

Dynamic Competitive Persuasion

Mark Whitmeyer*

December 24, 2023

Abstract

Two long-lived senders play a dynamic game of competitive persuasion. Each period, each provides information to a single short-lived receiver. When the senders also set prices, we unearth a folk theorem: if they are sufficiently patient, virtually any vector of feasible and individually rational payoffs can be sustained in a subgame perfect equilibrium. Without price-setting, there is a unique subgame perfect equilibrium. In it, patient senders provide less information—maximally patient ones none.

Keywords: Bayesian Persuasion, Information Design, Martingales, Dynamic Games

JEL Classifications: C72; C73; D15; D63

*Arizona State University. Email: mark.whitmeyer@gmail.com. I began this paper, then called "A Game of Martingales," at UT Austin and continued to work on it while I was at the University of Bonn. The DFG (project ID 390685813) supported me during this period. I am grateful to Isaac M. Sonin for suggesting this project and to V. Bhaskar, Raphael Boleslavsky, Andy Kleiner, Joseph Whitmeyer, Thomas Wiseman, and Kun Zhang for their feedback. I also thank the editor, Tilman Börgers, and two anonymous referees for their advice.

1 Introduction

How does competition over the long run shape information provision? Two long-lived senders interact over an infinite time horizon and compete each period through information design. Each period, each sender chooses a Blackwell experiment (or signal) about her independent binary type with the aim of attracting a sequence of short-lived receivers.

We analyze two cases. In the first, the senders compete through information alone, and each period's receiver selects the sender whose expected type is the highest. In the second case, the senders are sellers who also post prices each period. There, a receiver (consumer) purchases from a seller if and only if its expected type net of its price exceeds that of the other seller and the consumer's outside option (normalized to 0).

The two setups (with and without prices) are spartan, yet include the following essential ingredients. There is competition, as only one sender is selected each period, to the detriment of the other sender. There is a tension between immediate and distant incentives, as the short-run objective of besting one's opponent in the current period's competition is tempered by the long-run ramifications of the current period's persuasion.¹ There is also—when the senders also set prices—an extraction motive: there is no longer a fixed pie to split. Is the repeated interaction between senders to the receivers' benefit? How does the price instrument alter things? What is the role of the senders' patience?

The environment with prices captures competition between two duopolists who not only charge prices but provide information to consumers about their products. The setting without prices also corresponds to some stylized cases itself. Consider, for example, two universities competing to place a representative student from their graduating classes in a prestigious job. Each year, the schools dictate their grading policies to their professors—they commit to a signal about the hidden state, the quality of the student. Year-over-year, the pool of students that comprises each school's student body stays roughly the same, which generates the sort of gradual learning present in our model. A

¹The state is persistent and the senders' memories long, so any information provided today will affect play forever.

second environment that fits is a political one. Political opponents frequently run against each other multiple times throughout their careers and choose how much to reveal about their positions. Our stipulation that memories are long rings especially true here: there are quite a few infamous examples of past remarks (information revelation) coming back to wreak havoc on a candidate's campaign.

Information revelation (or risk-taking) comes at a cost: the history of experiments, and their realizations (and prices) is public, which leads to a (martingale) sequence of beliefs. When the senders do not post prices, we derive the unique subgame perfect equilibrium of the game. The game is formally similar to an all-pay auction; and, just as participants in such auctions typically choose continuous distributions over bids, the senders here choose continuous distributions over their "bids," the posteriors that their experiments produce. Moreover, in contrast to the all-pay auction, where winning bids are more costly the higher they are, the senders here do not care about the realized posteriors, *ceteris paribus*, so the senders choose distributions over posteriors that yield the other a *linear* payoff. As in an all-pay auction, this linearity ensures that the other is indifferent over the distribution she chooses.

In a single period game of competitive persuasion, [Au and Kawai \(2020\)](#) find that competition increases information provision: as the number of senders increases, each provides more and more information; in the limit, full information. Here, (absent prices) we discover that patience has the opposite effect: as senders become more patient, each provides less information. As they become maximally patient, the sender with the highest expected type provides no information. The other sender chooses a binary distribution with support on 0 and the mean of the high sender, then provides no information forever onward, herself.

Information provision is a blunt tool: today's posterior is tomorrow's prior and any revelation constrains what can be done forever onward. Moreover, the game is constant-sum, so there is no scope for cooperation; thus, the two senders' equilibrium values are unique. As a result, the analysis of the dynamic game can be reduced to the study of a single period interaction in which each sender's payoff is today's payoff plus the expected (discounted) value of the game for that sender. As noted above, these payoffs must be

linear in the realized posteriors and also yield the equilibrium values when evaluated at the vector of priors. The resulting pair of differential equations pins down the unique equilibrium distributions chosen by each sender.

Why do senders provide less information as they become more patient? The driving force behind this is the concavity of a sender's equilibrium value in her prior. These values (which are the senders' continuation values) become more important—and so payoffs become more concave—as the senders grow patient. A sender's equilibrium value is concave in her prior due to the fact that the richness of the strategy space—a sender can choose any Blackwell experiment—allows the disadvantaged sender (the one with the lower prior) to partially imitate the advantaged sender. This ensures diminishing returns to being advantaged: a disadvantaged sender's value is linear in her prior, and an advantaged sender's value is linear in the inverse of her prior, ergo strictly concave.

Allowing the senders to post prices grants them a great deal of flexibility and transforms the game into one that is no longer constant sum. They can adjust prices freely, which enables them to reward and punish each other as needed, essential to sustaining cooperation. The latitude granted by price setting engenders a folk theorem: for any vector of average payoffs that is in the set of feasible and strictly individually rational payoff vectors, there exists a subgame perfect equilibrium of the game that yields such payoffs provided the sellers are sufficiently patient. In contrast to the information-only game, in which long-run incentives dampen information provision, we obtain the result in the price-setting scenario by constructing equilibria in which the firms provide full information immediately before cooperating in the resulting repeated Bertrand game.

