

UCLA

UCLA Electronic Theses and Dissertations

Title

Essays in Information Economics

Permalink

<https://escholarship.org/uc/item/4jp9q4x3>

ISBN

9798247935254

Author

Zhai, Jiayin

Publication Date

2026-05-31

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Essays in Information Economics

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Economics

by

Jiayin Zhai

ABSTRACT OF THE DISSERTATION

Essays in Information Economics

by

Jiayin Zhai

Doctor of Philosophy in Economics

University of California, Los Angeles, 2026

Professor Moritz Meyer-ter-Vehn, Chair

This dissertation comprises three chapters that study how learning is shaped by attention constraints, peer pressure in social networks, and the information provider’s strategic incentives.

The first chapter studies how the incentive to prolong engagement affects the quality of information provided to a learning agent. In many advisory and digital environments, information providers benefit from keeping consumers engaged rather than from the actions consumers ultimately take. We develop a continuous-time principal-agent model in which the principal chooses the precision of a drift-diffusion signal process, and the agent decides when to stop learning and take an action. The main result is that optimal precision is U-shaped in the agent’s prior belief, reaching its minimum when the agent is most uncertain. The principal provides the least precise information to the agent with the highest underlying demand for information, because that agent’s willingness to remain engaged is greatest, giving the principal the most scope to dilute signal quality without triggering early exit. Despite receiving the worst information, the most uncertain agent also captures the largest consumer surplus. We develop two screening extensions that deliver a sharp contrast: when the agent’s prior is private information, all types rank contracts identically and any menu collapses to a single pooling contract; when the agent’s attention cost is private, single-

crossing holds and the principal can implement a fully separating menu with standard second-degree price discrimination features.

The second chapter studies how information-production costs shape persuasion in a dynamic setting. We study a continuous-time model in which the sender chooses the precision of a drift-diffusion signal process and the receiver decides when to stop learning and take an action. The sender directly benefits when the receiver takes a particular action. When information production is costless, the sender chooses maximum precision, and in the high-precision limit the receiver's payoff converges to the full-information benchmark. When information involves a fixed cost, the sender's decision is distorted on the extensive margin: she provides information only when the prior is sufficiently close to the receiver's indifference belief. When the marginal cost of precision is increasing, the distortion extends to the intensive margin: near the indifference belief, the receiver is relatively easy to persuade, so the sender reduces precision to save on production costs. Together, the two chapters show that the same continuous-time framework with the same instrument: signal precision, generates qualitatively different information distortions depending on the principal's objective. In both chapters, precision is low near the agent's indifference belief, but for opposite reasons: in the consultation model, the most uncertain agent is the most valuable to keep engaged, so the principal withholds precision to prolong learning; in the costly persuasion model, the most uncertain receiver is the easiest to persuade, so the sender economizes on precision to save on production costs. The two objectives also reverse the welfare implications: in the consultation model, the most uncertain agent captures the largest surplus despite receiving the worst information; in the persuasion model, the receiver near indifference may obtain a payoff strictly below the full-information benchmark precisely because the sender reduces information quality. The shape of optimal precision further reflects this contrast: U-shaped and symmetric around indifference in the consultation model, and decreasing toward indifference within the persuasion region in the costly persuasion model.

The third chapter studies how peer pressure in social networks shapes the evolution of social norms. We develop a network model in which agents repeatedly observe the behavior

of their neighbors and update their own actions based on both private preferences and locally perceived social norms. The chapter highlights a discontinuity in the role of private preferences. When agents place any positive weight on their own preferred action, initial actions have no effect on long-run behavior: equilibrium actions become weighted averages of agents' true preferences, and consensus generally cannot be reached. In contrast, when agents place zero weight on private preferences and simply copy their perceived social norms, the model becomes a standard DeGroot process, and actions converge to a uniform social norm determined by initial actions and network centrality, while true preferences no longer enter the long-run outcome. This discontinuity shows that even a small degree of attachment to private preferences can prevent complete social conformity, while pure peer-pressure updating can generate stable norms that diverge from what agents privately prefer.

The dissertation of Jiayin Zhai is approved.

Jay Lu

Ichiro Obara

Simon Board

Moritz Meyer-ter-Vehn, Committee Chair

University of California, Los Angeles

2026

To curiosity, and to never giving up.

TABLE OF CONTENTS

List of Figures	x
Acknowledgments	xi
Vita	xiii
1 The Optimal Level of Consultation	1
1.1 Introduction	1
1.2 Related Literature	4
1.3 Model	7
1.4 Agent's Problem: Optimal Stopping	8
1.5 Principal's Problem and Optimal Consultation Level	12
1.6 Extensions	18
1.6.1 Cost and Revenue: Consulting versus Platform	18
1.6.2 Screening with Heterogeneous Priors	21
1.6.3 Optimal Screening with Private Attention Costs	24
1.7 Conclusion	28
1.A Appendix to Chapter 1	31
1.A.1 Proof of Lemma 1.1	31
1.A.2 Derivation of the HJB Equation of the Agent	31
1.A.3 Existence, Uniqueness, and Monotonicity of the Optimal Boundary	33
1.A.4 Proof of Theorem 1.2	35
1.A.5 Proof of Theorem 1.3	36
References	40

2	The Optimal Experimentation Level in Costly Persuasion	42
2.1	Introduction	42
2.2	Related Literature	44
2.3	Model	46
2.3.1	Players and Payoffs	46
2.3.2	Static Persuasion Benchmark	47
2.3.3	Timing and Information	48
2.4	Receiver's Problem	49
2.5	Sender's Problem	52
2.6	Conclusion	59
2.A	Appendix to Chapter 2	61
2.A.1	Proofs for the Receiver's Problem	61
2.A.2	Proof of Proposition 2.3	63
2.A.3	High-Precision Limit and Receiver Learning Costs	65
	References	68
3	Evolving Social Norms on Networks	69
3.1	Introduction	69
3.2	Related Literature	72
3.3	Model	74
3.4	Benchmark: Pure Social Updating	77
3.5	Preference-Weighted Updating	78
3.6	Homophily and Group-Level Outcomes	80
3.7	Conclusion	82
3.A	Appendix to Chapter 3	84

3.A.1	Proof of Proposition 3.1	84
3.A.2	Proof of Proposition 3.3	84
3.A.3	Proof of Proposition 3.4	85
3.A.4	Mean-Field Group Equations	85
3.A.5	Asymmetric Within-Type Link Probabilities	86
References		87

LIST OF FIGURES

1.1	Agent value function under fixed precision	11
1.2	Optimal precision, stopping boundary, and payoffs	15
2.1	Static Bayesian Persuasion benchmark	47
2.2	Receiver’s value from optimal learning.	51
2.3	Sender’s expected payoff under different information-cost specifications.	54
2.4	Optimal information precision under different information-cost specifications.	56
2.5	Receiver payoff under optimal information policy and different cost specifications.	57
3.1	A simulated two-type homophilous network	76

ACKNOWLEDGMENTS

I am deeply grateful to my committee members, Moritz Meyer-ter-Vehn, Jay Lu, Ichiro Obara, and Simon Board, for their guidance, encouragement, and support throughout my graduate study. I am especially grateful to my advisor, Moritz Meyer-ter-Vehn, for his generous time, patience, and care. His commitment to my growth, his intellectual guidance, and his sense of responsibility helped me keep going and continue showing up even when things were difficult. I am also deeply thankful to Jay Lu for his encouragement and support throughout my Ph.D. years.

I also thank the faculty members in the theory group at UCLA Economics, including Alex Bloedel, Tomasz Sadzik, and Daniel Clark, for their insightful feedback and generous support at many stages of my research. I have benefited greatly from conversations with faculty members, classmates, and seminar participants at UCLA, and I am grateful for their comments, encouragement, and intellectual companionship.

I also thank the staff of the Department of Economics, especially Chiara Paz and Alice Wang, for their help and support throughout the program.

Graduate school has been challenging in ways I could and could not have anticipated, and I am fortunate to have been surrounded by friends who made these years meaningful and bright: Victoria Wang, Huihuang Zhu, Jay Yoo, Minghao Pan, Chang He, Jingyu Fan, for their friendship and inspiration, for those endless conversations and shared laughter, for the constant encouragement, and for sharing all the big and small things in life.

I am also grateful to my dear friends Yun-Pu Yang, Yi-Ting Lee, Yiping Wang, and Leroy Bondhus, for being my emotional support humans: for taking countless walks with me, for three-hour phone calls, for sharing food and joy, for working with me on Zoom, and for always being there for me. I am also grateful to the many friends not named here who filled this journey with memories, laughter, and joy.

I also want to thank my cat, Maggie, who came into my life at a particularly difficult time and brought me comfort, companionship, and joy. She reminded me daily that resting

is an essential part of life, and she never once doubted my ability to finish — largely because she had no idea what I was doing.

Finally, I want to thank my parents for their love and support. Thank you to my father, who gave me the courage to dream, and to my mother, who has always been my haven. This dissertation would not have been possible without them.

VITA

2013–2017	B.A. Economics, Minor in Mathematics, Peking University
2017–2019	M.Sc. Economics, Toulouse School of Economics
2019–2021	M.A. Economics, University of California, Los Angeles
2021	Research Assistant to Professor Jay Lu, UCLA
2020–2026	Teaching Assistant and Course Assistant, UCLA
2022	Excellence in Teaching Award, Microeconomic Theory, UCLA Department of Economics
2023	Research Assistant to Professor Manisha Shah, UCLA
2024	Instructor, Intermediate Microeconomic Theory, UCLA

CHAPTER 1

The Optimal Level of Consultation

1.1 Introduction

Advisors, consultants, and algorithmic recommendation systems often provide information to decision makers who face uncertainty before eventually taking an action. While these information providers may sometimes generate biased or incomplete information due to misaligned preferences or lack of commitment, in many environments they are largely indifferent to which action the decision maker ultimately takes. Instead, their primary objective is often to secure and sustain the decision maker's attention: a search platform may not care if a user eventually undergo surgery or purchase a product after browsing, but it wants the user to stay longer on the platform so their advertisement slots can be more valuable on the other side; a consultant may similarly earn revenue from the duration of the consultation rather than from the client's final decision, so she may have the same incentives to prolong the consultation period to increase billable hours.

In these settings, revenue depends not only on the accuracy of the information they deliver, but also on the amount of time the consumer remains engaged, and the information sellers face a trade-off between information quality and sustaining engagement: if the information provider reveals too much too quickly, the decision maker may resolve his uncertainty prematurely and leave. If the provider reveals too little, the decision maker may conclude that continued engagement is not worthwhile and exit. The provider must therefore choose the informativeness of the information process to keep the decision maker willing to continue learning, without resolving his uncertainty too quickly.

This paper studies how an information provider optimally manages signal precision in such environments. We develop a continuous-time principal-agent model in which the agent must choose between two actions whose payoffs depend on an unknown binary state. Before acting, the agent may purchase information from the principal at a constant flow cost. Information arrives through a Gaussian drift-diffusion process whose precision is chosen by the principal at time zero and committed to in advance. The principal receives no payoff from the agent's final action; her revenue is proportional only to the time the agent spends learning. We investigate how the optimal precision depends on the agent's prior belief.

The main result is that optimal information precision is U-shaped in the agent's prior belief, reaching its minimum when the agent is most uncertain, precisely when his prior is closest to indifference between the two actions. The principal provides the least precise information to the agent who has the highest underlying demand for information.

The intuition inverts the natural logic of information demand. An agent near indifference has the most to gain from learning, and has the highest willingness to remain engaged is high, the principal and exploit this feature can dilute signal precision without pushing him away. An agent with a more extreme prior is closer to acting immediately, and the principal must provide sharper information just to keep him at the table. Engagement maximization therefore reverses the expected relationship: the clients who need information most receive the worst of it.

Despite providing the least precise signal, the principal earns the most revenue from the most uncertain agent. Although low precision shrinks the continuation region, the agent's belief starts from the center of the belief space and must diffuse slowly to either boundary. The expected engagement time is longest at $p_0 = 1/2$, and this dominates the effect of low precision on the width of the stopping region. A symmetric result holds on the agent's side: the agent's surplus, i.e. the net value of participating in the information market relative to acting immediately, is also single-peaked at $p_0 = 1/2$. The most uncertain agent is simultaneously the most exploited and the most valuable, to both parties.

The model separates the agent’s flow cost of continued learning c from the principal’s flow revenue from engagement r , allowing it to cover two distinct market environments. In the consulting interpretation ($r = c$), the billing rate is a pure transfer and turns out to be entirely neutral: changing c leaves both parties’ payoffs unchanged, with the principal adjusting precision proportionally to maintain the same stopping boundaries. The billing rate functions as a unit of account rather than a source of distortion. In the platform interpretation ($r \neq c$), the two parameters move independently. An improvement in monetization technology (higher r) raises the platform’s revenue without affecting the information structure at all, while an increase in users’ attention cost (higher c) forces the platform to raise content precision to maintain engagement, eroding the revenue-to-cost ratio.

The paper develops two screening extensions that together deliver a dichotomy about when price discrimination is feasible in dynamic information markets.

When the agent’s prior belief is private information, screening is impossible. Because the agent’s stopping rule depends on precision and price only through their ratio, all types rank contracts identically regardless of their prior. Any incentive-compatible menu collapses to a single pooling contract, and the principal can only choose the overall precision level to balance revenue against participation. The prior determines how long an agent stays, but not how he ranks what he is offered.

When the agent’s attention cost is private information as in the consulting interpretation where $r = c$, screening is fully feasible. A lower-cost agent finds learning cheaper, stays engaged longer, and therefore gains more from any increase in signal precision. Single-crossing holds cleanly, and the principal can implement a tiered menu: high billing rate paired with high precision for low-cost clients, and low billing rate paired with low precision for high-cost clients. The optimal menu exhibits the standard features of second-degree price discrimination: efficiency at the bottom, downward distortion for higher-cost types, and a participation cutoff above which very high-cost clients are excluded. This tiered structure maps directly onto observed consulting and advisory market practices, where senior professionals provide sharper advice at higher hourly rates than junior ones.

The feasibility of price discrimination thus depends critically on the source of heterogeneity. Prior heterogeneity shifts the agent’s location in belief space without changing how he values precision relative to price; cost heterogeneity changes the effective marginal cost of learning and therefore generates the wedge the principal needs to separate types.

The paper proceeds as follows. Section 1.2 discusses related literature. Section 1.3 introduces the model. Section 1.4 characterizes the agent’s optimal stopping rule. Section 1.5 solves the principal’s problem, derives the U-shaped optimal precision, and establishes that both the principal’s value and the agent’s surplus are single-peaked at $p_0 = 1/2$. Section 1.6 develops three extensions: the consulting versus platform comparative statics, the impossibility of screening by prior, and the optimal screening menu under private attention costs. Section 1.7 concludes.

1.2 Related Literature

This paper relates to several strands of literature. First, this paper contributes to the literature on optimal stopping and sequential experimentation. The agent’s problem in this paper is an optimal stopping problem in the tradition of Wald (1947). The classical framework has been analyzed under drift-diffusion learning by Moscarini and Smith (2001), Fudenberg et al. (2015) and Ke et al. (2016), and extended to environments with endogenous experimentation by Zhong (2022). The paper also builds on Morris and Strack (2019), who show that in continuous-time drift-diffusion models the agent’s expected stopping time is equivalent to the expected change in the log-likelihood ratio. This equivalence allows dynamic experimentation problems to be recast as posterior-based optimization problems. In my model, because the principal’s revenue is proportional to expected stopping time, choosing precision is equivalent to choosing stopping posteriors. Unlike Morris and Strack (2019), who take precision as exogenous, this paper endogenizes information intensity and studies how a strategic principal selects precision to shape stopping behavior. Gonçalves (2024) shows that decision makers facing a speed-accuracy tradeoff choose stopping times generated by a drift-diffusion optimal-stopping rule in laboratory settings, and the resulting

decision structure mirrors the one faced by the agent in my model. This paper extends this tradition by embedding the stopping problem in a strategic setting, where the principal controls the experimentation intensity and the agent decides when to stop.

This paper connects the optimal stopping tradition to a growing literature in which a principal strategically controls an information process to influence an agent’s behavior or extract attention. The most closely related papers study principals who design information flows specifically to maximize engagement. [Koh and Sanguanmoo \(2022\)](#) develop a general theory of attention capture using arbitrary information structures, establishing that any achievable stopping-time distribution can be implemented by a structure that keeps the agent’s beliefs on a single deterministic path of increasing uncertainty. [Hébert and Zhong \(2022\)](#) show that the optimal engagement-maximizing strategy takes the form of a “dilution” strategy, which is a compound Poisson process where rare, belief-shifting signals are interspersed with silence, they show that the platform extracts all agent surplus. [Saeedi et al. \(2024\)](#) find that when the principal and agent hold different priors, the principal exploits this by revealing information in an order that caters to the agent’s bias, and they identify a speed-quality tradeoff when agent beliefs are private. The present paper is complementary but distinct: by restricting attention to a Gaussian diffusion framework in which the principal’s sole instrument is the initial precision, it delivers sharper comparative statics with the U-shaped relationship between optimal precision and the agent’s prior, and a clear screening result, namely that private beliefs collapse any menu to a single pooling contract because all types rank precision levels identically.

More broadly, this paper relates to the dynamic persuasion literature, in which a sender controls the information process to influence a receiver’s ultimate action. [Che et al. \(2023\)](#) study a setting in which the sender must maintain the receiver’s attention because persuasion takes time under flow-cost constraints; attention retention there arises because the sender compensates the receiver for listening, whereas in the present paper it arises endogenously from the principal’s choice of precision. [Henry and Ottaviani \(2019\)](#) reinterpret the Wald problem as a strategic communication problem with state-dependent Gaussian drift but fix

precision throughout; the present paper treats precision as the primary instrument and asks how its optimal level varies with the agent’s prior. [Orlov et al. \(2020\)](#) study a setting where waiting is itself a payoff-relevant action in a real-options environment, in contrast to the present paper where waiting serves as the medium through which the principal extracts revenue. [Gentzkow and Kamenica \(2014\)](#) extend [Kamenica and Gentzkow \(2011\)](#)’s concavification approach to costly signals with entropy-reduction costs; in my model, although information is delivered continuously, the principal’s problem can likewise be represented as a choice over stopping posteriors, preserving the posterior-based optimization logic. [Persson \(2018\)](#) and [Bloedel and Segal \(2018\)](#) study persuasion under information overload and rational inattention, respectively, where complexity or cognitive costs impose endogenous constraints on what the receiver processes.

