Information design through scarcity and social learning*

Alexei Parakhonyak† and Nick Vikander‡

February 2019

Abstract

We show that a firm may benefit from strategically creating scarcity for its product, in order to trigger herding behavior from consumers in situations where such behavior is otherwise unlikely. We consider a setting with social learning, where consumers observe sales from previous cohorts and update beliefs about product quality before making their purchase. Imposing a capacity constraint directly limits sales but also makes information coarser for consumers, who react favorably to a sell-out because they only infer that demand must exceed capacity. Neither large cohorts nor unbounded private signals guarantee efficient learning, because the firm acts strategically to influence the consumers' learning environment. Our results show that capacity constraints can serve as a practical tool for information design: if private signals are not too precise and capacity can be changed over time, then in large markets the firm's optimal choice of capacity delivers the same expected sales as the Bayesian persuasion solution.

JEL classifications: D21, D82, D83

Keywords: social learning, information design, capacity, Bayesian persuasion

---

*A previous version of this paper was circulated until the title 'Inducing Herding with Capacity Constraints'. The paper has benefited from discussions with Nemanja Antic, James Best, Itai Arieli, Vincent Crawford, Nicolas Drouhin, Andrea Galeotti, Ben Golub, Paul Heidhues, Stephan Lauermann, Ignacio Monzón, Steven Morris, David Myatt, Marek Pycia, Gosia Poplawska, Daniel Quigley, Larry Samuelson, Peter Norman Sørensen, Ran Spiegler, Bruno Strulovici, Balazs Szentes, Curtis Taylor, Marta Troya-Martinez, and participants of the Oxford Economic Theory Workshop, Game Theory Society World Congress (Maastricht 2016), EARIE (Lisbon, 2016), Transatlantic Theory Workshop 2016, Lancaster Game Theory Workshop 2016, International Industrial Organization Conference (Boston, 2017), the Royal Economic Society Annual Conference 2018, the XXXIII Jornadas de economia industrial 2018, the Midwest Economic Theory Conference 2018, as well as audiences in Cardiff, Gothenburg, and Copenhagen.

†University of Oxford, Department of Economics and Lincoln College. E-mail: aleksei.parakhonyak@economics.ox.ac.uk

‡University of Copenhagen, Department of Economics. E-mail: nick.vikander@econ.ku.dk

1 Introduction

This paper shows that firms may want to strategically create product scarcity, to influence consumer learning in such a way that can effectively persuade consumers to buy. As is standard in the literature on social learning (see foundational papers by Banerjee (1992) and Bikhchandani et al. (1992)), each consumer receives a noisy private signal about product quality, but also infers some information from observing earlier sales. To illustrate the mechanism at work, consider a cohort of ten consumers who visit a firm, where each buys if and only if her signal was good. A consumer who then arrives and observes initial sales of three may refuse to buy, even if her own signal was good, because she infers that only three out of the ten others had a good signal. But if the firm, without knowing product quality, had initially limited capacity to three sales per period, then the consumer would observe a sell-out, and only infer that there were *at least* three good signals. That may well convince the consumer to buy even if her own signal was bad, and trigger a positive purchase cascade.

Now suppose that in our example, one out of ten consumers per period is perfectly informed about quality, and that the firm is unconstrained. If quality is low, but say seven consumers in period 1 receive good signals, then everyone will buy in period 2 except for the informed consumer. The informed consumer's choice not to buy will perfectly reveal low quality and lead to zero sales in later periods. However, this same choice would not reveal any information if the firm had restricted capacity, because the informed consumer would be effectively pooled with those who were rationed.[1]

Thus, neither large cohort size (as in the 'guinea pigs' considered by Sgroi (2002)) nor unbounded signals (as in Smith and Sørensen (2000)) guarantees efficient social learning. The reason is that here, the consumers' learning environment is endogenous, and the firm can manipulate this environment by restricting capacity. The result is that consumers may fail to learn, in particular in situations where learning would hurt the firm.

The attractiveness of restricting capacity as a tool to manipulate learning can be understood

[1] In the spirit of Smith and Sørensen (2000), we use the word 'cascade' to refer to a situation where all consumers with boundedly informative signals take the same action regardless of their private information.

more broadly in terms of information design. A firm's choice of capacity affects the structure of consumers' public information, analogous to the way a sender determines the structure of a receiver's private signal in models of Bayesian persuasion. We show that in large markets, a firm's optimal capacity choice can deliver exactly the Bayesian persuasion outcome, and therefore result in optimal information provision. Bayesian persuasion mechanisms generally raise two practical concerns: how the signal space is determined and how a sender can commit to a particular signal structure. Neither concern is an issue in our setting, as both the signal space and commitment power follow naturally from the firm's choice of capacity.

