, 3$ map to k_0 with probabilities $1, s, 1-t, 0$ respectively; the optimal action is a. (2) The fundamentals $0, 1, 2, 3$ map to k_1 with probabilities $0, 1-s, t, 1$ respectively; the optimal action is b. Then

$$U = (x_0 + sx_1 + (1-t)x_2)h(a) + (y_0 + sy_1 + (1-t)y_2)h(1-a) + ...$$
$$+ ((1-s)x_1 + tx_2 + x_3)h(b) + ((1-s)y_1 + ty_2 + y_3)h(1-b).$$

We need to show that $s = t = 0$ is a suboptimal strategy. By the envelope theorem,

$$\frac{\partial U}{\partial s} = x_1(h(a) - h(b)) + y_1(h(1-a) - h(1-b)),$$
$$\frac{\partial U}{\partial t} = -x_2(h(a) - h(b)) - y_2(h(1-a) - h(1-b)).$$

As in Step 1.1, it now suffices to show that one of the two partial derivatives is positive at $s = 0$ or $t = 0$. Suppose both are nonpositive; then $\eta_2 \leq \frac{h(1-b)-h(1-a)}{h(a)-h(b)} \leq \eta_1$, contradicting the assumption that

$\eta_2 > \eta_1$.

Case (iii-c): $a = b$. Consider the same strategy as (iii-a). Then $\frac{\partial U}{\partial s}|_{s=0} = \frac{\partial U}{\partial t}|_{t=0} = 0$. The utility is $U|_{s=t=0} = (x_0 + x_1 + x_2 + x_3)h(a) + (y_0 + y_1 + y_2 + y_3)h(1-a)$. By the optimality of a and b, $\frac{y_0+y_2}{x_0+x_2} = \frac{h'(a)}{h'(1-a)} = \frac{h'(b)}{h'(1-b)} = \frac{y_1+y_3}{x_1+x_3}$, and hence $\frac{h'(a)}{h'(1-a)} = \frac{h'(b)}{h'(1-b)} = \frac{y_0+y_1+y_2+y_3}{x_0+x_1+x_2+x_3}$. Therefore, the alternative strategy with the partition $\{\{0,1,2,3\}, \varnothing\}$ yields the same optimal actions and utility as $s = t = 0$. However, it does not use both reports, so by Proposition 2 it is suboptimal. Hence the strategy at $s = t = 0$ is also suboptimal. This completes Step 1.2.

Step 1.3. Given that the solution to $P(N-1, 1)$ involves an ordered partition, prove that the solution of $P(N, 1)$ involves an ordered partition (for $N \geq 4$).

The proof is by contradiction. Suppose the optimal partition $\{B_0, B_1\}$ for $P(N, 1)$ is not ordered. We consider two cases.

Case (i): There exist neighboring fundamentals $i, i+1 \in B_0$ (or B_1; here we pick B_0 without loss of generality). In this case, let v be a fundamental and $(x_v, y_v) = (x_i + x_{i+1}, y_i + y_{i+1})$. Then $\Lambda(i) < \Lambda(v) < \Lambda(i+1)$. Consider the $P(N-1, 1)$ problem of partitioning $\{0, 1, ..., i-1, v, i+2, ..., N\}$. Since $\{B_0, B_1\}$ is optimal, by Lemma 2 the solution has to be $\{B_0', B_1'\}$ where $B_0' = B_0 \cup \{v\} \backslash \{i, i+1\}$ and $B_1' = B_1$. However, this is not an ordered partition, a contradiction.

Case (ii): There are no patterns as in case (i), i.e. $B_0 = \{0, 2, 4, ...\}$ and $B_1 = \{1, 3, 5, ...\}$. Let v be a fundamental and $(x_v, y_v) = (x_0 + x_2, y_0 + y_2)$. Then $\Lambda(v) < \Lambda(2)$. Consider the $P(N-1, 1)$ problem of partitioning $\{v, 1, 3, 4, ..., N\}$. Since $\{B_0, B_1\}$ is optimal, by Lemma 2 the solution has to be $\{B_0', B_1'\}$ with $B_0' = B_0 \cup \{v\} \backslash \{0, 2\}$ and $B_1' = B_1$. However, $\Lambda(v) < \Lambda(3) < \Lambda(4)$, so this is not an ordered partition, a contradiction.

By Steps 1.1, 1.2, and 1.3, $P(N, 1)$ has ordered-partition solutions.

Step 2. Prove that the solution to $P(N, n)$ involves an ordered partition.

The problem is to partition $\{0, 1, ..., N\}$ into $n+1$ sets. Denote the solution by $\{B_0, ..., B_n\}$. For any $0 \leq i, j \leq n$ $(i \neq j)$, set $B_i = \{K_1^i, K_2^i, ..., K_{m_i}^i\}$ and $B_j = \{K_1^j, K_2^j, ..., K_{m_j}^j\}$ and consider the problem $P(m_i + m_j, 1)$ with fundamental space $\{K_1^i, K_2^i, ..., K_{m_i}^i, K_1^j, K_2^j, ..., K_{m_j}^j\}$.

Since $\{B_0, ..., B_n\}$ is optimal, the solution has to be $\{B_i, B_j\}$ to avoid contradiction. By Step 1, the latter is an ordered partition. Therefore, for any two partition sets in $\{B_0, ..., B_n\}$, there is a cutoff such that the two partition sets lie on different sides of it. This implies that $\{B_0, ..., B_n\}$ is an ordered partition. □

B Theorem 1

Lemma 3 (properties of $\tilde{a}(K)$)**.** *Under Assumptions 3 and 4,*

(i) $\tilde{a}(K)$ *is well-defined, strictly increasing, and five times continuously differentiable on* $(\underline{K}, \bar{K})$*,*

(ii) for any n*,* $a_0^* < \tilde{a}(K_1^*) < a_1^* < \tilde{a}(K_2^*) < ... < \tilde{a}(K_n^*) < a_n^*$*, and*

(iii) $\mathcal{K}^*(n) \subset (\underline{K}, \bar{K}) \subset (\underline{K}^{(1)}, \bar{K}^{(0)})$ *for all* n*.*

Proof of Lemma 3. (i) The function $\tilde{a}(K)$ is implicitly determined by

$$\frac{f_1(K)}{f_0(K)} = \frac{(1-\pi)u_0 h'(1-a)}{\pi u_1 h'(a)}.$$

Because $\frac{\partial}{\partial a} h'(1-a)/h'(a) \neq 0$ on $(0,1)$, by the implicit function theorem, $\tilde{a}(K)$ is well-defined on $(\underline{K}, \bar{K})$ and inherits the smoothness of LHS minus RHS, so it is five times continuously differentiable. (By an analogous argument, $\tilde{K}(a)$, the inverse of $\tilde{a}(K)$ on $(0,1)$, is also well-defined and five times continuously differentiable.) By Assumption 4(iii), the left-hand side is strictly increasing in K. By Assumption 3, the right-hand side is strictly increasing in a. Hence $\tilde{a}(K)$ is strictly increasing.

