Contracting for Persuasion*

Anand Bharadwaj † Ganesh Iyer ‡

October 2, 2025

Abstract

A seller contracts with an intermediary (influencer) who can design and transmit information to persuade consumers to purchase the seller's product. The intermediary's utility is a combination of pecuniary payoffs arising from monetary incentives offered by the seller, as well as reputational payoffs arising from enhancing consumer welfare. We characterize the optimal contract for a given price to show how it depends on the relative valuations for monetary versus reputational payoffs. Using a duality-based approach, we solve for the sender-optimal information design for a general class of two-action Bayesian persuasion problems. The seller's optimal incentive is in the form of a share contract (sales commission) and it is non-monotonic in the influencer's relative preference for monetary payoffs over consumer welfare. The equilibrium commission incentive first increases and then decreases with the influencer's relative preference for monetary payoffs. With endogenous pricing, we show that pricing and persuasion are substitutes: i.e., the seller optimally sets the sales commission to induce maximal persuasion when the influencer's relative weight on monetary payoffs is sufficiently high, and no persuasion at all when it is low. We also show that this setting of pricing and persuasion is isomorphic to one in which the influencer instead functions as a data broker, who is informed about downstream consumers' privately held valuations for the product.

Keywords: Influencer marketing, information design, optimal contracting
JEL Classification: D82, D83, D86, M31

*We are grateful to Yuichiro Kamada, Sam Kapon, Chris Shannon, Steve Tadelis, and Jesse Yao for detailed comments and suggestions. We also thank various seminar participants, Ivan Balbuzanov, Dongkyu Chang, Liang Guo, Barry Hughes, Gustavo Manso, Stephen King, Kannan Sethuraman, and Steve Williams for their comments.

†Department of Economics, University of California, Berkeley. Email: anand_bharadwaj@berkeley.edu.
‡Haas School of Business, University of California, Berkeley. Email: giyer@berkeley.edu.

1 Introduction

From social media creators, to marketing agents, stock analysts, property appraisers, to restaurant critics and fashion influencers, persuasive agents are pervasive in the modern economy. In particular, the growth of the digital economy and social media platforms over the past two decades has contributed to a significant market for persuasive intermediaries. The online social media influencer economy itself is currently valued at around $250 billion globally.[1] Since the influential paper by Kamenica and Gentzkow (2011), substantial literature has developed on the economics of persuasion. Yet the role of persuasive agents as intermediaries and the design of optimal incentive contracts to induce intermediated persuasion have remained largely unexplored.

Many important examples exist in observed market scenarios in which persuasion occurs through intermediaries, that is, where a principal or a firm seeks to persuade by using an intermediary who sends information to influence receivers. Consider the following common examples:

- A firm seeks to sell its product to a consumer, and a product reviewer may provide information about the product's quality or suitability for the customer. Reviewers may be provided monetary incentives by the seller to influence purchases, and they must balance the pecuniary gains from sales commissions with the potential reputational harms arising from recommendations that are not in the interests of consumers.

- Credit rating agencies or securities analysts may have implicit or explicit conflicts of interest with the firms whose creditworthiness they are analyzing. For example, firms may steer future business to investment banks who rate their securities favorably. Again, the intermediary balances the rewards from greater persuasion with the concomitant risks to their credibility.

- Venture capitalists, such as Y Combinator, often take fixed-price equity stakes in star-

[1] "The $250 Billion Influencer Marketing Economy is Booming" (*Forbes*, 28 October 2024).

tups. The dual role of such intermediaries as screeners of project quality and as persuasive agents has been previously explored; see, for example, Mullainathan and Shleifer (2005) and Azarmsa and Cong (2020) for a discussion. Fundamentally, startups face a trade-off between selling more equity (and thereby providing stronger marketing incentives to the venture capitalist), versus retaining a greater fraction of equity but lowering the probability of a successful exit.

This trade-off between information provision and monetary incentives provided to influencers has also been a point of focus in recent policy debates. In essence, while influencers may provide economic benefits by facilitating the transmission of valuable information, thereby allowing consumers and decision makers to make better decisions, they may also have incentives to misrepresent information, especially when they have a financial interest in the products that they are advertising. These concerns have motivated the recent updates to the FTC guidelines, which require influencers to "clearly and conspicuously" disclose any material relationship to the brands they are endorsing. One might ask whether regulation of this type actually achieves its desired purpose, or might it indirectly cause harm by reducing incentives for efficient information transmission?

Motivated by these issues, we study the optimal design of contracts between a seller and a persuasive intermediary, who, in turn, transmits an information signal that influences the downstream purchase decision of a consumer. As Kotler et al. (2018) suggest, the trade-off between reach and cost (particularly in terms of lost revenue through sales commissions) is an important economic criterion in choosing an integrated marketing channel. Our model analyzes the mechanisms and the implications of this reach-cost trade-off within the Bayesian persuasion framework.

We consider a setup with three players: a seller who wants to sell a product of uncertain value to a buyer and an intermediary who can transmit information about the product's value to the buyer. The seller can design a contract that offers monetary incentives to the influencer to affect the information that is transmitted to the buyer. The value of the product

to the buyer (this could be thought of as the valuation of product quality or an idiosyncratic match value) is given by v, which is distributed on $[0, 1]$ according to the distribution F. The buyer will buy if $\mathbb{E}(v) \geq p$ and will not buy otherwise. We assume that the influencer has information about this value, but that it is not known by the buyer or to the seller. Consistent with the standard Bayesian persuasion setup, the influencer can commit to revealing a state-contingent signal to the buyer. The influencer cares about both her pecuniary payoff and her reputation with the buyer; that is, the payoff is a convex combination of the monetary compensation R and the buyer's *ex-post* utility. To this end, the influencer's *ex-post* utility for realized value v when the buyer purchases is given by $\alpha R + (1 - \alpha)(v - p)$ and is zero if the buyer does not purchase, where α represents the relative preference of the influencer for their monetary payoff vis-à-vis consumer utility. The parameter α has several real-world interpretations. For example, it may capture the relative preference for immediate versus long-term payoffs. It could also capture the lifecycle stage or scale of an influencer. In fact, it would be reasonable to posit that nascent or small-scale (micro) influencers exhibit a lower value of α and place greater weight on reputation building relative to their monetary compensation.

We first characterize the influencer's optimal persuasion scheme in response to a particular reward structure chosen by the seller. The analysis provides a general solution to the optimal persuasion scheme chosen by the influencer for the class of two-action persuasion problems, which has parallels to the structure of the optimal hypothesis test as per the Neyman-Pearson lemma. We then characterize the seller's optimal incentive contract for the influencer that is based on observable outcomes, namely, the buyer's purchase decision. Such a contract amounts to a flat transfer plus variable reward scheme (similar to a franchising contract), where the influencer is rewarded for a successful sale. The main analysis considers contracts under the usual limited liability assumption, which implies non-negative monetary payoffs for the influencer independent of whether a sale is realized or not. Indeed, without limited liability (and when the influencer places a sufficiently high weight on pecuniary pay-

offs), the seller can use the variable part of the contract to induce the maximum feasible amount of persuasion, while extracting surplus from the influencer up to their reservation profit using the fixed fee. But with limited liability, such a contract is infeasible, meaning that the reward offered by the seller will not only apply to the marginal sale, but also the inframarginal sales, creating a trade-off between increasing sales through persuasion and the influencer siphoning revenue through a higher equilibrium reward.

The first result is that the optimal variable reward varies non-monotonically with the influencer's relative preference for monetary incentives versus reputation, and we identify simple conditions under which R exhibits unimodal comparative statics. Initially, as α increases from zero, the influencer becomes more sensitive to monetary incentives. In response, the seller increases the variable incentive R. In turn, this "incentive effect" leads the influencer to commit to a scheme which increases the equilibrium amount of persuasion (i.e., the probability with which the buyer is rationally induced to purchase). The reward R continues to increase until a threshold of α at which the influencer generates the maximum feasible amount of persuasion. Beyond this threshold, as α increases, the seller actually finds it optimal to decrease the variable incentive even as the influencer's relative preference for monetary incentives increases. This is due to a "satiation effect" – because the influencer's persuasion scheme already generates the maximum amount of persuasion, any further increase in the preference for monetary incentives means that the seller can afford to offer progressively weaker incentives and yet induce the persuasion scheme that maximizes the probability of buyer purchase.

Next, we consider the scenario where the seller endogenously chooses and commits to a retail price prior to contracting with the influencer. This allows us to examine how incentives and persuasion interact in inducing the consumer to buy the product. In contrast to the case in which prices are fixed, the equilibrium contract involves the seller either opting for full persuasion or no persuasion at all. When α is sufficiently small and the influencer cares less about monetary payoffs compared to reputation, using persuasion to induce the buyer to

purchase is relatively more "expensive" for the seller than reducing the price. Consequently, the seller eschews the variable incentive (sets $R = 0$) and endogenously abstains from persuasion. However, when the influencer's preference for monetary payoffs is sufficiently high, it then becomes possible for the seller to use the variable incentive in the contract to induce the maximum amount of persuasion and to charge the maximum feasible price that extracts consumer surplus. Persuasion, in other words, becomes relatively less expensive than pricing. Thus, in our setup, the seller never uses partial persuasion, and in this sense, pricing and persuasion can be seen as substitutes in making the consumer buy the product.