In this sense, our paper brings out a close connection between Bayesian persuasion and social learning despite apparent differences between the two approaches. In the former, a sender directly influence agent learning by designing statistical experiments. In the latter, agents learn by observing each others' actions which may reveal their private information. And yet, in practice, a seller's simple choice to restrict capacity can end up persuading consumers, by marshalling the forces of social learning: consumers learn just as they would have, had the sender designed an optimal experiment.

Our focus on sell-outs and scarcity fits in with evidence that demand for some products seems to persistently far outstrip supply and that suggests a plausible cause is seller strategic behavior. Shortages for 'Beanie Babies' were accompanied by extremely high demand from 1996 to 1999. Debo et al. (2012) argue that scarcity for these toys was largely induced by the firm itself. Chicago magazine writes that scarcity was a marketing strategy chosen by Ty Warner, the toys' creator:

> The toys easily supplanted other fads such as Ninja Turtles and Cabbage Patch dolls partly because of Warner's strategy of deliberate scarcity. He rolled out each one — Spot the Dog, Squealer the Pig — in a limited quantity and then retired it.[2]

For high-end restaurants, tables at 'Noma' in Copenhagen remain notoriously difficult to come by, and reported waiting times for 'Damon Baehrel' or 'Club 33' in the United States

[2]See: https://www.chicagomag.com/Chicago-Magazine/May-2014/Ty-Warner/, accessed on February 7, 2019. See also: https://nypost.com/2015/02/22/how-the-beanie-baby-craze-was-concocted-then-crashed/, accessed on February 7, 2019.

range from ten to fourteen years. Music festivals, such as those in Roskilde and Reading, also often sell out well in advance, year after year, with tickets for Glastonbury 2016 selling out in just 30 minutes.[3] Sell-outs are also common in professional sports, where the Boston Red Sox enjoyed a sell-out streak that stretched from 2003 to 2013. In terms of performances, Courty and Pagliero (2012) argue that concert promoters believe that empty seats would reveal negative information to consumers, and therefore select venues and prices to make sellouts more likely. In a similar spirit, 'papering the house' is a common practice where some consumers quietly receive free tickets so as to fill empty seats. Our formal model does not capture all the richness of these particular situations, but echoes the same broad idea: consumers may react favourably to sellouts, and sellers may take into account when making their strategic choices.

Formally, we consider a model of social learning where cohorts of consumers arrive sequentially at a seller and observe sales from previous periods. Informed consumers perfectly know whether product quality is high or low, whereas uninformed consumers receive a private noisy signal. Consumers in each cohort simultaneously choose whether to buy or take their outside option and then leave the game. Initially, without knowing its quality, the seller can set a capacity constraint, where per period sales cannot exceed capacity. This constraint effectively coarsens the information of consumers since they cannot observe the extent of any excess demand.

We show that the seller may find it profitable to restrict capacity even though the constraint directly limits sales, because sell-outs drive up uninformed consumers' willingness to pay. The result can be a positive purchase cascade, where all consumers buy regardless of their signal, which in particular generates excess demand in later periods. This excess demand can actually benefit the seller, if quality is in fact low, by masking the actions of informed consumers who choose not to buy, thereby allowing the cascade to be maintained. Restricting capacity tends to help in situations that a priori seem grim, where the product likely has low quality.

Thus, while our formal analysis does not consider dynamic pricing, our framework provides a rationale as to why firms may not increase prices despite regular sellouts and likely excess demand. To maximize expected profits, in the presence of demand uncertainty, it would often

[3] See http://www.glastonburyfestivals.co.uk/glastonbury-2016-tickets-sell-out-in-30-minutes/, accessed on February 7, 2019.

seem reasonable to charge a price that is high enough to prevent sellouts from *always* occurring. Our analysis points out a downside to such a strategy, namely that a failure to sell out can reduce future profits by stopping a positive purchase cascade.

We then explore how restricting capacity performs relative to Bayesian persuasion, i.e. ex ante commitment to a rule that maps product quality to a binary purchase recommendation, where each consumer decides whether to buy based on the recommendation and on her private signal. We show that if private signals are relatively imprecise, then the optimal mechanism gives consumers a recommendation to buy whenever quality is high, and sometimes when quality is low. All consumers follow these recommendations regardless of their private signal.