1.1 Related Literature

There are a number of papers that study the problem of competitive persuasion (without prices) in a single period setting. They are [Boleslavsky and Cotton \(2015\)](#), who solve this problem for two players; [Albrecht \(2017\)](#), who sets the two-player competition in a political realm; [Au and Kawai \(2020\)](#), who extend the analysis to n players; and [Au and Kawai \(2019\)](#), who allow for correlated types. [Spiegler \(2006\)](#) solves a mathematically identical problem of firms competing for boundedly rational consumers.

Quite a few other works explore risk-taking contests in which agents with a fixed budget (or starting point) can “gamble” by choosing distributions on the positive portion of the real line. This problem is essentially equivalent to the competitive persuasion problem, with the minor difference that players choose distributions on $[0, \infty)$ rather than $[0, 1]$. [Wagman and Conitzer \(2012\)](#) and [Fang and Noe \(2016\)](#) focus on the martingale case, in which the expectations of the players’ distributions must be their budgets. In [Seel and Strack \(2013\)](#) the players choose when to stop Brownian motions with drift, and [Feng and Hobson \(2015\)](#) extend this to general diffusions.

The common theme throughout each competitive persuasion or gambling paper is the linearity of the equilibrium payoff function: players choose distributions that yield a payoff function for each other player that is linear in that player’s realization, leaving that player, herself, willing to choose an equilibrium distribution. As we discover in here, the necessity of the linear payoff persists. However, this payoff now includes the continuation payoff, which becomes increasingly important as players become more patient.

There are also a few recent investigations of dynamic information disclosure by a single sender. [Au \(2015\)](#) extends [Gentzkow and Kamenica \(2016\)](#) to such a setting: a sender wishes an agent to take one irreversible action and discloses information over time. Both players are long-lived, and as the players become maximally patient, full disclosure occurs virtually immediately. This contrasts starkly with our finding that maximally patient players are completely reticent. Also related is [Guo and Shmaya \(2018\)](#), who generalize [Au \(2015\)](#) and allow for correlated sender and receiver types.

A recent paper, [Koessler, Laclau, Renault, and Tomala \(2022\)](#), explores “long information design” in which persuaders disclose information at multiple stages about their (respective) states before the receiver chooses an action. As in this paper, they find that less information is disclosed when there are multiple stages of disclosure than in the one stage (single period) case. In addition, the equilibrium of this paper’s game coincides with the equilibrium in their product demonstration example when there is no deadline.

In contrast to the abundance of pure persuasion papers, only a few study competition in which sellers post prices and choose information structures. [Garcia \(2018\)](#), [Hwang, Kim, and Boleslavsky \(2018\)](#), and [Boleslavsky, Cotton, and Gurnani \(2017\)](#) all look at

variations of such a problem;² and [Au and Whitmeyer \(2020\)](#) and [Whitmeyer \(2020\)](#) investigate the effect of search frictions on this type of competition. The infinite horizon of this paper alters this competition drastically—just as long-run interaction and patience facilitate cooperation (and surplus extraction) in the standard model of Bertrand competition, so do they here.

2 The Game Without Prices

Time is discrete, indexed by the natural numbers, $t = 0, 1, 2, \dots$. There are two long-lived senders and a sequence of short-lived receivers, who have identical preferences. Each sender has a persistent binary type taking values $\{0, 1\}$ that is drawn prior to period 0 by nature and remains unchanged for the duration of the game. The senders' types are independent. x_0 (y_0) denotes the prior probability at period 0 that sender 1 (2) is value 1.

Each period, without knowing her (or the other sender's) type, each sender chooses a public signal about her type. Following the realizations of the two signals, that period's receiver selects the sender with the highest expected type (and randomizes fairly between the two if indifferent). We normalize the reward to a sender from being chosen to 1, and to 0 from not being chosen. The two senders discount the future by a common discount factor $\delta \in [0, 1)$. Each sender and receiver has perfect recall and observes the complete history of signals and signal realizations chosen by each sender.

Each sender has a compact metric space of signal realizations and each period t independently chooses a signal. As is now ubiquitous in this literature, we immediately reframe this problem as one of senders choosing Bayes-plausible distributions over posterior values: each period, given priors (x_t, y_t) , senders 1 and 2 choose random variables $X_t \sim F_t$ and $Y_t \sim G_t$ supported on $[0, 1]$ whose expected values are x_t and y_t , respectively.³

²The first two of these papers focus on symmetric settings and have the same timing as the stage-games of this paper: prices and information provision are determined simultaneously. In contrast, [Boleslavsky et al. \(2017\)](#) compare two different timings. In one, a firm chooses its information policy first (only one firm provides information in their duopoly model), which is observed before firms compete over prices. In the other, prices are chosen first, before the information policy is specified.

³It is without loss of generality to restrict the senders to pure strategies, since any randomization over

Naturally, a sender’s realized posterior in period t becomes the prior in period $t + 1$. \mathcal{F}_ω ($\omega = x, y$) denotes the set of Bayes-plausible distributions over posteriors.

Remark 1. It is without loss of generality to identify signals with Bayes-plausible distributions over posteriors.