Finally, we study the relationship of this problem with an ostensibly different but economically important setting – i.e., data brokerage. This setting refers to the case where consumers are privately informed about their valuations. The influencer is informed of consumer valuations and can therefore function as a *data broker* between the consumer and the firm. Suppose that contracts can be written conditional on the buyer's ultimate purchase decision and the influencer accordingly designs an optimal information transmission scheme. Interestingly, we show that this data brokerage problem is isomorphic to the pricing/persuasion problem with limited liability.

1.1 Related Research

Our paper contributes to the Bayesian persuasion literature pioneered by Kamenica and Gentzkow (2011) by considering the role of intermediaries who act as persuasive agents. This in turn enables us to establish the optimal contracts that a firm can design to govern the incentives of a persuasive intermediary. Thus, our analysis is distinct from the standard Bayesian persuasion literature with direct communication by a principal (sender) to persuade a receiver.

There is a related research stream on intermediated cheap talk and Bayesian persuasion, including Ambrus et al. (2013) and Corrao (2023). The focus of mediation in this research is the garbling of signals by intermediaries in a communication chain to influence a receiver.

There is no garbling of this type in our setup; rather, the principal (seller) must design a contract to induce an appropriate message policy from the intermediary in order to optimally alter the beliefs and influence the action of an agent/consumer. In this environment, we provide a general characterization of the structure of optimal contracts and the associated persuasion scheme of the intermediary.

The basic Bayesian persuasion framework has been extended in Kolotilin (2018) and Dworczak and Martini (2019), who develop a linear programming method to tackle persuasion problems and derive duality-based conditions for optimal signals when the sender's utility depends solely on the receiver's posterior expectation of the state. Since our analysis of intermediated persuasion focuses on a binary action (buy/no buy) setup, our analysis is related to Kolotilin (2015) and Kolotilin et al. (2017), which also characterize optimal persuasion schemes in binary action environments. Here we develop results pertaining to settings with strategic intermediaries where the sender's utility may depend on the action-state pair using a duality approach and show a connection to the Neyman-Pearson lemma, thereby providing a "sensitivity analysis" of the sender-optimal experiment in response to changes in their payoff structure. We also develop a simple and extensible proof strategy to demonstrate weak duality and attainability, leveraging the continuity of the constraint value function.

Several other papers have explored themes that may be seen as related to those we study here, especially with regard to how monetary incentives can lead to misleading information transmission. These include Inderst and Ottaviani (2009), who examine the inherent tension faced by marketing agents who have to increase sales while at the same time providing advice to consumers about the suitability/fit of the product for their needs. Krishna and Morgan (2008) examine the structure of optimal contracts in a principal-agent setup where the agent is the informed party. Also related is the work by Pei and Mayzlin (2022), which investigates the effects of affiliation relationships with influencers. The influencer in our model has a payoff structure similar to Durbin and Iyer (2009), in terms of the intermediary being

concerned with a combination of their pecuniary payoffs and their reputational concerns. However, because we study an environment with Bayesian persuasion, there is no scope for inference about the "corruptibility" of the influencer, since their commitment to a signal structure reveals this entirely.

The analysis is also related to the body of work on conditions for monotonicity and comparative statics for Bayesian persuasion mechanisms (as explored by Ivanov (2021), Mensch (2021) and Curello and Sinander (2022)), as well as classic research on monotone comparative statics (Milgrom and Shannon, 1994). In some sense, our results are related to both of these strands of the literature. Our findings about the optimal "bundling" of states in two-action Bayesian persuasion problems are thematically similar to the former body of work, and we show that pure monotonicity emerges as a special case of these results. However, given the setting of this paper, we focus our attention on providing a full characterization of the structure and comparative statics of persuasion schemes in a two-action environment. Later, when we study the comparative statics of the optimal *persuasion contract* with respect to the influencer's relative welfare weights on their pecuniary reward and on consumer surplus, we build on the monotone comparative statics literature. Here, some of the arguments need to be adapted to account for a contracting domain, and we derive quite general conditions under which we have *unimodal* comparative statics in such an environment.

The remainder of this paper will be structured as follows. Section 2 introduces the setup and derives the key results for the baseline model. Section 3 endogenizes prices and thereby explores the interplay between pricing and persuasion. Section 4 develops further extensions and discusses connections between the intermediated persuasion and data brokerage problems. Finally, Section 5 concludes the paper.

2 Model and Analysis

Consider a model with a seller who wants to sell a product of uncertain value to a buyer and an intermediary who is an influencer and who can transmit information about the product's value to the buyer. The seller offers the product to the buyer at price p, which is assumed to be fixed and exogenous for now, but that assumption will be relaxed later. The buyer's value (which can be thought of as an absolute quality or an idiosyncratic match value) is given by v, which is distributed on the state space $\Omega = [0, 1]$ according to a continuously differentiable distribution F with associated density function f. This value is not known by the buyer or to the seller (later, we analyze an alternative data brokerage setting where the value is privately known by the buyer), but it is known to the influencer. This informational advantage creates a market for persuasion facilitated by the influencer. In the spirit of the Bayesian persuasion setup, the influencer can commit *ex-ante* to communicate a state-contingent signal to the buyer, as in Kamenica and Gentzkow (2011). Here, a signal (interchangeably referred to as an experiment) denotes a mapping from Ω to the space of probability distributions over a set of messages M, that is, $\Delta(M)$.

The buyer's utility CS is quasilinear and is normalized at 0 if they do not purchase and is given by $v - p$ if they do purchase the product. Therefore, the buyer will buy (action B) if, in expectation, $v \geq p$ and will refrain from buying (action N) otherwise. As is standard in the literature, we focus on *sender-preferred perfect Bayesian equilibria* of the information transmission subgame.[2] When the buyer purchases, the influencer's payoff is given by $\alpha R + (1 - \alpha)CS$, where $\alpha \in [0, 1]$ parameterizes the relative preference weights that the influencer places on their pecuniary payoff (R) and the consumer surplus (CS).

[2]It is of course not always without loss of generality to consider sender-preferred equilibria, but it can be shown, using arguments from Lipnowski et al. (2024), that there is a unique equilibrium in the two-action Bayesian persuasion subgame that we study in this paper. Nevertheless, we invoke the sender-preferred equilibrium assumption to simplify the exposition and present it in a manner that would generalize to a broader class of persuasion problems. The assumptions in this paper broadly accord with the principle in both the information and mechanism design literature to focus on equilibria that maximize the designer's payoff. Indeed, we also implicitly assume that the equilibria in the full game (including the contract proposal) are optimal for the seller.

2.1 Characterization of Optimal Signals

Suppose that the seller offered no monetary compensation to the influencer. In such a case, the influencer will always report the true state, and the seller will only effect a sale in states above the price. In general, the seller can contract with the influencer to affect the information transmission and potentially alter the buyer's purchasing decision. Consider the standard contracting situation in which the seller can contract on the buyer's observable purchase decision, that is, contracts that have a reward R for the influencer if the buyer actually purchases the product and a fixed fee transfer T which is independent of purchase.

Prior to identifying what the optimal contract is for the seller, we must identify how the intermediary's persuasion scheme varies with the contractual terms. As in Kamenica and Gentzkow (2011), it suffices to consider $M = A \equiv \{B, N\}$; that is, we only need a message space as large as the action space and can restrict attention to straightforward signals that produce the equilibrium recommended action which the receiver always follows (this is essentially an analogue of the revelation principle for this setup).[3] Apart from characterizing the influencer's optimal signal in our setup, we will actually solve a more general two-action persuasion problem that subsumes the framework of this paper.

We proceed to derive a novel characterization of optimal signals in two-action Bayesian persuasion problems where the utility of the influencer (sender) is dependent on the state-action pair. In this more general persuasion problem, we still have a state space $[0, 1]$ and associated continuously differentiable distribution function F and a binary action space,

[3] As Kamenica and Gentzkow (2011) show, there is no loss of generality if we restrict attention to straightforward signals (see their Proposition 1). However, it may not be immediately obvious why it also suffices even in the presence of an intermediary. Now, contracts between the seller and the influencer essentially consist of an action-contingent transfer rule coupled with an incentive-compatible signal recommendation by the influencer. As long as the set of straightforward signals is contained within the space of candidate signals available to the influencer and the utilities of the players do not depend directly on the chosen signal, any non-straightforward signal recommendation can be replaced with the corresponding straightforward signal without loss of generality. In our setting, these conditions hold, since there are no explicit restrictions on the signal space and the utilities are solely dependent on the state and the action. Even if contracts were also allowed to depend on the influencer's message, any recommended signal with more than two messages could be replaced with an alternative signal solely consisting of the maximal-transfer messages inducing actions B and N respectively, preserving the validity of the argument.

where the actions are, without loss of generality, called "buy" (B) and "don't buy" (N). We assume that the sender (influencer) receives a payoff that varies depending on the action-state combination, that is, the sender's utility is $u(v, a)$, where v is the state and a is the receiver's chosen action. The receiver (buyer) chooses action B if their posterior expectation of the state exceeds a given cutoff $p < 1$ and chooses action N otherwise. We also assume that $p < \mathbb{E}(v)$, so that the receiver does not choose action B in the absence of information. We normalize $u(v, N) = 0$ for all v and also assume that $u(v, B)$ is bounded and continuous in v, so that we can write $u(v) \equiv u(v, B)$. Within the motivating framework of this paper, we interpret p as a price, but it could simply represent a cutoff for expectation-based persuasion of the receiver.