Using the optimal mechanism as a benchmark, we consider the optimal choice of capacity for a seller interested only in its informational impact. Here we abstract away from the cost of restricting capacity, by allowing the seller to remove its constraint following a sellout and assuming it is patient enough to only care about long-run expected sales. It turns out that if the market is large, then the optimal capacity constraint yields long-run sales that are asymptotically equivalent to those under the above-mentioned persuasion mechanism.

Our paper contributes to the literature on social learning with imperfect observability of past actions. Different work has assumed that agents can observe a random sample of actions that is anonymous (Banerjee and Fudenberg (2004), Smith and Sorensen (2013), Monzón and Rapp (2014), Monzón (2017)) or non-anonymous (Acemoglu et al. (2011), Lobel and Sadler (2015)), the aggregate total of all past actions (Callander and Hörner, 2009), the aggregate total of one particular action (Guarino et al. (2011), Herrera and Hörner (2013)), or only the choice of an agent's immediate predecessor (Çelen and Kariv, 2004). Unlike these papers, the information structure in our setting is endogenous, so consumers may fail to learn about low quality despite two features that the literature suggests should promote learning: multiple consumers who do not have access to social information (see Banerjee (1992), Sgroi (2002), Acemoglu et al. (2011), Smith and Sorensen (2013), Golub and Sadler (2017)); and unbounded private signals (see, e.g., Smith and Sørensen (2000), Banerjee and Fudenberg (2004)).

Our paper also relates to the recent literature on Bayesian persuasion pioneered by Kamenica and Gentzkow (2011). As our seller can only influence consumers' information through its choice of capacity, the paper complements work looking at a sender's choice from a restricted set of signal structures: Tsakas and Tsakas (2017) consider noise that distorts signal realizations, Perez-Richet and Skreta (2018) study sender manipulation of test results, and Ichihashi (2018) focuses on which signal-structure restrictions are optimal for the receiver. In terms of costly persuasion, Gentzkow and Kamenica (2016, 2017) and Mensch (2018) assume a direct cost associated with each experiment, whereas our cost of restricting capacity (i.e. foregone sales) is implicit and depends on consumer behavior. Other papers share our focus on dynamics (see Au (2015), Ely (2017), Renault et al. (2017), Best and Quigley (2017), Bizzotto et al. (2018), Orlov et al. (2018)), but none consider scarcity or social learning.[4]

A few recent papers consider information design as a way to influence social learning. Kremer et al. (2014) and Che and Hörner (2017) both look at settings where agents arrive sequentially, and show that a designer looking to maximize total surplus may prefer an information structure that is coarse, to encourage experimentation and promote learning. In contrast, we consider a firm looking to maximize profits, and show that it may choose a coarse information structure in order to limit learning.

Our results present a novel rationale for firms to strategically restrict capacity. Other work on such scarcity strategies has mainly focused on discouraging consumer strategic delay (DeGraba (1995), Nocke and Peitz (2007), Möller and Watanabe (2010)). In terms of scarcity and learning, Debo et al. (2012) show that a firm may reduce service speed so that consumers observe longer queues and infer higher quality. Stock and Balachander (2005) show consumers who observe that inventory is low may infer that earlier demand (and hence quality) was likely high, which can induce a firm to set low inventory. In contrast to our work, both papers assume the firm is privately informed about quality, and scarcity does not hide information from consumers in these settings, but instead may help reveal it.[5] The broader literature on how firms can influence

[4]The signal structure associated with a capacity constraint in our setting involves upper-tail censoring: revealing precise information about the state if news is sufficiently bad (i.e. demand below a threshold value), and coarse information otherwise. Dworczak and Martini (2017), Kolotilin et al. (2017), and Kolotilin and Zapechelnyuk (2018) all show that upper-tail censoring can at times be optimal in persuasion problems where the state is drawn from a continuous distribution.