This paper also contributes to the literature on markets for information. [Eső and Szentes \(2007\)](#) study a monopolistic seller of information who provides imperfect signals about a project’s value to a buyer whose action is contractible. The present paper departs from this setting in two respects: the principal’s payoff depends only on engagement rather than the agent’s action, and the action is not contractible, which is what gives rise to the precision-distortion result. [Bergemann et al. \(2018\)](#) analyze a data intermediary selling targeted information to advertisers with heterogeneous beliefs, and [Lehrer and Wang \(2024\)](#) characterize an optimal upfront transfer scheme for selling a fixed Blackwell experiment for repeated sampling. A broader treatment of this literature appears in [Bergemann and Bonatti \(2019\)](#).

The screening extension connects to the literature on second-degree price discrimination and mechanism design with single-crossing preferences (e.g., [Mussa and Rosen \(1978\)](#); [Maskin and Riley \(1984\)](#)). In standard models, different types rank contracts differently, creating the familiar wedge between adjacent types that gives rise to distortion and rent extraction. In contrast, in the present model all agent types rank contracts according to the same normalized precision-to-cost ratio. This eliminates the usual local incentive compatibility tradeoff: there is no crossing of marginal valuations across types, and hence no scope for separation. Any menu therefore collapses to a single contract that maximizes expected

revenue against the distribution of priors. This pooling outcome differs from canonical failure modes in screening, such as multidimensional types or limited commitment, because it arises from the structure of information valuation itself. The result also connects to the literature on selling information (e.g., [Bergemann et al. \(2015\)](#)), but unlike those models, heterogeneity in priors here does not generate differentiated pricing power.

1.3 Model

Terminal Payoffs. A principal (“she”) sells costly information to an agent (“he”) who must decide when to stop learning and take one of two actions $a \in \{A, B\}$. The unknown state is $\theta \in \{H, L\}$, with common prior $p_0 = \Pr(\theta = H)$. Each action yields state-dependent payoffs:

$$\begin{aligned} u(A, H) &= -l < 0, & u(A, L) &= l > 0, \\ u(B, H) &= h > 0, & u(B, L) &= -h < 0. \end{aligned}$$

The agent’s expected payoffs at belief p are

$$u_A(p) = l(1 - 2p), \quad u_B(p) = h(2p - 1).$$

The agent is indifferent between A and B at $p = \frac{1}{2}$, where both expected payoffs are zero. The principal obtains zero terminal payoff regardless of the action taken.

Timing. Time flows continuously from $t = 0$. There is no discounting for either player. At time 0, the principal commits to an information structure. At each instant, the agent decides whether to continue observing information or to stop and take action A or B . Once an action is taken, the game ends.

Information. Before choosing an action, the agent can observe informative signals about the state. The signal process $\{x_t\}_{t \geq 0}$ is a Brownian motion with state-dependent drift μ^θ , where $\mu^H = 1$ and $\mu^L = -1$. The principal chooses the precision parameter $n > 0$ at time 0. Equivalently, the signal process satisfies

$$dx_t^\theta = \mu^\theta dt + \frac{1}{\sqrt{n}} dW_t, \quad (1.1)$$

where W_t is a standard Brownian motion. A higher n corresponds to a more precise signal process.

Information Cost and Revenue. While the agent listens to the principal, he incurs a flow cost $c > 0$ per unit of time. This cost captures the price of continued access to information, the opportunity cost of attention, or the disutility of remaining engaged. The principal receives flow revenue $r > 0$ per unit of time while the agent remains engaged. In a standard consultation-fee interpretation, the agent's payment is a pure transfer to the principal, so $r = c$. In advertising-supported or platform environments, however, the agent's attention cost and the principal's monetized revenue from engagement need not coincide. If the agent stops at time T , the agent's total flow cost is cT , while the principal's total revenue is rT .

1.4 Agent's Problem: Optimal Stopping

Belief Dynamics. Given n , the signal process $\{x_t\}_{t \geq 0}$ induces a diffusion belief process $\{p_t\}_{t \geq 0}$. As in [Moscarini and Smith \(2001\)](#), or originally by Theorem 9.1 of [Liptser and Shiryaev \(2013\)](#), the agent's posterior belief $p_t = \Pr_t(\theta = H)$ satisfies

$$dp_t = 2p_t(1 - p_t)\sqrt{n} d\bar{W}_t, \quad (1.2)$$

where $\bar{W}_t$ is the observation-adapted Brownian innovation process.

Optimal Stopping Problem. The agent maximizes his expected terminal payoff net of information costs. Denote by $U(p_0; n)$ the value of the agent's optimal stopping problem when his prior is p_0 and the signal precision is n . Given precision n , the agent chooses a stopping time T to maximize

$$U(p_0; n) = \sup_T \mathbb{E} \left[u(p_T) - \int_0^T c \, dt \, \middle| \, p_0 \right], \quad (1.3)$$

$$dp_t = 2p_t(1 - p_t)\sqrt{n} \, d\bar{W}_t,$$

where $u(p_T) = \max\{u_A(p_T), u_B(p_T)\}$ is the payoff from taking the optimal terminal action at posterior p_T .

Since all information is ex ante identical and there is no discounting, it is without loss of generality to assume that the agent acquires information continuously until he stops. The agent chooses action A for sufficiently low beliefs, action B for sufficiently high beliefs, and continues learning on an interval $(\underline{p}, \bar{p})$.

In the continuation region, the Hamilton–Jacobi–Bellman equation is

$$U(p; n) = -c \, dt + \mathbb{E}[U(p + dp; n)]. \quad (1.4)$$

Using the generator of the belief process, this yields

$$2np^2(1 - p)^2 U''(p; n) = c, \quad (1.5)$$

or equivalently

$$U''(p; n) = \frac{c}{2np^2(1 - p)^2}. \quad (1.6)$$

The general solution is

$$U(p; n) = \frac{c}{2n} L(p) + a_1 p + a_2, \quad (1.7)$$

$$L(p) = p \ln \frac{p}{1 - p} + (1 - p) \ln \frac{1 - p}{p}. \quad (1.8)$$

The function $L(p)$ is nonnegative, strictly convex on $(0, 1)$, and symmetric around $1/2$.

The stopping posteriors and the constants in the value function are pinned down by value matching and smooth pasting:

$$U(\underline{p}; n) = u_A(\underline{p}) = l(1 - 2\underline{p}), \quad U'(\underline{p}; n) = u'_A(\underline{p}) = -2l, \quad (1.9)$$

$$U(\bar{p}; n) = u_B(\bar{p}) = h(2\bar{p} - 1), \quad U'(\bar{p}; n) = u'_B(\bar{p}) = 2h. \quad (1.10)$$

Lemma 1.1 *If the agent is indifferent between the two actions at $p = 1/2$, then for any $h, l > 0$,*

$$\bar{p} = 1 - \underline{p}.$$

The result follows by subtracting a linear term from the value function, which transforms the asymmetric payoff magnitudes h and l into a symmetric stopping problem around $1/2$. The proof is in Appendix [1.A.1](#).

Using Lemma [1.1](#) and subtracting the two derivative conditions gives

$$U'(\bar{p}; n) - U'(\underline{p}; n) = 2h - (-2l) = 2(h + l).$$

Since $U'(p; n) = \frac{c}{2n}L'(p) + a_1$ and $L'(\underline{p}) = -L'(\bar{p})$, the stopping boundary satisfies

$$\frac{c}{n}L'(\bar{p}) = 2(h + l).$$

Proposition 1.1 *For signal precision n , agent flow cost c , and payoff magnitudes $h, l > 0$, the stopping boundaries satisfy*

$$L'(\bar{p}) = -L'(\underline{p}) = \frac{2n(h + l)}{c}, \quad (1.11)$$

with $\bar{p} = 1 - \underline{p}$.

Proposition 1.2 *Let $(\underline{p}, \bar{p})$ be defined by Proposition [1.1](#). Then:*

- as the agent's flow cost c decreases, the experimentation region widens;
- as signal precision n increases, the experimentation region widens;
- as the decision stakes $h + l$ increase, the experimentation region widens.

Proof 1.1 By Proposition 1.1,

$$L'(\bar{p}) = \frac{2n(h+l)}{c}.$$

On $(1/2, 1)$, the function L is strictly convex and symmetric around $1/2$, so L' is strictly increasing and satisfies $L'(1/2) = 0$. Hence $\bar{p}$ is increasing in n , increasing in $h + l$, and decreasing in c . Since $\underline{p} = 1 - \bar{p}$, the continuation interval $(\underline{p}, \bar{p})$ widens when $\bar{p}$ increases. This proves the stated comparative statics.




Figure 1.1: Agent value function under fixed precision

Notes: The red solid curves show the agent's value function under a fixed signal precision n . The black dotted lines show the expected payoff with no information. The green dashed line is the full information benchmark. The dotted vertical lines mark the stopping posteriors $\underline{p}$ and $\bar{p}$. Panel (a) considers the case in which both actions are risky. Panel (b) considers the case in which action A is safe and yields zero payoff, while action B is risky. In both cases, the continuation region is symmetric around the indifference belief.

For a fixed precision n , the stopping posteriors $\{\underline{p}, \bar{p}\}$ are independent of the prior p_0 . If $p_0 \in (\underline{p}, \bar{p})$, the agent continues learning until his belief exits the experimentation region. If $p_0 \notin (\underline{p}, \bar{p})$, the agent stops immediately and acts on his prior.

The stopping rule depends on n and c only through the normalized precision n/c . This observation is useful below: the agent's incentives are governed by precision per unit flow cost, while the principal's payoff is governed by the revenue she earns per unit of engagement.

1.5 Principal's Problem and Optimal Consultation Level

The principal observes the agent's prior p_0 and chooses the precision n to maximize expected engagement revenue. Anticipating the agent's optimal stopping rule, the principal solves

$$\begin{aligned} V(p_0) &= \sup_{n>0} r \mathbb{E}[T \mid p_0, n], \\ \text{s.t.} \quad T &\in \arg \max_{\tau} \mathbb{E}[u(p_{\tau}) - c\tau \mid p_0, n]. \end{aligned} \tag{1.12}$$

The agent's cost is governed by c , while the principal's revenue is governed by r .

[Morris and Strack \(2019\)](#) show the equivalence between Wald's sequential sampling cost and the log-likelihood cost. In the present setting,

$$\begin{aligned} \mathbb{E}[T \mid p_0, n] &= \frac{1}{2n} \mathbb{E}[L(p_T) - L(p_0) \mid p_0, n] \\ &= \frac{1}{2n} [L(\bar{p}(n)) - L(p_0)], \end{aligned} \tag{1.13}$$

where the last equality applies to $p_0 \geq 1/2$ and follows from the symmetry $\bar{p}(n) = 1 - \underline{p}(n)$ and the symmetry of L : $L(\underline{p}) = L(\bar{p})$. The analogous expression for $p_0 \leq 1/2$ uses the lower stopping boundary $\underline{p}(n)$.

This representation makes the principal's choice transparent. Using the boundary condition in (1.11), choosing a precision n is equivalent to choosing the induced stopping posterior $\bar{p}$. For $p_0 \geq 1/2$, the principal's problem reduces to:

$$\sup_{\bar{p} \in (p_0, 1)} \frac{r}{c} (h + l) \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}. \tag{1.14}$$

Equivalently, for $p_0 \leq 1/2$, the principal chooses a lower stopping posterior satisfying $\underline{p} < p_0$:

$$\sup_{\underline{p} \in (0, p_0)} -\frac{r}{c}(h+l) \frac{L(\underline{p}) - L(p_0)}{L'(\underline{p})}. \quad (1.15)$$

This representation shows that the principal's problem of choosing n based on p_0 is equivalent to targeting the posterior $\bar{p}(p_0)$ or $\underline{p}(p_0)$ at which the agent stops, balancing informativeness with length of engagement. We can trace how optimal $\bar{p}$ or $\underline{p}$ changes in p_0 , and use the monotonicity of $\bar{p}$ or $\underline{p}$ and n , further establishing how the optimal consultation level n changes in p_0 :

Theorem 1.1 *For each prior $p_0 \in (0, 1)$, the principal's problem admits a unique optimal consultation level $n^*(p_0)$. The function $n^*(p_0)$ is symmetric around $1/2$, strictly decreases on $(0, 1/2)$, strictly increases on $(1/2, 1)$, and reaches its unique minimum at $p_0 = 1/2$.*

Proof 1.2 *See Appendix [1.A.3](#).*

The principal provides the least precise information to the agent with the most unfocused prior belief. This result inverts the natural intuition about information demand: an agent near indifference has the strongest underlying motive to learn, as each signal is most valuable to him. Here the agent with the highest underlying demand for information is also the agent whose has the highest willingness to remain engaged, and the principal exploits precisely this. Because the uncertain agent's willingness to remain engaged is high, the principal can dilute precision without triggering early exit. As the prior moves away from $1/2$, the agent is more confident and closer to taking an action. To keep such an agent engaged, the principal must provide a more precise signal. Engagement maximization thus reverses the expected relationship: the clients most in need of information receive the worst of it.

To see the mechanism in terms of the principal's optimization, recall that the principal's payoff from a given stopping posterior $\bar{p}$ is proportional to $[L(\bar{p}) - L(p_0)]/L'(\bar{p})$. The numerator is the informational distance the agent must travel; the denominator, through the boundary condition, reflects the cost of that precision. For an agent near $p_0 = 1/2$, the


numerator is large and the marginal cost of sustaining engagement is low, so the principal can afford to set a wide continuation region and extract a long expected stopping time with relatively little precision. For an agent near $p_0 = 0$ or $p_0 = 1$, the numerator is small and the agent is close to acting immediately, so the principal must offer a tight, highly precise signal just to keep him at the table.

The proof sketch is as follows. First we focus on the case when $p_0 \in [1/2, 1)$, and let the principal choose the optimal consultation level n to target at the belief upper bound $\bar{p}$. We showed earlier that the principal's problem is essentially choosing $\bar{p}$ to maximize the expected engagement, and we want to know how the optimal $\bar{p}$ changes in p_0 . We can use either implicit function theorem, or Topkis' theorem as in [Milgrom and Shannon \(1994\)](#) to show that $\bar{p}^*$ strictly increases in p_0 for $p_0 \in [1/2, 1)$. Then since $\bar{p}$ strictly increases in n , we can show that $n^*(p_0)$ strictly increases in p_0 for $p_0 \in [1/2, 1)$. When $p_0 \in (0, 1/2)$, we can focus on the optimization problem in terms of $\underline{p}$. Similarly, we can show that $\underline{p}^*$ strictly increases in p_0 for $p_0 \in (0, 1/2)$, and because $\underline{p}$ strictly decreases in n , it follows that $n^*(p_0)$ strictly decreases in p_0 for $p_0 \in (0, 1/2)$. Therefore, the optimal $n^*(p_0)$ reaches the lowest value when $p_0 = 1/2$.


Here the principal earns revenue equal to the flow fee r multiplied by the expected duration of learning. We know that the agent with more unfocused prior belief gets lower information precision, and the experimentation region is the smallest. Although for the same experimentation region, it takes longer for agents with more unfocused prior to reach the stopping boundaries; with the personalized information quality $n^*(p_0)$ and the resulting experimentation region $(\underline{p}(p_0), \bar{p}(p_0))$, it's unclear how the expected learning time, as well as the principal's expected revenue changes with p_0 . Using the envelope theorem, we obtain the following result:

Theorem 1.2 *Under the optimal consultation level $n^*(p_0)$, the principal's value function $V(p_0)$ is symmetric around $1/2$ and strictly single-peaked, attaining its unique maximum at $p_0 = 1/2$.*


Proof 1.3 *See Appendix [1.A.4](#).*




(a) Optimal Consultation Level n^*



(b) Optimal Stopping Boundary $\bar{p}^*$



(c) Principal's Value Function



(d) Agent's Value Function

Figure 1.2: Optimal precision, stopping boundary, and payoffs

Notes: The figure summarizes the principal's optimal information policy and the induced payoffs as functions of the agent's prior p_0 . Panel (a) shows that optimal precision is U-shaped and reaches its minimum at the indifference prior $p_0 = 1/2$. Panel (b) shows the corresponding upper stopping boundary. Panel (c) shows that the principal's value is single-peaked at $p_0 = 1/2$. Panel (d) shows the agent's value under the principal's optimal policy, and the red dashed line shows the agent's value with no information. 1.3 shows that the agent's surplus is maximized at $p_0 = 1/2$, i.e. the distance between the two curves.

Although the principal offers the least precise signal at $p_0 = 1/2$, she earns the most from this agent. Low precision shrinks the continuation region, but the agent starts from the center of the belief space and must diffuse slowly to either boundary before stopping. We find the expected time to exit is highest when the prior is most central, even under low precision. The principal exploits the uncertain agent most aggressively and profits the most from him.

Together, the two theorems describe the principal's complete strategy. Precision is set low to extract engagement from the uncertain agent; the agent's expected stopping time is longest where precision is lowest; and the most uncertain agent is simultaneously the most exploited and the most valuable. The separation between c and r clarifies what drives each side of this picture. The agent's stopping rule is governed entirely by c . The principal's payoff is governed by r . In a pure consultation-fee setting these coincide, but the model applies equally to advertising-supported platforms and advisory markets where they need not. The ratio r/c scales the principal's revenue and leaves the U-shaped precision profile intact: the shape of the distortion is a property of the engagement motive itself, not of how attention is monetized.

Define the agent's surplus as the value of participating in the information market relative to acting immediately on the prior:

$$S(p_0) := U(p_0; n^*(p_0)) - u(p_0), \quad (1.16)$$

where $u(p_0) = \max\{u_A(p_0), u_B(p_0)\}$ is the agent's payoff from stopping without any information. The surplus $S(p_0)$ measures how much the agent benefits from entering the information market, net of all flow costs paid during learning.

Theorem 1.3 *Under the optimal consultation level $n^*(p_0)$, the agent's surplus $S(p_0)$ is symmetric around $1/2$ and strictly single-peaked, attaining its unique maximum at $p_0 = 1/2$. Moreover, $S(p_0) \rightarrow 0$ as $p_0 \rightarrow 0$ or $p_0 \rightarrow 1$.*

Proof 1.4 *See Appendix [1.A.5](#).*

Theorem 1.3 mirrors Theorem 1.2 on the agent's side of the market. Both payoffs are single-peaked at $p_0 = 1/2$, yet for distinct reasons. The principal earns the most from the uncertain agent because his belief takes the longest to reach either stopping boundary, generating the greatest expected engagement time. The agent gains the most because his prior is farthest from any action threshold, so the option to learn and act on better information is most valuable exactly when uncertainty is highest.

The two results together paint a coherent picture of the information market at the optimal policy. At $p_0 = 1/2$, the principal sets the lowest precision and the agent pays the lowest per-unit-time learning cost, yet both players reach their peak payoffs. The mechanism is the same for each: the agent's belief has the most ground to cover before stopping, and this translates into the longest engagement spell on the principal's side and the largest option value of learning on the agent's side. Diluted precision does not undo the agent's gain because even a slow signal is worth waiting for when uncertainty is greatest.