In this two-action environment, the choice of an optimal signal structure is equivalent to choosing a measurable function $\beta : [0, 1] \to [0, 1]$, where $\beta(v)$ denotes the probability that the "buy" message will be transmitted for a given state realization v, so that the "don't buy" signal will be transmitted with complementary probability $1 - \beta(v)$. Given that the optimal experiment must respect Bayes plausibility, the problem reduces to finding the function $\beta(v)$ that corresponds to the optimal signal. The sender attains a payoff of $u(v)$ when the receiver chooses action B and receives payoff zero otherwise. The influencer's optimization problem is as follows:

$$\begin{aligned} &\max_{\beta(.)} \int_0^1 u(v)\beta(v)\, dF(v) \\ \text{subject to} \quad &\int_0^1 (v - p)\beta(v)\, dF(v) \geq 0 \\ &\beta(v) \in [0, 1] \ \forall \ v \in [0, 1] \end{aligned}$$

Therefore, the Lagrangian is given by the integral of $(u(v) + \lambda(v - p))\beta(v)$ over $[0, 1]$. We then invoke arguments based on Lagrangian duality to show that the optimal $\beta(v)$ follows a zero-one structure; that is, it takes the value of 1 over a set of states S containing values of v that make the integrand of the dual function $u(v) + \lambda(v - p)$ positive for some $\lambda \geq 0$, and

it takes the value of 0 in all other states (with the possible exception of the boundary states v such that $u(v) + \lambda(v - p) = 0$, for which we may employ a probabilistic tie-breaking rule with an interior value of β if this set has positive measure). Formally, these arguments can be summarized in the following proposition, where we have defined $S(\lambda)$ as the set of states v such that $u(v) + \lambda(v - p) > 0$.

Proposition 1. *In this general two-action Bayesian persuasion problem, there exists an optimal persuasion scheme for the sender under which they send message B with probability 1 for the set of states $S(\lambda^*)$, which consists of all states v such that $u(v) + \lambda^*(v - p) > 0$, where $\lambda^* = \inf_{\lambda \geq 0} \int_{S(\lambda)} (v - p)\, dF(v) \geq 0$, and message N with probability 1 for all states v such that $u(v) + \lambda^*(v - p) < 0$.*

Proof. See the Appendix. □

We begin with a sketch of the analysis strategy to facilitate understanding of the results. First, we can specify the primal optimization problem for the influencer as follows:

$$\max_{\beta(.)} \min_{\lambda \geq 0} \int_0^1 (u(v) + \lambda(v - p))\beta(v)\, dF(v)$$

The associated Lagrangian dual problem is:

$$\min_{\lambda \geq 0} \max_{\beta(.)} \int_0^1 (u(v) + \lambda(v - p))\beta(v)\, dF(v)$$

The optimizer of the dual problem depends on the Lagrange multiplier λ. By weak duality, the dual is an upper bound for the solution to the primal optimization problem. We show that there exists some value of λ such that we can construct a corresponding solution for the dual for which the constraint exactly binds. Ignoring issues related to the tiebreaking rule, this essentially amounts to showing that the value of the constraint is continuous in λ, which yields the result of the intermediate value theorem. If we have a "jump" at λ^* in the constraint value function, we can handle this by having the sender randomize over messages for the boundary states. But such an optimizer must also be feasible in the primal problem, which, combined with weak duality, verifies

that it solves the primal problem.

Several points are worth noting here. First, we can generically ignore the tiebreaking issue whenever the boundary set has zero measure (this will be satisfied if $u(v)$ is not an affine function of v with negative slope over an interval with positive measure). Second, if $u(v)$ is increasing in v, it follows that a monotone pooling rule, where $S = [q, 1]$ for some $q \in [0, 1]$ is optimal as indicated in the proposition. Moreover, if $u(v) \geq 0$ for all v, the result simply amounts to including all states where $v \geq p$ and ordering the remaining states by $\frac{u(v)}{p-v}$ in descending order, pooling such states up to the point where the constraint exactly binds. This latter special case illustrates the connection of this proposition to the Neyman-Pearson lemma, which is also, at its heart, an allocation problem of this sort, but where the values $u(v)$ and costs $p - v$ are replaced with the likelihoods under the alternative and null hypotheses respectively, since the Neyman-Pearson optimization problem corresponds to maximizing the statistical power of a test for a given size.[4] Much like Berliant (1984), our result illustrates the application of the optimization-based ideas in the Neyman-Pearson lemma to the context of strategic communication and develops an analysis strategy that allows duality to be applied in an infinite-dimensional space.[5]

In short, the alignment condition for monotonicity is simply that $u(v)$ increases in v, whereby the sender would choose a monotone message scheme. This is intuitive: if the cheapest messages to send (that is, the ones that are most aligned with the downstream receiver's preferences) are

[4]The Neyman-Pearson lemma (Neyman and Pearson, 1933) characterizes optimal hypothesis tests and can be stated as follows: Let P_0 and P_1 be candidate probability distributions for a random variable X with respective densities $p_0(x)$ and $p_1(x)$. Let a *test* denote a function $\phi : \Omega \to [0, 1]$. We say a test has *significance level* or *size* α if $\mathbb{E}_0[\phi(X)] \leq \alpha$, and that a test is a *uniformly most powerful* (UMP) test with size α if it maximizes $\mathbb{E}_1[\phi(X)]$ among the class of tests with size α. The Neyman-Pearson lemma states that a test ϕ is UMP of size α if and only if it satisfies

$$\phi(x) = \begin{cases} 1 & \text{if } p_1(x) > kp_0(x), \\ 0 & \text{if } p_1(x) < kp_0(x), \end{cases}$$

for some constant k, and $\mathbb{E}_0[\phi(X)] \leq \alpha$. In this paper, we develop an analogous duality-based proof strategy for the study of Bayesian persuasion problems.

[5]An interesting question is whether similar results of this type would hold beyond the two-action environment that we study. The primary impediment to this generalization is that there does not appear to be an obvious way to prove attainability of the Lagrange dual function within the feasible space for the primal problem and to tractably characterize the associated Lagrange multipliers. Accordingly, it is also worth noting that there is no generalization of the Neyman-Pearson lemma that extends to the case of more than two hypotheses. However, the two-action assumption covers a range of practically relevant situations, including the binary purchase decision of a consumer that we study in this paper.

also the most preferred by the sender, the chosen persuasion scheme must indeed be monotone. The proposition shows that we can generalize this notion of monotonicity to a setting without $u(v)$ being increasing in v, whereby we might wish to include a "bad" state v if it slackens the Bayes plausibility constraint.

Let us return after that digression to the characterization of how the influencer's optimal persuasion scheme varies with the terms of the contract, assuming that the influencer does accept the contract. Since the influencer's utility function $\alpha R + (1-\alpha)(v-p)$ is monotonically increasing in v, the aforementioned monotonicity condition is satisfied, and the seller optimally chooses $s = [q, 1]$, where we call q the *persuasion threshold* and $1 - F(q)$ the *persuasion quantity*. Since the influencer always transmits the "buy" signal when it is optimal for the buyer to purchase, we can write $q = p - y$, where $y \geq 0$. Moreover, the maximum quantity of persuasion is achieved for the largest y such that Bayes plausibility is still respected; i.e., y is the maximum value such that $\mathbb{E}(v|v \geq p - y) \geq p$. In response to a given contract, the influencer chooses their optimal threshold $p - x$, where $x \leq y$, so as to equate the weighted marginal monetary reward with the weighted harm to the consumer, whereby $(1-\alpha)x = \alpha R$ (assuming that $(1-\alpha)y \geq \alpha R$, since otherwise the influencer will just choose $x = y$). Hence, $x = \min\left(\frac{R\alpha}{1-\alpha}, y\right)$.