Proof. Any sequence of signals induces a sequence of distributions over posterior beliefs. Consequently, we need only to argue that we do not “lose” equilibria by working directly with the sequence of beliefs. But this is nearly immediate: the game at hand is constant-sum, and in the proof of Proposition 2.1, we derive the two senders’ equilibrium values, which depend only on the vector of prior means. This means that the dynamic game can be reduced to a single-period game, in which each sender’s payoff is, moreover, posterior-separable—as today’s payoff, and the discounted continuation payoff (which is the expected discounted value of the game) depend only on the realized posterior today. This posterior-separability implies the remark. ■

2.1 Analysis of the Game without Prices

We begin by characterizing each sender’s equilibrium payoff. In the single period game of Boleslavsky and Cotton (2015), the sender with the higher mean (WLOG sender 1) obtains payoff $1 - \frac{y_0}{2x_0}$ and sender 2 gets $\frac{y_0}{2x_0}$. Here, as we discover, in any subgame perfect equilibrium of the infinite-horizon game, these are the two senders’ average payoffs.

Proposition 2.1. *In any subgame perfect equilibrium of the infinite horizon game, the two senders’ average payoffs are their payoffs of the one-shot game.*

The proof of this result may be found in the appendix but here is a sketch. It is clear that sender 2’s equilibrium payoff must be at least $\frac{1}{1-\delta} \left(\frac{y_0}{2x_0} \right)$: given any strategy by sender 1, sender 2 could guarantee herself this payoff by choosing a compound distribution in which she realizes 0 or x_0 , then, following the latter realization, mimics sender 1 thenceforth. However, it is *a priori* unclear whether sender 2 can be limited to just this payoff—the proposition verifies that this is indeed the case: we construct a strategy for

distributions is strategically equivalent to the resulting marginalized distribution.

sender 1 that limits sender 2 to this value. As the game is constant-sum, this pins down the unique equilibrium payoffs for both senders.

Proposition 2.1 allows us to reduce the dynamic game to a single-period game. In any period, given any distribution H_2 chosen by sender 2, sender 1 solves

$$\sup_{F \in \mathcal{F}_x} \left\{ \int_0^1 \Gamma_1(w) dF(w) \right\},$$

where

$$\Gamma_1(w) := H_2(w) + \frac{\delta}{1-\delta} \int_0^w \left(1 - \frac{z}{2w}\right) dH_2(z) + \frac{\delta}{1-\delta} \int_w^1 \frac{w}{2z} dH_2(z).$$

From here the rest of our analysis is easy. Boleslavsky and Cotton (2015) (as do, e.g., Fang and Noe (2016) and Au and Kawai (2020)) identify that the equilibrium payoff function must have a linear form.⁴ There can be no jumps—corresponding to interior atoms—in a sender’s payoff, since in her best response she prefers to place a mass just above the jump in her payoff; which, in turn, means that her opponent is not best responding. The support of the two senders’ equilibrium distributions must be the same (or else one would prefer to spread out the measure on the subset uninhabited by her competitor), nor can there be gaps in the support except possibly at the top (or else senders would want to place mass points). Similar arguments preclude Γ_1 (or sender 2’s payoff, Γ_2) from being strictly convex or concave.

Given the linear form, it remains to solve the differential equation(s) $\Gamma_i = \lambda_i + \eta_i w$ ($i = 1, 2$), accompanied by the conditions $\int_0^1 w dF(w) = x$, $\int_0^1 w dG(w) = y$, and that F and G are cdfs supported on (a subset of) the unit interval. This pins down the unique equilibrium distributions at every history. Defining $\mu := \left(1 + \sqrt{1-\delta}\right)^{-1}$, the sender with the higher current mean (WLOG sender 1) chooses a continuous, atomless, distribution on $\left[0, \frac{x}{\mu}\right]$ if $x \leq \mu$ and a distribution with a continuous, atomless, portion on $\frac{1-x}{(1-\delta)\mu}$ and a mass point on 1 if $x \geq \mu$. sender 2 partially mimics sender 1: she chooses sender 1’s distribution with probability y/x and places a mass point on 0 with its complement (which is needed in order to satisfy the Bayes-plausibility condition). Summing things up, and leaving the details of the distributions and payoff functions to the appendix,

⁴Refer to either Lemma 4.1 in Boleslavsky and Cotton (2015), the supplementary appendix in Fang and Noe (2016) or Theorem 1 of Au and Kawai (2020).

Theorem 2.2. *There is a unique subgame-perfect equilibrium of the game. It is Markovian, in which the state variable (x, y) is the current vector of priors. Each sender chooses a distribution that generates a payoff function for her opponent that is linear in the opponent's posterior.*

With Proposition 2.1 in hand, Theorem 2.2 is easy, and the economics behind the result is straightforward. In short, it is a standard “matching pennies” argument: each sender's continuous distribution must give the other a linear payoff so that she, herself, is willing to choose a continuous distribution. The necessity of this linear form then makes it clear why the proposition holds: the senders are indifferent over any distribution supported on $\left[0, \frac{x}{\mu}\right]$, two of which are the degenerate distribution supported on $\{x\}$ (for sender 1) and the distribution for sender 2 that (partially) mimics this, i.e., the distribution supported on $\{0, x\}$ with mean y .

As senders become more patient, the relative importance of besting one's opponent decreases. Indeed, consider first the immediate payoff (whose relative magnitude wanes with patience): for a fixed realization of a sender's opponent, z , it is a step function that jumps from 0 to 1 as soon as the sender's posterior exceeds z . On the other hand, the long-run payoff exhibits no such jump: for a fixed realization z it is linearly increasing in the sender's posterior on $[0, z)$, then strictly concave on $[z, 1]$; and is, therefore, concave on the entire unit interval. Moreover, this payoff function is not only continuous on $(0, 1]$ but differentiable. These observations suggest the next result: today's payoff encourages information provision, as each sender wants to provide a marginally higher posterior than her opponent; whereas the continuation payoff discourages such disclosure, due to its inherent concavity. Consequently, as the latter becomes more important, senders should provide less information.