(ii) Given a partition interval (K_1, K_2), its optimal action a^* satisfies

$$\frac{\Pr(K \in (K_1, K_2) | \theta = 1)}{\Pr(K \in (K_1, K_2) | \theta = 0)} = \frac{(1-\pi)u_0 h'(1-a)}{\pi u_1 h'(a)}.$$

The right-hand side is strictly increasing in a. The left-hand side lies in $\left(\frac{f_{\theta=1}(K_1)}{f_{\theta=0}(K_1)}, \frac{f_{\theta=1}(K_2)}{f_{\theta=0}(K_2)}\right)$ by Assumption 4(iii). Hence, $\tilde{a}(K_1) < a^* < \tilde{a}(K_2)$. This result applies to all partition sets.

(iii) The optimal action a_0^* for $(-\infty, K_1^*)$ satisfies $a_0^* > 0$, so $\tilde{a}(K_1^*) > 0$, i.e., $K_1^* > \underline{K}$. Similarly,

for $(K_n^*, +\infty)$, $a_n^* < 1$, so $\tilde{a}(K_n^*) < 1$, i.e., $K_n^* < \bar{K}$. □

Lemma 4 (cutoffs are dense in the limit)**.** *Under Assumptions 3 and 4,*

$$(i) \qquad \lim_{n\to\infty} \max_{K_i^*, K_{i+1}^* \in \mathcal{K}^*(n)} |K_{i+1}^* - K_i^*| = 0, \text{ and}$$

$$(ii) \qquad \lim_{n\to\infty} K_1^* = \underline{K}, \quad \lim_{n\to\infty} K_n^* = \bar{K}.$$

Proof of Lemma 4. Let $v(a, K_1, K_2) := \Pr(K \in (K_1, K_2))E[u(a;\theta)|K \in (K_1, K_2)]$, $v^*(K_1, K_2) := \max_a v(a, K_1, K_2)$, and denote the optimal action on (K_1, K_2) by $a(K_1, K_2)$.

The proof is by contradiction. For Lemma 4(i), suppose there exists $\delta > 0$ such that for a subsequence $\{n_j\}$,

$$\max_{K_i^*, K_{i+1}^* \in \mathcal{K}^*(n_j)} |K_{i+1}^* - K_i^*| > \delta.$$

By a slight abuse of notation, let (K_i^*, K_{i+1}^*) denote the longest interval in the partition given by $\mathcal{K}^*(n_j)$. Using the cutoff set $\mathcal{K}^*(n_j) \cup \{(K_{i+1}^* + K_i^*)/2\}$ improves $\mathcal{K}^*(n_j)$ by at least $w := \min_{K\in(\underline{K}, \bar{K}-\delta)} v^*(K, K+\delta/2) + v^*(K+\delta/2, K+\delta) - v^*(K, K+\delta)$. By Assumptions 3 and 4(iii), $a(K, K+\delta/2)$ and $a(K+\delta/2, K+\delta)$ cannot coincide for any K. Hence $w > 0$.

For Lemma 4(ii), I present the proof for $\underline{K}$. Suppose there exists $K' > \underline{K}$ such that for a subsequence $\{n_j\}$, $K_1^* > K'$. Define K^{**} as $(\underline{K} + K')/2$ if $\underline{K} > -\infty$ and $K' - 1$ if $\underline{K} = -\infty$. Using the cutoff set $\mathcal{K}^*(n_j) \cup \{K^{**}\}$ improves $\mathcal{K}^*(n_j)$ by at least $w' := \inf_{n_j} v^*(\underline{K}, K^{**}) + v^*(K^{**}, K_1^*) - v^*(\underline{K}, K_1^*) \geq (v^*(\underline{K}, K^{**}) - v(a(\underline{K}, K_1^*), \underline{K}, K^{**})) + \inf_{n_j}\{v^*(K^{**}, K_1^*) - v(a(\underline{K}, K_1^*), K^{**}, K_1^*)\}$. Here $w' > 0$ because by Assumptions 3 and 4(iii), $a(\underline{K}, K^{**})$ and $a(\underline{K}, K_1^*)$ cannot coincide.

Let $\hat{\mathcal{K}}_1(n_1), \hat{\mathcal{K}}_2(n_2) \subset (\underline{K}, \bar{K})$ be any two cutoff sets satisfying $\hat{\mathcal{K}}_1(n_1) \subset \hat{\mathcal{K}}_2(n_2)$, with expected utility levels $\hat{U}_1$ and $\hat{U}_2$. Obviously, $\hat{U}_1 \leq \hat{U}_2$, $\hat{U}_2 \leq M := h(1)\max\{u_0, u_1\} < +\infty$, and $\hat{U}_1 \geq m := \max_a \pi u_1 h(a) + (1-\pi)u_0 h(1-a) > -\infty$. Hence $\hat{U}_2 - \hat{U}_1 \leq M - m$. Under n_j, let $\hat{\mathcal{K}}_2(n_2)$ be $\mathcal{K}^*(n_j)$ and let $\hat{\mathcal{K}}_1(n_1)$ be $\{K_i^* \in \mathcal{K}^*(n_j) | i \text{ is even}\}$. Then, setting $K_0^* := \underline{K}$ and $K_{n+1}^* := \bar{K}$, we get $\hat{U}_2 - \hat{U}_1 = \sum_{t=1,\ldots,\lfloor (n_j+1)/2 \rfloor} \nu_t$, where $\nu_t = v^*(K_{2t-2}^*, K_{2t-1}^*) + v^*(K_{2t-1}^*, K_{2t}^*) - v^*(K_{2t-2}^*, K_{2t}^*)$ for $1 \leq t < \lfloor (n_j+1)/2 \rfloor$ and $\nu_t = v^*(K_{2t-2}^*, K_{2t-1}^*) + \ldots + v^*(K_{n_j}^*, \bar{K}) - v^*(K_{2t-2}^*, \bar{K})$ for

$t = \lfloor (n_j + 1)/2 \rfloor$. Each ν_t is nonnegative, and since $M - m$ is the upper bound for their sum, the three smallest values of ν_t, denoted by $\nu_{\min 1}$, $\nu_{\min 2}$, and $\nu_{\min 3}$, respectively, have their sum $\nu_{\min 1} + \nu_{\min 2} + \nu_{\min 3}$ converge to zero as $n_j \to \infty$. Therefore, given w or w', there exists $\tilde{n}_j$ such that $\nu_{\min 1} + \nu_{\min 2} + \nu_{\min 3}$ is less than w or w'. Let $\nu_{\min 1}$, $\nu_{\min 2}$, and $\nu_{\min 3}$ be achieved at $\tilde{t}_{\min 1}$, $\tilde{t}_{\min 2}$, and $\tilde{t}_{\min 3}$, respectively. Thus, let $\tilde{t}$ be any of the three that does not coincide with K_i^* or K_{i+1}^*, the cutoff set $\mathcal{K}^*(\tilde{n}_j) \cup \{(K_i^* + K_{i+1}^*)/2\} \backslash \{K_{2\tilde{t}-1}^*\}$ improves on $\mathcal{K}^*(\tilde{n}_j)$ by at least $w - \min_t \nu_t > 0$, a contradiction that proves Lemma 4(i). Let $\tilde{t}$ be any of the three that does not equal one. The cutoff set $\mathcal{K}^*(\tilde{n}_j) \cup \{K^{**}\} \backslash \{K_{2\tilde{t}-1}^*\}$ improves on $\mathcal{K}^*(\tilde{n}_j)$ by at least $w' - \min_t \nu_t > 0$, a contradiction that proves Lemma 4(ii). □