2.2 Optimal Contracts and Comparative Statics

Given that we have now characterized how the persuasion scheme (experiment) chosen by the influencer varies in response to any contract, we can return to the seller's problem of optimally designing the contract with the influencer.[6] A useful benchmark is the setting where the influencer's persuasion scheme is directly contractible. In that case, to induce a given value of x on the part of the influencer, the seller must compensate them for the scaled lost consumer surplus. The pointwise compensation payment for a given negative consumer surplus state is $\frac{1-\alpha}{\alpha}(p - v)$, so that the total payment is $\int_{p-x}^{p} \frac{1-\alpha}{\alpha}(p - v)\, dF(v)$. Since the marginal value of an additional sale is p, the seller optimally chooses $x = \min\left\{x \;\middle|\; 0 \leq x \leq y \text{ and } p \geq \frac{1-\alpha}{\alpha}x\right\}$. Thus, the seller optimally induces $x = y$

[6] The influencer's outside option if they decline the offered contract is to choose a fully revealing experiment, which would yield the influencer the payoff of $(1-\alpha)(v-p)$ in the states with positive consumer surplus. Thus, any acceptable contract must respect the participation constraint of giving the influencer at least the payoff $(1-\alpha)\int_p^1 (v-p)\, dF(v)$.

if $p \geq \frac{1-\alpha}{\alpha} y$, for which a sufficient but not necessary condition turns out to be $\alpha > \frac{1}{2}$ (that is, the influencer values their pecuniary payoff more than the welfare of the consumer). Otherwise, the seller sets $x = \frac{\alpha}{1-\alpha} p$.

Let us now consider the main analysis where the seller can contract only on the observable action or the purchase decision of the buyer (for example, as in actual commission contracts). We consider two sets of assumptions for this contract. In the first case, which we term a *franchising contract*, we have an unrestricted purchase-contingent contract, with a flat transfer T in conjunction with a commission or reward R when the buyer purchases. In the case of unlimited liability if negative values of T are permissible (that is, the influencer can be charged a franchising fee by the seller), it is possible to use the contract to replicate any direct contract on the influencer's persuasion scheme. The contract structure consists of the minimal reward R that incentivizes the influencer to choose the designated cutoff $p - x$. Hence, R is set so that $\alpha R + (1 - \alpha)(-x) = 0$; that is, $R = \frac{1-\alpha}{\alpha} x$, and the baseline transfer T is the flat fee that makes the influencer exactly indifferent between engaging in the designated level of persuasion and not accepting the proposed contract. That is, $\alpha \left[T + R(1 - F(p - x))\right] = (1 - \alpha) \int_{p-x}^{p} (p - v)\, dF(v)$, so that $T = \int_{p-x}^{p} \frac{1-\alpha}{\alpha} (p - v)\, dF(v) - R(1 - F(p - x))$.

Next, consider the case of *limited liability* for the influencer (i.e., the influencer must receive a non-negative monetary transfer irrespective of whether the buyer purchases or not). In that case, the reward not only affects the marginal "unit", but also the inframarginal sales, engendering a price-quantity trade-off. We begin by observing that it suffices to consider flat-reward contracts, that is, contracts that pay zero when no sale is made and pay a non-negative reward R otherwise. We can characterize the optimal contract, that is, the seller's optimal choice of R, as follows. First, note that there is no benefit to the seller from setting a reward above the smallest value R^u that induces $x = y$ ($\frac{(1-\alpha)y}{\alpha}$), so that the relevant domain for R is $\left[0, \frac{(1-\alpha)y}{\alpha}\right]$.

For a given choice of R in the range $\left[0, \frac{(1-\alpha)y}{\alpha}\right]$, the seller gets payoff

$$\pi(\alpha, R) = \left(1 - F\left(p - \frac{R\alpha}{1-\alpha}\right)\right)(p - R)$$

Note that the values of this objective function at the left and right endpoints are $p(1 - F(p))$ and

$(1 - F(p - y))\left(p - \frac{(1-\alpha)y}{\alpha}\right)$ respectively. Thus, the seller's task is to choose R to maximize their payoff function over this domain. A natural and interesting question is how the seller's optimal choice of R varies with respect to the influencer's preference weight α. The mechanism that governs the optimal choice is a combination of two effects: as α increases and the weight on the pecuniary payoff increases, the influencer becomes more responsive to monetary incentives. In contrast, as α falls, stronger monetary incentives are needed to convince the influencer to adopt a lower standard (i.e., choose a buy message for lower valuations). Some insight can also be gained by looking at the extremal values of α: if $\alpha = 0$, the influencer does not care about their monetary payoff and will act fully in the buyer's interest, while if $\alpha = 1$, the influencer only cares about their monetary payoff and will engage in the maximum amount of feasible persuasion even for an infinitesimally small reward.

As such, we would like to identify conditions under which we will have simple and intuitive comparative statics. Based on the informal discussion in the last paragraph, it is natural to posit that under fairly general conditions, the optimal commission will exhibit unimodal comparative statics in α; that is, it will first increase for sufficiently small α and decrease for sufficiently large α. We proceed by proving four lemmas invoking increasingly restrictive conditions that jointly permit a formalization of this intuition, noting that subsets of these assumptions imply some, but not all, of the overall results. To simplify the seller's payoff function, let $k = \frac{\alpha}{1-\alpha}$. Since k is a monotonically increasing function of α, we can rewrite the payoff function in terms of k and R as:

$$\pi(k, R) = (1 - F(p - kR))\,(p - R).$$

Differentiating the payoff function with respect to R, we find that the derivative is $\pi_R(k, R) = -(1 - F(p - kR)) + k(p - R)f(p - kR)$. Indeed, $\pi_R(k, 0) = F(p) - 1 + kpf(p)$, so that $\pi_R(k, 0) < 0$ for $k = 0$ and $\pi_R(k, 0) > 0$ for sufficiently large k.

Thus, we observe that when k is sufficiently large ($k > k_2$ for some k_2), $\pi_R(k, 0) > 0$. Hence, when $k > k_2$, it follows by continuity of π_R that $\pi_R(k, R) > 0$ for $R \in [0, R_1(k)]$, where $R_1(k) > 0$. However, the issue is the potential dependence of R_1 on k. Now, with an additional regularity condition of $\pi_{RR}(k, R)$ being uniformly bounded from below in k (a unidirectional analog of a

standard sufficiency condition for equicontinuity), it follows by the mean value theorem that such a uniform R_1 (agnostic of k) indeed exists. Now, let us call the seller's reward scheme *maxed out* if the seller optimally chooses $R = \frac{y}{k}$. Since $\frac{y}{k}$ is strictly decreasing in k and tends to 0 as k tends to ∞, we have the following lemma:

Lemma 1. *With the regularity condition that $\pi_{RR}(k, R)$ is uniformly bounded from below in k, the seller optimally chooses to* max out*, that is, they select $R = \frac{y}{k}$, when k is sufficiently large.*

Now, we make the standard assumption that $\pi_R(k, R)$ is strictly decreasing in R (i.e., π is strictly concave in R), which ensures a unique optimal R. Then, it follows that choosing $R = 0$ is optimal if and only if $\pi_R(k, 0) \leq 0$. Thus, we have the following lemma:

Lemma 2. *If $\pi_R(k, R)$ is strictly decreasing in R, the seller chooses $R = 0$ for $k \in [0, k_1]$, where $k_1 > 0$.*

We have now shown that the schedule of R is flat at 0 for $k \in [0, k_1]$ and decreasing when k is sufficiently large. Can we generalize these conclusions to fully characterize the comparative statics of R with respect to k? To do so, we need a further assumption, namely, that π_R is increasing in k. Indeed, with this assumption, the seller's payoff function (over the relevant domain) satisfies the increasing differences condition in Milgrom and Shannon (1994). Ignoring the domain restriction, the maximizer must be increasing in k. The issue, however, is that the right endpoint of the domain is strictly decreasing in k. We begin by proving a lemma that states that if the seller opts to max out for $k = k^*$, then the seller also chooses to max out for all $k > k^*$. The intuition is that the derivative for the higher value of k dominates the derivative for the lower value of k for the compressed domain, so the right endpoint must be the optimal value of R in the case with the higher value of k.

Lemma 3. *If π_R is strictly decreasing in R and increasing in k and the seller opts to max out for $k = k^*$, then the seller also chooses to max out for all $k > k^*$.*

Proof. If the seller opts to max out for $k = k^*$, then $\pi_R\left(\frac{y}{k}\right) \geq 0$. Hence, for $k' > k^*$, $\frac{y}{k'} < \frac{y}{k^*}$, so that $\pi_R\left(\frac{y}{k^*}\right) \geq 0$, which in turn implies that the seller also opts to max out for $k = k'$. □

Therefore, from the lemma, we know that beyond a certain cutoff value k^* of k, the optimal reward scheme has the "maxing out" property, and that the choice of R will therefore be strictly decreasing in k over this region. Combining this with analogues of the monotone comparative statics results in Milgrom and Shannon (1994) for the case of $k < k^*$, we find that the optimal value of R is increasing in k over this latter region. Thus, with these assumptions, the intuitive trade-off that we discussed earlier between greater incentivization and satiation becomes clear-cut, generating this unimodal pattern in R. The remaining results are formally stated and proven below (with the previously invoked assumptions that π_R is strictly decreasing in R and increasing in k).