The previous paragraph's intuition was mostly mechanical. It argues that the decrease in information provision as senders become more patient is driven by the increased importance of the continuation value, which is concave. But why is it concave? The fact that at equilibrium, the disadvantaged sender does a mimicking strategy—whereby its distribution corresponds to an initial binary distribution with support on 0 and the advantaged sender's mean, the latter of which is followed by an exact duplicate of the advantaged sender's distribution—means that the disadvantaged sender's equilibrium payoff is lin-

ear in its mean. Then, the constant-sum aspect of the game means that the advantaged sender's payoff is linear in the inverse of her mean; and is, therefore, concave. In short, the ability of the disadvantaged sender to imitate means that being disadvantaged is not so bad and that there are diminishing returns to being advantaged.

Proposition 2.3. *As senders become more patient they provide less information. That is, if $\hat{\delta} \geq \delta$ then the equilibrium distributions for discount factor δ are mean-preserving spreads of those distributions for discount factor $\hat{\delta}$.*

As the senders become maximally patient, the advantaged sender becomes completely uninformative and the disadvantaged sender's distribution converges to a binary distribution where one of the realizations leaves the receiver just willing to select her and the other realization reveals that her type is 0. To wit,

Corollary 2.4. *As $\delta \rightarrow 1$, the equilibrium distributions converge to the following: sender 1 chooses the degenerate distribution with support $\{x\}$ and sender 2 chooses the binary distribution with support $\{0, x\}$.*

In the limit, today's payoff is negligible, and all that matters is the continuation payoff, which is concave for any fixed realization of a sender's opponent. Consequently, the advantaged sender prefers not to release any information, whereas the disadvantaged sender's payoff is linear on $[0, x]$, leaving her willing to choose the specified binary distribution.

As noted in the introduction, the equilibrium described in Corollary 2.4 coincides with the equilibrium in the product demonstration application (with no deadline) of [Koessler et al. \(2022\)](#). This equivalence is implied by their Lemmas 9 and 10 (which follow from results in [Laraki and Renault \(2019\)](#)), which establishes that their demonstration game without a deadline can be approximated by a discounted (splitting) game as senders become maximally patient.

An earlier version of this paper focuses on a finite-horizon (T -period) model.⁵ In the unique subgame-perfect equilibrium of that game, obtained via backward induction, the

⁵Please refer to v4 of this paper on ArXiv.

senders' strategies and equilibrium values are the natural analogs of their equilibrium strategies and values here. Notably, as the horizon grows long ($T \uparrow \infty$), the equilibrium converges to the one identified in Theorem 2.2.

3 The Game With Prices

We carry over several of the assumptions from the game without prices. Namely, time is discrete and the game is played between two long-lived senders, who wish to entice a sequence of short-lived receivers. Now, the two senders are sellers and the receivers consumers. Each period each seller publicly chooses a Blackwell experiment *and* posts a price. Given a price p and realized value from a seller's good z , a consumer's utility from purchasing from that seller is $z - p$. For simplicity, we stipulate that the consumer's outside option is 0.

No longer do we limit the analysis to binary types: initially, at time $t = 0$, the values of seller 1's and seller 2's goods to the consumer are independent random variables $X_0 \sim F_0$ and $Y_0 \sim G_0$, respectively, that are each supported on subsets of the unit interval. In principle, this infinite-horizon game of price and information competition is a very difficult problem: each period's information provision decision by a firm determines tomorrow's prior, and because of this the firms' decisions do not reduce to choosing mean-preserving contractions of the prior, as they would if this were a single-period problem. However, as it turns out we can bypass these difficulties entirely.

Before doing so, let us reflect on a few things. If this were a market with just one seller, the optimal pricing and information policy for the (now) monopolist would be to never provide any information and charge a price equal to the good's expected quality. Trade is efficient—since the consumer always buys—and the seller extracts all of the surplus. Furthermore, if this monopolist were extremely patient, it would be almost optimal for her to immediately provide full information, then set a price equal to the consumer's value—as the discount factor approaches 1 the seller would approximate her optimal payoff.

On the other hand, with two sellers, it is not efficient for the sellers to provide no information. With two sellers, efficiency requires not only that the consumer always purchase

from a seller (since her outside option is 0) but also that the consumer purchase from the seller with the higher value good. Accordingly, it is efficient for both sellers to provide full information; and, therefore, if the sellers are extremely patient, it is approximately optimal for the sellers to provide full information immediately and for the higher value seller to sell to the consumer at a price equal to the consumer's value.

The (near-)optimality of immediate information revelation suggests an approach in the repeated game between the sellers. We should have the sellers immediately (in period 0) provide full-information-maximizing the available surplus-then cooperate over the remainder of the interaction in splitting the surplus. As it turns out, this also helps us in constructing the equilibrium: from $t = 1$ on, (on path) there is no longer uncertainty and commencing in any such subgame, the sellers are now in a standard repeated game. Moreover, it is easier to tackle unilateral deviations in the initial period, since the other (non-deviating) seller will have a much simpler strategy space from $t = 1$ on. Thenceforth, her prior is degenerate and therefore she can only affect payoffs through the price instrument.