[5]Vikander (2018) considers an informed firm that may limit capacity to influence consumer beliefs, but

consumer learning has mainly focused on pricing (Welch (1992), Bose et al. (2006), Bose et al. (2008), Sayedi (2018)), rather than scarcity.[6]

We describe our model in Section 2. Section 3 present preliminary results regarding consumer learning under restricted capacity, along with our key technical lemma, and Section 4 is devoted to the baseline model where all consumers are uninformed. Section 5 asymptotically compares the informational impact of capacity constraints to Bayesian persuasion. We extend the baseline model in Section 6 to consider informed consumers, and Section 7 then concludes. All the proofs are presented in Appendix A. Appendix B contains more general analysis of our deterministic and stochastic models. Appendix C presents our results on pricing in the presence of capacity constraints.

2 Model

Suppose there is a product or service of unknown quality and two possible states of the world, $\Omega = \{G, B\}$. In state G, quality is good and each consumer who buys obtains $u_G = 1$. In state B, quality is bad and each consumer who buys obtains $u_B = 0$. A consumer who does not buy gets reservation utility $r \in (0, 1)$.

The actual state is not initially known, neither to the seller nor to consumers. A priori beliefs of all players are that $P(G) \equiv \beta$ and $P(B) = 1-\beta$. In each period there are $2n$ potential buyers, who are either *informed* or *uninformed*. Before making her purchase decision, each informed consumer receives a signal that reveals the state for sure. Each uninformed consumer receives a noisy private signal, $s \in \{g, b\}$, where $P(g|G) = P(b|B) \equiv \alpha \in (1/2, 1)$. By $\alpha < 1$, our signals are *boundedly informative*. We also assume that a consumer without further information would follow her signal, i.e. $P(G|s = g) > r > P(G|s = b)$. We model the number of informed consumers in two ways. In the **deterministic setting**, we will assume there is a fixed number m of informed consumers in every period, where $1 \leq m \leq n$. In the **stochastic setting**, we

assumes bounded rationality and social image concerns.

[6]See also Gill and Sgroi (2008) on product testing, Gill and Sgroi (2012) on choice of reviewers, and Aoyagi (2010), Liu and Schiraldi (2012), and Bhalla (2013) on simultaneous versus sequential product launch.

will assume that each of the $2n$ consumers is informed with probability $\varepsilon > 0$.[7] We will also consider a **baseline** where all consumers are uninformed, which would correspond to imposing $m = 0$ or $\varepsilon = 0$ in the deterministic or stochastic settings.

At the start of the game, $t = -1$, the seller can set a *capacity constraint* $K \leq 2n$. This capacity choice is irreversible and limits potential sales in each period (i.e. how many consumers can buy), which cannot exceed capacity. The state of the world is realized at $t = 0$, so the constraint itself does not reveal any information.[8]

In each period $t \in [1, \infty)$, $2n$ consumers arrive and observe both capacity and total sales from the consumers in previous cohorts. That is, consumers do not directly observe quantity demanded in each period, but only quantity sold. Notice that a sell-out in period $t' < t$, where sales equal capacity, need not imply that demand precisely equaled capacity.

Each of the $2n$ consumers who arrive in period t choose whether they want to buy based on (i) prior beliefs about quality, β, (ii) a private signal, s, and (iii) observed sales from the previous cohorts.[9] The decision of informed consumers is solely determined by their signal. If more than K consumers want to buy in period t, a random selection of them are served, and the remaining consumers use their outside option. All period-t consumers then leave the market forever, and a new cohort of $2n$ consumers arrives in period $t + 1$.

The seller receives a fixed profit per consumer who buys, normalized to 1. He discounts profits in future periods with a factor of δ, and sets K so as to maximize expected discounted future profits.

[7]Our approach of modeling unboundedly informative signals, through the presence of fully informed consumers, differs from the more common approach of assuming continuous signals, and dramatically helps with tractability.

[8]Parsa et al. (2005) document that about 60% of new restaurants fail within three years, which suggests that their owners had imprecise information about quality when opening and setting capacity. Ultimately, we require that the seller and consumers holds the same prior, as is common in the literature on social learning, see, e.g., Bose et al. (2006), Bose et al. (2008) and Bhalla (2013).

[9]In our setting it does not matter how long a history is observed, provided that consumers observe at least sales from the two previous periods.