At the extremes, $S(p_0) \rightarrow 0$ for a simple reason. As p_0 approaches 0 or 1, the agent is already close to the optimal action without any information. The principal responds by raising precision to retain him at all, but the continuation region narrows and the agent stops almost immediately, paying negligible flow cost and gaining negligible informational value. The net surplus vanishes.

This result also has a market interpretation. Even though the principal provides the worst information quality to the most uncertain agent, that agent still captures the largest consumer surplus. In this sense, the exploitation and the benefit are two sides of the same force: high willingness to remain engaged allows the principal to dilute quality, but it also ensures the agent derives substantial value from what he learns.

1.6 Extensions

1.6.1 Cost and Revenue: Consulting versus Platform

The model separates two parameters that are often conflated: the agent's flow cost of continued learning c , and the principal's flow revenue from engagement r . How these two relate to each other determines the economic interpretation of the model and the comparative statics of the equilibrium. We develop two interpretations in turn.

Scenario 1: Consulting ($r = c$). In a consultation-fee setting, the agent's flow payment is a direct transfer to the principal. The billing rate is both what the agent pays and what the principal receives, so $r = c$. The ratio $r/c = 1$ and the principal's objective reduces to

$$(h + l) \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})},$$

which depends on neither c nor r . The optimal stopping posterior $\bar{p}^*(p_0)$ is therefore determined entirely by the belief space and the payoff stakes $h + l$, with no role for the billing rate.

This delivers the following comparative statics. Suppose the billing rate c changes, perhaps because the client's outside option improves, or because market competition shifts the prevailing consultation fee.

Proposition 1.3 *In the consulting interpretation ($r = c$), for fixed p_0 , the following hold as c changes:*

1. *The optimal stopping posterior $\bar{p}^*(p_0)$ is unchanged.*
2. *The optimal precision $n^*(p_0)$ scales proportionally with c : $n^*(p_0) \propto c$.*
3. *The agent's surplus $S(p_0)$ is unchanged.*
4. *The principal's value $V(p_0)$ is unchanged.*

Proof 1.5 *The optimal stopping posterior satisfies $R(\bar{p}^*) = L(p_0)$, which is independent of c , so $\bar{p}^*$ is invariant. Since $L'(\bar{p}^*) = 2n^*(h+l)/c$ and $\bar{p}^*$ is fixed, n^* scales proportionally with c . The agent's surplus from Theorem 1.3 is*

$$S(p_0) = (h+l) \left[\frac{L(\bar{p}^*) - L(p_0)}{L'(\bar{p}^*)} + \bar{p}^* - p_0 \right],$$

which depends on c only through $\bar{p}^$; since $\bar{p}^*$ is invariant, so is $S(p_0)$. Finally, $V(p_0) = (r/c)(h+l)[L(\bar{p}^*) - L(p_0)]/L'(\bar{p}^*) = (h+l)[L(\bar{p}^*) - L(p_0)]/L'(\bar{p}^*)$, which is also invariant to c .*

The last two bullet points together say that in the consulting interpretation, the billing rate is completely neutral: changing c affects neither party's payoff. The only equilibrium response is a proportional adjustment in precision. When the billing rate rises, the principal delivers sharper information to maintain the same continuation region, shortening the agent's engagement and keeping the principal's total revenue unchanged. The billing rate functions only scales the flow of payments between the two parties without affecting the underlying economics of the interaction.

This neutrality result has a practical implication for the screening extension in Section 1.6.3. When $r = c$, the billing rate is both the instrument of screening and the source of distortion. The principal uses the billing rate to sort clients by their private attention cost, and the neutrality of c at the aggregate level does not carry over to the cross-sectional menu: different billing rates induce different precisions and hence different continuation regions, which is exactly what generates the separating equilibrium.

Scenario 2: Platform ($r \neq c$). In a platform or advertising-supported setting, the agent's cost of attention and the principal's revenue from engagement are distinct. The agent incurs attention cost c , which could be the opportunity cost of time spent on the platform, or the disutility of processing information, while the principal monetizes engagement at rate r through advertising, data collection, or subscription fees. The two parameters move independently.

The optimal stopping posterior remains $\bar{p}^*(p_0)$, determined by $R(\bar{p}^*) = L(p_0)$ independently of both c and r . The platform's equilibrium is therefore characterized by the following comparative statics.

Proposition 1.4 *In the platform interpretation ($r \neq c$), for fixed p_0 , the following hold:*

1. **Change in c (attention cost):** $n^*(p_0)$ scales proportionally with c ; $S(p_0)$ is unchanged; $V(p_0)$ scales with r/c and therefore decreases.
2. **Change in r (monetization rate):** $n^*(p_0)$ is unchanged; $S(p_0)$ is unchanged; $V(p_0)$ scales with r/c and therefore increases.
3. **Proportional change in both:** if r and c change by the same factor, $n^*(p_0)$ scales with c , $S(p_0)$ is unchanged, and $V(p_0)$ is unchanged.

Proof 1.6 *All three results follow from $\bar{p}^*$ being invariant to c and r , $n^* \propto c$ from the boundary condition, $S(p_0)$ depending only on $\bar{p}^*$, and $V(p_0) = (r/c)(h + l)[L(\bar{p}^*) - L(p_0)]/L'(\bar{p}^*)$ scaling with r/c .*

The platform results reveal a clean separation between two sources of value. The information structure, the stopping boundaries, the agent's continuation region, and the distribution of engagement periods, depends only on c . The monetization rate r scales the principal's revenue without touching the information structure at all. A platform that improves its advertising technology (higher r) earns more from the same engagement behavior without adjusting content quality. A platform facing more distracted users (higher c) must raise content precision to maintain the same engagement spell, but earns less per unit of attention as the ratio r/c falls.

Comparison. The consulting and platform interpretations differ in one fundamental respect. In the consulting case, $r = c$ and the billing rate is neutral: it adjusts without consequence because higher fees are exactly offset by higher precision, leaving engagement

and payoffs unchanged. In the platform case, r and c are decoupled, and the platform faces genuine tradeoffs: a higher attention cost c erodes the ratio r/c and reduces platform revenue, while an improvement in monetization r raises revenue without distorting the information environment. The consulting model is appropriate when the principal's revenue derives solely from the client's fee; the platform model is appropriate when the principal monetizes attention through a channel that is independent of the agent's direct cost of engagement, such as advertising, data, subscriptions.

1.6.2 Screening with Heterogeneous Priors

We now consider a setting in which the agent's prior belief p_0 is private information. The principal observes only a known distribution F over types $p_0 \in (0, 1)$ and may offer a menu of contracts to elicit the agent's type.

Contracts and menus. A contract specifies a precision level and a flow price. Denote a contract by (n, r) , where $r > 0$ is the flow price the agent pays per unit of time while learning. Let the menu be

$$\mathcal{M} = \{(n(p), r(p))\}_{p \in (0,1)}.$$

A type- p_0 agent selects a contract from $\mathcal{M}$ and then solves the optimal stopping problem induced by that contract. Under contract (n, r) , the agent's value is

$$U(p_0; n, r) = \sup_T \mathbb{E} \left[u(p_T) - \int_0^T r \, dt \mid p_0, n \right].$$

The principal's revenue from a type- p_0 agent who selects contract (n, r) is $r \mathbb{E}[T \mid p_0, n]$.

The key observation is that the agent's stopping problem depends on n and r only through the ratio $\eta = n/r$. To see this, note that the boundary condition from Proposition 1.1 becomes

$$L'(\bar{p}) = -L'(\underline{p}) = 2\eta(h + l),$$

so the stopping boundaries, and hence the agent's entire continuation strategy, are pinned down by η alone. A higher η means more precision per unit flow price: the agent learns faster relative to what he pays, the continuation region widens, and he stays engaged longer. The agent's value $U(p_0; n, r)$ is therefore strictly increasing in η for every prior p_0 .

This single-dimensional sufficiency is what destroys the scope for screening. In a standard second-degree price discrimination problem, different types rank bundles differently: a high-demand type is willing to pay more for quality, while a low-demand type is not, and this creates the wedge the principal exploits to separate them.

Here, however, all types agree on the ranking of contracts: any type strictly prefers a contract with a higher η to one with a lower η , regardless of his prior. Formally, for any two contracts (n_i, r_i) and (n_j, r_j) :

$$\frac{n_i}{r_i} > \frac{n_j}{r_j} \implies U(p_0; n_i, r_i) \geq U(p_0; n_j, r_j) \quad \forall p_0 \in (0, 1).$$

There is no crossing of preferences across types: all agents want the same contract, and no menu can exploit type-specific differences in willingness to pay.

Corollary 1.1 *Let the principal offer any menu of contracts $\mathcal{M} = \{(n(p), r(p))\}$ to agents with private priors $p_0 \in (0, 1)$. Then:*

1. *all types $p_0 \in (0, 1)$ weakly prefer the contract with the highest normalized precision $\eta = n/r$;*
2. *no incentive-compatible menu can induce different participating types to choose different normalized precision levels;*
3. *any incentive-compatible menu is payoff-equivalent, from the agents' perspective, to a single pooling contract in normalized precision.*

The proof follows directly from the observations above. If any incentive-compatible menu contained two contracts with $\eta_i > \eta_j$, all types would strictly prefer contract i , violating

incentive compatibility. Hence all participating types must be offered the same η , and the menu collapses to a single pooling contract.

The source of this failure differs from the usual reasons screening breaks down. In standard models with multidimensional types, limited commitment, or ex-post participation constraints, screening fails because the principal's instruments are insufficient to separate types who have different preferences over contracts. Here, the failure is more fundamental: different priors do not generate different preferences over contracts at all. Although a type near $p_0 = 1/2$ will remain engaged longer under any given contract, and thus generate more revenue for the principal, he does not value a high- η contract any differently from a type near $p_0 = 0$ or 1 . The prior determines how long the agent stays, but not how he ranks what he is offered. Because the ranking is common across types, there is no local incentive compatibility tradeoff to exploit, and no distortion can be used to relax a binding incentive constraint.

Since any incentive-compatible mechanism reduces to a single normalized precision η^* , the principal's problem simplifies to choosing η^* to maximize expected revenue against the distribution F , subject to each type's individual rationality constraint. The participation constraint binds for types near $p_0 = 0$ or 1 : these agents have little underlying demand for information, and even a modest flow price r may not justify continued learning. As η^* decreases (lower precision per unit price), the participation region shrinks further toward the center of the type space. The principal therefore faces a classic extensive-margin tradeoff: raising η^* expands participation but tightens the revenue extracted per engaged agent, while lowering η^* concentrates revenue on the most uncertain agents who remain willing to pay.

The optimal contract thus targets agents with intermediate priors. Those with sufficiently extreme beliefs are already close to acting without information, and will find the contract unattractive and opt out. The pooling outcome means the principal cannot tailor precision to the agent's belief even when she knows the distribution F : the same contract is offered to all, and participation self-selects on uncertainty.

This result clarifies the limits of personalization in information markets. A legal advisor who cannot observe how confident a client already is cannot offer a lower billing rate per hour for a more efficient service to the more uncertain client, because all clients would prefer the more precise option regardless. A digital platform that cannot contractually condition its information density on users' prior beliefs will find that any attempt to offer differentiated information products collapses: more targeted recommendations for a higher subscription fee, etc., since every user prefers the higher-quality product. The only tool the principal retains is the overall level of the normalized precision, which she sets to balance revenue against the cost of excluding agents with extreme priors.

1.6.3 Optimal Screening with Private Attention Costs

We now return to the consultation-fee interpretation, in which the agent's flow cost of attention and the principal's flow revenue coincide: $r = c$. The agent's private type is his flow cost of attention c , distributed according to a continuously differentiable distribution G on $[\underline{c}, \bar{c}]$ with density $g > 0$. The prior $p_0 \in (0, 1)$ is commonly known. The principal does not observe c directly but may offer a menu of consultation tiers to elicit it.

A consultation tier specifies a precision level and a billing rate. Because $r = c$ in the consultation-fee interpretation, the natural instrument is a menu of pairs (n, c) , where c is simultaneously the billing rate charged to the agent and the agent's flow cost of continued engagement. Formally, the principal offers a menu

$$\mathcal{M} = \{(n(\tilde{c}), \tilde{c})\}_{\tilde{c} \in [\underline{c}, \bar{c}]},$$

and a type- c agent selects a tier $\tilde{c}$ from $\mathcal{M}$. Under the selected tier $(n(\tilde{c}), \tilde{c})$, the agent faces billing rate $\tilde{c}$ per unit of time but his stopping rule is governed by his private cost c . Specifically, he solves

$$U(p_0; n(\tilde{c}), c) = \sup_T \mathbb{E} \left[u(p_T) - \int_0^T c \, dt \, \middle| \, p_0, n(\tilde{c}) \right],$$

and his net payoff from selecting tier $\tilde{c}$ is

$$V(c, \tilde{c}) := U(p_0; n(\tilde{c}), c) - \tilde{c} \cdot \mathbb{E}[T \mid p_0, n(\tilde{c}), c],$$

where the second term reflects the total billing cost at rate $\tilde{c}$ over the agent's engagement spell. The principal's revenue from a type- c agent who selects tier $\tilde{c}$ is $\tilde{c} \cdot \mathbb{E}[T \mid p_0, n(\tilde{c}), c]$.

The key difference from Section 1.6.2 is that c enters the agent's problem in two distinct ways: through the stopping rule (via $\eta = n/c$) and through the billing cost (via $\tilde{c} \cdot \mathbb{E}[T]$). These two channels have opposing effects on the agent's preference for precision, but the first dominates.

To see this, consider the agent's marginal rate of substitution between precision n and billing rate $\tilde{c}$ at a truthful report $\tilde{c} = c$:

$$\text{MRS}(c) := \frac{\partial U / \partial n}{\mathbb{E}[T \mid p_0, n, c]} = \frac{(1/c) \cdot \partial U / \partial \eta}{\mathbb{E}[T \mid p_0, n, c]}.$$

This measures how much additional billing rate the agent would accept in exchange for a unit increase in precision, keeping his net payoff constant. The MRS is strictly decreasing in c : a lower-cost agent gains more from any increase in precision (through the $1/c$ factor in the numerator) while his longer expected engagement time amplifies the gain relative to the billing cost in the denominator.

Lemma 1.2 *For any two tiers with $n_1 > n_2$ and $c_1 < c_2$, the net gain from choosing tier 1 over tier 2,*

$$\Delta(c) := V(c, c_1) - V(c, c_2) = [U(p_0; n_1, c) - U(p_0; n_2, c)] - [c_1 \mathbb{E}[T \mid n_1, c] - c_2 \mathbb{E}[T \mid n_2, c]],$$

is strictly decreasing in c . Hence lower- c agents have strictly higher marginal willingness to pay for precision, and the single-crossing property holds with the ordering: lower c is the high type.

Proof 1.7 Differentiate $\Delta(c)$ with respect to c , holding the tier choices fixed. Using $\eta_k = n_k/c$ for $k = 1, 2$:

$$\begin{aligned}\frac{d}{dc}[U(p_0; n_k, c)] &= \frac{\partial U}{\partial \eta} \cdot \frac{\partial \eta_k}{\partial c} = -\frac{n_k}{c^2} \frac{\partial U}{\partial \eta} \Big|_{\eta_k}, \\ \frac{d}{dc}[c_k \mathbb{E}[T \mid n_k, c]] &= c_k \frac{d \mathbb{E}[T]}{dc} = -\frac{c_k n_k}{c^2} \frac{d \mathbb{E}[T]}{d\eta} \Big|_{\eta_k}.\end{aligned}$$

Therefore

$$\Delta'(c) = -\frac{1}{c^2} \sum_{k=1}^2 (-1)^{k+1} \left[\frac{\partial U}{\partial \eta} - c_k \frac{d \mathbb{E}[T]}{d\eta} \right]_{\eta_k}.$$

At a truthful report, the agent's first-order condition for stopping gives $\partial U / \partial \eta > c_k d \mathbb{E}[T] / d\eta$ when precision is below the efficient level, and the inequality is strict and larger for lower η (higher c). Since $\eta_1 > \eta_2$, the bracketed term is larger at η_2 than at η_1 , making $\Delta'(c) < 0$.

By Lemma 1.2, the screening problem has the standard structure of a quality-price discrimination model (Mussa and Rosen, 1978). We parameterize the menu by the precision schedule $n(c)$ and appeal to the standard approach of replacing global incentive compatibility with local incentive compatibility plus monotonicity.

A type- c agent who reports $\tilde{c}$ obtains $V(c, \tilde{c})$. Incentive compatibility requires truthful reporting to be optimal. By the envelope theorem, the agent's equilibrium rent $R(c) := V(c, c)$ satisfies

$$R'(c) = \frac{\partial V(c, \tilde{c})}{\partial c} \Big|_{\tilde{c}=c} = \frac{\partial U}{\partial c} \Big|_{n(c)} - \mathbb{E}[T \mid p_0, n(c), c] < 0,$$

where the inequality follows because both the value gain from lower cost ($\partial U / \partial c < 0$ since lower c means lower flow cost and higher η , so this term contributes positively to $-R'(c)$) and the billing saving ($-\mathbb{E}[T] < 0$ contributes positively to $-R'(c)$) work in the same direction. Rents are strictly decreasing in c : the lowest-cost agent earns the highest rent.

Integrating by parts in the standard way, the principal's expected revenue can be written as

$$\int_{\underline{c}}^{\bar{c}} \left[c \mathbb{E}[T \mid p_0, n(c), c] + \frac{G(c)}{g(c)} R'(c) \right] g(c) dc - R(\bar{c}),$$

where $R(\bar{c}) = 0$ at the optimum (the individual rationality constraint binds for the highest-cost participating type). Substituting $R'(c)$ and maximizing pointwise yields the following characterization.