Given distributions over valuations F and G , the maximal surplus is

$$S_{F,G} := \mathbb{E}[\max\{X, Y\}] .$$

Understanding $V_{i,min} = V_{i,min}(F, G)$ and $V_{i,max} = V_{i,max}(F, G)$ for $i = 1, 2$, we define

$$V_{1,min} := \int_0^1 \int_0^x (x - y) dG(y) dF(x), \quad V_{2,min} := \int_0^1 \int_0^y (y - x) dF(x) dG(y) ,$$

$$V_{1,max} := \int_0^1 x dF(x), \quad V_{2,max} := \int_0^1 y dG(y) ,$$

and

$$V(F, G) := \{(v_1, v_2) \in [0, 1]^2 : V_{1,min} < v_1 < V_{1,max} \ \& \ V_{2,min} < v_2 < V_{2,max} \ \& \ v_1 + v_2 < S_{F,G}\} .$$

$V_{i,min}$ is the average payoff a seller i can (approximately) guarantee herself when she is very patient. $V_{i,max}$ is seller i 's monopoly payoff, which is obviously the maximal payoff seller i can obtain. $V(F, G)$ is our analog of the feasible and individually rational set.

Theorem 3.1. *For any $(v_1, v_2) \in V(F_0, G_0)$ there exists a discount factor $\underline{\delta} < 1$ such that for any $\delta \geq \underline{\delta}$ there is a Subgame Perfect Equilibrium of the game in which the average payoff to seller i is v_i .*

The detailed construction and proof of this result may be found in the appendix. We assemble the equilibria by having both sellers immediately provide full information, then cooperate in one of the ensuing subgames.

Several more comments regarding Theorem 3.1 are in order. The theorem is an “anything goes” result for surplus splits, provided senders are sufficiently patient. It is not a complete characterization of equilibrium behavior. Indeed, it is easy to see that behavior other than the immediate full revelation we specified can be sustained at equilibrium. For example, provided senders are sufficiently patient, there are a variety of equilibria in which they provide less than full information in the first k periods then provide full information in the $(k + 1)$ th. In contrast, Theorem 2.2 is an exact characterization of the unique equilibrium strategies for arbitrary δ .

The reason why we can obtain such a sharp characterization when there are no prices but instead must settle for a folk theorem when there are prices is because the former game is constant-sum and the latter game is not; and because the senders in the latter game have aspects of their action sets that are not persistent. Obviously, the constant-sum aspect of the game without prices pins down the unique values for the two senders, negating any multiplicity. With prices, the set of feasible and individually rational payoffs is no longer a singleton. Moreover, today’s prices do not affect future prices, so senders have the requisite flexibility to cooperate.

3.1 One-Shot Interaction

Furthermore, *both* the infinite horizon (and patience) and the price instrument are needed to generate the equilibrium multiplicity described in Theorem 3.1. [Au and Whitmeyer \(2020\)](#) study a single-period, binary-value, directed search model in which sellers compete by providing information and setting prices. Proposition 3.7 of that paper reveals that there is a unique symmetric equilibrium in which sellers provide full information

and randomize over prices. The one-shot version of this paper’s game is a special case of [Au and Whitmeyer](#)’s setting, when the search cost is 0. [Hwang et al. \(2018\)](#) also identify a unique equilibrium in a game with price-setting and information provision. There, the distribution over posteriors has an alternating concave-linear form, where the nonlinearities arise due to constraints by the prior, and the prices are deterministic.

4 Discussion

Absent prices, the senders in this paper face a challenging problem. Each period, they are tasked with designing an experiment to entice a receiver, yet it and its realizations are fully observable and so affect the future beliefs of each player. These dynamic concerns engender caution. As senders become more patient, they choose distributions that are more and more *uninformative*.

Even though the game is infinite horizon, there is no scope for cooperation: any distribution that yields a sender a payoff that is not linear in her posterior gives her the opportunity to take (permanent) advantage of her opponent, and the constant-sum nature of the game means there can be no give and take between the senders. However, once price-setting is introduced, the permanence of information provision impedes cooperation no more. In fact, the sellers can extract maximal surplus by providing information immediately, then cooperating to eliminate the consumers’ rents.

Our construction of equilibria when prices are present seems realistic. That is, we achieve cooperation by having firms reveal information immediately, then proceed with a price-setting game. Anecdotally, this seems to accurately reflect the typical behavior in a product’s life cycle—there is an initial flurry of advertising for a new product, then little thereafter. Our result in the persuasion-only game also seems to correspond nicely to outcomes in the applications: for instance, the caginess of politicians is well-known, and schools, especially those whose input quality is highly esteemed, grade coarsely and/or provide little information about their students’ performance.⁶

⁶A stark example of this is prohibition of grade disclosure by many top MBA programs—Chicago (Booth), Columbia, UC Berkeley (Haas), Stanford GSB, and UPenn (Wharton).

References

- B.C. Albrecht. Political persuasion. *Mimeo*, February 2017.
- Pak Hung Au. Dynamic information disclosure. *The RAND Journal of Economics*, 46(4): 791–823, 2015.
- Pak Hung Au and Keiichi Kawai. Competitive disclosure of correlated information. *Economic Theory*, Jan 2019.
- Pak Hung Au and Keiichi Kawai. Competitive information disclosure by multiple senders. *Games and Economic Behavior*, 119:56 – 78, 2020.
- Pak Hung Au and Mark Whitmeyer. Attraction via prices and information. *Mimeo*, November 2020.
- Raphael Boleslavsky and Christopher Cotton. Grading standards and education quality. *American Economic Journal: Microeconomics*, 7(2):248–79, 4 2015.
- Raphael Boleslavsky, Christopher S Cotton, and Haresh Gurnani. Demonstrations and price competition in new product release. *Management Science*, 63(6):2016–2026, 2017.
- Dawei Fang and Thomas H. Noe. Skewing the odds: Taking risks for rank-based rewards. *Mimeo*, March 2016.
- Han Feng and David Hobson. Gambling in contests modelled with diffusions. *Decisions in Economics and Finance*, 38(1):21–37, 2015.
- Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.
- Daniel Garcia. Competitive information design. *Mimeo*, 2018.
- Matthew Gentzkow and Emir Kamenica. A rothschild-stiglitz approach to bayesian persuasion. *American Economic Review*, 106(5):597–601, 5 2016.
- Y. Guo and E. Shmaya. The value of multistage persuasion. *Mimeo*, December 2018.