Proof of Theorem 1. By Lemma 4, it suffices to show that the functions

$$\hat{s}_n(K) = \begin{cases} s_n(K_i^*) + \frac{s_n(K_{i+1}^*) - s_n(K_i^*)}{K_{i+1}^* - K_i^*}(K - K_i^*) \text{ if } K \in [K_i^*, K_{i+1}^*), \\ \frac{s_n(K_1^*)}{K_1^* - \underline{\underline{K}}}(K - \underline{\underline{K}}) \text{ if } K \in (\underline{\underline{K}}, K_1^*), \quad 0 \text{ if } K \leq \underline{\underline{K}}, \quad 1 \text{ if } K \geq K_n^* \end{cases}$$

converge, where $\underline{\underline{K}} := \max\{\underline{K}, K_1^* - 1\}$. $\hat{s}_n(K)$ are absolutely continuous cumulative distribution functions. By the theorem in Scheffé (1947), it suffices to show their densities,

$$d_n(a) = \begin{cases} \frac{s_n(K_{i+1}^*) - s_n(K_i^*)}{K_{i+1}^* - K_i^*} \text{ if } K \in [K_i^*, K_{i+1}^*), \\ \frac{s_n(K_1^*)}{K_1^* - \underline{\underline{K}}} \text{ if } K \in (\underline{\underline{K}}, K_1^*), \quad 0 \text{ otherwise}, \end{cases}$$

converge pointwise to some limiting density almost everywhere. We proceed in two steps.

Step 1. Find the target density. By Lemma 3, we can define $\tilde{K} : (0, 1) \to (\underline{K}, \bar{K})$ as the inverse function of $\tilde{a}(K)$ on $(\underline{K}, \bar{K})$. To simplify notation, let $t_1(K) = h(\tilde{a}(K))$, $t_0 = h(1 - \tilde{a}(K))$, $I_i = K_{i+1}^* - K_i^*$, and $J_i = \tilde{K}(a_i^*) - \tilde{K}(a_{i-1}^*)$. Then the first-order conditions for a_i^* and K_i^* are respectively

$$\frac{F_1(K_{i+1}^*) - F_1(K_i^*)}{F_0(K_{i+1}^*) - F_0(K_i^*)} = \frac{F_1'(\tilde{K}(a_i^*))}{F_0'(\tilde{K}(a_i^*))}, \tag{B.1}$$

$$\frac{t_1(\tilde{K}(a_i^*)) - t_1(\tilde{K}(a_{i-1}^*))}{t_0(\tilde{K}(a_i^*)) - t_0(\tilde{K}(a_{i-1}^*))} = \frac{t_1'(K_i^*)}{t_0'(K_i^*)}. \tag{B.2}$$

In (B.1), Taylor-expanding $F_1(K_{i+1}^*)$, $F_1(K_i^*)$, $F_0(K_{i+1}^*)$, and $F_0(K_i^*)$ to fourth order at $\bar{K}_i := (K_{i+1}^* + K_i^*)/2$ and $F_0(K_{i+1}^*)$ and $F_0(K_i^*)$ to first order at $\bar{K}_i$, we get

$$\begin{aligned}\tilde{K}(a_i^*) - \bar{K}_i =& \frac{1}{24}\frac{F_0'''(\bar{K}_i)F_1'(\bar{K}_i) - F_1'''(\bar{K}_i)F_0'(\bar{K}_i)}{F_1'(\bar{K}_i)F_0''(\bar{K}_i) - F_0'(\bar{K}_i)F_1''(\bar{K}_i)} I_i^2 + \frac{R_1 + R_2 + R_3 + R_4 + R_5 + R_6}{F_1'(\bar{K}_i)F_0''(\bar{K}_i) - F_0'(\bar{K}_i)F_1''(\bar{K}_i)} \\ :=& \Gamma(\bar{K}_i) I_i^2 + Rem_i^a, \qquad \text{where}\end{aligned} \tag{B.3}$$

$$\begin{aligned}
R_1 &= C_1(F_0'(\bar{K}_i)F_1'''(K_d) - F_1'(\bar{K}_i)F_0'''(K_d')) && \times (\tilde{K}(a_i^*) - \bar{K}_i)^2, \\
R_2 &= C_2(F_0'''(\bar{K}_i)F_1''(\bar{K}_i) - F_1'''(\bar{K}_i)F_0''(\bar{K}_i) && \times I_i^2(\tilde{K}(a_i^*) - \bar{K}_i), \\
R_3 &= C_3(F_0'''(\bar{K}_i)F_1'''(K_d) - F_1'''(\bar{K}_i)F_0'''(K_d')) && \times I_i^2(\tilde{K}(a_i^*) - \bar{K}_i)^2, \\
R_4 &= C_4((F_0^{(5)}(K_c'') - F_0^{(5)}(K_c'''))F_1'(\bar{K}_i) - (F_1^{(5)}(K_c) - F_1^{(5)}(K_c'))F_0'(\bar{K}_i)) && \times I_i^4, \\
R_5 &= C_5((F_0^{(5)}(K_c'') - F_0^{(5)}(K_c'''))F_1''(\bar{K}_i) - (F_1^{(5)}(K_c) - F_1^{(5)}(K_c'))F_0''(\bar{K}_i)) && \times I_i^4(\tilde{K}(a_i^*) - \bar{K}_i), \\
R_6 &= C_6((F_0^{(5)}(K_c'') - F_0^{(5)}(K_c'''))F_1'''(K_d) - (F_1^{(5)}(K_c) - F_1^{(5)}(K_c'))F_0'''(K_d')) && \times I_i^4(\tilde{K}(a_i^*) - \bar{K}_i)^2.
\end{aligned}$$

In (B.2), Taylor-expanding $t_1(\tilde{K}(a_i^*))$, $t_1(\tilde{K}(a_{i-1}^*))$, $t_0(\tilde{K}(a_i^*))$, and $t_0(\tilde{K}(a_{i-1}^*))$ to fourth order at $\bar{\bar{K}}_i := (\tilde{K}(a_i^*) + \tilde{K}(a_{i-1}^*))/2$ and $t_1'(K_i^*)$ and $t_0'(K_i^*)$ to first order at $\bar{\bar{K}}_i$, we get

$$\begin{aligned}K_i^* - \bar{\bar{K}}_i =& \frac{1}{24}\frac{t_0'''(\bar{\bar{K}}_i)t_1'(\bar{\bar{K}}_i) - t_1'''(\bar{\bar{K}}_i)t_0'(\bar{\bar{K}}_i)}{t_1'(\bar{\bar{K}}_i)t_0''(\bar{\bar{K}}_i) - t_0'(\bar{\bar{K}}_i)t_1''(\bar{\bar{K}}_i)} J_i^2 + \frac{S_1 + S_2 + S_3 + S_4 + S_5 + S_6}{t_1'(\bar{\bar{K}}_i)t_0''(\bar{\bar{K}}_i) - t_0'(\bar{\bar{K}}_i)t_1''(\bar{\bar{K}}_i)} \\ :=& T(\bar{\bar{K}}_i) J_i^2 + Rem_i^K, \qquad \text{where}\end{aligned} \tag{B.4}$$