3 Preliminary Results

3.1 Consumer Behaviour

Define $Q^i_\omega(j)$ as the probability of j good signals from $2n$ consumers in period 1, conditional on the state being $\omega \in \{G, B\}$, in either the deterministic (denote $i = det$), stochastic (denote $i = sto$), or baseline (denote $i = base$) setting. In the deterministic setting, m consumers receive a correct signal for sure, and the remaining $2n-m$ consumers each receive a correct signal with probability α, which implies

$$Q^{det}_G(j) = \binom{2n-m}{j-m} \alpha^{j-m}(1-\alpha)^{2n-j}, \tag{1}$$

if $j \geq m$, and $Q^{det}_G(j) = 0$ if $j < m$, along with

$$Q^{det}_B(j) = \binom{2n-m}{j} \alpha^{2n-m-j}(1-\alpha)^{j}, \tag{2}$$

if $j \leq 2n-m$ and $Q^{det}_B(j) = 0$ if $j > 2n-m$.

In the stochastic setting, each consumer receives a correct signal with probability $\epsilon+(1-\epsilon)\alpha$, where ϵ is the probability of being informed. This implies:

$$Q^{sto}_G(j) = \binom{2n}{j} \left[1-(\alpha+\varepsilon-\alpha\varepsilon)\right]^{2n-j} (\alpha+\varepsilon-\alpha\varepsilon)^j, \tag{3}$$

$$Q^{sto}_B(j) = \binom{2n}{j} \left[1-(\alpha+\varepsilon-\alpha\varepsilon)\right]^{j} (\alpha+\varepsilon-\alpha\varepsilon)^{2n-j}. \tag{4}$$

The baseline probabilities $Q^{base}_\omega(j)$ can be obtained by setting $m = \varepsilon = 0$ in the above expressions. We also introduce an unconditional probability of having j good signals

$$Q^i(j) = \beta Q^i_G(j) + (1-\beta)Q^i_B(j), \qquad j \in \{base, det, sto\}$$

We start by showing that $Q^i_\omega(j)$ exhibits the following four properties.

Lemma 1. *In the baseline, deterministic, and stochastic settings, probabilities* $Q^i_\omega(j)$, $\omega \in$

$\{G, B\}$, $i \in \{base, det, sto\}$, *satisfy the following conditions:*

(i) $\frac{Q^i_B(j)}{Q^i_G(j)}$ *is non-increasing in* j.

(ii) $\frac{Q^i_B(n)}{Q^i_G(n)} = 1$, $\frac{Q^i_B(n-1)}{Q^i_G(n-1)} \geq \left(\frac{\alpha}{1-\alpha}\right)^2$ *and* $\frac{Q^i_B(n+1)}{Q^i_G(n+1)} \leq \left(\frac{1-\alpha}{\alpha}\right)^2$.

(iii) $Q^i_B(j) = Q^i_G(2n - j)$ *for all* $j \leq 2n$.

(iv) $Q^i_G(j) > Q^i_G(2n - j)$ *if* $j \geq n + 1$.

That is, (i) more good signals means the good state is more likely, (ii) having an equal number of good and bad signals is just as likely in either state, and the smallest difference from an equal number is sufficiently informative, (iii) j good signals in the bad state is as likely as j bad signals in the good state, and (iv) in the good state, $j \geq n + 1$ good signals is more likely than j bad ones.

In what follows we assume that consumers follow their private signals in the absence of other information, i.e. $P(G|s = g) > r > P(G|s = b)$, or

$$\frac{\beta\alpha}{\beta\alpha + (1-\beta)(1-\alpha)} > r > \frac{\beta(1-\alpha)}{\beta(1-\alpha) + (1-\beta)\alpha}. \tag{5}$$

Now we formulate the optimal behaviour for a consumer in period t facing an unconstrained seller, following a sequence of sales $(S_1, \ldots, S_{t-1})$.

Lemma 2. *Suppose the seller is unconstrained and consider an uninformed consumer A acting in period* $t \geq 2$. *Then in the baseline, deterministic, and stochastic settings:*

1. *If* $t = 2$ *or* $t > 2$ *and* $S_\tau = n$ *for all* $\tau \leq t - 1$:

   (a) if $S_{t-1} > n$ *then A buys regardless of her own signal;*

   (b) if $S_{t-1} = n$ *then A follows her own signal;*

   (c) if $S_{t-1} < n$ *then A does not buy regardless of her own signal.*

2. *If* $t > 2$ *and* $\max_{\tau \leq t-2} S_\tau > n$:

(a) if $S_{t-1} = 2n$, then A buys regardless of her own signal;

(b) if $S_{t-1} < 2n$ then A does not buy regardless of her own signal.

3. *If $t > 2$ and $\max_{\tau \leq t-2} S_\tau < n$:*

(a) if $S_{t-1} > 0$, then A buys regardless of her own signal;

(b) if $S_{t-1} = 0$ then A does not buy regardless of her own signal.