Proposition 1.5 *Suppose G satisfies the standard regularity condition that $c + G(c)/g(c)$ is strictly increasing. Then the principal's optimal menu $\mathcal{M}^* = \{(n^*(c), c)\}$ has the following properties:*

1. *Efficiency at the bottom: the lowest-cost type $\underline{c}$ receives the efficient precision $n^*(\underline{c})$ that maximizes total surplus at that type.*
2. *Downward distortion: all higher-cost types $c > \underline{c}$ receive a precision strictly below the efficient level. The precision schedule $n^*(c)$ is strictly decreasing in c .*
3. *Participation cutoff: there exists a threshold $\tilde{c} \leq \bar{c}$ such that types with $c > \tilde{c}$ are excluded and offered no contract. The cutoff satisfies $R(\tilde{c}) = 0$.*
4. *Rent gradient: the equilibrium rent $R(c)$ is strictly decreasing in c , with $R(\tilde{c}) = 0$ at the participation cutoff.*

Proof 1.8 *Pointwise maximization of the virtual surplus integrand over $n(c)$ yields the first-order condition*

$$\frac{\partial}{\partial n} \left[c \mathbb{E}[T] + \frac{G(c)}{g(c)} R'(c) \right] = 0.$$

Since $R'(c) < 0$ and $G(c)/g(c) > 0$, the virtual surplus includes a negative correction to the social optimum, implying the optimal $n^(c)$ is below the efficient level for all $c > \underline{c}$. Monotonicity $n^{*'}(c) < 0$ follows from the regularity assumption. The participation cutoff $\tilde{c}$ is determined by the condition that the virtual surplus is non-negative; types above $\tilde{c}$ are excluded because the rent required to serve them exceeds the revenue they generate.*

The optimal menu has the structure of a tiered consulting schedule. The principal offers a ladder of service tiers: high billing rate paired with high precision for low-cost clients, and low billing rate paired with low precision for high-cost clients. Low-cost clients, i.e. those who find deliberation and continued engagement cheap, are willing to pay high billing rates because the value of precise advice is large relative to their cost of staying. High-cost clients, those with high opportunity costs of time, prefer cheaper, lower-quality tiers because their engagement is brief regardless of precision.

The downward distortion in precision for high-cost types serves the standard rent-extraction motive: by degrading service quality for high-cost clients, the principal makes it unattractive for low-cost clients to mimic them, allowing the principal to charge low-cost clients a higher billing rate without losing them to a cheaper tier. The participation cutoff reflects the extensive margin: very high-cost clients generate so little engagement revenue that even at a low billing rate they are not worth serving, and the principal excludes them entirely.

This tiered structure maps naturally onto observed market practices. Law firms and consulting firms offer senior partners at high hourly rates and junior associates at lower rates, with more experienced professionals providing sharper and more precise advice. Investment banks offer research tiers of varying depth and frequency at different price points. The model provides a micro-foundation for this tiering directly from the engagement motive and the private cost of client attention, without requiring any exogenous quality differences across service providers.

1.7 Conclusion

This paper studies how an information seller optimally chooses the precision of a continuous-time signal when her revenue depends on how long the agent remains engaged rather than on which action he ultimately takes. The analysis delivers three sets of results.

The central result is that optimal precision is U-shaped in the agent’s prior belief, reaching its minimum at the indifference prior $p_0 = 1/2$. The principal provides the least precise information to the most uncertain agent, also the one with the highest underlying demand for information, because that agent’s willingness to remain engaged is greatest and gives the principal the most scope to slow down learning.

We also find that distortion does not harm the uncertain agent in net terms. The agent’s surplus is also single-peaked at $p_0 = 1/2$: despite receiving the lowest quality information, the most uncertain agent captures the largest gain from participating in the information market. His prior is farthest from any action threshold, so the option value of learning is highest exactly where uncertainty is greatest. The principal’s value and the agent’s surplus therefore peak at the same prior, for the same underlying reason: starting from the center of the belief space, the agent’s belief must travel the long to either stopping boundary slowly, generating the greatest engagement time and the greatest informational gain simultaneously. The exploitation and the benefit are two sides of the same force.

The model separates the agent’s flow cost c from the principal’s flow revenue r , and the two market interpretations differ sharply in their comparative statics. In the consulting case ($r = c$), the billing rate is entirely neutral: a change in c is fully absorbed by a proportional adjustment in precision, leaving both the agent’s surplus and the principal’s value unchanged. The billing rate is a unit of account, not a source of distortion. In the platform case ($r \neq c$), the two parameters are decoupled. An improvement in monetization technology raises platform revenue without touching the information structure, while an increase in users’ attention cost forces the platform to raise content quality to maintain engagement, eroding the revenue-to-cost ratio.

The two screening extensions deliver a contrast result about when price discrimination is feasible in dynamic information markets. When the agent’s prior is private, all types rank contracts identically because the prior affects engagement duration but not the marginal value of precision. No menu can separate them, and the principal is forced to a single pooling contract. When the agent’s attention cost is private, single-crossing holds because

lower-cost types gain more from precision and are willing to pay higher billing rates to obtain it. The principal can implement a fully separating menu, with the standard features of second-degree price discrimination: efficiency at the bottom, downward distortion for higher-cost types, and exclusion of the highest-cost types.

The source of the difference is the role each form of heterogeneity plays in the agent's stopping problem. Prior heterogeneity shifts the agent's location in belief space without changing how he ranks precision levels relative to price. Cost heterogeneity changes the effective marginal cost of continued learning and therefore generates the type-specific marginal valuations that screening requires. A platform or advisor who can observe behavioral proxies for attention cost, including session length, return frequency, response latency, can implement the full range of price discrimination. One who can only infer the agent's prior belief gains nothing from offering differentiated tiers.

The analysis shows that information distortions can arise from the engagement motive alone, without preference misalignment over the agent's final action. This mechanism helps explain why information provision may be slowest precisely when the decision maker's need is greatest, and why tiered service structures in consulting and advisory markets emerge naturally from the economics of attention

APPENDICES

1.A Appendix to Chapter 1

1.A.1 Proof of Lemma 1.1

Define

$$\tilde{U}(p; n) = U(p; n) - (h - l) \left(p - \frac{1}{2} \right).$$

Since the subtracted term is linear, $\tilde{U}''(p; n) = U''(p; n)$. The transformed terminal payoffs are

$$\begin{aligned} \tilde{u}_B(p) &= u_B(p) - (h - l) \left(p - \frac{1}{2} \right) = (h + l) \left(p - \frac{1}{2} \right), \\ \tilde{u}_A(p) &= u_A(p) - (h - l) \left(p - \frac{1}{2} \right) = -(h + l) \left(p - \frac{1}{2} \right). \end{aligned}$$

The transformed problem is symmetric around $1/2$: the belief diffusion depends on $p(1 - p)$, and the transformed payoff obstacles are mirror images. Hence the optimal continuation region is symmetric around $1/2$, so

$$\bar{p} - \frac{1}{2} = \frac{1}{2} - \underline{p}.$$

Therefore,

$$\bar{p} = 1 - \underline{p}.$$

1.A.2 Derivation of the HJB Equation of the Agent

The agent's value function under precision level n is

$$U(p_0; n) = \sup_T \mathbb{E} \left[u(p_T) - \int_0^T c \, dt \mid p_0 \right], \quad dp_t = 2p_t(1 - p_t)\sqrt{n} \, d\tilde{W}_t,$$

where $u(p_T)$ is the expected payoff from taking the optimal static action at belief p_T .

Over a small interval $[0, \Delta]$, the agent either stops immediately or continues for duration Δ :

$$U(p; n) = \max \left\{ u(p), \mathbb{E}[-c\Delta + U(p_\Delta; n) \mid p_0 = p] \right\}.$$

By Itô's lemma, the generator of p_t satisfies

$$\mathcal{L}f(p) = \frac{1}{2}\sigma^2(p)f''(p) = 2np^2(1-p)^2f''(p),$$

so that

$$\mathbb{E}[U(p_\Delta; n) \mid p_0 = p] = U(p; n) + 2np^2(1-p)^2U''(p; n)\Delta + o(\Delta).$$

Substituting and simplifying,

$$U(p; n) = \max \left\{ u(p), U(p; n) + (2np^2(1-p)^2U''(p; n) - c)\Delta + o(\Delta) \right\}.$$

Taking the limit as $\Delta \rightarrow 0$, we obtain the HJB equation:

$$0 = \max \left\{ u(p) - U(p; n), 2np^2(1-p)^2U''(p; n) - c \right\}. \quad (\text{HJB-U})$$

In the stopping region, $U(p; n) = u(p)$ and $2np^2(1-p)^2U''(p; n) - c \leq 0$. In the continuation region, $U(p; n) > u(p)$ and

$$2np^2(1-p)^2U''(p; n) = c.$$

Defining the value of experimentation $V(p; n) := U(p; n) - u(p) \geq 0$, we obtain

$$0 = \max \left\{ -V(p; n), 2np^2(1-p)^2(V''(p; n) + u''(p)) - c \right\}.$$

Since $u(p)$ is piecewise linear and $u''(p) = 0$ except at the indifference point, the HJB simplifies away from the kink to

$$0 = \max \left\{ -V(p; n), 2np^2(1-p)^2V''(p; n) - c \right\}. \quad (\text{HJB-V})$$

This characterizes the agent's optimal stopping rule in the absence of discounting.

1.A.3 Existence, Uniqueness, and Monotonicity of the Optimal Boundary

Consider first the case $p_0 \geq 1/2$. The principal's problem can be written as

$$\max_{\bar{p} \in (p_0, 1)} f(\bar{p}, p_0) := \frac{r}{c}(h + l) \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}.$$

The positive multiplicative constant $\frac{r}{c}(h + l)$ does not affect the maximizer, so define

$$g(\bar{p}, p_0) := \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}.$$

The function g is positive on $(p_0, 1)$ and satisfies

$$\lim_{\bar{p} \downarrow p_0} g(\bar{p}, p_0) = 0, \quad \lim_{\bar{p} \uparrow 1} g(\bar{p}, p_0) = 0.$$

Thus any maximizer is interior.

The first-order condition is

$$g_{\bar{p}}(\bar{p}, p_0) = 0 \iff L'(\bar{p})^2 - (L(\bar{p}) - L(p_0))L''(\bar{p}) = 0.$$

Equivalently,

$$R(\bar{p}) = L(p_0), \quad R(p) := L(p) - \frac{L'(p)^2}{L''(p)}.$$

We now show that this equation has a unique solution in $(p_0, 1)$.

Using

$$L(p) = (2p - 1) \ln \frac{p}{1 - p},$$

we have

$$L'(p) = 2 \ln \frac{p}{1 - p} + \frac{2p - 1}{p(1 - p)}, \quad L''(p) = \frac{1}{p^2(1 - p)^2}.$$

Differentiating R gives

$$R'(p) = L'(p) \left(\frac{L'(p)L'''(p)}{[L''(p)]^2} - 1 \right).$$

Since

$$L'''(p) = \frac{2(2p-1)}{p^3(1-p)^3},$$

the term in parentheses can be written as

$$2(2p-1)p(1-p)L'(p) - 1.$$

Let $t = 2p - 1 \in (0, 1)$. Then this term becomes

$$A(t) = t(1-t^2) \ln \frac{1+t}{1-t} + 2t^2 - 1.$$

We claim that $A(t)$ has exactly one zero on $(0, 1)$. First, $A(0) = -1$ and $\lim_{t \uparrow 1} A(t) = 1$. A direct calculation shows that $A'(t) > 0$ on $(0, 1/\sqrt{3})$, so A crosses zero at most once on this interval. For $t \geq 1/\sqrt{3}$, the inequality

$$\ln \frac{1+t}{1-t} \geq 2t$$

implies

$$A(t) \geq 2t^2(1-t^2) + 2t^2 - 1 = 4t^2 - 2t^4 - 1 > 0.$$

Therefore $A(t)$ has exactly one zero on $(0, 1)$. Hence $R'(p)$ changes sign exactly once on $(1/2, 1)$: R first decreases and then increases.

Since

$$R(p_0) = L(p_0) - \frac{L'(p_0)^2}{L''(p_0)} < L(p_0)$$

for $p_0 > 1/2$, while

$$\lim_{p \uparrow 1} R(p) = +\infty,$$

the equation $R(\bar{p}) = L(p_0)$ has a unique solution in $(p_0, 1)$. Hence the principal's problem has a unique interior maximizer, denoted by $\bar{p}^*(p_0)$.

To establish monotonicity, observe that

$$g_{p_0}(\bar{p}, p_0) = -\frac{L'(p_0)}{L'(\bar{p})}.$$

Differentiating with respect to $\bar{p}$ gives

$$g_{\bar{p}p_0}(\bar{p}, p_0) = \frac{L'(p_0)L''(\bar{p})}{[L'(\bar{p})]^2}.$$

For $p_0 > 1/2$ and $\bar{p} > 1/2$, this expression is strictly positive. Hence g has strictly increasing differences in $(\bar{p}, p_0)$. Since the maximizer is unique, the monotone comparative statics theorem implies that $\bar{p}^*(p_0)$ is strictly increasing in p_0 on $(1/2, 1)$.

Finally, the optimal consultation level is pinned down by the agent's stopping-boundary condition

$$L'(\bar{p}^*(p_0)) = \frac{2n^*(p_0)(h+l)}{c}.$$

Since L' is strictly increasing on $(1/2, 1)$, the strict increase of $\bar{p}^*(p_0)$ implies that $n^*(p_0)$ is strictly increasing on $(1/2, 1)$. By symmetry, $n^*(p_0)$ is strictly decreasing on $(0, 1/2)$ and reaches its unique minimum at $p_0 = 1/2$.

1.A.4 Proof of Theorem 1.2

Fix any prior $p_0 \in (1/2, 1)$ and write the principal's value as

$$V(p_0) := \sup_{\bar{p} \in (p_0, 1)} f(\bar{p}, p_0), \quad f(\bar{p}, p_0) := \frac{r}{c}(h+l) \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}.$$

By Appendix 1.A.3, the maximizer $\bar{p}^*(p_0)$ is unique and interior. Therefore,

$$V(p_0) = f(\bar{p}^*(p_0), p_0).$$

By the envelope theorem,

$$V'(p_0) = \frac{\partial f}{\partial p_0}(\bar{p}^*(p_0), p_0).$$

Holding $\bar{p}$ fixed, differentiation yields

$$\frac{\partial f}{\partial p_0}(\bar{p}, p_0) = -\frac{r}{c}(h+l) \frac{L'(p_0)}{L'(\bar{p})}.$$

Thus

$$V'(p_0) = -\frac{r}{c}(h+l) \frac{L'(p_0)}{L'(\bar{p}^*(p_0))}.$$

Since $p_0 > 1/2$ and $\bar{p}^*(p_0) > p_0 > 1/2$, both $L'(p_0)$ and $L'(\bar{p}^*(p_0))$ are positive. Therefore,

$$V'(p_0) < 0 \quad \text{for all } p_0 \in (1/2, 1).$$

By symmetry of the problem around $1/2$, the same argument implies

$$V'(p_0) > 0 \quad \text{for all } p_0 \in (0, 1/2).$$

Combining these inequalities with the symmetry $V(p_0) = V(1 - p_0)$, we obtain that V is strictly single-peaked and attains its unique maximum at $p_0 = 1/2$.

1.A.5 Proof of Theorem 1.3

Step 1: Closed form for $S(p_0)$. Fix $p_0 \in (1/2, 1)$; the case $p_0 \in (0, 1/2)$ follows by symmetry. In the continuation region, the agent's value function is

$$U(p_0; n^*) = \frac{c}{2n^*} [L(\bar{p}^*) - L(p_0)] + (h-l)p_0 + a_2,$$

where $a_1 = h - l$ is pinned down by the symmetry of stopping boundaries (Lemma 1.1), and a_2 is determined by value-matching at $\bar{p}^*$:

$$a_2 = (h + l)\bar{p}^* - h - \frac{c}{2n^*}L(\bar{p}^*).$$

Substituting and using the boundary condition $\frac{c}{2n^*} = \frac{h+l}{L'(\bar{p}^*)}$ from Proposition 1.1:

$$U(p_0; n^*) = (h + l) \left[\frac{L(\bar{p}^*) - L(p_0)}{L'(\bar{p}^*)} + \bar{p}^* \right] + (h - l)p_0 - h.$$

Since $u(p_0) = h(2p_0 - 1)$ for $p_0 > 1/2$, the surplus is

$$S(p_0) = (h + l) \left[\frac{L(\bar{p}^*) - L(p_0)}{L'(\bar{p}^*)} + \bar{p}^* - p_0 \right], \quad (1.A.1)$$

where $\bar{p}^* = \bar{p}^*(p_0)$ is the optimal stopping boundary characterized in Appendix 1.A.3. Define $W(p_0) := S(p_0)/(h + l)$ for convenience.

Step 2: Symmetry and boundary behavior. By symmetry of the full problem around $1/2$, $W(p_0) = W(1 - p_0)$, so $p_0 = 1/2$ is a critical point of W .

As $p_0 \rightarrow 1^-$, the principal must set $n^*(p_0) \rightarrow \infty$ to retain the agent at all (Theorem 1.1), so the continuation region shrinks to a point and $\bar{p}^*(p_0) \rightarrow 1$ and $\underline{p}^*(p_0) \rightarrow p_0$. Hence $L(\bar{p}^*) - L(p_0) \rightarrow L(1) - L(1) = 0$ and $\bar{p}^* - p_0 \rightarrow 0$, giving $S(p_0) \rightarrow 0$. The same holds at $p_0 \rightarrow 0$ by symmetry.

Step 3: $W'(p_0) < 0$ for $p_0 \in (1/2, 1)$. Differentiating (1.A.1) with respect to p_0 and applying the principal's first-order condition $g_{\bar{p}}(\bar{p}^*, p_0) = 0$ (i.e. $L'(\bar{p}^*)^2 - [L(\bar{p}^*) - L(p_0)]L''(\bar{p}^*) = 0$), the terms involving $d\bar{p}^*/dp_0$ cancel, leaving

$$W'(p_0) = \frac{d\bar{p}^*}{dp_0} - \frac{L'(p_0)}{L'(\bar{p}^*)} - 1. \quad (1.A.2)$$

From the implicit function theorem applied to the first-order condition $R(\bar{p}^*) = L(p_0)$, where $R(p) := L(p) - L'(p)^2/L''(p)$:

$$\frac{d\bar{p}^*}{dp_0} = \frac{L'(p_0)}{R'(\bar{p}^*)}. \quad (1.A.3)$$

Substituting (1.A.3) into (1.A.2):

$$W'(p_0) = L'(p_0) \left[\frac{1}{R'(\bar{p}^*)} - \frac{1}{L'(\bar{p}^*)} \right] - 1. \quad (1.A.4)$$

We show this is strictly negative by considering the two regions identified in Appendix 1.A.3, where it is established that R' changes sign exactly once on $(1/2, 1)$.

Case (i): $R'(\bar{p}^*) < 0$. Then $1/R'(\bar{p}^*) < 0$ and $1/L'(\bar{p}^*) > 0$, so the bracket in (1.A.4) is negative. Since $L'(p_0) \geq 0$ for $p_0 \geq 1/2$, the first term is non-positive, giving $W'(p_0) \leq -1 < 0$.