$$\begin{aligned}
S_1 &= C_1(t_0'(\bar{\bar{K}}_i)t_1'''(K_{dd}) - t_1'(\bar{\bar{K}}_i)t_0'''(K_{dd}')) && \times (K_i^* - \bar{\bar{K}}_i)^2, \\
S_2 &= C_2(t_0'''(\bar{\bar{K}}_i)t_1''(\bar{\bar{K}}_i) - t_1'''(\bar{\bar{K}}_i)t_0''(\bar{\bar{K}}_i)) && \times J_i^2(K_i^* - \bar{\bar{K}}_i), \\
S_3 &= C_3(t_0'''(\bar{\bar{K}}_i)t_1'''(K_{dd}) - t_1'''(\bar{\bar{K}}_i)t_0'''(K_{dd}')) && \times J_i^2(K_i^* - \bar{\bar{K}}_i)^2, \\
S_4 &= C_4((t_0^{(5)}(K_{cc}'') - t_0^{(5)}(K_{cc}'''))t_1'(\bar{\bar{K}}_i) - (t_1^{(5)}(K_{cc}) - t_1^{(5)}(K_{cc}'))t_0'(\bar{\bar{K}}_i)) && \times J_i^4,
\end{aligned}$$

$$S_5 = C_5((t_0^{(5)}(K_{cc}'') - t_0^{(5)}(K_{cc}'''))t_1''(\bar{\bar{K}}_i) - (t_1^{(5)}(K_{cc}) - t_1^{(5)}(K_{cc}'))t_0''(\bar{\bar{K}}_i)) \quad \times J_i^4(K_i^* - \bar{\bar{K}}_i),$$

$$S_6 = C_6((t_0^{(5)}(K_{cc}'') - t_0^{(5)}(K_{cc}'''))t_1'''(K_{dd}) - (t_1^{(5)}(K_{cc}) - t_1^{(5)}(K_{cc}'))t_0'''(K_{dd}')) \quad \times J_i^4(K_i^* - \bar{\bar{K}}_i)^2.$$

Here, K_d and K_d' (between $\tilde{K}(a_i^*)$ and $\bar{K}_i$) are in the remainders of $F_1'(\tilde{K}(a_i^*))$ and $F_0'(\tilde{K}(a_i^*))$; K_c and K_c'' (between K_{i+1}^* and $\bar{K}_i$) are in the remainders of $F_1(K_{i+1}^*)$ and $F_0(K_{i+1}^*)$; and K_c' and K_c''' (between $\bar{K}_i$ and $\tilde{K}(a_i^*)$) are in the remainders of $F_1(K_i^*)$ and $F_0(K_i^*)$. We omit the subscripts i in this notation. Differentiation is possible by the smoothness in Assumption 4. The denominator $F_1'(\bar{K}_i)F_0''(\bar{K}_i) - F_0'(\bar{K}_i)F_1''(\bar{K}_i)$ is positive when $\bar{K}_i \in (\underline{K}, \bar{K})$, since $(F_1'(K)/F_0'(K))' > 0$ by Assumption 4.

Also, K_{dd} and K_{dd}' (between K_i^* and $\bar{\bar{K}}_i$) are in the remainders of $t_1'(K_i^*)$ and $t_0'(K_i^*)$; K_{cc} and K_{cc}'' (between $\tilde{K}(a_i^*)$ and $\bar{\bar{K}}_i$) are in the remainders of $t_1(\tilde{K}(a_i^*))$ and $t_0(\tilde{K}(a_i^*))$; and K_{cc}' and K_{cc}''' (between $\bar{\bar{K}}_i$ and $\tilde{K}(a_{i-1}^*)$) are in the remainders of $t_1(\tilde{K}(a_{i-1}^*))$ and $t_0(\tilde{K}(a_{i-1}^*))$. Again we omit the subscripts i. Differentiation is possible by the smoothness in Assumption 3 and Lemma 3(i), and $t_1'(\bar{\bar{K}}_i)t_0''(\bar{\bar{K}}_i) - t_0'(\bar{\bar{K}}_i)t_1''(\bar{\bar{K}}_i) > 0$ when $\bar{\bar{K}}_i \in (\underline{K}, \bar{K})$ by Assumption 3. Thus,

$$I_i - I_{i-1} = -2\Gamma(\bar{K}_i)I_i^2 - 2\Gamma(\bar{K}_{i-1})I_{i-1}^2 - 4T(\bar{\bar{K}}_i)J_i^2 - 2Rem_i^a - 2Rem_{i-1}^a - 4Rem_i^K, \quad \text{(B.5)}$$

$$J_i - I_{i-1} = -2\Gamma(\bar{K}_{i-1})I_{i-1}^2 - 2T(\bar{\bar{K}}_i)J_i^2 - 2Rem_{i-1}^a - 2Rem_i^K. \quad \text{(B.6)}$$

To find the limit of $d_n(x)$, we first investigate $d_n(x)/d_n(y)$ on $(\underline{K}, \bar{K})$. Choose any closed interval $[K_L, K_H] \subset (\underline{K}, \bar{K})$ and examine $d_n(x)/d_n(y)$, for any $x, y \in [K_L, K_H]$ $(x < y)$. Define $i_z := \max\{i : K_i^* \leq z\}$. By Lemma 4, $\lim_{n\to\infty} K_1^* = \underline{K}$, so for a large n,

$$d_n(x) = \frac{s_n(K_{i_x+1}^*) - s_n(K_{i_x}^*)}{K_{i_x+1}^* - K_{i_x}^*} = \frac{1/n}{I_{i_x}}.$$

Thus, for any large n,

$$\frac{d_n(x)}{d_n(y)} = \frac{(1/n)/I_{i_x}}{(1/n)/I_{i_y}} = \frac{I_{i_y}}{I_{i_y-1}} \times ... \times \frac{I_{i_x+1}}{I_{i_x}} = \exp\left(\ln\left(\frac{I_{i_y}}{I_{i_y-1}}\right) + ... + \ln\left(\frac{I_{i_x+1}}{I_{i_x}}\right)\right). \tag{B.7}$$

The first-order Taylor expansion of $\ln\left(\frac{I_i}{I_{i-1}}\right)$ at 1 gives $\ln\left(\frac{I_i}{I_{i-1}}\right) = \frac{I_i}{I_{i-1}} - 1 - \frac{1}{2t_i}(\frac{I_i}{I_{i-1}} - 1)^2$, where t_i is between $\frac{I_i}{I_{i-1}}$ and 1, so Eq. (B.7) becomes

$$\frac{d_n(x)}{d_n(y)} = \exp\left(\sum_{i=i_x+1}^{i_y}(\frac{I_i}{I_{i-1}} - 1)\right) / \exp\left(\sum_{i=i_x+1}^{i_y}\frac{1}{2t_i}(\frac{I_i}{I_{i-1}} - 1)^2\right). \tag{B.8}$$

By Eqs. (B.5) and (B.6), we have

$$\frac{I_i}{I_{i-1}} = 1 - 2\Gamma(\bar{K}_i)I_i\frac{I_i}{I_{i-1}} - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 4T(\bar{\bar{K}}_i)J_i\frac{J_i}{I_{i-1}} - \frac{2Rem_i^a + 2Rem_{i-1}^a + 4Rem_i^K}{I_{i-1}}, \tag{B.9}$$

$$\frac{J_i}{I_{i-1}} = 1 - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 2T(\bar{\bar{K}}_i)J_i\frac{J_i}{I_{i-1}} - \frac{2Rem_{i-1}^a + 2Rem_i^K}{I_{i-1}}. \tag{B.10}$$

Now let us discuss Eq. (B.8). First, we show three useful results. Results 1 and 3 are with proofs, and Result 2 is a direct corollary of Result 1.