Following the literature on social learning, Lemma 2 describes optimal behavior on the equilibrium path; we do not specify consumers' beliefs and best responses for histories which cannot arise from equilibrium behaviour. As each consumer's payoff does not depend on the actions of those who follow, this approach is not restrictive. The idea of Lemma 2 is that initial sales of at least $n + 1$ out of $2n$ is sufficiently good news to trigger a positive purchase cascade where all uninformed consumers buy. This cascade continues unless (or until) there is a period with fewer than $2n$ sales, which would reveal that an informed consumer chose not to buy, and that the state was actually bad. This would in turn trigger a negative cascade. The story is similar for sales of at most $n - 1$ in period 1, which triggers a negative cascade where uninformed consumers do not buy, unless (or until) there is a period with strictly positive sales.

Now we look into consumer behaviour when the seller restricts capacity. It follows from Lemma 2 that a capacity $K \geq n + 1$ limits sales compared to being unconstrained, but does not increase the probability of starting a positive cascade. As such, we will focus on levels of capacity $K \leq n$ in our analysis. Moreover, setting capacity very low means that multiple sell-outs may be necessary to trigger a cascade: say, selling out a restaurant of one seat given one hundred potential customers has very little impact on consumers' beliefs. However, as Lemma 3 will establish, consumers become increasingly optimistic after each sell-out.

Consider a consumer in cohort $l + 1$ who receives a bad signal, realizes there were sell-outs in all l previous periods, and believes that all earlier consumers followed their private signals. Let $\gamma(l, K)$ denote this consumer's belief that the state is good. Then:

$$\gamma(l,K) = \frac{P(G \cap b, l \text{ sell-outs})}{P(G \cap b, l \text{ sell-outs}) + P(B \cap b, l \text{ sell-outs})} =$$

$$\frac{\beta(1-\alpha)\left[\sum_{j=K}^{2n} Q_G^i(j)\right]^l}{\beta(1-\alpha)\left[\sum_{j=K}^{2n} Q_G^i(j)\right]^l + (1-\beta)\alpha\left[\sum_{j=K}^{2n} Q_B^i(j)\right]^l} =$$

$$\frac{1}{1 + \frac{1-\beta}{\beta}\frac{\alpha}{1-\alpha}\left[\frac{\sum_{j=K}^{2n} Q_B^i(j)}{\sum_{j=K}^{2n} Q_G^i(j)}\right]^l} \quad (6)$$

where $i \in \{base, det, sto\}$.

Lemma 3. *In the baseline, deterministic, and stochastic settings, for all* $1 \leq K \leq n$*, consumer beliefs* $\gamma(l,K)$ *are increasing in* l*, with* $\lim_{l\to\infty} \gamma(l,K) = 1$.

As beliefs $\gamma(l,K)$ are increasing in l, a sufficiently long sequence of sell-outs will eventually lead to a positive purchase cascade. We use the results of Lemma 3 for our derivation of general profit functions and analysis of social learning in the long run. There we illustrate that a seller might sometimes prefer to set a low capacity that delays the start of a positive cascade. That being said, we now show that for all parameter values satisfying our assumptions, there is always a value of the outside option r such that a cascade is triggered by a single sell-out.

Lemma 4. *In the baseline, deterministic, and stochastic settings, for all* $K \leq n$ *there exists* $r \in (0,1)$ *such that* $\gamma(1,K) > r > \gamma(0,K) = P(G|s = b)$.

We will apply Lemma 4 for establishing our main existence result, i.e. that for some parameter values the seller prefers to restrict capacity.

3.2 Key Lemma

Our aim is to show that there is a range of parameters in all settings such that the seller wants to restrict its capacity. Moreover, we show that each value of $K \leq n$ can sometimes lead to higher profits than remaining unconstrained. We start with the following preliminary result, which is essential for proving our main theorems in Sections 4 and 6. Let $\mathcal{Q} = \{Q(j),\ 0 \le j \le 2n\}$ be a discrete probability measure: $Q(j) \in [0,1]$ and $\sum_{j=0}^{2n} Q(j) = 1$. This measure can correspond to that in any one of our settings described in Section 3.1, as well as in some alternative ones, but at this point we are agnostic about the actual probabilistic model, so we do not use model superscripts. Define

$$\pi_u^{\mathcal{Q}} = \frac{1}{1-\delta Q(n)} \left(\sum_{j=0}^{2n} jQ(j) + \frac{2n\delta}{1-\delta} \sum_{j=n+1}^{2n} Q(j) \right), \tag{7}$$

and

$$\pi_c^{\mathcal{Q}}(K) = \sum_{j=0}^{2n} \min\{j, K\} Q(j) + \frac{K\delta}{1-\delta} \sum_{j=K}^{2n} Q(j), \tag{8}$$

with $\delta \in (0,1)$.