Ilwoo Hwang, Kyungmin Kim, and Raphael Boleslavsky. Competitive advertising and pricing. *Mimeo*, 2018.

F. Koessler, M Laclau, J Renault, and T Tomala. Long information design. *Theoretical Economics*, 17(2):883–927, December 2022.

R. Laraki and J. Renault. Acyclic gambling games. *Mathematics of Operations Research*, *Forthcoming*, 2019.

Christian Seel and Philipp Strack. Gambling in contests. *Journal of Economic Theory*, 148(5):2033–2048, 2013.

Ran Spiegler. Competition over agents with boundedly rational expectations. *Theoretical Economics*, 1(2):207–231, 2006.

Liad Wagman and Vincent Conitzer. Choosing fair lotteries to defeat the competition. *International Journal of Game Theory*, 41(1):91–129, 2012.

Mark Whitmeyer. Persuasion produces the (diamond) paradox. *Mimeo*, 2020.

A Omitted Proofs

A.1 Proposition 2.1 Proof

Proof. The proof of this result is constructive. Recall the definition

$$\mu = \mu(\delta) := \frac{1}{1 + \sqrt{1 - \delta}}.$$

We also introduce the cdfs (with l and h representing “low” and “high,” respectively and $\xi := \max\{x, y\}$).

$$G_l^\xi(w) = \left(\frac{\mu w}{\xi}\right)^{\frac{1}{\sqrt{1-\delta}}} \quad \text{on} \quad \left[0, \frac{\xi}{\mu}\right], \quad (\text{A1})$$

and

$$G_h^\xi(w) = \begin{cases} (1-a) \left(\frac{(1-\xi\delta)\mu w}{(1-\xi)}\right)^{\frac{1}{\sqrt{1-\delta}}} & \text{if } 0 \leq w < \frac{(1-\xi)}{(1-\xi\delta)\mu} \\ 1-a, & \text{if } \frac{(1-\xi)}{(1-\xi\delta)\mu} \leq w < 1, \\ 1, & \text{if } w \geq 1 \end{cases} \quad (\text{A2})$$

where

$$a = a(\xi) := 2 - \frac{1}{\xi} - \left(\frac{\delta}{1-\delta} \right) \frac{(1-\xi)^2}{\xi}.$$

Let sender 2 play the following Markovian strategy. Each period she chooses distribution G (though henceforth we suppress the subscript) defined as follows. When $x \leq y$, $G = G_l^y$ if $y \leq \mu$ and $G = G_h^y$ if $y > \mu$. When $y < x$, $G = \frac{x}{y}G_l^x + 1 - \frac{x}{y}$ if $x \leq \mu$ and $G = \frac{x}{y}G_h^x + 1 - \frac{x}{y}$ if $x > \mu$.

Given this, sender 1 faces a Markov decision problem. Her payoff is (recursively):

$$V(x, y) = \sup_{F \in \mathcal{F}_x} \left\{ \int_0^1 G(w) dF(w) + \delta \int_0^1 \int_0^1 V(w, z) g(z) dz dF(w) \right\}.$$

Let us guess that the following strategy is an optimal strategy: if $x \geq y$, sender 1 chooses the degenerate distribution with support on $\{x\}$; and if $x < y$, she chooses a binary distribution with support $\{0, y\}$. A direct substitution and some algebra yields sender 1 a continuation payoff of $\frac{1}{1-\delta} \left(1 - \frac{y}{2x}\right)$ for realizations $x \geq y$ and $\frac{1}{1-\delta} \frac{x}{2y}$ for realizations $y > x$. Appealing to the one-shot deviation principle, sender 1 solves

$$\sup_{F \in \mathcal{F}_x} \left\{ \int_0^1 \Phi(w) dF(x) \right\},$$

where

$$\Phi(w) := G(w) + \frac{\delta}{1-\delta} \int_0^w \left(1 - \frac{z}{2w}\right) g(z) dz + \frac{\delta}{1-\delta} \int_w^1 \frac{w}{2z} g(z) dz.$$

Substituting in for G , we have when $\mu \geq x \geq y$,

$$\Phi(w) = \begin{cases} \frac{1}{1-\delta} \left[1 + y \frac{w-2x}{2x^2} \right], & \text{if } 0 \leq w < \frac{x}{\mu} \\ \frac{1}{1-\delta} \left[1 - \frac{y}{2x} (1 - \sqrt{1-\delta}) \right], & \text{if } \frac{x}{\mu} \leq w \leq 1 \end{cases},$$

when $\mu \geq y \geq x$,

$$\Phi(w) = \begin{cases} \frac{1}{1-\delta} \left[\frac{w}{2y} \right], & \text{if } 0 \leq w < \frac{y}{\mu} \\ \frac{1}{1-\delta} \left[\frac{1}{2\mu} \right], & \text{if } \frac{y}{\mu} \leq w \leq 1 \end{cases},$$