Case (ii): $R'(\bar{p}^*) > 0$. Both terms in the bracket are positive, so we must bound their difference. Setting $t = 2\bar{p}^* - 1 \in (0, 1)$ and using the explicit derivatives

$$L'(p) = 2 \ln \frac{p}{1-p} + \frac{2p-1}{p(1-p)}, \quad L''(p) = \frac{1}{p^2(1-p)^2}, \quad L'''(p) = \frac{2(2p-1)}{p^3(1-p)^3},$$

a direct computation gives

$$R'(\bar{p}^*) = L'(\bar{p}^*) \cdot [2(2\bar{p}^* - 1)\bar{p}^*(1 - \bar{p}^*)L'(\bar{p}^*) - 1] = L'(\bar{p}^*) \cdot A(t),$$

where $A(t) = t(1 - t^2) \ln \frac{1+t}{1-t} + 2t^2 - 1$ is the function analyzed in Appendix 1.A.3. In Case (ii), $A(t) > 0$, so

$$\frac{1}{R'(\bar{p}^*)} - \frac{1}{L'(\bar{p}^*)} = \frac{1}{L'(\bar{p}^*)} \left[\frac{1}{A(t)} - 1 \right] = \frac{1 - A(t)}{L'(\bar{p}^*) \cdot A(t)}.$$

Substituting into (1.A.4) and using (1.A.3):

$$W'(p_0) = \frac{L'(p_0)}{L'(\bar{p}^*)} \cdot \frac{1 - A(t)}{A(t)} - 1.$$

The monotone comparative statics established in [Appendix 1.A.3](#) imply $\bar{p}^*(p_0) > p_0 > 1/2$ for $p_0 \in (1/2, 1)$, and L' is strictly increasing on $(1/2, 1)$, so $L'(p_0) < L'(\bar{p}^*)$, giving $L'(p_0)/L'(\bar{p}^*) < 1$. Since $A(t) \in (0, 1]$ in Case (ii) (the zero of A marks the boundary between the two cases), we have $(1 - A(t))/A(t) \geq 0$, so the first term is in $[0, 1)$. Therefore

$$W'(p_0) < 1 \cdot 1 - 1 = 0.$$

Combining both cases, $W'(p_0) < 0$ for all $p_0 \in (1/2, 1)$. By symmetry, $W'(p_0) > 0$ for all $p_0 \in (0, 1/2)$. Together with $W(p_0) = W(1 - p_0)$ and $W(p_0) \rightarrow 0$ at the boundaries, the surplus $S(p_0) = (h + l)W(p_0)$ is strictly single-peaked and attains its unique maximum at $p_0 = 1/2$.

References

- Bergemann, Dirk and Alessandro Bonatti (2019) “Markets for information: An introduction,” *Annual Review of Economics*, 11 (1), 85–107.
- Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin (2018) “The design and price of information,” *American economic review*, 108 (1), 1–48.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2015) “The limits of price discrimination,” *American Economic Review*, 105 (3), 921–957.
- Bloedel, Alexander W and Ilya R Segal (2018) “Persuasion with rational inattention,” *Available at SSRN 3164033*.
- Che, Yeon-Koo, Kyungmin Kim, and Konrad Mierendorff (2023) “Keeping the listener engaged: a dynamic model of bayesian persuasion,” *Journal of Political Economy*, 131 (7), 000–000.
- Eső, Péter and Balázs Szentes (2007) “The price of advice,” *The Rand Journal of Economics*, 38 (4), 863–880.
- Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki (2015) “Stochastic choice and optimal sequential sampling,” *arXiv preprint arXiv:1505.03342*.
- Gentzkow, Matthew and Emir Kamenica (2014) “Costly persuasion,” *American Economic Review*, 104 (5), 457–462.
- Gonçalves, Duarte (2024) “Speed, Accuracy, and Complexity,” *arXiv preprint arXiv:2403.11240*.
- Hébert, Benjamin and Weijie Zhong (2022) “Engagement maximization,” *arXiv preprint arXiv:2207.00685*.
- Henry, Emeric and Marco Ottaviani (2019) “Research and the approval process: The organization of persuasion,” *American Economic Review*, 109 (3), 911–955.
- Kamenica, Emir and Matthew Gentzkow (2011) “Bayesian persuasion,” *American Economic Review*, 101 (6), 2590–2615.
- Ke, T Tony, Zuo-Jun Max Shen, and J Miguel Villas-Boas (2016) “Search for information on multiple products,” *Management Science*, 62 (12), 3576–3603.
- Koh, Andrew and Sivakorn Sanguanmoo (2022) “Attention capture,” *arXiv preprint arXiv:2209.05570*.
- Lehrer, Ehud and Tao Wang (2024) “The value of information in stopping problems,” *Economic Theory*, 78 (2), 619–648.

- Liptser, Robert S and Albert N Shiryaev (2013) *Statistics of random processes: I. General theory*, 5: Springer Science & Business Media.
- Maskin, Eric and John Riley (1984) “Monopoly with incomplete information,” *The RAND Journal of Economics*, 15 (2), 171–196.
- Milgrom, Paul and Chris Shannon (1994) “Monotone comparative statics,” *Econometrica: Journal of the Econometric Society*, 157–180.
- Morris, Stephen and Philipp Strack (2019) “The wald problem and the relation of sequential sampling and ex-ante information costs,” *Available at SSRN 2991567*.
- Moscarini, Giuseppe and Lones Smith (2001) “The optimal level of experimentation,” *Econometrica*, 69 (6), 1629–1644.
- Mussa, Michael and Sherwin Rosen (1978) “Monopoly and product quality,” *Journal of Economic theory*, 18 (2), 301–317.
- Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov (2020) “Persuading the principal to wait,” *Journal of Political Economy*, 128 (7), 2542–2578.
- Persson, Petra (2018) “Attention manipulation and information overload,” *Behavioural Public Policy*, 2 (1), 78–106.
- Saeedi, Maryam, Yikang Shen, and Ali Shourideh (2024) “Getting the agent to wait,” *arXiv preprint arXiv*, 2407 (5).
- Wald, Abraham (1947) “Foundations of a general theory of sequential decision functions,” *Econometrica, Journal of the Econometric Society*, 279–313.
- Zhong, Weijie (2022) “Optimal dynamic information acquisition,” *Econometrica*, 90 (4), 1537–1582.

CHAPTER 2

The Optimal Experimentation Level in Costly Persuasion

2.1 Introduction

When one individual receives information carefully designed to swing his action, he might still be willing to follow the guidance, even when he is aware of the possible misalignment of interests. This seemingly surprising result follows from the obedience principle in the Bayesian persuasion model of [Kamenica and Gentzkow \(2011\)](#): when the information designer, or sender, can publicly commit to an information structure, and when communication is costless and instantaneous, more information is valuable to the receiver because it allows him to make more accurate inferences about the state of the world and take the corresponding action.

However, in practice persuasion is rarely frictionless. Consider a seller pitching a product to a potential buyer. Collecting information and tailoring the demonstration often takes time and effort. Likewise, the consumer must pay attention in order to process the information, which is costly. Similarly, consider a pharmaceutical company applying for FDA approval for a new drug. Designing and conducting experiments is costly for the drug company, and the FDA must also invest time and manpower to evaluate the relevant evidence. Persuasion is therefore often costly to both parties. The sender must design and collect information, so her cost may vary with the difficulty of information production and the quality of the signal. The receiver, on the other hand, must decide how long to pay attention to the information before taking an action.

This paper studies a dynamic persuasion problem in which the sender chooses the precision of an information process and the receiver chooses when to stop learning. The sender wants to induce action B , regardless of the state, while the receiver wants to match his action to the state. The sender can commit ex ante to the precision of a drift-diffusion information process. After observing this choice, the receiver decides how long to listen before choosing one of two actions. The sender faces an information-production cost that may depend on the chosen precision, while the receiver pays a constant flow cost of attention while processing the information.

The sender's information choice is nontrivial only when the prior satisfies $p_0 < 1/2$. In this region, the receiver would choose action A without additional information, while the sender wants to induce action B . When $p_0 \geq 1/2$, the receiver's default action is already B , so persuasion is unnecessary and the sender provides no information. The analysis therefore focuses on the persuasion region $p_0 < 1/2$, where information can change the receiver's action.

The key feature of the model is that the sender does not directly choose a distribution of posterior beliefs. Instead, she chooses the precision of a signal process. Given this precision, the receiver solves an optimal stopping problem and stops learning once his belief reaches one of two endogenous boundaries. The sender's precision choice therefore affects persuasion by changing the receiver's stopping boundaries and the probability that the receiver's posterior reaches the sender-favorable action region.

The main result is that, within the persuasion region, the sender's optimal experimentation level depends critically on the cost of information production. When information is costless for the sender, she chooses the most precise feasible information process because doing so maximizes the probability that the receiver's belief reaches the sender-favorable action region. In the high-precision limit, the receiver learns almost instantly, his expected learning cost vanishes, and his payoff converges to the full-information benchmark.

When information production involves a fixed cost, the sender's decision is distorted on the extensive margin. She provides information only when the prior is sufficiently close to

the receiver’s indifference belief. Conditional on providing information, however, she still chooses the most precise feasible signal. When the marginal cost of precision is increasing, the sender’s decision is also distorted on the intensive margin. Near the receiver’s indifference belief, the receiver is already relatively easy to persuade, so the marginal benefit of additional precision is small. The sender may therefore choose a lower-precision signal, even though the receiver would prefer more accurate information.

The mechanism is simple. A more precise signal expands the receiver’s learning region and increases the probability that the posterior eventually reaches the sender-favorable stopping boundary. However, precision is valuable to the sender only insofar as it increases the probability of inducing action B . Once the receiver is close to being persuaded, additional precision has limited marginal benefit for the sender. As a result, costly persuasion can generate both no information provision for unfavorable priors and low-precision information provision for relatively favorable priors.

The rest of the chapter is organized as follows. Section 2.2 discusses the related literature. Section 2.3 introduces the model. Section 2.4 characterizes the receiver’s optimal stopping problem. Section 2.5 studies the sender’s optimal choice of information precision. Section 2.6 concludes.

2.2 Related Literature

This paper relates to several strands of literature. First, it contributes to the literature on Bayesian persuasion and information design by studying persuasion in a dynamic environment with costly information production. [Kamenica and Gentzkow \(2011\)](#) show that when a sender can commit to an information structure, the sender’s problem can be solved by choosing a distribution of posterior beliefs subject to Bayes plausibility. [Kamenica \(2019\)](#) and [Bergemann and Morris \(2019\)](#) provide broader surveys of the information-design approach. The present paper keeps the commitment assumption, but restricts the sender’s choice to the precision of a drift-diffusion information process. The sender therefore can-

not directly choose an arbitrary distribution over posterior beliefs; instead, she chooses an information technology that determines the receiver’s endogenous stopping boundaries.

The closest paper is [Che et al. \(2023\)](#), who study a dynamic model of Bayesian persuasion in which information takes time and is costly for the sender to generate and for the receiver to process. They derive the belief dynamics under the optimal information structure and show that the sender must leave rents to the receiver to compensate him for his listening costs. They also show that a range of outcomes between the sender-optimal static persuasion outcome in [Kamenica and Gentzkow \(2011\)](#) and full revelation can arise in Markov perfect equilibria as persuasion costs vanish. This paper instead focuses on a commitment environment in which the sender chooses the precision of a drift-diffusion signal process at the outset. This restriction makes it possible to connect the sender’s problem directly to the receiver’s optimal stopping boundaries and to characterize how the sender’s information-production cost shapes the optimal level of experimentation.

This paper is also related to [Henry and Ottaviani \(2019\)](#), who study research and approval as a dynamic persuasion problem motivated by drug regulation. In their model, evidence accumulates over time, and the allocation of stopping authority affects approval standards and welfare. They compare institutional arrangements in which the sender controls experimentation, the receiver can request additional evidence, or the receiver can commit to an approval standard. In contrast, the present paper focuses on a different margin of strategic persuasion: the sender chooses the precision of the evidence-generating process, while the receiver retains the stopping decision. The question is therefore not who decides when evidence is sufficient, but how precise the evidence-generating process should be.

Second, the receiver’s problem builds on the literature on sequential experimentation and optimal stopping. [Wald \(1947\)](#) provides the classical foundation for sequential decision problems. [Moscarini and Smith \(2001\)](#) characterize the optimal level of experimentation in a single-agent environment. In this paper, the receiver’s learning problem is a standard optimal stopping problem, but the experimentation technology is chosen strategically by the sender. The model can therefore be interpreted as a strategic decomposition of the optimal

experimentation problem: the sender chooses the information intensity, while the receiver chooses when to stop learning.

Finally, this paper relates to the literature on costly communication, rational inattention, and costly information acquisition. [Gentzkow and Kamenica \(2014\)](#) extend the concavification approach in [Kamenica and Gentzkow \(2011\)](#) to costly signals with an entropy-reduction cost. [Bloedel and Segal \(2018\)](#) study persuasion when the receiver faces rational inattention costs, so information disclosure serves both to persuade the receiver and to manipulate his attention. [Persson \(2018\)](#) studies costly communication when the receiver may face information overload as the sender complexifies information. More recently, [Matyskova and Montes \(2023\)](#) study Bayesian persuasion when the receiver can acquire independent information at a posterior-separable cost, and [Wei \(2021\)](#) studies persuasion under costly learning. These papers show that information costs can substantially change the sender’s optimal disclosure policy. The present paper differs by modeling the receiver’s cost as a flow cost of processing a sender-chosen diffusion process and by focusing on how the sender’s cost of precision determines the optimal level of experimentation.

2.3 Model

2.3.1 Players and Payoffs

There are two players: a sender, denoted by she, and a receiver, denoted by he. There is an unknown payoff-relevant state $\theta \in \{H, L\}$, with common prior $p_0 = \Pr(\theta = H)$. The receiver eventually chooses an action $a \in \{A, B\}$. His payoff is

$$\begin{aligned} u(A, L) &= \pi_A, & u(A, H) &= -\pi_A, \\ u(B, L) &= -\pi_B, & u(B, H) &= \pi_B. \end{aligned}$$

The receiver wants to match action A with the low state and action B with the high state. The sender's payoff is state independent:

$$v(a, \theta) = v \cdot \mathbf{1}\{a = B\}.$$

Thus, the sender wants to induce action B regardless of the state.

Given belief p , the receiver's payoff from stopping and choosing the better action is

$$u(p) = \max\{\pi_A(1 - 2p), \pi_B(2p - 1)\}.$$

The payoff normalization implies that the receiver's indifference belief is $\hat{p} = 1/2$. This symmetry is useful because it pins down the receiver's stopping boundaries around the same belief that separates the two actions. The assumption is not without loss of generality, but it keeps the analysis focused on the sender's choice of information precision.

2.3.2 Static Persuasion Benchmark




Figure 2.1: Static Bayesian Persuasion benchmark

Notes: The red solid curves show expected payoffs under the sender-optimal static information structure. The blue squares show expected payoffs under the prior belief p_0 without additional information. The green dashed lines show expected payoffs under full revelation.

It is useful to begin with the static Bayesian persuasion benchmark in [Kamenica and Gentzkow \(2011\)](#). Suppose persuasion is costless for both players and the sender can commit

to a public information structure. When $p_0 < 1/2$, the receiver's default action is A , so the sender wants to design information that pushes the receiver's posterior above $1/2$ with the highest possible probability. Bayes plausibility requires the posterior mean to equal p_0 . The sender-optimal static experiment therefore splits the prior between 0 and $1/2$, inducing action B with probability $2p_0$, so the sender maximizes the probability of inducing $\hat{p}$ as well as action B . The resulting payoff is shown in the left panel of Figure 2.1. The receiver, on the other hand, receives no rents from persuasion. He gets the same payoff as in the situation where no information is provided, as shown in the right panel of Figure 2.1.

The dynamic model below imposes a different constraint. The sender does not choose arbitrary posterior distributions. Instead, she chooses the precision of a signal process, and the receiver chooses when to stop observing it. The relevant posterior distribution is therefore generated by the receiver's stopping boundaries.

2.3.3 Timing and Information

In the dynamic persuasion problem, time starts continuously from 0. At time 0, the sender chooses the precision $n > 0$ of the signal process. Conditional on the state, the signal evolves as

$$dx_t^\theta = \mu^\theta dt + \frac{1}{\sqrt{n}} dW_t,$$

where $\mu^H = 1$, $\mu^L = -1$, and W_t is a standard Brownian motion. The parameter n measures the precision, or quality, of information. A higher n makes the signal more informative about the state.

The sender pays a one-time information-production cost $k(n)$. We assume that $k(n)$ is a continuous non-decreasing function: the more precise the signal is, the more costly it is. We also assume that $k(0) = 0$, $k(n) \leq v/2$ to avoid the trivial situation that the information is too expensive for the sender.

At each time instance, the receiver can decide whether to observe a piece of information to update his belief at a constant cost $c \geq 0$, then he can choose to take action A or B , or to wait to listen to the sender's next piece of information. Once the receiver acts, the game ends. The sender receives v if the receiver chooses B and 0 otherwise, net of the cost $k(n)$. The constant flow cost c implies the attention cost doesn't depend on the signal quality: the receiver needs to incur this attention cost as long as he decides to listen to the signal, regardless of the information context.

The assumption that the sender chooses a constant precision at time 0 restricts the class of feasible information policies. The restriction is useful because it gives a direct link between the sender's precision choice and the receiver's optimal stopping boundaries. Throughout the analysis, we focus on the case with no discounting. This keeps the model focused on the tradeoff between information precision and the receiver's flow cost of attention.

2.4 Receiver's Problem

Fix the sender's choice of precision n . The receiver's posterior belief $p_t = \Pr(\theta = H \mid x_s, s \leq t)$ evolves according to

$$dp_t = 2p_t(1 - p_t)\sqrt{n} d\bar{W}_t,$$

where $\bar{W}_t$ is the innovation process. The belief process is a martingale. Since future signal increments are ex ante identical, the receiver's problem is an optimal stopping problem: he chooses how long to listen, and then chooses the action that is optimal given his posterior belief.

Let $U(p; n)$ denote the receiver's value from entering the learning process at belief p when the signal precision is n . The receiver pays a flow attention cost $c > 0$ while learning.

His problem is

$$U(p_0; n) = \sup_T \mathbb{E} [u(p_T) - cT \mid p_0],$$

subject to the belief dynamics above. The receiver stops at a lower boundary $\underline{p}$ and an upper boundary $\bar{p}$. At $\underline{p}$, he chooses action A ; at $\bar{p}$, he chooses action B .

On the continuation region $(\underline{p}, \bar{p})$, the Hamilton-Jacobi-Bellman equation is

$$0 = -c + 2n(p(1-p))^2 U''(p; n).$$

And the stopping boundaries satisfy value matching and smooth pasting:

$$\begin{aligned} U(\underline{p}; n) &= \pi_A(1 - 2\underline{p}), & U(\bar{p}; n) &= \pi_B(2\bar{p} - 1), \\ U'(\underline{p}; n) &= -2\pi_A, & U'(\bar{p}; n) &= 2\pi_B. \end{aligned}$$

Rearranging the HJB equation, we get

$$U''(p; n) = \frac{c}{2np^2(1-p)^2}.$$

Thus the receiver's value is convex on the continuation region. Lower attention costs or higher signal precision make continuation more attractive and expand the learning region.