Lemma 4. *If* $k^* = \inf_k \left\{R(k) = \frac{v^*}{k}\right\}$*, then the seller's optimal choice of* R *is decreasing for* $k > k^*$*.*

Proof. This follows trivially from Lemma 3, since the seller opts to *max out* for all $k > k^*$ and the right endpoint $\frac{y}{k}$ is strictly decreasing in k. □

We can then combine these lemmas to derive the following proposition, which shows that the optimal commission is indeed unimodal in α.

Proposition 2. *The seller's optimal choice of* R *is increasing for* $k \leq k^*$ *and decreasing for* $k \geq k^*$*.*

Proof. See the Appendix. □

This proposition establishes the main result pertaining to the optimal contract. The optimal choice of the purchase-contingent reward R first increases and then decreases with the influencer's relative preference for monetary incentives. We can illustrate the mechanisms underlying this result using the example of F being the uniform distribution on $[0, 1]$. In the uniform case, $\pi_R(k, R) = p - kR - 1 + k(p - R) = (k+1)p - 2kR - 1$, which is strictly decreasing in R. Moreover, $\pi_{Rk}(k, R) = (p - 2R)$. The solution for $\pi_R(k, R) = 0$ is $R = \frac{(k+1)p-1}{2k} < \frac{p}{2}$, so the optimal R is always less than $\frac{p}{2}$, implying that π_R is increasing in k. We demonstrate the content of the results for the uniform distribution with $p = \frac{1}{2}$. Figure 1 indeed confirms the non-monotonic (unimodal) relationship between R and α. Initially, when α is sufficiently small the influencer is not sensitive to monetary incentives, and in response the firm finds it optimal to offer no incentives (i.e., $R = 0$). As the influencer becomes increasingly more responsive to monetary incentives, the variable incentive with α. This is driven by an "incentive effect", which induces the influencer to commit

to an information transmission scheme that progressively increases the amount of persuasion (or equivalently, the probability with which the buyer is rationally induced to purchase as determined by Bayes plausibility). The variable incentive R continues to increase until a threshold value of α, at which point the influencer generates the maximum feasible amount of persuasion. Interestingly, beyond this threshold, the seller actually decreases the variable incentive even as α increases. This is a "satiation effect" or, in other words, a "maxing out" effect – because the influencer's persuasion scheme already generates the maximum amount of persuasion, any further increase in their relative preference for monetary incentives permits the seller to offer progressively weaker incentives and yet maximize the probability of buyer purchase.

Revisiting some of the interpretations for the parameter α that we discussed earlier, these theoretical results provide some potentially interesting empirical corollaries. First, influencers with high reputational concerns (for example, those at the early stages of their lifecycle) may not take any commissions. Second, the finding from Fainmesser and Galeotti (2021) that micro-influencers typically earn higher per-post commissions relative to macro-influencers can be explained within our model by the former having an intermediate value of α and the latter exhibiting a sufficiently large value of α to place them in the satiation region. Naturally, this intuition is not fully precise, since the relationship is unimodal. However, our theoretical results can be seen as consistent with the anecdotal observation that nano-influencers (that is, influencers with a very small follower base) are often compensated through free products or flat per-post fees, rather than through a commission model.

Given Proposition 2, another implication is that the consumer surplus is decreasing in α, since the cutoff state is monotonically decreasing in α. Moreover, the seller's profit is increasing in α (for this, we do not even require the restrictions for the proposition, since for any given R the probability of purchase increases with α).

Figure 2 shows that the probability of purchase (sales quantity) is monotonically increasing in α and eventually reaches the maximum feasible level that satisfies Bayes plausibility in the "maxing out" region. The seller's profit is also monotonically increasing in α, which is intuitive as the seller benefits from the influencer's greater responsiveness to monetary incentives. In contrast, the consumer surplus decreases monotonically in α as the persuasion cutoff moves to the left. Echoing

the unimodality of the optimal R, the influencer's monetary payoff increases until the satiation threshold is reached and decreases thereafter.

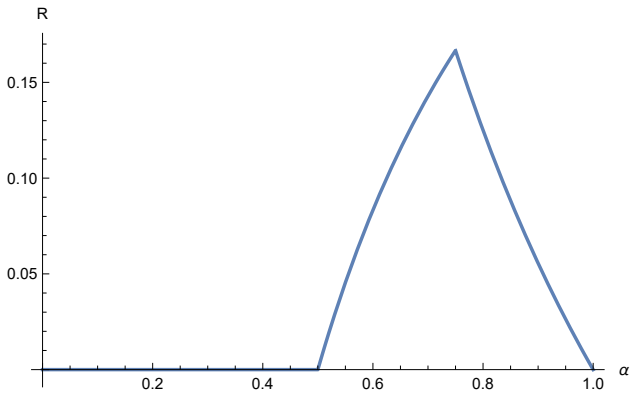

Figure 1: Illustration of unimodality of optimal R versus α

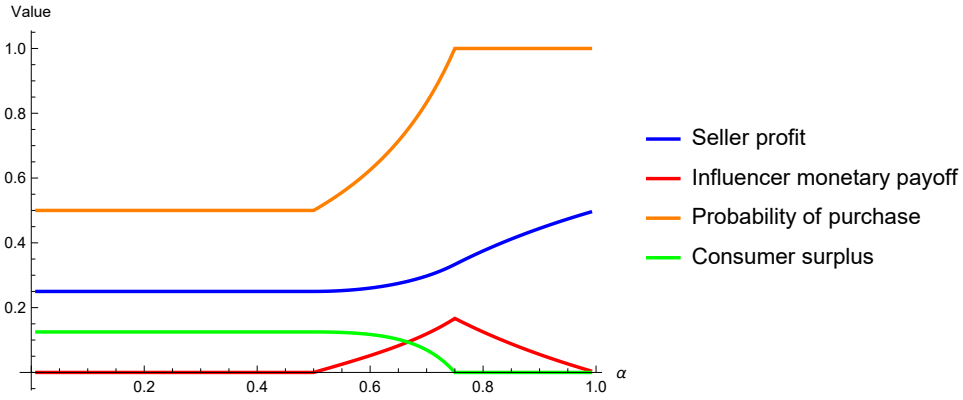

Figure 2: Comparative statics of seller profit, influencer monetary payoff, probability of purchase and consumer surplus with respect to α

3 Endogenous Pricing

As discussed in Kamenica and Gentzkow (2011), there are two general ways in which a principal/sender can motivate an agent/receiver to take a desired action – incentives, which alter the receiver's marginal payoffs, and persuasion, which alters the receiver's beliefs. We now allow the seller to endogenously choose a retail price p prior to engaging in contract negotiations with the influencer. This enables us to examine the potential interplay between incentives and persuasion.

We can ask whether these two mechanisms are substitutes or complements and examine their effects on seller profits and consumer welfare.

A useful simplification is to think of the seller as choosing their desired probability of purchase (quantity) first and then selecting the combination of price and persuasion appropriately to feasibly enforce this quantity. For a chosen marginal state cutoff of q, if the influencer were maximally persuasive, the effective price that could be charged would be $p = \mathbb{E}(v|v \geq q)$, which we call the *grossed-up price*. We assume that the seller's outside option is to sell at price p and only realize sales in the states of the world above p.

Suppose that there is no limited liability constraint. Then, as in our analysis of the baseline setup in the previous section, the seller can set R and T to compensate the influencer for the (scaled) lost utility from reduced consumer surplus induced by persuasion. Analogous to our discussion in Section 2.2, it follows that $\alpha \geq \frac{1}{2}$ is a sufficient but not necessary condition for the seller to induce maximal persuasion from the influencer. If $\alpha \geq \frac{1}{2}$, the seller nets the profit from the effective price $p = \mathbb{E}(v|v \geq q)$ with quantity $1 - F(q)$, but must pay a "silence bribe" of the foregone consumer surplus associated with this price (scaled by the factor $k \equiv \frac{1-\alpha}{\alpha}$). When the seller chooses price p and induces full persuasion, or, equivalently, sells quantity q at the retail price p, the lost consumer surplus is (by symmetry) equal to $\int_p^1 (v - p)\, dF(v)$. Thus, if we make q the variable of choice, our objective function is

$$U = \int_q^1 v\, dF(v) - k \int_{p(q)}^1 (v - p(q))\, dF(v)$$

Moreover, $p(q)$ is governed by the implicit function $\int_q^1 (v - p(q))\, dF(v) = 0$. Now, taking the first-order condition with respect to q, we find that $\frac{dU}{dq} = -qf(q) - k \int_p^1 (-dF(v)) \frac{dp}{dq} = 0$. This is non-trivial to solve in closed form, since $\frac{dp}{dq}$ does not have an obvious closed-form representation, but the key insight is that $q = 0$ is no longer optimal, as the first term in the derivative is zero, while the second term is strictly positive, which is different to what we would have in the absence of the influencer.