when $x \geq \max\{\mu, y\}$,

$$\Phi(w) = \begin{cases} \frac{1}{1-\delta} \left[1 + y \frac{w-2x}{2x^2} \right], & \text{if } 0 \leq w < \frac{(1-x)}{(1-x\delta)\mu} \\ \frac{1}{1-\delta} \left[1 + \frac{y \left(\frac{(1+\sqrt{1-\delta})(1-x)}{1-\delta x} - 2x \right)}{2x^2} \right], & \text{if } \frac{(1-x)}{(1-x\delta)\mu} \leq w < 1 \\ \frac{1}{1-\delta} \left[\frac{y}{x} \left(1 - \frac{a(x)}{2} - \delta \frac{x-a(x)}{2} \right) + 1 - \frac{y}{x} \right], & \text{if } 1 \leq w \end{cases},$$

and when $y \geq \max\{\mu, x\}$,

$$\Phi(w) = \begin{cases} \frac{1}{1-\delta} \left[\frac{w}{2y} \right], & \text{if } 0 \leq w < \frac{(1-y)}{(1-y\delta)\mu} \\ \frac{1}{1-\delta} \left[\frac{(1-y)}{2y(1-y\delta)\mu} \right], & \text{if } \frac{(1-y)}{(1-y\delta)\mu} \leq w < 1 \\ \frac{1}{1-\delta} \left[1 - \frac{a(y)}{2} - \delta \frac{y-a(y)}{2} \right], & \text{if } 1 \leq w \end{cases}$$

Evidently, our specified strategy is indeed optimal. Thus, sender 2 can always guarantee herself the payoffs $\frac{1}{1-\delta} \left(1 - \frac{x_0}{2y_0} \right)$ when $y_0 \geq x_0$ and $\frac{1}{1-\delta} \left(\frac{y_0}{2x_0} \right)$ when $y_0 \leq x_0$. Since sender 1 can do likewise, we conclude that the unique equilibrium payoff vector is as stated. ■

A.2 Theorem 2.2 Proof

Proof. The necessity of the linear payoff at equilibrium follows from [Boleslavsky and Cotton \(2015\)](#). Sufficiency is immediate by concavification. Finally, it is straightforward to verify that the distribution G (Expressions [A1](#) and [A2](#)) specified in [Proposition 2.1](#)'s proof is the unique solution to the differential equation that arises when specifying the linear payoff and imposing the various conditions (Bayes-plausibility and support). ■

A.3 Proposition 2.3 Proof

Proof. Fix means $y \leq x$. Let G be the equilibrium distribution chosen by seller 1 for discount factor δ and \hat{G} be its analog for discount factor $\hat{\delta} > \delta$. Let $x \leq \mu(\delta)$. It suffices to show that G crosses \hat{G} on $(0, 1)$ once from above on $\left(0, \frac{x}{\mu}\right)$. Obviously $G(0) = \hat{G}(0)$. Moreover, $\frac{x}{\mu} > \frac{x}{\hat{\mu}}$ and so G and \hat{G} intersect at least once on $\left(0, \frac{x}{\hat{\mu}}\right)$. It remains to show that this intersection point is unique. Replacing $\sqrt{1-\delta}$ and $\sqrt{1-\hat{\delta}}$ with β and $\hat{\beta}$, respectively, on $\left[0, x(1+\hat{\beta})\right]$, we define

$$T(w) := G(w) - \hat{G}(w) = \left(\frac{w}{x(1+\beta)} \right)^{\frac{1}{\beta}} - \left(\frac{w}{x(1+\hat{\beta})} \right)^{\frac{1}{\hat{\beta}}}.$$

Directly,

$$T'(w) = \frac{1}{w} \left[\frac{1}{\beta} \left(\frac{w}{x(1+\beta)} \right)^{\frac{1}{\beta}} - \frac{1}{\hat{\beta}} \left(\frac{w}{x(1+\hat{\beta})} \right)^{\frac{1}{\hat{\beta}}} \right].$$

Pick some $a \in (0, x(1 + \hat{\beta}))$ at which point $T(a) = 0$ (we argued above that such a point exists). Since $\beta > \hat{\beta}$, we have

$$0 = \frac{T(a)}{\hat{\beta}a} = \frac{1}{\hat{\beta}a} \left[\left(\frac{a}{x(1 + \beta)} \right)^{\frac{1}{\beta}} - \left(\frac{a}{x(1 + \hat{\beta})} \right)^{\frac{1}{\beta}} \right] > \frac{1}{a} \left[\frac{1}{\beta} \left(\frac{a}{x(1 + \beta)} \right)^{\frac{1}{\beta}} - \frac{1}{\hat{\beta}} \left(\frac{a}{x(1 + \hat{\beta})} \right)^{\frac{1}{\beta}} \right] = T'(a) .$$

Accordingly, whenever T crosses the axis on $(0, x(1 + \hat{\beta}))$, it must do so from above, which can therefore happen only once. The same steps can be followed to show that G crosses \hat{G} once from above when $x \geq \mu(\delta)$. ■

A.4 Theorem 3.1 Proof

Proof. At any t , the state vector is the pair of then-current prior distributions (F_t, G_t) . The following remark is useful. Suppose at some point each seller's quality has been fully revealed. Denote this period t^* —to clarify, this is the first period t such that the prior held by a consumer for each seller is degenerate.⁷ In that case, the state vector from period t^* on is (δ_a, δ_b) , for some $(a, b) \in [0, 1]^2$, where WLOG $a \geq b$. For any such vector, the maximal average total payoff for the two sellers is a . Then, in the parlance of [Fudenberg and Maskin \(1986\)](#) (henceforth FM86), the feasible and strictly individually rational set is

$$R(a, b) := \{x, y \in [0, 1]^2 : a - b < x \leq a \text{ \& } 0 < y \leq a - x\} .$$

By Theorem 1 in [FM86](#),

Remark 2. In the subgame commencing at period t^* in which the value realizations for sellers 1 and 2 are a and $b (\leq a)$, respectively, for any $(x, y) \in R(a, b)$ there exists a discount factor $\underline{\delta}$ such that for any $\delta \geq \underline{\delta}$ there is a subgame perfect equilibrium in which the vector of average payoffs for the two firms is (x, y) .