Proposition 2.1 *For fixed payoffs, the receiver's continuation region expands when the flow cost c decreases or when the signal precision n increases. In each case, the receiver's value $U(p; n)$ weakly increases.*

Moreover, the HJB equation admits the general solution

$$U(p; n) = \frac{c}{2n} \left[p \ln \frac{p}{1-p} + (1-p) \ln \frac{1-p}{p} \right] + A_1 p + A_2,$$




Figure 2.2: Receiver's value from optimal learning.

Notes: The red solid curves show the receiver's expected payoffs with no information. The blue dashed curve shows the receiver's value from optimal learning under the baseline parameters. The blue dots mark the baseline stopping boundaries $\underline{p}$ and $\bar{p}$. In panel (b), the orange dotted curve shows the receiver's value when the flow cost c is lower or the signal precision n is higher. The orange dots mark the expanded stopping boundaries $\underline{p}'$ and $\bar{p}'$.

where A_1 and A_2 are constants. The two value-matching conditions and two smooth-pasting conditions can pin down $(\underline{p}, \bar{p}, A_1, A_2)$. The payoff symmetry also implies that the receiver's stopping region should be symmetric around the indifference belief.

Lemma 2.1 *For every n and c , the receiver's stopping boundaries satisfy*

$$\bar{p} = 1 - \underline{p}.$$

The symmetry follows from the payoff normalization around $1/2$ and the fact that the learning technology depends on p only through $p(1 - p)$. The formal proof is in Appendix 2.A.1.

Using Lemma 2.1, smooth pasting implies $A_1 = \pi_B - \pi_A$. The value-matching conditions then imply

$$\begin{aligned} \ln\left(\frac{1 - \underline{p}}{\underline{p}}\right) + \frac{1 - 2\underline{p}}{2\underline{p}(1 - \underline{p})} &= \frac{n(\pi_A + \pi_B)}{c}, \\ \ln\left(\frac{\bar{p}}{1 - \bar{p}}\right) + \frac{2\bar{p} - 1}{2\bar{p}(1 - \bar{p})} &= \frac{n(\pi_A + \pi_B)}{c}. \end{aligned}$$

Equivalently, the lower boundary is pinned down by the first equation together with $\bar{p} = 1 - \underline{p}$. The stopping region expands when the stakes are larger, the signal is more precise, or attention is less costly.

Proposition 2.2 *An increase in π_A or π_B expands the receiver's continuation region and weakly increases the receiver's value.*

2.5 Sender's Problem

The sender chooses information precision n at time 0. Given n , the receiver chooses the optimal stopping time T characterized in Section 2.4. The sender obtains payoff v whenever the receiver chooses action B , and pays an information-production cost $k(n)$. Her problem is therefore

$$\sup_n v \Pr(a = B \mid n) - k(n).$$

Since the receiver chooses action B when the posterior reaches the upper stopping boundary $\bar{p}(n)$, the sender's payoff is proportional to the probability that the receiver's posterior reaches $\bar{p}(n)$.

If $p_0 \geq 1/2$, the receiver's default action is already B , so persuasion is unnecessary. The interesting case is $p_0 < 1/2$, where the receiver would choose A without additional information. In this region, the sender chooses the precision of the signal to maximize the probability of reaching the upper boundary, net of the production cost.

Lemma 2.1 makes the sender's problem tractable. By choosing n , the sender indirectly chooses the receiver's stopping boundaries $(\underline{p}(n), \bar{p}(n))$, with $\bar{p}(n) = 1 - \underline{p}(n)$. Since beliefs are martingales, the probability $\mu(n)$ of reaching the upper boundary satisfies

$$p_0 = (1 - \mu(n))\underline{p}(n) + \mu(n)\bar{p}(n).$$

Therefore,

$$\mu(n) = \frac{p_0 - \underline{p}(n)}{\bar{p}(n) - \underline{p}(n)} = \frac{p_0 - \underline{p}(n)}{1 - 2\underline{p}(n)}. \quad (2.1)$$

The sender's problem can therefore be written as

$$\sup_n v \frac{p_0 - \underline{p}(n)}{1 - 2\underline{p}(n)} - k(n). \quad (2.2)$$

Since higher precision expands the receiver's continuation region, $\underline{p}(n)$ decreases in n . It follows from equation (2.1) that $\mu(n)$ increases in n : more precise information increases the probability that the posterior reaches the sender-favorable boundary.


The following proposition summarizes the sender's optimal information choice.

Proposition 2.3 *Suppose $p_0 < 1/2$.*


1. *If $k(n) \equiv 0$, the sender chooses the maximum feasible precision.*
2. *If $k(n) \equiv k > 0$, there exists a cutoff $p^* \in (0, 1/2)$ such that the sender chooses the maximum feasible precision for $p_0 \in [p^*, 1/2)$ and chooses no information for $p_0 < p^*$.*
3. *If $k'(n) > 0$, there exists a cutoff $p^* \in (0, 1/2)$ such that the sender chooses no information for $p_0 < p^*$ and chooses positive information for $p_0 \in [p^*, 1/2)$. Within the persuasion region, the optimal precision decreases as p_0 approaches $1/2$.*

See the proof is in Appendix 2.A.2. Figure 2.3 illustrates the proposition.


The proposition shows how information-production costs shape the sender's optimal information choice. When information is costless, the sender only cares about maximizing the probability $\mu(n)$ that the receiver's posterior reaches the upper stopping boundary. Since higher precision expands the receiver's continuation region and raises $\mu(n)$, the sender chooses the maximum feasible precision. In the limit, the information process approximates full revelation and maximizes the probability of inducing action B within the class of Brownian experiments generated by the receiver's stopping rule.




(a) Costless information: $k(n) = 0$



(b) Fixed information cost: $k(n) = k > 0$



(c) Comparing two precision levels



(d) Increasing information cost: $k'(n) > 0$

Figure 2.3: Sender's expected payoff under different information-cost specifications. *Notes:* The red solid curve shows the sender's expected payoff under the optimal information choice. The green dashed lines are full-revelation payoff benchmarks. The blue dotted lines, when present, show payoffs generated by fixed finite-precision Brownian experiments with receiver stopping boundaries $(\underline{p}, \bar{p})$. Panel (a) gives the costless benchmark. Panel (b) shows the fixed-cost case, where p^* is the cutoff prior below which the sender does not provide information. Panel (c) compares two feasible precision levels $n_1 > n_2$, with $k(n_1) > k(n_2)$. Although n_1 generates a higher persuasion probability, n_2 can yield a higher net payoff near the receiver's indifference belief. Panel (d) shows the upper envelope induced by an increasing information-production cost.

With a fixed information-production cost, the sender's intensive-margin incentive is unchanged. Conditional on providing information, she still chooses the maximum feasible precision because doing so maximizes the persuasion probability $\mu(n)$. The fixed cost affects only the extensive margin. If the prior is sufficiently unfavorable, the expected persuasion gain is too small to cover the fixed cost, and the sender provides no information. If the prior is sufficiently close to $1/2$, persuasion is worthwhile, and the sender provides the most precise feasible signal. Thus, a fixed cost generates a cutoff p^* , but does not distort the sender's precision choice conditional on providing information at all.

When the cost of information increases with precision, the sender faces a tradeoff: higher precision raises the probability of reaching the sender-favorable boundary, but it also raises the cost of producing information. The marginal benefit of precision depends on the prior. To see this, compare two information structures as in Figure 2.3. The first, n_1 , is sufficiently precise to approximate full revelation. The second, $n_2 < n_1$, induces stopping boundaries

$$(1/2 - \varepsilon, 1/2 + \varepsilon)$$

for some $\varepsilon \in (0, 1/2)$. Since n_1 is more precise than n_2 , it is also more costly. The additional persuasion probability generated by n_1 rather than n_2 is

$$\begin{aligned} \mu_{\text{full}} - \mu_{\varepsilon} &= p_0 - \frac{p_0 - (1/2 - \varepsilon)}{2\varepsilon} \\ &= \left(\frac{1}{2\varepsilon} - 1\right) \left(\frac{1}{2} - p_0\right). \end{aligned}$$

For any fixed $\varepsilon \in (0, 1/2)$, this gain converges to zero as p_0 approaches $1/2$. Hence, near the receiver's indifference belief, the additional persuasion benefit from moving from n_2 to n_1 is small. If the cost savings from using n_2 are large enough, the sender strictly prefers the lower-precision experiment.

This logic extends to the continuous choice of precision. For sufficiently unfavorable priors, persuasion is too costly and the sender provides no information. At the cutoff prior p^* , the sender begins to provide information. As the prior moves closer to $1/2$, the receiver

becomes easier to persuade, so the sender can reduce precision while still inducing action B with sufficiently high probability. Therefore, within the persuasion region, the optimal precision decreases as p_0 approaches the receiver's indifference belief.






Figure 2.4: Optimal information precision under different information-cost specifications.

Notes: The red curves show the sender's optimal precision choice $n^*(p_0)$. The horizontal axis is the receiver's prior p_0 . Panel (a) shows the costless case, in which the sender chooses the maximum feasible precision $\bar{n}$ whenever information is needed for persuasion. Panel (b) shows the fixed-cost case: p^* is the cutoff prior below which the sender provides no information, while for $p_0 \in [p^*, 1/2]$ the sender chooses $\bar{n}$. Panel (c) compares two feasible precision levels $n_1 > n_2$, with $k(n_1) > k(n_2)$; near the receiver's indifference prior, the lower precision n_2 can yield a higher net payoff. Panel (d) shows the case with increasing marginal cost of precision. The sender provides no information for $p_0 < p^*$, and within the persuasion region the optimal precision decreases as p_0 approaches $1/2$.


Proposition 2.3 characterizes the sender's optimal information choice, and this choice also determines the receiver's expected payoff. Figure 2.5 illustrates the receiver's payoff induced by the sender's policy under the same four information-cost specifications.




(a) Costless information: $k(n) = 0$



(b) Fixed information cost: $k(n) = k > 0$



(c) Comparing two precision levels



(d) Increasing information cost: $k'(n) > 0$

Figure 2.5: Receiver payoff under optimal information policy and different cost specifications.

Notes: The red curves show the receiver's payoff under the sender's chosen information policy. The black dotted lines show the receiver's payoff from no information. The green dashed lines show the receiver's payoff under full information. The horizontal axis is the receiver's prior p_0 . Panel (a) shows the costless case, in which the sender chooses the maximum feasible precision $\bar{n}$; as $\bar{n} \rightarrow \infty$, the receiver's payoff converges to the full-information benchmark. Panel (b) shows the fixed-cost case: p^* is the cutoff prior below which the sender provides no information. Panel (c) compares two feasible precision levels $n_1 > n_2$, with $k(n_1) > k(n_2)$; near the receiver's indifference prior, the sender may choose the lower-precision experiment because it is less costly, even though it gives the receiver a lower learning value. Panel (d) shows the case with increasing marginal cost of precision, where the sender reduces precision as the prior approaches 1/2.

Whenever the sender provides an informative signal, the receiver's payoff lies weakly above the no-information payoff. This is because the receiver can always stop immediately and take the same action he would have taken without information. At the same time, unless the signal becomes arbitrarily precise, the receiver's payoff remains below the full-information benchmark because learning takes time and the receiver pays a flow cost while processing the signal.

Panel (a) shows the costless-information benchmark. When the sender's information-production cost is zero, the sender chooses the maximum feasible precision. As precision becomes arbitrarily large, the receiver learns almost instantly and his expected learning cost vanishes. Hence, the receiver's payoff converges to the full-information payoff. In this limiting case, the sender's optimal Brownian experiment approximates full revelation: the receiver obtains the value of knowing the state, while the sender maximizes the probability that the receiver takes action B . Appendix 2.A.3 provides the formal argument.

Panel (b) shows how a fixed information-production cost affects the receiver through the extensive margin of information provision. When the prior is below the cutoff p^* , persuasion is not profitable for the sender, so no information is provided and the receiver obtains his no-information payoff. Once the prior is sufficiently close to the receiver's indifference belief, persuasion becomes worthwhile for the sender. The sender then provides the most precise feasible signal, and the receiver benefits from learning. T

Panels (c) and (d) show that when precision is costly, the receiver's payoff is also shaped by the sender's intensive-margin choice of information quality. A more precise signal gives the receiver a higher learning value, but it is more costly for the sender to produce. Near the receiver's indifference belief, the marginal persuasion gain from additional precision is small because the receiver is already relatively easy to persuade. The sender may therefore choose a lower-precision experiment, even though the receiver would prefer more precise information. As a result, the receiver's payoff can remain strictly below the full-information benchmark in the region where the sender provides information.

2.6 Conclusion

This chapter studies a dynamic persuasion problem in which the sender chooses the precision of an information process and the receiver chooses when to stop learning. The model links Bayesian persuasion to optimal experimentation: the sender does not directly choose a distribution of posterior beliefs, but instead chooses the precision of a drift-diffusion signal process. Given this choice, the receiver solves an optimal stopping problem. The sender's precision choice therefore determines the receiver's stopping boundaries and, through them, the probability that the receiver eventually takes the sender's preferred action.

The sender's information choice is nontrivial only when the prior satisfies $p_0 < 1/2$. In this region, the receiver would choose action A without additional information, while the sender wants to induce action B . When $p_0 \geq 1/2$, the receiver's default action is already B , so persuasion is unnecessary and the sender provides no information. The main analysis therefore focuses on the region $p_0 < 1/2$, where information can change the receiver's action.

The main result is that, within this persuasion region, the sender's optimal experimentation level depends critically on the cost of producing information. When information production is costless, the sender chooses the most precise feasible signal. In the high-precision limit, the receiver learns almost instantly, his expected learning cost vanishes, and his payoff converges to the full-information benchmark. Thus, even though the sender is biased toward action B , costless precision makes the induced outcome approximate full revelation within the Brownian experimentation environment.

When information production is costly, this full-information logic breaks down. With a fixed information-production cost, the sender's decision is distorted on the extensive margin: she provides information only when the prior is sufficiently close to the receiver's indifference belief. Conditional on providing information, however, she still chooses the most precise feasible signal. When the marginal cost of precision is increasing, the sender is also distorted on the intensive margin. Near the receiver's indifference belief, the receiver is already relatively easy to persuade, so the marginal gain from additional precision is small. The sender

may therefore choose a lower-precision signal, even though the receiver would prefer more accurate information.

The analysis highlights a simple force in costly persuasion. More precise information raises the probability that the receiver's posterior reaches the sender-favorable stopping boundary, but precision is valuable to the sender only when it increases this persuasion probability. Once the receiver is close to being persuaded, additional precision has limited marginal benefit for the sender. As a result, costly persuasion can generate both no information provision for unfavorable priors and low-precision information provision for relatively favorable priors.

This mechanism also clarifies how parametric dynamic persuasion differs from the static Bayesian persuasion benchmark. In the static model with flexible information design, the sender chooses posterior beliefs subject only to Bayes plausibility. In the dynamic model, posterior beliefs are generated by a costly learning process and by the receiver's endogenous stopping rule. The sender's problem can therefore be interpreted as choosing stopping boundaries indirectly through the precision of the signal process. This provides a tractable way to study persuasion when information is gradual, attention is costly, and the quality of evidence is itself the only strategic choice.

APPENDICES

2.A Appendix to Chapter 2

2.A.1 Proofs for the Receiver's Problem

Proof .1 (Proof of Proposition 2.1) *Fix the payoff primitives. The receiver's continuation value solves*

$$0 = -c + 2n (p(1 - p))^2 U''(p; n)$$

on the continuation region, together with value matching and smooth pasting at the stopping boundaries.

A decrease in c lowers the flow cost of continuing. For any stopping time T , the payoff

$$\mathbb{E} [u(p_T) - cT \mid p_0]$$

is weakly higher when c is lower. Taking the supremum over stopping times implies that the value function is weakly higher at every belief. Since the stopping payoff $u(p)$ is unchanged, the set of beliefs at which continuation is optimal weakly expands.

An increase in n makes the signal more precise. Equivalently, the posterior process becomes a time change of the same martingale: higher n allows the receiver to reach any given posterior threshold more quickly in expectation. Since the receiver pays a flow cost while learning, this raises the value of continuation relative to stopping. The value function therefore weakly increases, and the continuation region weakly expands.

Hence the receiver's continuation region expands when c decreases or n increases, and the receiver's value weakly increases in each case.

Proof .2 (Proof of Lemma 2.1) Consider the transformed value function

$$\tilde{U}(p; n) = U(p; n) - (\pi_B - \pi_A) \left(p - \frac{1}{2} \right).$$

The transformation subtracts a linear term, so

$$\tilde{U}''(p; n) = U''(p; n).$$

It therefore leaves the learning technology and the curvature of the continuation value unchanged.

The transformed terminal payoff from choosing B is

$$\begin{aligned} \tilde{u}_B(p) &= \pi_B(2p - 1) - (\pi_B - \pi_A) \left(p - \frac{1}{2} \right) \\ &= (\pi_A + \pi_B) \left(p - \frac{1}{2} \right), \end{aligned}$$

while the transformed terminal payoff from choosing A is

$$\begin{aligned} \tilde{u}_A(p) &= \pi_A(1 - 2p) - (\pi_B - \pi_A) \left(p - \frac{1}{2} \right) \\ &= -(\pi_A + \pi_B) \left(p - \frac{1}{2} \right). \end{aligned}$$

Thus, after subtracting a linear term, the two stopping payoff lines are mirror images around $p = 1/2$. The learning technology is also symmetric around $1/2$, because the diffusion coefficient depends on p only through $p(1 - p)$. Therefore the transformed optimal stopping problem is symmetric around $1/2$. The stopping boundaries must satisfy

$$\bar{p} = 1 - \underline{p}.$$

Proof .3 (Proof of Proposition 2.2) *The receiver's value on the continuation region satisfies*

$$U''(p; n) = \frac{c}{2np^2(1-p)^2}.$$

Integrating twice gives

$$U(p; n) = \frac{c}{2n} \left[p \ln \frac{p}{1-p} + (1-p) \ln \frac{1-p}{p} \right] + A_1 p + A_2.$$

The constants A_1, A_2 and the stopping boundaries are determined by value matching and smooth pasting. By Lemma 2.1, $\bar{p} = 1 - \underline{p}$. Smooth pasting then implies

$$A_1 = \pi_B - \pi_A.$$

Substituting this expression into the value-matching conditions yields

$$\begin{aligned} \ln \left(\frac{\underline{p}}{1-\underline{p}} \right) + \frac{1-2\underline{p}}{2\underline{p}(1-\underline{p})} &= \frac{n(\pi_A + \pi_B)}{c}, \\ \ln \left(\frac{\bar{p}}{1-\bar{p}} \right) - \frac{2\bar{p}-1}{2\bar{p}(1-\bar{p})} &= -\frac{n(\pi_A + \pi_B)}{c}. \end{aligned}$$

The right-hand side of the lower-boundary equation is increasing in n , increasing in $\pi_A + \pi_B$, and decreasing in c . Therefore, higher stakes, higher precision, or lower attention cost expand the receiver's continuation region. The receiver's value weakly increases because each of these changes makes continuation more valuable relative to stopping.