Expressions (7) and (8) give seller profits assuming that cascades start upon a single sell-out and, once started, continue forever. This is true in our baseline, but not in the deterministic or stochastic setting. The following Lemma shows that if probabilities $\mathcal{Q}$ satisfy certain properties, the seller prefers to be constrained, provided that cascades continue forever. Later, in the proofs of the corresponding theorems, we will show that (i) parameters of the model can be chosen such that $Q^i(j)$, $i \in \{base, det, sto\}$ satisfy the properties required by Lemma 5, and (ii) the assumption that cascades continue forever understates the seller's incentive to restrict capacity for $i \in \{det, sto\}$.

Lemma 5. *Consider a sequence of probability measures* $\{\mathcal{Q}^s\}_{s=0}^{\infty}$ *such that* $\lim_{s\to\infty} \frac{Q^s(n)}{\sum_{j=n+1}^{2n} Q^s(j)} = \infty$*. Then, there exist* δ *and* T_0*, such that for all* $s > T_0$*,* $n > 1$*,* $K \le n$

$$\pi_u^{\mathcal{Q}^s} < \pi_c^{\mathcal{Q}^s}(K)$$

Suppose furthermore that $\lim_{s\to\infty} Q^s(n) = 0$*. Then, for any* $\delta \in \left(\sqrt{\frac{2n-K}{2n}}, 1\right)$*, there exists* T_1 *such that for all* $s > T_1$ *we have* $\pi_u^{\mathcal{Q}^s} < \pi_c^{\mathcal{Q}^s}(K)$*.*

4 Baseline Setting

4.1 Main Results

We start our analysis with the baseline and look into the seller's incentives to restrict capacity. We define a *correct cascade* as a situation where all consumers buy regardless of their signal and the state is good, or where they all refuse to buy regardless of their signal and the state is bad. We define an *incorrect cascade* as a situation where all consumers refuse to buy and the state is good, or where they all buy and the state is bad.

Profits for an unconstrained seller are

$$\pi_u^{base} = \sum_{j=0}^{2n} jQ^{base}(j) + \delta Q^{base}(n)\pi_u + \delta \sum_{j=n+1}^{2n} Q^{base}(j)\frac{2n}{1-\delta}, \tag{9}$$

with $Q^{base}(j)$ being probabilities in the baseline case. Period-1 sales of strictly more (less) than n immediately trigger a positive (negative) cascade, which continues in all later periods.

According to Lemma 4, there are values of r for which a single sell-out generates a cascade. For such values of r, profits for a seller with capacity constraint $K \leq n$ are

$$\pi_c^{base}(K) = \sum_{j=0}^{2n} \min\{j, K\}Q^{base}(j) + \delta \sum_{j=K}^{2n} Q^{base}(j)\frac{K}{1-\delta}. \tag{10}$$

Notice that (9) and (10) correspond to (7) and (8), if we set $Q(j) = Q^{base}(j)$ in the latter two expressions.

Theorem 1. *For any $n > 1$, $K \leq n$ and $\delta > \sqrt{\frac{2n-K}{2n}}$, there are $(\alpha, \beta) \in (1/2, 1) \times (0, 1/2)$ and $r > 0$ for which the seller can increase its profits above the unconstrained level by restricting capacity to K.*

The seller faces a tradeoff, as restricting capacity limits sales in each period, but also increases the probability of a positive cascade. A key difference between (9) and (10) is that below average period-1 demand, between K and n, can only trigger a positive cascade if the seller is constrained. This effect can make it profitable to restrict capacity, in particular when a positive cascade occurs after a single sell-out. Intuitively, having set a capacity constraint will tend to help the seller when the state turns out to be bad, because the constraint can hide information from consumers. The seller effectively engages in obfuscation, which increases the probability of an incorrect positive cascade (where all consumers buy despite the state being bad) and decreases the probability of an incorrect negative cascade (where no consumer buys despite the state being good).