Given this, we construct the equilibrium by having the firms provide full information in the first period ($t = 0$) and charge $p = 0$, then play the appropriate subgame perfect equilibrium in the (from period 1 on) *repeated*—that is, no longer dynamic—game. This

⁷Shortly we will stipulate that $t^* = 1$ on path, i.e., it is only initially ($t = 0$) that a consumer's belief about a seller may be nondegenerate.

construction evidently yields the set $V(F_0, G_0)$. It remains to tackle deviations and specify strategies at subgames unreachable via unilateral deviations.

We can divide the remainder into the following three cases.

- **Case 1:** histories in which the state vector at some $t \geq 1$ is (δ_a, δ_b) for some $(a, b) \in [0, 1]^2$, and at least one sender deviated at some previous point.
- **Case 2:** histories in which the state vector at some $t \geq 1$ is (F, G) , where at least one of F or G is not degenerate, and in which both senders deviated previously.
- **Case 3:** histories in which the state vector at some $t \geq 1$ is (F, δ_b) ($b \in [0, 1]$), where F is not degenerate, and in which only sender 1 deviated at some previous point.

Now let us specify strategies in each case. We impose that at least one of F_0 or G_0 is not degenerate (or else [FM86](#)'s folk theorem yields the result immediately), in which case $t^* \geq 1$. Define $p_i(t)$ to be 1 plus the number of times seller i charged a non-zero price before period t .

Case 1: WLOG $a \geq b$. We impose that if both senders deviated at some preceding $t' < t^*$ they play a SPE from t^* on in which the vector of average payoffs is $\left(a - b + \frac{\eta_1^{p_1(t^*)}}{t^*}, \frac{\eta_2^{p_2(t^*)}}{t^*}\right)$ for some small $1 \gg \eta_1, \eta_2 > 0$. If only one sender deviated at some previous point, we stipulate that the vector of average payoffs from t^* on is $\left(a - b + \frac{\varepsilon^{p_1(t^*)}}{t^*}, b - \frac{\varepsilon^{p_1(t^*)}}{t^*}\right)$, if it was sender 1 who deviated, and $\left(a - \frac{\varepsilon^{p_2(t^*)}}{t^*}, \frac{\varepsilon^{p_2(t^*)}}{t^*}\right)$, if it was sender 2 who deviated, for some small $1 \gg \varepsilon > 0$. [FM86](#) guarantees that such SPE exist (starting from period t^*) provided δ is sufficiently large.

Case 2: We specify that both senders charge price $p = 0$ and provide full information. Consequently, the state in $t + 1$ is (δ_a, δ_b) (for some $a, b \in [0, 1]$), which is case 1, and so the vector of average payoffs from then on is either $\left(a - b + \frac{\eta_1^{p_1(t)}}{t+1}, \frac{\eta_2^{p_2(t)}}{t+1}\right)$ or $\left(\frac{\eta_1^{p_1(t)}}{t+1}, b - a + \frac{\eta_2^{p_2(t)}}{t+1}\right)$, depending on whether a or b is higher, and where $1 \gg \eta_1, \eta_2 > 0$ are small. On path, $p_i(t) = p_i(t + 1)$ for $i = 1, 2$.

The two senders' average payoffs are

$$u_1 := \delta \left\{ \frac{\eta_1^{p_1(t)}}{t+1} + \int_0^1 \int_b^1 (a-b) dF_t(a) dG_t(b) \right\},$$

and

$$u_2 := \delta \left\{ \frac{\eta_2^{p_2(t)}}{t+1} + \int_0^1 \int_0^b (b-a) dF_t(a) dG_t(b) \right\}.$$

Appealing to the one-shot deviation principle, suppose sender 1 deviates by charging $p' \in (0, 1]$ and choosing some signal.^a sender 1's payoff is bounded above by

$$(1-\delta) \int_0^1 \int_b^1 (a-b) dF_t(a) dG_t(b) + \delta \left\{ \frac{\eta_1^{p_1(t)+1}}{t+1} + \int_0^1 \int_b^1 (a-b) dF_t(a) dG_t(b) \right\},$$

which is less than u_1 for all sufficiently large δ . The same logic implies that sender 2 cannot deviate profitably to some $p'' > 0$. Any immediate gain will vanish as δ approaches 1 and her payoff in the cooperative portion will be discretely lower.

^aIt is clear that sender 1 cannot profit by a deviation in which she continues to charge price 0. This merely delays the cooperative (repeated-game) portion of the dynamic game.

Case 3: This is analogous to the previous case: we specify that both senders charge $p = 0$ and that seller 1 provides full information. The remainder proceeds in a kindred manner *mutatis mutandis*.

Finally, we need to check incentives in period 0. WLOG we do so for seller 1. Recall that seller 2 is providing full information. On path, seller 1 obtains some $v_1 \in (V_{1,min}, V_{1,max})$. If it deviates at $t = 0$, for any $\rho > 0$ there exists a $\hat{\delta} < 1$ such that for all $\delta \geq \hat{\delta}$, seller 1's average payoff is bounded above by

$$\int_0^1 \int_y^1 (x-y) dF_0(x) dG_0 + \rho.$$

Observing $V_{1,min}(F_0, G_0) = \int_0^1 \int_y^1 (x-y) dF_0(x) dG_0$ allows us to conclude the result. ■