2.A.2 Proof of Proposition 2.3

Proof .4 *Suppose $p_0 < 1/2$. For a given precision n , the receiver stops at boundaries $(\underline{p}(n), \bar{p}(n))$, with $\bar{p}(n) = 1 - \underline{p}(n)$. Since beliefs are martingales, the probability that the*

posterior reaches the upper boundary is

$$\mu(n) = \frac{p_0 - \underline{p}(n)}{\bar{p}(n) - \underline{p}(n)} = \frac{p_0 - \underline{p}(n)}{1 - 2\underline{p}(n)}.$$

By Proposition 2.1, higher precision expands the receiver's continuation region, so $\underline{p}(n)$ is decreasing in n . Since

$$\frac{\partial}{\partial \underline{p}} \left[\frac{p_0 - \underline{p}}{1 - 2\underline{p}} \right] = \frac{2p_0 - 1}{(1 - 2\underline{p})^2} < 0$$

for $p_0 < 1/2$, it follows that $\mu(n)$ is increasing in n .

First suppose $k(n) \equiv 0$. The sender's objective is $v\mu(n)$. Since $\mu(n)$ is increasing in n , the sender chooses the maximum feasible precision.

Second suppose $k(n) \equiv k > 0$. Conditional on providing information, the fixed cost does not affect the marginal benefit of precision. Therefore, the sender again chooses the maximum feasible precision. Let $\bar{n}$ denote the maximum feasible precision, and let

$$\mu(\bar{n}; p_0)$$

be the probability of reaching the upper boundary under this precision. Providing information is optimal if and only if

$$v\mu(\bar{n}; p_0) - k \geq 0.$$

The persuasion probability is increasing in p_0 , so there exists a cutoff $p^* \in (0, 1/2)$ such that the sender provides information for $p_0 \in [p^*, 1/2)$ and provides no information for $p_0 < p^*$.

Third suppose $k'(n) > 0$. The sender's objective is

$$v\mu(n; p_0) - k(n).$$

For sufficiently unfavorable priors, the maximum attainable persuasion probability is too small relative to the cost of information production, so the sender chooses no information. This gives a cutoff $p^* \in (0, 1/2)$ below which the sender does not provide information.

Within the persuasion region, the sender trades off persuasion probability against information-production cost. To see why the optimal precision can decrease as p_0 approaches $1/2$, compare two information structures. The first is sufficiently precise to approximate full revelation. The second induces stopping boundaries

$$(1/2 - \varepsilon, 1/2 + \varepsilon)$$

for some $\varepsilon \in (0, 1/2)$. The gain in persuasion probability from full revelation relative to the second experiment is

$$\begin{aligned} \mu_{\text{full}} - \mu_{\varepsilon} &= p_0 - \frac{p_0 - (1/2 - \varepsilon)}{2\varepsilon} \\ &= \left(\frac{1}{2\varepsilon} - 1 \right) \left(\frac{1}{2} - p_0 \right). \end{aligned}$$

This gain converges to zero as p_0 approaches $1/2$. Therefore, near the receiver's indifference belief, the sender can reduce precision while losing only a small amount of persuasion probability. If the cost saving from reducing precision exceeds this small probability loss, the sender strictly prefers the less precise signal. Hence, within the persuasion region, optimal precision decreases as p_0 approaches $1/2$.

2.A.3 High-Precision Limit and Receiver Learning Costs

This subsection clarifies the limiting case in which the sender has no information-production cost and can choose arbitrarily high precision. The receiver still pays a flow cost c while observing the signal. However, this learning cost becomes negligible as the precision of the signal becomes arbitrarily large.

Fix a precision level n , and let the receiver stop at boundaries

$$(p(n), \bar{p}(n)).$$

The receiver's total expected learning cost is

$$C_R(p_0; n) = c\mathbb{E}[T \mid p_0, n].$$

Using the log-likelihood representation of expected stopping time, this can be written as

$$C_R(p_0; n) = \frac{c}{2n} [\mathbb{E}L(p_T) - L(p_0)],$$

where

$$L(p) = p \log \frac{p}{1-p} + (1-p) \log \frac{1-p}{p}.$$

Under the symmetric stopping rule, $\bar{p}(n) = 1 - p(n)$, and therefore

$$L(\bar{p}(n)) = L(p(n)).$$

Hence, when p_0 lies in the continuation region,

$$C_R(p_0; n) = \frac{c}{2n} [L(\bar{p}(n)) - L(p_0)].$$

As $n \rightarrow \infty$, the receiver's stopping boundaries converge to the extreme posteriors:

$$p(n) \rightarrow 0, \quad \bar{p}(n) \rightarrow 1.$$

To see the rate, recall that the receiver's stopping payoff has slopes

$$u'_A(p) = -2\pi_A, \quad u'_B(p) = 2\pi_B.$$

Combining smooth pasting with the receiver's value function implies the boundary condition

$$L'(\bar{p}(n)) = \frac{2n(\pi_A + \pi_B)}{c}.$$

Near $p = 1$, $L'(\bar{p})$ diverges at rate $1/(1 - \bar{p})$. Therefore,

$$1 - \bar{p}(n) = O\left(\frac{1}{n}\right).$$

It follows that

$$L(\bar{p}(n)) = O(\log n).$$

Thus,

$$C_R(p_0; n) = \frac{c}{2n} [L(\bar{p}(n)) - L(p_0)] \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Intuitively, higher precision makes the receiver's belief move faster. Although the receiver becomes willing to wait for more extreme posterior beliefs, the informational distance to those beliefs grows only logarithmically, while the speed of learning grows linearly in n . Therefore, the receiver's total expected learning cost vanishes in the high-precision limit.

This observation clarifies the costless-information benchmark. If the sender has no information-production cost and can choose arbitrarily high precision, she chooses the maximum feasible precision. In the limit, the receiver learns almost instantly, and his payoff converges to the full-information payoff. For $p_0 < 1/2$, the sender's probability of inducing action B converges to the full-information probability p_0 . Thus, in the costless benchmark, the sender's optimal Brownian information policy approximates full revelation.

References

- Bergemann, Dirk and Stephen Morris (2019) “Information design: A unified perspective,” *Journal of Economic Literature*, 57 (1), 44–95.
- Bloedel, Alexander W and Ilya R Segal (2018) “Persuasion with rational inattention,” *Available at SSRN 3164033*.
- Che, Yeon-Koo, Kyungmin Kim, and Konrad Mierendorff (2023) “Keeping the listener engaged: a dynamic model of bayesian persuasion,” *Journal of Political Economy*, 131 (7), 000–000.
- Gentzkow, Matthew and Emir Kamenica (2014) “Costly persuasion,” *American Economic Review*, 104 (5), 457–462.
- Henry, Emeric and Marco Ottaviani (2019) “Research and the approval process: The organization of persuasion,” *American Economic Review*, 109 (3), 911–955.
- Kamenica, Emir (2019) “Bayesian persuasion and information design,” *Annual Review of Economics*, 11 (1), 249–272.
- Kamenica, Emir and Matthew Gentzkow (2011) “Bayesian persuasion,” *American Economic Review*, 101 (6), 2590–2615.
- Matyskova, Ludmila and Alfonso Montes (2023) “Bayesian persuasion with costly information acquisition,” *Journal of Economic Theory*, 211, 105678.
- Moscarini, Giuseppe and Lones Smith (2001) “The optimal level of experimentation,” *Econometrica*, 69 (6), 1629–1644.
- Persson, Petra (2018) “Attention manipulation and information overload,” *Behavioural Public Policy*, 2 (1), 78–106.
- Wald, Abraham (1947) “Foundations of a general theory of sequential decision functions,” *Econometrica, Journal of the Econometric Society*, 279–313.
- Wei, Dong (2021) “Persuasion under costly learning,” *Journal of Mathematical Economics*, 94, 102451.

CHAPTER 3

Evolving Social Norms on Networks

3.1 Introduction

Private incentives and social norms jointly shape many economic decisions. Individuals may privately prefer one action but choose another because they observe that others around them behave differently, or because they believe that others expect them to conform. At the same time, individual actions together also shape the evolution of social norms: when many people repeatedly choose a behavior, the behavior becomes more visible, more expected, and potentially more persistent. For example, recycling behavior is shaped not only by private costs and environmental preferences, but also by whether recycling appears to be common and expected in one's community; as more individuals recycle, the behavior becomes increasingly visible and can further reinforce recycling as a local norm ([Cialdini et al., 1990](#)). This feedback between expectations and behavior is especially important in networks, where individuals observe only a limited and local subset of society as their perceived norms.

This paper develops a simple network model of norm dynamics in which private preferences, observable actions, and perceived norms may diverge. Agents differ in their privately preferred actions, but they observe only the past actions of their neighbors. The model compares two ways in which observed behavior can affect future behavior. In the benchmark case, agents update only by averaging locally observed actions, as in a standard DeGroot model. In the main case, agents instead combine their own preferred action with the average action observed in their neighborhood. This comparison highlights a basic tension. Pure social updating can generate consensus, but the resulting consensus need not reflect the distribution of private preferences. Preference-weighted updating preserves the role of private

preferences, but it generally leads to persistent heterogeneity rather than a common social norm.

The model is motivated by three characteristics of social norms. First, social norms often generate a positive feedback loop between behavior and expectations: the more widely a behavior is practiced, the more it becomes expected, which in turn reinforces adherence (Benhabib et al., 2010). Second, norms can display local conformity and global diversity: behavior may be relatively uniform within a group but differ substantially across groups. Third, social norms can be stable over time, even when individual actions are occasionally perturbed, because repeated conformity may absorb small shocks. At the same time, stability does not imply that a norm reflects private preferences: it may persist because individuals observe public behavior but not the preferences that generate it.

This is closely related to pluralistic ignorance, a social psychological phenomenon in which individuals misperceive others' private attitudes and therefore conform to a behavior that many privately oppose.¹ At a large scale, citizens may appear to support a political regime because opposition is rarely expressed publicly, even though many privately oppose the status quo (Kuran, 1991). In everyday social settings, students may drink more than they would like because each person mistakenly believes that others are more comfortable with heavy drinking than they actually are (Prentice and Miller, 1993).

In both cases, individual behavior reinforces the perceived norm: repeated public conformity makes the behavior appear more widely accepted, which can further encourage others to conform, but the discrepancy between private preferences and observed behavior can make a social norm appear more stable than it is: once individuals learn that others share their private views, the prevailing behavior may unravel quickly. A more recent literature Bursztyn et al. (2020), study young married men in Saudi Arabia and show that many privately support their wives working outside the home, but underestimate the extent to which other men share this view. When informed about the true level of support, men become more

¹Katz et al. (1931) describe pluralistic ignorance as a situation in which “most people privately hold an opinion, but they incorrectly believe that most other people hold the contrary opinion, and therefore end up acting against their own view.”

willing to help their wives search for jobs. This example illustrates an important distinction that motivates this paper: observed behavior can be a poor proxy for private preferences, especially when individuals observe behavior locally and do not observe the preferences that generate it.

The model formalizes these ideas in a two-type random network. Agents of the same type have the same privately preferred action and are more likely to be linked to each other. Types therefore determine agents' preferred actions, and shape the local neighborhoods from which agents observe behavior. This structure allows the model to study how homophily affects the relationship between private preferences, local perceived norms, and long-run behavior.

The analysis first studies a benchmark with pure social updating. In this case, agents update only by averaging the past actions of their neighbors, as in a standard DeGroot model. Under connectedness and aperiodicity of the network, actions converge to a consensus. The consensus is a degree-weighted average of initial actions, so it depends on initial behavior and network position rather than directly on private preferences. The benchmark therefore captures a case in which repeated social averaging can generate a stable common norm, even though the resulting norm need not reflect what agents privately prefer.

The main model then introduces private preferences into the updating rule. Each period, an agent chooses a weighted average of her privately preferred action and the average action observed in her neighborhood. The neighborhood average can be interpreted as a local perceived norm. Under this rule, initial actions do not affect long-run behavior. Instead, the limiting action of each agent is determined by the distribution of private preferences and the structure of the network. Because agents continue to place positive weight on their own preferred actions, consensus generally does not arise when preferred actions are heterogeneous.

The comparison highlights a discontinuity between pure imitation and preference-weighted updating. When agents put all weight on observed behavior, local averaging can produce a common long-run norm. When agents put even an arbitrarily small positive weight on their

own preferred actions, long-run behavior remains tied to private preferences and generally stays heterogeneous. Homophily further affects the steady state by changing the composition of agents' neighborhoods. Stronger same-type interaction makes local perceived norms more type-specific and can sustain greater differences across groups.

The chapter proceeds as follows. Section 3.2 discusses related literature on social norms, network learning, contagion, and pluralistic ignorance. Section 3.3 introduces the network environment. Section 3.4 studies the benchmark case in which agents update only by averaging locally observed behavior. Section 3.5 introduces private preferences into the updating rule and characterizes the resulting steady state. Section 3.6 studies how homophily and type composition shape group-level outcomes. Section 3.7 concludes.

3.2 Related Literature

This paper relates to several strands of literature. First, it contributes to the literature on social norms and norm dynamics. Social norms are often understood as patterns of behavior sustained by expectations, coordination motives, sanctions, or a desire to conform. [Young \(2015\)](#) reviews how norms can emerge and persist through evolutionary dynamics, while [Benhabib et al. \(2010\)](#) emphasize the broader role of social interactions in shaping economic behavior. [Cialdini et al. \(1990\)](#) emphasize that observed behavior can shape individuals' perceptions of what is socially typical. This perspective is useful because agents in this paper respond to locally observed behavior as perceived social norm, which may or may not reveal underlying private preferences.

Second, this paper builds on the literature on social learning in networks. In Bayesian sequential-learning models, agents observe private signals and the actions of previous agents, and the main question is whether information is eventually aggregated ([Banerjee, 1992](#); [Bikhchandani et al., 1992](#); [Smith and Sørensen, 2000](#)). The updating rule in this paper is more closely related to the heuristic DeGroot model ([DeGroot, 1974](#)), in which agents repeatedly update beliefs or actions by averaging the opinions of their neighbors. [DeMarzo](#)

et al. (2003) provide a microfoundation for repeated averaging through a model of persuasion bias, where agents repeatedly communicate but fail to fully account for repetition in the information they receive. Golub and Jackson (2010) study when naive learning in networks aggregates information, and Golub and Jackson (2012) show how homophily affects the speed and accuracy of learning and best-response dynamics. Acemoglu et al. (2010) study the spread of misinformation in social networks and show how influential or forceful agents can interfere with information aggregation. In contrast to both Bayesian and non-Bayesian social-learning models, this paper does not assume that agents are trying to learn a common true state. Instead, agents may have different private preferred actions. The central issue is therefore not whether society learns the truth, but whether observed behavior comes to reflect private preferences.

Third, the paper is related to models of contagion, cascades, and threshold behavior. Granovetter (1978) develops threshold models in which individuals adopt a behavior when sufficiently many others have adopted it. Morris (2000) studies contagion in local interaction systems and characterizes when a behavior initially adopted by a small set of agents can spread through the population. Watts (2002) studies global cascades on random networks and shows how network structure affects the possibility of large-scale diffusion. The present paper studies a different but related mechanism. Agents choose continuous actions and respond linearly to locally observed behavior while continuing to place weight on their private preferred actions. The focus is therefore not on whether a small seed can generate a cascade, but on how local social influence and private preferences jointly shape the long-run distribution of behavior.

Finally, the paper connects to the literature on pluralistic ignorance and misperceived social norms. Pluralistic ignorance arises when individuals systematically misperceive the private attitudes or behaviors of others. Katz et al. (1931) provide an early discussion of this phenomenon, while Miller (2023) reviews its origins and consequences across social contexts. Bursztyn et al. (2020) provide economic evidence that correcting misperceived norms can change behavior. Smerdon et al. (2020) distinguish between bad norms sustained by coordi-

nation incentives and bad norms sustained by uncertainty about others' preferences, finding support for the pluralistic ignorance channel. The model in this paper captures a simple version of this mechanism: when agents update only from observed behavior, the resulting social norm may persist even when it differs from the distribution of private preferences.

3.3 Model

There is a finite set of agents $N = \{1, \dots, n\}$. Each agent i has a type

$$\theta_i \in \{0, 1\},$$

which determines her privately preferred action. A type-0 agent prefers action 0, while a type-1 agent prefers action 1. Let

$$N_k = \{i \in N : \theta_i = k\}, \quad n_k = |N_k|, \quad k \in \{0, 1\}.$$

The share of type-1 agents is

$$\lambda = \frac{n_1}{n}.$$

Thus $1 - \lambda$ is the share of type-0 agents. The type vector $\theta = (\theta_i)_{i \in N}$ is fixed throughout the analysis.

Network. Agents interact through an undirected network A . The matrix A is an $n \times n$ adjacency matrix, where $A_{ij} = 1$ if agents i and j are linked and $A_{ij} = 0$ otherwise. Since the network is undirected, $A_{ij} = A_{ji}$. Let

$$d_i(A) = \sum_j A_{ij}$$

denote agent i 's degree, and let

$$D(A) = \sum_i d_i(A)$$

denote the total degree in the network.

Given a realized network A , define the row-normalized interaction matrix $T(A)$ by

$$T_{ij}(A) = \frac{A_{ij}}{d_i(A)}.$$

Thus $T_{ij}(A)$ is the weight that agent i places on agent j 's past action when forming her perceived local norm. The matrix $T(A)$ is row-stochastic. When there is no ambiguity, we write T instead of $T(A)$.

The network is generated by a two-type random network. Conditional on the type vector θ , links are independently drawn according to

$$\mathbf{P} = \begin{bmatrix} p & q \\ q & p \end{bmatrix}, \quad p > q.$$

Thus two agents of the same type are linked with probability p , while two agents of different types are linked with probability q . The assumption $p > q$ captures homophily. For simplicity, the two types have the same within-type linking probability, so the analysis isolates the effects of homophily and type composition rather than type-specific differences in network density.²

Figure 3.1 illustrates one realization of such a network.

Local observation. Agents observe past actions in their own neighborhoods. Given T , agent i 's perceived local norm at time $t - 1$ is

$$\bar{x}_i(t - 1) = \sum_j T_{ij} x_j(t - 1).$$

This is a local object: it is the average action observed by agent i , not the population average.