Result 1. We have $\max_i |I_i/I_{i-1} - 1| \to 0$ and $\max_i |J_i/I_{i-1} - 1| \to 0$ as $n \to \infty$.

Proof of Result 1. For brevity, let $X_i = I_i/I_{i-1}$ and $Y_i = J_i/I_{i-1}$. Equations (B.9) and (B.10) can be rearranged into the following system:

$$X_i = 1 - 2\Gamma(\bar{K}_i)I_iX_i - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 4T(\bar{\bar{K}}_i)J_iY_i - 2\frac{Rem_i^a/I_{i-1}}{X_i}X_i - 2Rem_{i-1}^a/I_{i-1} - \frac{4Rem_i^K/I_{i-1}}{Y_i}Y_i,$$

$$Y_i = 1 - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 2T(\bar{\bar{K}}_i)J_iY_i - 2Rem_{i-1}^a/I_{i-1} - \frac{2Rem_i^K/I_{i-1}}{Y_i}Y_i.$$

Solving this system, we obtain

$$X_i = \frac{1 - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 2Rem_{i-1}^a/I_{i-1} - (4T(\bar{\bar{K}}_i)J_i + \frac{4Rem_i^K/I_{i-1}}{Y_i})Y_i}{1 + 2\Gamma(\bar{K}_i)I_i + 2\frac{Rem_i^a/I_{i-1}}{X_i}}, \tag{B.11}$$

$$Y_i = \frac{1 - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 2Rem^a_{i-1}/I_{i-1}}{1 + 2T(\bar{\bar{K}}_i)J_i + \frac{2Rem^K_i/I_{i-1}}{Y_i}}. \tag{B.12}$$

Here, Γ, T and all functions appearing in Rem^a and Rem^K that consist of higher-order derivatives of F and t are continuous on $[K_L, K_H]$ by the continuous differentiability assumption, and are therefore bounded. Let $M > 0$ be a uniform upper bound on their absolute values.

Notice that in Rem^a_{i-1} we have $|\tilde{K}(a^*_{i-1}) - \bar{K}_{i-1}| \le I_{i-1}$, since $\tilde{K}(a^*_{i-1}), \bar{K}_{i-1} \in [K^*_{i-1}, K^*_i]$ by Lemma 3(ii). Hence, by the triangle inequality, $|Rem^a_{i-1}| \le M(I^2_{i-1}+I^3_{i-1}+I^4_{i-1}+I^4_{i-1}+I^5_{i-1}+I^6_{i-1})$, so $|Rem^a_{i-1}/I_{i-1}| \le M(I_{i-1} + I^2_{i-1} + 2I^3_{i-1} + I^4_{i-1} + I^5_{i-1})$. Since $\max_i I_{i-1} \to 0$ in Lemma 4, $\max_i |Rem^a_{i-1}/I_{i-1}| \to 0$.

Also, in Rem^K_i we have $|K^*_i - \bar{\bar{K}}_i| \le J_i$, since $\tilde{K}^*_i, \bar{\bar{K}}_i \in [\tilde{K}(a^*_{i-1}), \tilde{K}(a^*_i)]$. Hence $|Rem^K_i| \le M(J^2_i + J^3_i + J^4_i + J^4_i + J^5_i + J^6_i)$, and so $|(Rem^K_i/I_{i-1})/Y_i| = |Rem^K_i/J_i| \le M(J_i + J^2_i + 2J^3_i + J^4_i + J^5_i)$. Since $\max_i J_i \to 0$ in Lemma 4, $\max_i |(Rem^K_i/I_{i-1})/Y_i| \to 0$.

We conclude that in Eq. (B.12), $|\Gamma|, |T| \le M$, and all indexed terms uniformly converge to 0. Thus, $Y_i \to 1$ uniformly on $[K_L, K_H]$.

Furthermore, in Rem^a_i we have $|\tilde{K}(a^*_i) - \bar{K}_i| \le I_i$, since $\tilde{K}(a^*_i), \bar{K}_i \in [K^*_i, K^*_{i+1}]$. Hence $|Rem^a_i| \le M(I^2_i + I^3_i + I^4_i + I^4_i + I^5_i + I^6_i)$, and so $|(Rem^a_i/I_{i-1})/X_i| = |Rem^a_i/I_i| \le M(I_i + I^2_i + 2I^3_i + I^4_i + I^5_i)$. Since $\max_i I_i \to 0$ in Lemma 4, $\max_i |(Rem^a_i/I_{i-1})/X_i| \to 0$.

Therefore, in Eq. (B.11), $|\Gamma|, |T| \le M$, Y_i uniformly converges to 1, and all other indexed terms uniformly converge to 0. Thus, $X_i \to 1$ uniformly on $[K_L, K_H]$. □

Result 2. There exists $\eta \in (0, 1)$ such that $I_i/I_{i-1}, J_i/I_{i-1} \in [1 - \eta, 1 + \eta]$ for any i and large n.

Definition. Let $(a_m)_{m=1}^M$, $(b_m)_{m=1}^M$, and $(c_m)_{m=1}^M$ be vectors of nonnegative integers and let $z = \min(a_m + b_m + c_m)_{m=1}^M$. Let $o_i(z)$ denote $\sum_{m=1}^M I^{a_m}_{i-1} I^{b_m}_i J^{c_m}_i$.

Result 3. We have $\sum_{i=i_x+1}^{i_y} o_i(z) \to 0$ for $x, y \in [K_L, K_H]$ if $z \ge 2$.

Proof of Result 3. Let $a, b, c \ge 0$ be integers. It suffices to show that $\sum_{i_x+1}^{i_y} I^a_{i-1} I^b_i J^c_i \to 0$ for $x, y \in [K_L, K_H]$ if $a + b + c \ge 2$. Without loss of generality, let $a > 1$; then $0 \le \sum_{i_x+1}^{i_y} I^a_{i-1} I^b_i J^c_i \le$

$(\sum_{i_x+1}^{i_y} I_{i-1})(\max_i I_{i-1})^{a-1}(\max_i I_i)^b(\max_i J_i)^c = (y-x)(\max_i I_{i-1})^{a-1}$

$(\max_i I_i)^b(\max_i J_i)^c \to 0$. This proof also works for $b > 1$ or $c > 1$. □

To understand the term $I_i/I_{i-1} - 1$ in Eq. (B.8), substitute Eq. (B.10) into the right-hand side of Eq. (B.9), and subtract 1 from both sides to get

$$\frac{I_i}{I_{i-1}} - 1 = -2\Gamma(\bar{K}_i)I_i - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 4T(\bar{\bar{K}}_i)J_i + Residual_i, \tag{B.13}$$

where $Residual_i =$

$$\begin{aligned}
&- (2Rem_i^a + 2Rem_{i-1}^a + 4Rem_i^K)/I_{i-1} \\
&- 2\Gamma(\bar{K}_i)I_i\left(-2\Gamma(\bar{K}_i)I_i\frac{I_i}{I_{i-1}} - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 4T(\bar{\bar{K}}_i)J_i\frac{J_i}{I_{i-1}} - \frac{2Rem_i^a + 2Rem_{i-1}^a + 4Rem_i^K}{I_{i-1}}\right) \\
&- 4T(\bar{\bar{K}}_i)J_i\left(-2\Gamma(\bar{K}_{i-1})I_{i-1} - 2T(\bar{\bar{K}}_i)J_i\frac{J_i}{I_{i-1}} - \frac{2Rem_{i-1}^a + 2Rem_i^K}{I_{i-1}}\right).
\end{aligned}$$