Let us now focus on the setting with limited liability. Suppose that the seller chooses the marginal state cutoff q and some retail price p, which is less than or equal to $\mathbb{E}(v|v \geq q)$. The question is what reward R is needed to support price p. The reward R solves $\alpha R + (1-\alpha)(q-p) = 0$,

so that $R = \frac{1-\alpha}{\alpha}(p-q) = k(p-q)$, where k is defined here as $\frac{1-\alpha}{\alpha}$. Hence, we find that the seller's effective price is $p - R = p - k(p-q) = (1-k)p + kq$. As such, the effective price earned by the seller is a convex combination of the state cutoff and retail price, and we see that if $\alpha < \frac{1}{2}$ $(k > 1)$, the seller will eschew persuasion, setting $R = 0$ and $p = q$, and if $\alpha \geq \frac{1}{2}$ $(k \leq 1)$, the seller will use the maximum amount of feasible persuasion (that is, they will charge the grossed-up price to the customer). The intuition is straightforward: if the influencer has a higher relative weight on their pecuniary payoff, then persuasion is cheaper than pricing, while the converse holds if they have a higher relative weight on consumer surplus. It is in this sense that we can think of price reductions and persuasion as instrumental substitutes; to attain a given q, the seller can cut the price and reduce their reliance on persuasion, or they can offer a higher price in conjunction with a greater reward to induce more persuasion.[7]

Thus, if $k > 1$, the seller's problem reduces to the standard monopoly pricing problem, which we term the *persuasion-free baseline*, and they will earn the optimal fixed price monopoly profit. Now, note that the revenue with the grossed-up price (excluding the payment to the influencer) is equal to the monopoly profit with perfect price discrimination (full surplus extraction) for values above q, so that for $k \leq 1$, the seller's profit is $kq(1-F(q)) + (1-k)\int_q^1 v\, dF(v)$. Taking the derivative with respect to q and using Leibniz's rule, we find that it is $k(1-F(q)) - kqf(q) - (1-k)qf(q) = k(1-F(q)-qf(q)) - (1-k)qf(q)$. The key insight is that the derivative is strictly lower than the monopoly profit derivative, and we can write the condition for the derivative to be zero as $(1-k)q + k\left(q - \frac{1-F(q)}{f(q)}\right) = 0$, which is a k-weighted virtual value condition using the language of Myerson (1981). Assuming that the virtual value function is increasing, that is, regularity in the sense of Myerson (1981) is satisfied, the k-weighted virtual value function is also increasing, implying that the optimal q makes this weighted virtual value function equal to zero. Therefore, the seller's optimal q is increasing in k (decreasing in α) and is no larger than that chosen in the persuasion-free baseline. Consumer surplus is reduced to zero when $k \leq 1$ $(\alpha \geq \frac{1}{2})$, with the residual surplus

[7]There is some subtlety in how we interpret substitutability or complementarity between pricing and persuasion. What we are saying here is that the seller can choose between lower prices and higher rewards to attain a given probability of purchase or state cutoff q. It is not true, however, that the seller's optimal R is increasing in p when p is exogenously given. To see this, observe that there is no scope for persuasion either when $p = 0$ (the buyer will always purchase) or $p = 1$ (due to Bayes plausibility), while persuasion is evidently feasible for interior values of p. Hence, there will generally be a non-monotonic relationship between p and R. Nevertheless, fixing q elucidates the *functional* substitutability of pricing and persuasion.

shared between the seller and the influencer. Thus, in the case with flexible prices, the presence of influencers whose utility functions are more strongly weighted toward consumer surplus is beneficial to consumer surplus, since for $\alpha \geq \frac{1}{2}$, consumer surplus falls to zero, while otherwise, consumer surplus is positive. It is worthwhile to contrast this with the dynamics of consumer surplus in the fixed price environment shown in Figure 2. There, consumer surplus remains positive even for an interval of $\alpha \geq \frac{1}{2}$, while with a flexible price, consumer surplus drops to zero for all $\alpha \geq \frac{1}{2}$. It is unsurprising that the consumer is worse off when the seller has an additional instrument for surplus extraction. We can summarize these results in the following proposition:

Proposition 3. *Assuming that contracts must respect limited liability, if $\alpha \geq \frac{1}{2}$ (or equivalently, $k \leq 1$), the seller uses persuasion; otherwise, the seller abstains from persuasion and only uses pricing, charging the optimal static monopoly price. Moreover, the seller never opts for partial persuasion, and whenever they use persuasion, consumer surplus falls to zero and, if the distribution F satisfies Myersonian regularity, the sales quantity (the probability of purchase) is increasing in α and is at least as large as that in the persuasion-free baseline.*

4 Extensions

4.1 Data brokerage

We now compare the above analysis of pricing and persuasion with what appears to be an ostensibly different problem, namely, that of a *data brokerage*. Specifically, the influencer might be in a position to act as a data provider for the firms. The existing literature typically considers a data intermediary who can sell information that enables firms to better target their consumers (Bergemann and Bonatti, 2015). However, our data brokerage setup is distinct in that the influencer plays the role not only of a data provider, but also of a persuasion intermediary for consumers. To capture this data brokerage problem, assume that the buyer is privately informed of their value v, and in addition, that the influencer also observes this value.

Consider first the case where the buyer's ultimate purchase decision is contractible. In this case, the seller will structure the base transfer T and the purchase-contingent commission R to the

influencer, as well as the downstream price p to the buyer, in order to incentivize the influencer to truthfully report the buyer's value. The buyer purchases with probability one if $v \geq p$ but does not purchase otherwise. Now, the seller could commit to setting a price precisely equal to the influencer's report. The problem that remains is to prevent the influencer from undertaking "downside deviations," that is, reporting a lower value of the state to raise consumer surplus. Suppose that the true value is v and the influencer considers a deviation to report $v' < v$, assuming that the seller sets the price equal to the influencer's report. Then, we require $\alpha R(v) \geq \alpha R(v') + (1-\alpha)(v - v')$ for incentive compatibility to hold. Therefore, $R(v) - R(v') \geq \frac{1-\alpha}{\alpha}(v - v')$. Now, we will proceed to show that if $\alpha \geq \frac{1}{2}$, the seller does not sell to any type below a cutoff q and optimally sells to the remaining types $v \geq q$ at price v. Hence, R is optimally set so that $R(v) = \frac{1-\alpha}{\alpha}(v - q)$ for $v \geq q$ and $R(v) = 0$ otherwise, giving the seller payoff

$$U = \int_q^1 v \, dF(v) - k \int_q^1 (v - q) \, dF(v)$$

Let us again assume that F satisfies Myersonian regularity. Now, the derivative $\frac{dU}{dq}$ is $-qf(q) + k(1 - F(q))$, which we set to zero for the first-order condition. Hence, $q - k\frac{1-F(q)}{f(q)} = kq + (1-k)q - k\frac{1-F(q)}{f(q)} = k\left(q - \frac{1-F(q)}{f(q)}\right) + (1-k)q = 0$. But this is precisely the same weighted virtual value condition that we derived in the pricing/persuasion problem with limited liability. We summarize and prove these results in the following proposition:

Proposition 4. *Consider the data brokerage environment, where the buyer's ultimate purchase decision is contractible, $\alpha \geq \frac{1}{2}$, and F satisfies regularity. In equilibrium, the seller induces a cutoff state q such that for all $v \geq q$ they choose $R(v) = \frac{1-\alpha}{\alpha}(v - q)$ if a sale is realized and a price equal to the state, whereas for $v < q$, the seller chooses $R(v) = 0$ and does not make a sale. The optimal base transfer is $T = 0$. The influencer recommends purchase when $v \geq q$ and no purchase otherwise. Finally, the equilibrium q solves the zero weighted virtual type condition $k\left(q - \frac{1-F(q)}{f(q)}\right) + (1-k)q = 0$.*

Proof. See the Appendix. □

It is interesting that the first-order conditions for the data brokerage problem and the pricing/persuasion contracting problem with the limited liability constraint are exactly the same. In-

deed, these settings are ostensibly quite different, but there are clearly very close links between the two problems. Thus, the analysis uncovers an underlying connection between two economically important phenomena in contemporary influencer and consumer information markets. Intuitively, the limited liability condition in the Bayesian persuasion problem with endogenous pricing is isomorphic to the individual rationality constraint in the data brokerage problem.

Indeed, it is also instructive to compare this to another scenario where the buyer's ultimate purchase decision is not contractible. This is closer in spirit to the setup in Krishna and Morgan (2008). In this environment, the seller can set transfers based on the influencer's report, but in essence cannot block "upside deviations". Can the seller incentivize truthful reporting of the state without committing to a fixed price? Suppose we have states v_1 and v_2, transfers $t(v_1)$ and $t(v_2)$ and prices $p(v_1)$ and $p(v_2)$. If the buyer is willing to buy in state v_2 at $p(v_2)$, then the incentive compatibility condition becomes $\alpha t(v_2) + (1-\alpha)(v_2 - p(v_2)) \geq \alpha t(v_1) + (1-\alpha)(v_2 - p(v_1))$, so that $t(v_2) - t(v_1) \geq \frac{1-\alpha}{\alpha}(p(v_2) - p(v_1))$. If the buyer is willing to buy in state v_1 at $p(v_1)$ and at $p(v_2)$, the resulting incentive compatibility condition is $\alpha t(v_1) + (1-\alpha)(v_1 - p(v_1)) \geq \alpha t(v_2) + (1-\alpha)(v_1 - p(v_2))$, so that $t(v_2) - t(v_1) \leq \frac{(1-\alpha)}{\alpha}(p(v_2) - p(v_1))$. The problem for the influencer is that if the state is 0, then there would always be zero consumer surplus, but if any other state gave a positive transfer, then the influencer would deviate to reporting that state. Hence, all states must give the same transfer and have the same associated price, and the seller is forced to simply set the monopoly price and earn the static monopoly profit. In fact, this is exactly what emerges in the Bayesian persuasion setting when neither the signal scheme nor the buyer's ultimate purchase decision are contractible, implying that the seller eschews any contract with the influencer and simply sets the monopoly price.