²Allowing the within-type linking probabilities to differ across types would not affect the convergence results below. It would change the group-level comparative statics by allowing the two types to differ in within-type network density.




Figure 3.1: A simulated two-type homophilous network

Notes: Nodes denote agents, shapes or colors denote types, and links are more likely between agents of the same type than between agents of different types.

Preferences and updating. Time is discrete. In each period t , agent i chooses an action $x_i(t) \in [0, 1]$. The action may be interpreted as the intensity with which the agent adopts a behavior associated with type 1.

Agent i trades off her preferred action θ_i against her perceived local norm $\bar{x}_i(t-1)$. Specifically, she solves

$$\max_{x_i(t)} -\alpha (x_i(t) - \bar{x}_i(t-1))^2 - (1 - \alpha) (x_i(t) - \theta_i)^2,$$

where $\alpha \in (0, 1)$. The parameter α measures the strength of social pressure. The optimal action is

$$x_i(t) = \alpha \bar{x}_i(t-1) + (1 - \alpha) \theta_i.$$

Substituting for $\bar{x}_i(t-1)$, the updating rule is

$$x_i(t) = \alpha \sum_j T_{ij} x_j(t-1) + (1 - \alpha) \theta_i.$$

In vector form,

$$\mathbf{x}(t) = \alpha T \mathbf{x}(t-1) + (1-\alpha)\theta.$$

3.4 Benchmark: Pure Social Updating

Before studying the preference-weighted updating rule, it is useful to consider a benchmark in which behavior evolves only through social averaging. In this benchmark, agents update by averaging the past actions of their neighbors:

$$\mathbf{x}(t) = T \mathbf{x}(t-1).$$

By iteration,

$$\mathbf{x}(t) = T^t \mathbf{x}(0).$$

Long-run behavior therefore depends on the initial action vector and the structure of the network, but not directly on the type vector unless types are already reflected in initial actions.

Proposition 3.1 *Suppose A is connected and aperiodic. Under the updating rule*

$$\mathbf{x}(t) = T \mathbf{x}(t-1),$$

actions converge to consensus. In particular,

$$\lim_{t \rightarrow \infty} x_i(t) = \sum_j \frac{d_j(A)}{D(A)} x_j(0)$$

for every agent i .

The proof is standard and is given in Appendix 3.A. The result shows that pure social updating aggregates local observations into a common long-run action. Since the network is

undirected, the consensus is a degree-weighted average of initial actions. Agents with higher degree receive larger weight in the limiting norm.

This benchmark highlights one way in which a stable social norm may emerge from repeated observation and imitation. If agents update only by averaging observed behavior, then the network can generate consensus even though the resulting norm need not reflect what agents privately prefer. Types matter only to the extent that they affect initial behavior or network position.

3.5 Preference-Weighted Updating

We now return to the main updating rule,

$$\mathbf{x}(t) = \alpha T \mathbf{x}(t-1) + (1-\alpha)\theta.$$

Unlike the pure social-updating benchmark, agents continue to place positive weight on their privately preferred actions. The analysis in this section is conditional on a realized network A , and hence on the induced interaction matrix T .

Proposition 3.2 *For any row-stochastic interaction matrix T and any $\alpha \in (0, 1)$, the sequence $\{\mathbf{x}(t)\}_{t \geq 0}$ converges to*

$$\mathbf{x} = (1-\alpha)(I - \alpha T)^{-1}\theta.$$

The limiting action vector is independent of the initial action vector $\mathbf{x}(0)$.

Proof 3.1 *Iterating the updating rule gives*

$$\mathbf{x}(t) = (\alpha T)^t \mathbf{x}(0) + \sum_{s=0}^{t-1} (\alpha T)^s (1-\alpha)\theta.$$

Since T is row-stochastic, its spectral radius is one. Hence the spectral radius of αT is $\alpha < 1$. Therefore $(\alpha T)^t \rightarrow 0$, and

$$\sum_{s=0}^{\infty} (\alpha T)^s = (I - \alpha T)^{-1}.$$

Taking limits yields the result.

Define

$$M = (1 - \alpha)(I - \alpha T)^{-1}.$$

Then $\mathbf{x} = M\theta$. The matrix M is the long-run influence matrix: M_{ij} measures the weight that agent i 's limiting action places on agent j 's type.

Proposition 3.3 *The matrix M is row-stochastic. Hence each limiting action is a weighted average of agents' types:*

$$x_i = \sum_j M_{ij} \theta_j.$$

The proof follows from the Neumann-series representation of M and is given in Appendix [3.A](#).

The result has two implications. First, initial actions do not affect long-run behavior. Temporary shocks fade because each agent places positive weight on her own type in every period. Second, limiting actions aggregate types through the network. Direct and indirect paths both matter, but paths of length s are discounted by α^s .

Proposition 3.4 *For $n > 1$, there is no interaction matrix T such that the preference-weighted updating process leads to consensus for every type vector $\theta \in [0, 1]^n$.*

The proof is given in Appendix [3.A](#).

Preference-weighted updating therefore generates partial conformity rather than consensus. Social pressure moves agents toward their perceived local norms, but persistent type heterogeneity remains because agents continue to place positive weight on their own preferred actions.

The contrast with the benchmark is sharp. Under pure social updating, long-run behavior is a degree-weighted average of initial actions. Under preference-weighted updating,

$$\mathbf{x} = (1 - \alpha)(I - \alpha T)^{-1}\theta.$$

Initial actions do not matter in the long run, while the type vector matters directly. Preference-weighted updating therefore preserves a connection between long-run behavior and underlying preferences, but it generally prevents consensus.

3.6 Homophily and Group-Level Outcomes

The previous section characterized equilibrium behavior for any realized interaction matrix T . We now return to the two-type homophilous network and describe how homophily and type composition affect group-level behavior.

Let

$$\bar{x}_0 = \frac{1}{n_0} \sum_{i \in N_0} x_i, \quad \bar{x}_1 = \frac{1}{n_1} \sum_{i \in N_1} x_i$$

denote the average limiting action among type-0 and type-1 agents. Since type-0 agents prefer action 0 and type-1 agents prefer action 1, the difference

$$\bar{x}_1 - \bar{x}_0$$

measures behavioral separation across types.

Mean-field approximation. To build intuition, consider a large-network approximation in which each agent's neighborhood composition is close to its expectation. Under the two-type random network, a type-0 agent's expected share of type-0 neighbors is

$$\frac{n_0 p}{n_0 p + n_1 q} = \frac{(1 - \lambda)p}{(1 - \lambda)p + \lambda q},$$

while her expected share of type-1 neighbors is

$$\frac{n_1 q}{n_0 p + n_1 q} = \frac{\lambda q}{(1 - \lambda)p + \lambda q}.$$

Similarly, a type-1 agent's expected share of type-0 neighbors is

$$\frac{n_0 q}{n_0 q + n_1 p} = \frac{(1 - \lambda)q}{(1 - \lambda)q + \lambda p},$$

and her expected share of type-1 neighbors is

$$\frac{n_1 p}{n_0 q + n_1 p} = \frac{\lambda p}{(1 - \lambda)q + \lambda p}.$$

Under this approximation, group-level equilibrium behavior is described by

$$\bar{x}_0 = \alpha \left(\frac{(1 - \lambda)p}{(1 - \lambda)p + \lambda q} \bar{x}_0 + \frac{\lambda q}{(1 - \lambda)p + \lambda q} \bar{x}_1 \right),$$

and

$$\bar{x}_1 = \alpha \left(\frac{(1 - \lambda)q}{(1 - \lambda)q + \lambda p} \bar{x}_0 + \frac{\lambda p}{(1 - \lambda)q + \lambda p} \bar{x}_1 \right) + (1 - \alpha).$$

The first equation reflects the fact that type-0 agents prefer action 0, while the second reflects the fact that type-1 agents prefer action 1. The derivation is given in [Appendix 3.A](#).

Comparative statics. The system highlights three forces.

First, α measures the strength of social pressure. When α is close to zero, agents place little weight on the perceived local norm, and equilibrium behavior remains close to the type vector:

$$\mathbf{x} \approx \boldsymbol{\theta}.$$

As α rises, agents respond more strongly to local behavior. Social pressure therefore moves agents away from their preferred actions and toward their neighborhood averages.

Second, the gap $p - q$ measures the strength of homophily. When $p - q$ is large, agents are mostly exposed to others of the same type. Type-0 agents then observe behavior close to action 0, while type-1 agents observe behavior close to action 1. Stronger homophily therefore tends to increase behavioral separation, as measured by $\bar{x}_1 - \bar{x}_0$. When p is close to q , neighborhoods are less segregated, and agents from different types face more similar local norms.

Third, λ measures the prevalence of type-1 agents. As λ increases, type-1 behavior becomes more common in the population and hence more common in mixed neighborhoods. This tends to raise the local norm faced by type-0 agents and to keep type-1 agents closer to their preferred action. Thus type composition affects behavior through the local social environments agents observe.

These comparative statics are qualitative, but they show how the random-network structure enters the model. Homophily determines how different local norms are across types; type composition determines which behavior is more prevalent in mixed neighborhoods; and α determines how strongly agents respond to these local norms.

3.7 Conclusion

This paper develops a simple network model of social norm dynamics when observed behavior and private preferences may diverge. Agents have privately preferred actions, observe past actions in their local neighborhoods, and update their behavior in response to both private preferences and perceived local norms. The model compares two updating rules: pure social updating, in which agents only average locally observed behavior, and preference-weighted updating, in which agents also continue to place weight on their own preferred actions.

Under pure social updating, behavior converges to a common norm under standard connectedness and aperiodicity conditions. The limiting norm is a degree-weighted average of initial actions, so it depends on initial behavior and network position rather than directly on the distribution of private preferences. Under preference-weighted updating, by contrast,

the limiting action vector is

$$\mathbf{x} = (1 - \alpha)(I - \alpha T)^{-1}\theta.$$

Initial actions do not matter in the long run, while private preferences matter directly. Preference-weighted updating therefore preserves a connection between long-run behavior and underlying preferences, but it generally prevents consensus when preferences are heterogeneous.

The model also explores the role of homophily. When same-type links are more likely than cross-type links, agents of different types face systematically different local norms. Stronger homophily makes local perceived norms more type-specific and tends to sustain behavioral separation across groups. Type composition also matters because the prevalence of each type affects the behavior observed in mixed neighborhoods.

These results are useful for thinking about pluralistic ignorance and misperceived social norms. If individuals observe actions but not the preferences behind them, public behavior can be a poor proxy for private attitudes. A stable norm may persist not because most agents privately endorse it, but because individuals respond to the behavior observed in their local social environments. The model shows how this link between public behavior and private preferences depends critically on the updating rule: pure imitation can produce consensus while severing the connection to preferences, whereas preference-weighted updating keeps behavior tied to preferences but preserves disagreement.

One interesting extension would be to allow agents to differ not only in their privately preferred actions, but also in their exposure to external information sources. Such an extension would add a third force to the updating rule: in addition to the average behavior observed in the agent's neighborhood and the agent's private preferred action, behavior would also respond to the agent's external information environment. This would distinguish social influence transmitted through the network from information exposure that is external to the network but heterogeneous across agents. It would also provide a way to study how local social norms interact with media exposure, public information, or targeted interventions in shaping long-run behavior.

APPENDICES

3.A Appendix to Chapter 3

3.A.1 Proof of Proposition 3.1

Since A is undirected, connected, and aperiodic, the Markov chain with transition matrix T has a unique stationary distribution. For an undirected network, the stationary weight of node j is proportional to its degree:

$$\pi_j = \frac{d_j(A)}{D(A)}.$$

Therefore T^t converges to a matrix whose rows are all equal to π . Hence

$$\lim_{t \rightarrow \infty} \mathbf{x}(t) = \lim_{t \rightarrow \infty} T^t \mathbf{x}(0) = \left(\sum_j \pi_j x_j(0) \right) \mathbf{1}.$$

Substituting $\pi_j = d_j(A)/D(A)$ gives

$$\lim_{t \rightarrow \infty} x_i(t) = \sum_j \frac{d_j(A)}{D(A)} x_j(0)$$

for every agent i .

3.A.2 Proof of Proposition 3.3

Using the Neumann series,

$$M = (1 - \alpha) \sum_{s=0}^{\infty} \alpha^s T^s.$$

Since T is row-stochastic, each T^s is row-stochastic. Thus M is nonnegative, and each row sums to

$$(1 - \alpha) \sum_{s=0}^{\infty} \alpha^s = 1.$$

Hence M is row-stochastic.

3.A.3 Proof of Proposition 3.4

Suppose there exists an interaction matrix T such that the preference-weighted updating process leads to consensus for every type vector $\theta \in [0, 1]^n$. Then $M\theta$ must be proportional to $\mathbf{1}$ for every θ . This requires all rows of M to be identical. Hence M has rank one and is singular.

But

$$M = (1 - \alpha)(I - \alpha T)^{-1}.$$

Since T is row-stochastic, the spectral radius of T is one. Because $\alpha \in (0, 1)$, the spectral radius of αT is less than one, so $I - \alpha T$ is invertible. Therefore M is invertible, a contradiction.

3.A.4 Mean-Field Group Equations

This appendix derives the group-level equations used in Section 3.6. Consider a large two-type random network with type-1 share $\lambda = n_1/n$. Conditional on being type 0, an agent's expected number of type-0 links is proportional to $n_0 p$, while her expected number of type-1 links is proportional to $n_1 q$. Hence her expected share of type-0 neighbors is

$$\frac{n_0 p}{n_0 p + n_1 q} = \frac{(1 - \lambda)p}{(1 - \lambda)p + \lambda q},$$

and her expected share of type-1 neighbors is

$$\frac{n_1 q}{n_0 p + n_1 q} = \frac{\lambda q}{(1 - \lambda)p + \lambda q}.$$

Similarly, conditional on being type 1, the expected share of type-0 neighbors is

$$\frac{n_0 q}{n_0 q + n_1 p} = \frac{(1 - \lambda)q}{(1 - \lambda)q + \lambda p},$$

and the expected share of type-1 neighbors is

$$\frac{n_1 p}{n_0 q + n_1 p} = \frac{\lambda p}{(1 - \lambda)q + \lambda p}.$$

Replacing individual neighborhood averages with these expected shares gives

$$\bar{x}_0 = \alpha \left(\frac{(1 - \lambda)p}{(1 - \lambda)p + \lambda q} \bar{x}_0 + \frac{\lambda q}{(1 - \lambda)p + \lambda q} \bar{x}_1 \right),$$

and

$$\bar{x}_1 = \alpha \left(\frac{(1 - \lambda)q}{(1 - \lambda)q + \lambda p} \bar{x}_0 + \frac{\lambda p}{(1 - \lambda)q + \lambda p} \bar{x}_1 \right) + (1 - \alpha).$$

3.A.5 Asymmetric Within-Type Link Probabilities

The main text assumes that both types have the same within-type linking probability p . If instead

$$\mathbf{P} = \begin{bmatrix} p & q \\ q & r \end{bmatrix}, \quad p > q, \quad r > q,$$

the convergence results are unchanged. The only change is in the group-level approximation.

The type-0 equation becomes

$$\bar{x}_0 = \alpha \left(\frac{(1 - \lambda)p}{(1 - \lambda)p + \lambda q} \bar{x}_0 + \frac{\lambda q}{(1 - \lambda)p + \lambda q} \bar{x}_1 \right),$$

while the type-1 equation becomes

$$\bar{x}_1 = \alpha \left(\frac{(1 - \lambda)q}{(1 - \lambda)q + \lambda r} \bar{x}_0 + \frac{\lambda r}{(1 - \lambda)q + \lambda r} \bar{x}_1 \right) + (1 - \alpha).$$

Thus allowing $p \neq r$ introduces type-specific network density, but does not change the main distinction between preference-weighted updating and pure social updating.

References

- Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi (2010) “Spread of (Mis)information in Social Networks,” *Games and Economic Behavior*, 70 (2), 194–227.
- Banerjee, Abhijit V (1992) “A simple model of herd behavior,” *The quarterly journal of economics*, 107 (3), 797–817.
- Benhabib, Jess, Alberto Bisin, and Matthew O Jackson (2010) *Handbook of social economics*: Elsevier.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992) “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of political Economy*, 100 (5), 992–1026.
- Bursztyn, Leonardo, Alessandra L González, and David Yanagizawa-Drott (2020) “Misperceived social norms: Women working outside the home in saudi arabia,” *American economic review*, 110 (10), 2997–3029.
- Cialdini, Robert B., Raymond R. Reno, and Carl A. Kallgren (1990) “A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places,” *Journal of Personality and Social Psychology*, 58 (6), 1015–1026.
- DeGroot, Morris H (1974) “Reaching a consensus,” *Journal of the American Statistical association*, 69 (345), 118–121.
- DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel (2003) “Persuasion bias, social influence, and unidimensional opinions,” *The Quarterly journal of economics*, 118 (3), 909–968.
- Golub, Benjamin and Matthew O Jackson (2010) “Naive learning in social networks and the wisdom of crowds,” *American Economic Journal: Microeconomics*, 2 (1), 112–49.
- (2012) “How homophily affects the speed of learning and best-response dynamics,” *The Quarterly Journal of Economics*, 127 (3), 1287–1338.
- Granovetter, Mark (1978) “Threshold Models of Collective Behavior,” *American Journal of Sociology*, 83 (6), 1420–1443.
- Katz, Daniel, Floyd Henry Allport, and Margaret Babcock Jenness (1931) “Students’ attitudes; a report of the Syracuse University reaction study..”
- Kuran, Timur (1991) “Now Out of Never: The Element of Surprise in the East European Revolution of 1989,” *World Politics*, 44 (1), 7–48.
- Miller, Dale T. (2023) “A Century of Pluralistic Ignorance: What We Have Learned About Its Origins, Forms, and Consequences,” *Frontiers in Social Psychology*, 1, 1260896.

- Morris, Stephen (2000) “Contagion,” *The Review of Economic Studies*, 67 (1), 57–78.
- Prentice, Deborah A. and Dale T. Miller (1993) “Pluralistic Ignorance and Alcohol Use on Campus: Some Consequences of Misperceiving the Social Norm,” *Journal of Personality and Social Psychology*, 64 (2), 243–256.
- Smerdon, David, Theo Offerman, and Uri Gneezy (2020) ““Everybody’s Doing It”: On the Persistence of Bad Social Norms,” *Experimental Economics*, 23 (2), 392–420.
- Smith, Lones and Peter Sørensen (2000) “Pathological outcomes of observational learning,” *Econometrica*, 68 (2), 371–398.
- Watts, Duncan J. (2002) “A Simple Model of Global Cascades on Random Networks,” *Proceedings of the National Academy of Sciences*, 99 (9), 5766–5771.
- Young, H. Peyton (2015) “The Evolution of Social Norms,” *Annual Review of Economics*, 7 (1), 359–387.