In $Residual_i$, we first examine Rem_i^a. Since $|\tilde{K}(a_i^*) - \bar{K}_i| \leq I_i$ (by the proof of Result 1), each of R_2, R_3, R_4, R_5, R_6 divided by its denominator should have absolute value at most a constant times I_i^3, I_i^4, I_i^4, I_i^5, I_i^6, respectively. Since $|R_1|$ is at most a constant times I_i^2, by Eq. (B.4), $R_1 = C_1(F_0'(\bar{K}_i)F_1'''(K_d) - F_1'(\bar{K}_i)F_0'''(K_d'))(\Gamma(\bar{K}_i)I_i^2 + Rem_i^a)^2$. The terms here have absolute values at most a constant times I_i^4, $I_i^2 Rem_i^a$ (which is at most a constant times $I_i^2(I_i^2+I_i^3+I_i^4+I_i^4+I_i^5+I_i^6)$), and $(Rem_i^a)^2$ (which is at most a constant times $I_i^4+...+I_i^{12}$). Thus, $|R_1|$ is at most a constant times I_i^4 plus higher-order terms. Hence, $|Rem_i^a| \leq \frac{M_1}{1+\eta}(I_i^3 + ... + I_i^{12})$, for a constant M_1. By Result 2, for large n, $|\frac{Rem_i^a}{I_{i-1}}| \leq \frac{M_1}{1+\eta}\frac{I_i}{I_{i-1}}(I_i^2+...+I_i^{11}) \leq M_1 o_i(2)$. Similarly, $|\frac{Rem_{i-1}^a}{I_{i-1}}| \leq M_2(I_{i-1}^2+...+I_{i-1}^{11}) \leq M_2 o_i(2)$.

By analogous arguments, $|Rem_i^K| \leq \frac{M_3}{1+\eta}(J_i^3 + ... + J_i^{12})$, where M_3 is a constant. For large n, $|\frac{Rem_i^K}{I_{i-1}}| \leq \frac{M_3}{1+\eta}\frac{J_i}{I_{i-1}}(J_i^2 + ... + J_i^{11}) \leq M_3 o_i(2)$.

Hence, for large n such that $0 < 1 - \eta \leq I_i/I_{i-1}, J_i/I_{i-1} \leq 1 + \eta$ and a constant M_4,

$$|Residual_i| \leq M_4 o_i(2) \tag{B.14}$$

on $[K_L, K_H]$. By Result 3,

$$\sum_{i=i_x+1}^{i_y} Residual_i \to 0, \tag{B.15}$$

since $0 \leq |\sum_{i=i_x+1}^{i_y} Residual_i| \leq \sum_{i=i_x+1}^{i_y} |Residual_i| \leq \sum_{i=i_x+1}^{i_y} M_4 o_i(2) \to 0$. Squaring Eq. (B.13) and using Eq. (B.14), for large n and some constant M_5 we get $|\frac{I_i}{I_{i-1}} - 1|^2 \leq M_5 o_i(2)$. Hence for the denominator of Eq. (B.8) we have

$$\exp\left(\sum_{i=i_x+1}^{i_y} \frac{1}{2t_i}(\frac{I_i}{I_{i-1}} - 1)^2\right) \to 1, \tag{B.16}$$

because, with t_i between 1 and I_i/I_{i-1} and hence between $1-\eta$ and $1+\eta$ by Result 2,

$$0 \leq |\sum_{i=i_x+1}^{i_y} \frac{1}{2t_i}(\frac{I_i}{I_{i-1}} - 1)^2| \leq \frac{1}{2(1-\eta)} \sum_{i=i_x+1}^{i_y} (\frac{I_i}{I_{i-1}} - 1)^2 \leq \frac{1}{2(1-\eta)} \sum_{i=i_x+1}^{i_y} M_5 o_i(2) \to 0.$$

Now, substituting Eq. (B.13) in Eq. (B.8), we can rewrite Eq. (B.8) as

$$\begin{aligned}\frac{d_n(x)}{d_n(y)} &= \exp\left(\sum_{i=i_x+1}^{i_y} (-2\Gamma(\bar{K}_i)I_i - 2\Gamma(\bar{K}_{i-1})I_{i-1} - 4T(\bar{\bar{K}}_i)J_i)\right) \times Rest(n;[x,y]) \\ &= \exp\left(RS(-4\Gamma;[x,y]) + RS(-4T;[x,y])\right) \times Rest(n;[x,y]),\end{aligned} \tag{B.17}$$

where $Rest(n;[x,y]) \to 1$ by Eq. (B.15) and Eq. (B.16), and $RS(f,[x,y])$ denotes a Riemann sum of f on $[x,y]$. Because $\int_x^y f'(a)/f(a)da = \ln(|f(a)|)|_x^y$, when $x, y \in [K_L, K_H]$,

$$\frac{d_n(x)}{d_n(y)} \to \exp\left(-4\int_x^y \Gamma(K)dK - 4\int_x^y T(K)dK\right)$$

$$\begin{aligned}&= \left|\frac{F_1'(x)F_0''(x) - F_0'(x)F_1''(x)}{F_1'(y)F_0''(y) - F_0'(y)F_1''(y)}\right|^{\frac{1}{6}} \left|\frac{t_1'(x)t_0''(x) - t_0'(x)t_1''(x)}{t_1'(y)t_0''(y) - t_0'(y)t_1''(y)}\right|^{\frac{1}{6}} \\ &= \left|\frac{F_1'(x)F_0''(x) - F_0'(x)F_1''(x)}{F_1'(y)F_0''(y) - F_0'(y)F_1''(y)}\right|^{\frac{1}{6}} \left|\frac{h'(\tilde{a}(x))h''(1-\tilde{a}(x)) + h'(1-\tilde{a}(x))h''(\tilde{a}(x))}{h'(\tilde{a}(y))h''(1-\tilde{a}(y)) + h'(1-\tilde{a}(y))h''(\tilde{a}(y))}\right|^{\frac{1}{6}} \left|\frac{\tilde{a}'(x)}{\tilde{a}'(y)}\right|^{\frac{1}{2}}.\end{aligned}$$

Under Assumption 3, $|h'(\tilde{a}(x))h''(1-\tilde{a}(x)) + h'(1-\tilde{a}(x))h''(\tilde{a}(x))| = -(h'(\tilde{a}(x))h''(1-\tilde{a}(x)) + h'(1-\tilde{a}(x))h''(\tilde{a}(x)))$. By the monotone likelihood ratio property in Assumption 4, $|F_1'(x)F_0''(x) -$

$F_0'(x)F_1''(x)| = F_1''(x)F_0'(x) - F_0''(x)F_1'(x)$. By the integrability assumption in Theorem 1, the denominator as a function of y is integrable on $(\underline{K}, \bar{K})$. Let $m(y) > 0$ denote the denominator scaled by its integral on $(\underline{K}, \bar{K})$, so that $\int_{\underline{K}}^{\bar{K}} m(y)dy = 1$. Then, pointwise for $x, y \in [K_L, K_H]$,

$$\lim_{n\to\infty} \frac{d_n(x)}{d_n(y)} = \frac{m(x)}{m(y)}. \tag{B.18}$$

For any $x, y \in (\underline{K}, \bar{K})$, Eq. (B.18) holds because there always exist K_L, K_H such that $x, y \in [K_L, K_H]$.