4.2 Imperfectly informed influencer

Thus far, we have made the assumption that the influencer is fully informed about the state of the world. We now briefly explore the implications of noise in the influencer's information. Rather than pursuing the analysis in full depth, we highlight some of the key economic intuitions, which also illustrate the dual role of the influencer as a facilitator of sales and a protector of consumer welfare in our model.

Suppose that with probability ρ, the influencer is uninformed of the state and receives a signal that is randomly drawn from the prior distribution F, and that with complementary probability $1-\rho$, they are perfectly informed of the state. A nice corollary of this setup is that the distribution over the signals received by the influencer is precisely the same as the prior distribution F. Moreover, it is also straightforward to see that the influencer's optimal experiment again involves conflating a set of states to the right of a cutoff x. However, the key difference is that the pointwise contribution of a state to the consumer surplus in the event that the buyer purchases is now given by $\rho\mu+(1-\rho)v-p$, where $\mu = \mathbb{E}(v)$, instead of $v - p$.

The impact of imperfect information on the role of the influencer is most fruitfully analyzed in the environment with endogenous prices. Here, we apply the method used in Section 3, where we set a marginal state cutoff q and examine what the feasible values are for the retail price p. Using the conditions for zero pointwise consumer surplus and Bayes plausibility, we find that for a given q, we must have $p \in [\rho\mu+(1-\rho)q, \rho\mu+(1-\rho)\mathbb{E}(v|v \geq q)]$. Moreover, the value of R that induces a given p is $R = k(p - \rho\mu - (1-\rho)q)$, where $k = \frac{\alpha}{1-\alpha}$. Therefore, adapting the arguments of Section 3, we we can show that the seller employs the maximum quantity of feasible persuasion if $\alpha \geq \frac{1}{2}$ ($k \leq 1$) and does not use persuasion otherwise.

Thus, without persuasion, the seller's profit is given by $[\rho\mu + (1-\rho)q)(1 - F(q)]$, while with full persuasion, the seller's profit is $k(\rho\mu+(1-\rho)q)(1-F(q))+(1-k)\left[\rho\mu + (1-\rho)\mathbb{E}(v|v \geq q)\right]$. The latter expression simplifies to $\rho\mu(1-F(q))+(1-\rho)\left[kq(1 - F(q) + (1-k)\int_q^1 dF(v)\right]$. Maintaining the assumption of F satisfying Myersonian regularity, the first-order condition for the optimal q becomes $q - \frac{1-F(q)}{f(q)} = -\frac{\rho\mu}{1-\rho}$ without persuasion and $(1-k)q + k\left(q - \frac{1-F(q)}{f(q)}\right) = -\frac{\rho\mu}{1-\rho}$ with full persuasion. Note that if these first-order conditions do not bind with equality for any q, the optimal choice of q will be $q = 0$. In short, we find that the seller's optimal choice of q is decreasing in ρ; that is, the probability of purchase is increasing in ρ. Indeed, this is an intuitive finding, in that when $\rho = 1$, the influencer is entirely uninformed, and the seller can simply choose $p = \mu$ and extract the entire surplus. This result indicates the dual role of the influencer in our environment: for a given price, their presence may permit a greater sales volume and thereby benefit the seller, but more generally, their alignment with consumer welfare in the absence of pecuniary incentives impedes rent extraction by the seller.

5 Summary and Conclusion

A significant amount of persuasion of consumers, especially in the modern digital economy, occurs through intermediaries such as social media and influencers, analysts, and online recommendation agents. In this paper, we study the design of contracts by a seller to optimally govern the persuasion policies of an intermediary who influences the purchase decision of a buyer. The influencer's preferences are a combination of their utility for their pecuniary payoffs and consumer surplus, which may be seen as representing a basic trade-off faced by influencers in a broad range of settings between immediate monetary payoffs from persuading users and the value of maintaining future reputational payoffs.

The optimal contract responds to the influencer's taste for pecuniary incentives versus consumer welfare or reputational concerns. It consists of a variable reward R that varies non-monotonically with the influencer's relative preference for monetary incentives versus consumer welfare, and we identify simple conditions under which R exhibits unimodal comparative statics. This points to an underlying trade-off in the contract design between an incentive effect and a satiation effect. As the relative preference for monetary incentives increases, the seller increases the variable reward, inducing the influencer to increase the amount of persuasion and thereby the probability of purchase. The reward R continues to increase until a threshold at which the influencer generates the maximum possible amount of persuasion. Beyond this threshold, the seller actually finds it optimal to decrease the variable incentive even as the influencer's relative preference for monetary incentives increases. Because the influencer already generates the maximum amount of persuasion, the seller can afford to offer progressively weaker incentives and yet induce the persuasion scheme that maximizes the probability of buyer purchase.

We analyze an important aspect of the problem by considering the scenario where the seller endogenously chooses and commits to a retail price prior to contracting with the influencer. This allows us to examine the basic question of how incentives and persuasion interact in this setup to influence the consumer to buy the product. The important insight is that the equilibrium contract involves the seller either opting for full persuasion or no persuasion at all. Partial persuasion is never used, and in this sense, persuasion and endogenous pricing can be seen as instrumental

substitutes for the seller in attaining a given probability of the ultimate purchase. In other words, the seller may rely more on price reductions while reducing the incentives for persuasion, or they may increase the influencer's commission to induce more persuasion but charge higher prices to the buyer. The former is optimal when the influencer's relative preference for monetary incentives is small, whereas the latter is optimal when the influencer has relatively higher reputational concerns compared to monetary payoffs.

Some natural extensions to the intermediated persuasion problem presented in this paper could be explored in future work. One could consider the possibility of multiple "types" of influencers, characterized by their value of α. If the influencer types were uncertain to consumers, then in dynamic settings with reputational concerns, this could lead to interesting trade-offs between preserving reputation and accepting short-run monetary rewards. Another interesting aspect that could be explored is the effect of relaxing the various commitment assumptions. Indeed, there are two sources of commitment: the first is on the part of the seller (principal) in terms of their contract with the influencer and their pricing decision. The second is in terms of the influencer's commitment to their chosen signal structure/experiment. Finally, it might be worthwhile to explore the effects of competition between sellers for a common persuasion intermediary's recommendation, especially in relation to how their optimal contracts would be structured in a competitive environment and the concomitant effects on firm profits and consumer welfare.

References

Ambrus, A., E. M. Azevedo, and Y. Kamada (2013): "Hierarchical cheap talk," *Theoretical Economics*, 8, 233–261.

Azarmsa, E. and L. W. Cong (2020): "Persuasion in relationship finance," *Journal of Financial Economics*, 138, 818–837.

Bergemann, D. and A. Bonatti (2015): "Selling Cookies," *American Economic Journal: Microeconomics*, 7, 259–294.

Berliant, M. (1984): "A characterization of the demand for land," *Journal of Economic Theory*, 33, 289–300.

Corrao, R. (2023): "Mediation Markets: The Case of Soft Information," Working Paper.

Curello, G. and L. Sinander (2022): "The comparative statics of persuasion," *arXiv preprint arXiv:2204.07474*.

Durbin, E. and G. Iyer (2009): "Corruptible advice," *American Economic Journal: Microeconomics*, 1, 220–242.

Dworczak, P. and G. Martini (2019): "The simple economics of optimal persuasion," *Journal of Political Economy*, 127, 1993–2048.

Fainmesser, I. P. and A. Galeotti (2021): "The market for online influence," *American Economic Journal: Microeconomics*, 13, 332–372.

Inderst, R. and M. Ottaviani (2009): "Misselling through agents," *American Economic Review*, 99, 883–908.

Ivanov, M. (2021): "Optimal monotone signals in Bayesian persuasion mechanisms," *Economic Theory*, 72, 955–1000.

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

Kolotilin, A. (2015): "Experimental design to persuade," *Games and Economic Behavior*, 90, 215–226.

——— (2018): "Optimal information disclosure: A linear programming approach," *Theoretical Economics*, 13, 607–635.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a privately informed receiver," *Econometrica*, 85, 1949–1964.

Kotler, P., K. L. Keller, S. H. Ang, C. T. Tan, and S. M. Leong (2018): *Marketing management: an Asian perspective*, Pearson London.