Step 2. Prove convergence.

To show that for any $y \in (\underline{K}, \bar{K})$,

$$\lim_{n\to\infty} d_n(y) = m(y), \tag{B.19}$$

notice that by Eq. (B.18),

$$\begin{aligned}\frac{1}{m(y)} &= \int_{\underline{K}}^{\bar{K}} \frac{m(x)}{m(y)}dx = \int_{\underline{K}}^{\bar{K}} \lim_{n\to\infty} \frac{d_n(x)}{d_n(y)}dx \\ &\le \liminf_{n\to\infty} \int_{\underline{K}}^{\bar{K}} \frac{d_n(x)}{d_n(y)}dx = \liminf_{n\to\infty} \frac{1}{d_n(y)} = \frac{1}{\limsup_{n\to\infty} d_n(y)},\end{aligned}$$

where the inequality is by Fatou's lemma and the subsequent equality comes from the fact that $\int_{\underline{K}}^{\bar{K}} d_n(x)dx = 1$ for any n. Hence,

$$m(y) \ge \limsup_{n\to\infty} d_n(y) > 0. \tag{B.20}$$

Therefore $\limsup_{n\to\infty} d_n(y)$ and $d_n(y)$ are dominated by $m(y)$, which is integrable on $(\underline{K}, \bar{K})$. By the Fatou–Lebesgue theorem,

$$1 = \limsup_{n\to\infty} 1 = \limsup_{n\to\infty} \int_{\underline{K}}^{\bar{K}} d_n(y)dy \le \int_{\underline{K}}^{\bar{K}} \limsup_{n\to\infty} d_n(y)dy \le \int_{\underline{K}}^{\bar{K}} m(y)dy = 1,$$

so $\int_{\underline{K}}^{\bar{K}} \limsup_{n\to\infty} d_n(y)dy = \int_{\underline{K}}^{\bar{K}} m(y)dy$. Combined with Eq. (B.20), this implies that, almost

everywhere,

$$\limsup_{n\to\infty} d_n(y) = m(y). \tag{B.21}$$

Next, I show by contradiction that $\liminf_{n\to\infty} d_n(y) = m(y)$. Suppose otherwise; then by Eq. (B.21), there exists $y_0 \in (\underline{K}, \bar{K})$ such that $\liminf_{n\to\infty} d_n(y_0) = m(y_0) - \delta \in [0, m(y_0))$. Hence there is a convergent subsequence $\{d_{n_k}(y_0)\}$ such that $\lim_{k\to\infty} d_{n_k}(y_0) = m(y_0) - \delta$. By Eq. (B.18), for any $x \in (\underline{K}, \bar{K})$,

$$\lim_{k\to\infty} d_{n_k}(x) = \frac{m(y_0) - \delta}{m(y_0)} m(x).$$

Then, on the one hand,

$$\int_{\underline{K}}^{\bar{K}} \lim_{k\to\infty} d_{n_k}(x)dx = \frac{m(y_0) - \delta}{m(y_0)},$$

but on the other hand, because $d_{n_k}(x) \le \limsup_{k\to\infty} d_{n_k}(x) \le \limsup_{n\to\infty} d_n(x) \le m(x)$, by the dominated convergence theorem we have

$$\int_{\underline{K}}^{\bar{K}} \lim_{k\to\infty} d_{n_k}(x)dx = \lim_{k\to\infty} \int_{\underline{K}}^{\bar{K}} d_{n_k}(x)dx = \lim_{k\to\infty} 1 = 1,$$

a contradiction. This proves Eq. (B.19). Hence we have $s(K) := s_n(K) \to \int_{\underline{K}}^{K} m(y)dy$ as discussed in the beginning of the proof. Therefore, $s'(K) = m(K) \propto \kappa_H(K)^{\frac{1}{6}} \kappa_F(K)^{\frac{1}{6}}$. □

C Propositions 4, 5 and 6

Proof of Proposition 5. First we show the symmetry of $\tilde{a}'(K)$. By the definition of $\tilde{a}(K)$,

$$\frac{h'(1-a)}{h'(a)} = \frac{u_1\pi}{u_0(1-\pi)} \exp\left(\frac{2\mu}{\sigma^2}K\right). \tag{C.1}$$

Hence,

$$\tilde{K}(a) = \frac{\sigma^2}{2\mu}\left(\ln h'(1-a) - \ln h'(a) - \ln\frac{u_1\pi}{u_0(1-\pi)}\right) \tag{C.2}$$

for $a \in (0,1)$, and $\tilde{K}(a) = 2K_{1/2} - \tilde{K}(1-a)$, implying $\tilde{K}(a)$ is symmetric about $(\frac{1}{2}, K_{1/2})$. Thus, $\tilde{a}(K)$ is symmetric about $(K_{1/2}, \frac{1}{2})$ on $(\underline{K}, \bar{K})$. This symmetry obviously holds outside $(\underline{K}, \bar{K})$. The symmetry of $\tilde{a}(K)$ implies the symmetry of $\tilde{a}'(K)$ about $K = K_{1/2}$.

Next, define

$$\kappa_h(K) := -h'(1-\tilde{a}(K))h''(\tilde{a}(K)) - h''(1-\tilde{a}(K))h'(\tilde{a}(K))$$

so that $\kappa_H(K) = \kappa_h(K)\tilde{a}'(K)^3$. It suffices to show the symmetry of $\kappa_h(K)$. Evaluating $\kappa_h(K)$ at $K_{1/2}+\delta$ and $K_{1/2}-\delta$ with $\tilde{a}(K_{1/2}+\delta)+\tilde{a}(K_{1/2}-\delta) = 1$, we get $\kappa_h(K_{1/2}+\delta) = \kappa_h(K_{1/2}-\delta)$. □

Proof of Proposition 6. (i) Since the bell-shaped $\kappa_H(K)^{\frac{1}{6}}$, i.e., $\kappa_h(K)^{\frac{1}{6}}\tilde{a}'(K)^{\frac{1}{2}}$ is symmetric about $K_{1/2}$, its value is higher if K is closer to $K_{1/2}$. Then, for K such that $KK_{1/2} > 0$, $s'(K) = \exp\left(-\frac{K^2}{6\sigma^2}\right)\kappa_h(K)^{\frac{1}{6}}\tilde{a}'(K)^{\frac{1}{2}} \geq \exp\left(-\frac{K^2}{6\sigma^2}\right)\kappa_h(-K)^{\frac{1}{6}}\tilde{a}'(-K)^{\frac{1}{2}} = s'(-K)$. The inequality holds because $K_{1/2}$ and K are closer together than $K_{1/2}$ and $-K$.

(i*) By arguments analogous to those used for (i), the value of the U-shaped $\kappa_h(K)^{\frac{1}{6}}\tilde{a}'(K)^{\frac{1}{2}}$ is lower if K is closer to $K_{1/2}$. Then, for K such that $KK_{1/2} > 0$, we have $s'(K) = \exp\left(-\frac{K^2}{6\sigma^2}\right)\kappa_h(K)^{\frac{1}{6}}\tilde{a}'(K)^{\frac{1}{2}} \leq \exp\left(-\frac{K^2}{6\sigma^2}\right)\kappa_h(-K)^{\frac{1}{6}}\tilde{a}'(-K)^{\frac{1}{2}} = s'(-K)$.