Krishna, V. and J. Morgan (2008): "Contracting for information under imperfect commitment," *The RAND Journal of Economics*, 39, 905–925.

Lipnowski, E., D. Ravid, and D. Shishkin (2024): "Perfect Bayesian Persuasion," *arXiv preprint arXiv:2402.06765*.

Mensch, J. (2021): "Monotone persuasion," *Games and Economic Behavior*, 130, 521–542.

Milgrom, P. and C. Shannon (1994): "Monotone comparative statics," *Econometrica: Journal of the Econometric Society*, 157–180.

Mullainathan, S. and A. Shleifer (2005): "Persuasion in Finance," NBER Working Paper 11838, National Bureau of Economic Research.

Myerson, R. B. (1981): "Optimal auction design," *Mathematics of Operations Research*, 6, 58–73.

Neyman, J. and E. S. Pearson (1933): "IX. On the problem of the most efficient tests of statistical hypotheses," *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231, 289–337.

Pei, A. and D. Mayzlin (2022): "Influencing social media influencers through affiliation," *Marketing Science*, 41, 593–615.

6 Appendix

Proof of Proposition 1. The primal problem can be written as:

$$\max_{\beta(.)} \min_{\lambda \geq 0} \int_0^1 (u(v) + \lambda(v-p))\beta(v)\, dF(v)$$

The associated Lagrangian dual problem is:

$$\min_{\lambda \geq 0} \max_{\beta(.)} \int_0^1 (u(v) + \lambda(v-p))\beta(v)\, dF(v)$$

We will *a priori* assume that the solutions to the primal and the dual problems exist, and then show as part of the proof that existence is indeed guaranteed. We know that the Lagrangian dual function $\max_{\beta(.)} \int_0^1 (u(v) + \lambda(v-p))\beta(v)\, dF(v)$ is an upper bound for the optimal value of the objective function in the primal problem for any $\lambda \geq 0$ by weak duality. To show this, observe that the maximizer $\beta(.)^*$ in the primal problem by definition satisfies the constraint in the original constrained optimization problem, which implies that $\max_{\beta(.)} \int_0^1 (u(v) + \lambda(v-p))\beta(v)\, dF(v) \geq \int_0^1 (u(v) + \lambda(v-p))\beta^*(v)\, dF(v) \geq \int_0^1 u(v)\beta^*(v)\, dF(v)$, as required.

It is straightforward to see that for a given value of λ, the optimizer of the dual problem follows the hypothesized "zero-one" structure, where $\beta(v) = 1$ for $v \in S(\lambda)$ and $\beta(v) = 0$ otherwise. Now, consider how the optimal β in the dual problem varies with λ. Indeed, we have four cases for different values of v, in terms of whether these values of v are included in S:

- $u(v) \geq 0$ and $v - p \geq 0$. Such values of v are always assigned $\beta(v) = 1$, since they have a negative "shadow price" and a positive reward.
- $u(v) \geq 0$ and $v - p < 0$. Such values of v have a positive reward and a positive shadow price, so are initially assigned $\beta(v) = 1$ when λ is small (that is, up to a cutoff value of λ), but are subsequently assigned $\beta(v) = 0$.
- $u(v) < 0$ and $v - p < 0$. Such values of v are always assigned $\beta(v) = 0$, since they have a positive shadow price and a negative reward.
- $u(v) < 0$ and $v - p \geq 0$. Such values of v have a negative reward and a negative shadow price, so are only useful to slacken the constraint. Thus, they are assigned $\beta(v) = 0$ up to cutoff value of λ, and are assigned $\beta(v) = 1$ for λ larger than this cutoff.

Hence, the value of the primal objective function for the Lagrange dual optimizer falls monotonically as λ increases and that the constraint value $\int_{S(\lambda)} (v-p)\, dF(v)$ increases (favorably). Moreover, by boundedness of u, the constraint value function must be non-negative for sufficiently large λ, implying that $\lambda^* = \inf_{\lambda \geq 0} \int_{S(\lambda)} (v-p)\, dF(v) \geq 0$ is well defined. Now, if the Bayes plausibility constraint was not violated at $\lambda = 0$, we would immediately be able to choose the associated S as our optimal set in the primal problem (that is, set $\beta(v) = 1$

for $v \in S$ and $\beta(v) = 0$ otherwise) and would be done. If, instead, the Bayes plausibility constraint was violated at $\lambda = 0$, we can proceed to partition the boundary states $B^* = \{ v : u(v)+\lambda^*(v-p) = 0 \}$ into $B^+ = B \cap \{ v : v-p \geq 0 \}$ and $B^- = B \cap \{ v : v-p < 0 \}$.

Now, it is easy to see that $0 \geq \lim_{\lambda \uparrow \lambda^*} \int_{S(\lambda)} (v-p)\, dF(v) = \int_{S(\lambda^*)} (v-p)\, dF(v) + \int_{B^-} (v-p)\, dF(v)$ and $0 \leq \lim_{\lambda \downarrow \lambda^*} \int_{S(\lambda)} (v-p)\, dF(v) = \int_{S(\lambda^*)} (v-p)\, dF(v) + \int_{B^+} (v-p)\, dF(v)$ (these limits exist from the monotonicity and boundedness of the constraint value function). If $\int_{S(\lambda^*)} (v-p)\, dF(v) + \int_{B^-} (v-p)\, dF(v) = \int_{B^+} (v-p)\, dF(v)$, we can directly choose $\beta(v) = 1$ for all $v \in B^*$ to obtain a dual optimizer for which the constraint exactly binds. Otherwise, we may select $\beta(v) = 1$ for all $v \in B^+$ and $\beta(v) = -\frac{\int_{S(\lambda)} (v-p)\, dF(v) = \int_{S(\lambda^*)} (v-p)\, dF(v) + \int_{B^+} (v-p)\, dF(v)}{\int_{S(\lambda^*)} (v-p)\, dF(v) + \int_{B^-} (v-p)\, dF(v)}$ for all $v \in B^-$, again yielding a dual optimizer at which the constraint exactly binds.

But we know that the Lagrange dual function is an upper bound for the objective function and that it is exactly equal to the objective for this value of λ. Thus, in all cases, the optimal value of λ, λ^*, is given by $\lambda^* = \inf_{\lambda \geq 0} \int_{S(\lambda)} (v-p)\, dF(v) \geq 0$. Moreover, the sender's optimal signal coincides with this solution to the dual problem for the given value of λ^*; that is, we know that $\beta(v) = 1$ for $v \in S(\lambda^*)$ and $\beta(v) = 0$ for all v such that $u(v) + \lambda^*(v-p) < 0$., establishing the validity of the proposition.

□

Proof of Proposition 2. First, we prove the result with strict inequality signs. Consider $k_1 < k_2 < k^*$. Let $R_2 = R(k_2)$, $R_1 = R(k_1)$ and observe that R_2 is feasible when $k = k_1$. Suppose for the sake of contradiction that $R_2 < R_1$ (so that R_1 is an interior point in the domain and $\pi_R(R_1, k_1) = 0$). Then $\pi_R(R_2, k_2) > \pi_R(R_1, k_2) > \pi_R(R_1, k_1) = 0$. But then for $k = k_2$, π is strictly increasing at $R = R_2$, contradicting the optimality of this choice and proving the result (since the second part has already been shown). Moreover, by Berge's theorem, since the constraints and objective function are continuous, we have that $R(k)$ is a continuous function, implying that the strict inequality signs can be changed to weak inequality signs and completing the proof.

□

Proof of Proposition 4. Let the lowest type to which the seller wishes to sell be q.[8] We know from the earlier discussion that for any type $q' > q$ to which the seller plans to sell, we need $R(q') - R(q) \geq \frac{1-\alpha}{\alpha}(p(q') - p(q)) = k(p(q') - p(q))$. Hence, we have that $p(q') - R(q') \leq p(q') - k(p(q') - p(q)) - R(q)$. Therefore, when $\alpha \geq \frac{1}{2}$ (that is, $k \leq 1$), $p(q') - R(q')$ is maximized when $p(q')$ is maximal, i.e. when $p(q') = q'$.

Thus, the seller cannot do better than to set $R(q) = 0$ and $T = 0$ (to make the individual rationality constraint bind), as well as to assign $p(q) = q$ and $R(q') = k(q' - q)$ for any $q' \geq q$, since this attains the upper bound of profit for fixed q. Moreover, it is straightforward to show that this mechanism is globally incentive compatible and individually rational. Thus, the seller's payoff is identical to what we derived in the pricing/persuasion problem with limited liability, namely, $U = \int_q^1 v\, dF(v) - k \int_q^1 (v - q)\, dF(v)$, and the associated first-order condition is the given weighted virtual type condition, given that F satisfies regularity. □

[8] If this does not exist, we can still show that $p(q') = q'$ for every type q' to which the seller plans to sell. Now, letting $q = \inf q'$, we can show that $R(q) \to 0$ as $q' \to q$. Hence, we can freely alter the mechanism to make the seller also sell to type q where $R(q) = 0$ and $p(q) = q$ without affecting any incentive constraints.