(ii) Differentiating both sides of Eq. (C.1) with respect to K, we get

$$\tilde{a}' = \frac{2\mu}{\sigma^2}\frac{h'(\tilde{a})^2}{\kappa_h(K)}\frac{u_1\pi}{u_0(1-\pi)}\exp\left(\frac{2\mu}{\sigma^2}K\right).$$

Differentiating the multiplicative inverse of both sides, we get

$$\tilde{a}' = \frac{2\mu}{\sigma^2}\frac{h'(1-\tilde{a})^2}{\kappa_h(K)} \Bigg/ \left(\frac{u_1\pi}{u_0(1-\pi)}\exp\left(\frac{2\mu}{\sigma^2}K\right)\right).$$

The product of these two equations gives $\tilde{a}'^2$. Its square root is

$$\tilde{a}' = \frac{2\mu}{\sigma^2}\frac{h'(\tilde{a})h'(1-\tilde{a})}{\kappa_h(K)},$$

implying that

$$(\kappa_h^{\frac{1}{6}}(\tilde{a}')^{\frac{1}{2}})^6 = \kappa_h(\tilde{a}')^3 = (\frac{2\mu}{\sigma^2})^3 \frac{h'(\tilde{a})^3 h'(1-\tilde{a})^3}{\kappa_h(K)^2} = (\frac{2\mu}{\sigma^2})^3 \frac{h'(\tilde{a})^3 h'(1-\tilde{a})^3}{(-h'(1-\tilde{a})h''(\tilde{a}) - h''(1-\tilde{a})h'(\tilde{a}))^2}.$$

Hence, $\kappa_h^{\frac{1}{6}}(\tilde{a}')^{\frac{1}{2}}$ is increasing (decreasing) if and only if

$$\frac{d}{d\tilde{a}} \left(\frac{h'(\tilde{a})^3 h'(1-\tilde{a})^3}{(-h'(1-\tilde{a})h''(\tilde{a}) - h''(1-\tilde{a})h'(\tilde{a}))^2} \right) \times \frac{da}{dK} \geq (\leq)\ 0,$$

i.e.,

$$\frac{d}{da} \left(\frac{h'(a)^3 h'(1-a)^3}{(-h'(1-a)h''(a) - h''(1-a)h'(a))^2} \right) \geq (\leq)\ 0. \tag{C.3}$$

For $\kappa_h^{\frac{1}{6}}(\tilde{a}')^{\frac{1}{2}}$ to be bell-shaped in K, this needs to hold if and only if $a < \frac{1}{2}$.

(iii) A sufficient condition for Eq. (C.3) ($\geq$) to hold is that both $\frac{h''(a)}{h'(a)}$ and $\frac{h'''(a)}{h'(a)}$ be decreasing in a. If so, then for $a < \frac{1}{2}$ we have $a < 1-a$, and hence

$$3\left(\frac{h''(a)}{h'(a)} - \frac{h''(1-a)}{h'(1-a)}\right)\kappa_h + 2\left(\frac{h'''(a)}{h'(a)} - \frac{h'''(1-a)}{h'(1-a)}\right) h'(a)h'(1-a) \geq 0.$$

This implies that Eq. (C.3) ($\geq$) holds. □

Proof of Proposition 7. To show Proposition 7, we only need to show that

$$\lim_{\sigma \to 0} (s(r) - s(-r)) = 1$$

for any $r > 0$.

Consider the case when $K_{1/2} > 0$ without loss of generality. Since $K_{1/2} \to 0$ as $\sigma \to 0$, we have $K_{1/2} < \frac{r}{2}$ for small σ values. By Proposition 5, this means $\kappa_H(K)^{\frac{1}{6}} \geq \kappa_H(r)^{\frac{1}{6}}$ for $K \in (0, r)$ and $\kappa_H(K)^{\frac{1}{6}} \leq \kappa_H(r)^{\frac{1}{6}}$ for $K > r$. Also, it is obvious that $\kappa_H(K)^{\frac{1}{6}} \geq \kappa_H(-r)^{\frac{1}{6}}$ for $K \in (-r, 0)$ and

$\kappa_H(K)^{\frac{1}{6}} \leq \kappa_H(r)^{\frac{1}{6}}$ for $K < -r$. Therefore, we have

$$\frac{1-s(r)}{s(r)-s(0)} = \frac{\int_r^{+\infty} \kappa_H(K)^{\frac{1}{6}} \kappa_F(K)^{\frac{1}{6}} dK}{\int_0^r \kappa_H(K)^{\frac{1}{6}} \kappa_F(K)^{\frac{1}{6}} dK} \leq \frac{\kappa_H(r)^{\frac{1}{6}} \int_r^{+\infty} \kappa_F(K)^{\frac{1}{6}} dK}{\kappa_H(r)^{\frac{1}{6}} \int_0^r \kappa_F(K)^{\frac{1}{6}} dK} = \frac{\int_r^{+\infty} \kappa_F(K)^{\frac{1}{6}} dK}{\int_0^r \kappa_F(K)^{\frac{1}{6}} dK},$$

$$\frac{s(-r)-0}{s(0)-s(-r)} = \frac{\int_{-\infty}^{-r} \kappa_H(K)^{\frac{1}{6}} \kappa_F(K)^{\frac{1}{6}} dK}{\int_{-r}^0 \kappa_H(K)^{\frac{1}{6}} \kappa_F(K)^{\frac{1}{6}} dK} \leq \frac{\kappa_H(-r)^{\frac{1}{6}} \int_{-\infty}^{-r} \kappa_F(K)^{\frac{1}{6}} dK}{\kappa_H(-r)^{\frac{1}{6}} \int_{-r}^0 \kappa_F(K)^{\frac{1}{6}} dK} = \frac{\int_{-\infty}^{-r} \kappa_F(K)^{\frac{1}{6}} dK}{\int_{-r}^0 \kappa_F(K)^{\frac{1}{6}} dK}.$$

We know $\kappa_F(K)^{\frac{1}{6}}$ is a $N(0, 3\sigma^2)$ density, so

$$\frac{\int_r^{+\infty} \kappa_F(K)^{\frac{1}{6}} dK}{\int_0^r \kappa_F(K)^{\frac{1}{6}} dK} = \frac{\int_{-\infty}^{-r} \kappa_F(K)^{\frac{1}{6}} dK}{\int_{-r}^0 \kappa_F(K)^{\frac{1}{6}} dK},$$

and hence, as $\sigma \to 0$,

$$0 \leq \frac{1-s(r)+s(-r)-0}{s(r)-s(-r)} \leq \frac{\int_r^{+\infty} \kappa_F(K)^{\frac{1}{6}} dK}{\int_0^r \kappa_F(K)^{\frac{1}{6}} dK} + \frac{\int_{-\infty}^{-r} \kappa_F(K)^{\frac{1}{6}} dK}{\int_{-r}^0 \kappa_F(K)^{\frac{1}{6}} dK} \to 0.$$

This means $s(r) - s(-r) \to 1$ and hence proves the result. □