Having shown that the seller may sometimes find it profitable to restrict capacity, we now say something about the conditions under which this can occur. Consistent with the above intuition, the seller will only restrict capacity if it is sufficiently likely that the state is bad.

Proposition 1. *Suppose that $\beta \geq 1/2$. Then for all parameter values, such that consumers in the first period follow their signals, the seller prefers not to restrict capacity: $\pi_u^{base} > \pi_c^{base}(K)$ for all $K \leq 2n$.*

When the state is good, initial sales are likely to be sufficiently high to trigger a positive cascade regardless of whether the seller restricts capacity, and an unconstrained seller then enjoys higher sales. The proof shows that for $\beta \geq \frac{1}{2}$, the good state is sufficiently likely that expected sales for an unconstrained seller exceed n per period. This is more than the seller could possibly earn per period by restricting capacity to $K \leq n$, which are the only capacity levels that can increase the probability of a positive cascade. The result of Proposition 1 does not rely on the outside option satisfying Lemma 4, i.e. it holds regardless of whether a single sellout is sufficient to trigger a cascade, or if multiple sellouts are necessary (see Section 4.2).[10]

The sufficient condition provided in Proposition 1 is simple and intuitive, but the necessary condition is much harder to obtain, so we investigate this question numerically. Figure 1(a) shows that the region where restricting capacity can be optimal is larger when there are many consumers. Intuitively, an increase in the number of consumers makes period-1 demand more informative, so that an unconstrained seller is more likely to suffer a negative cascade in the bad state. Restricting capacity then becomes more attractive, as a way to hide unfavourable

[10]A result similar to Proposition 1 can be obtained in the deterministic and stochastic models. The proof is available upon request.

Figure 1: Areas where capacity restriction can be optimal.

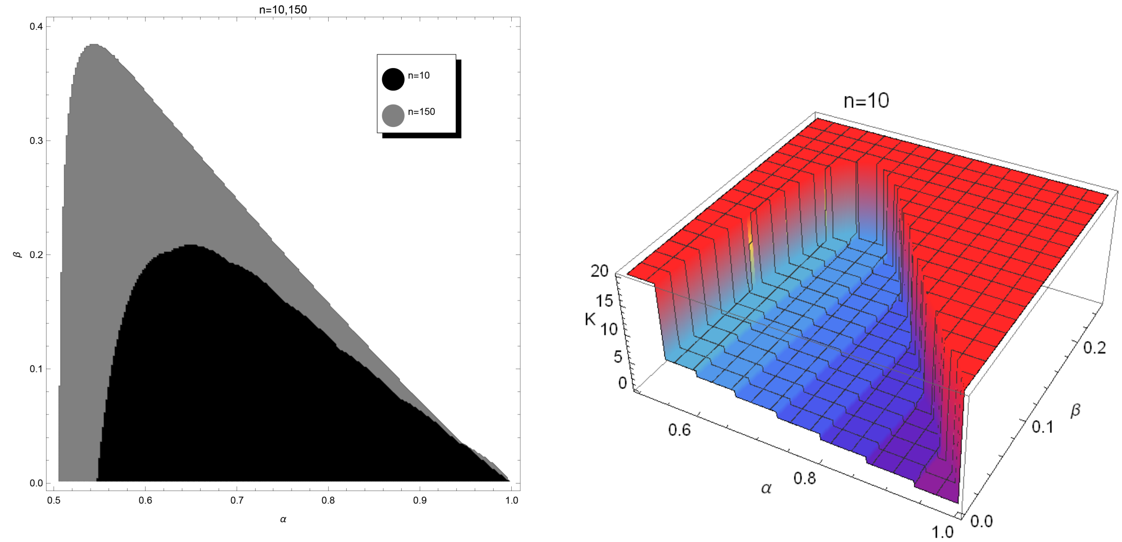

(a) Baseline: area where restricting capacity is optimal for some r, δ

(b) Baseline: optimal K

information from consumers.

Figure 1(a) also shows an interesting interaction between signal precision and the prior. When the signal is very imprecise ($\alpha \approx 1/2$), at least half of the consumers will likely buy, regardless of the state, so the seller will remain unconstrained and bet on generating a positive cascade. As signal precision increases, an unconstrained seller becomes increasingly likely to experience a negative cascade when the state is bad, so the seller will restrict capacity for sufficiently small β. A further increase in signal precision makes restricting capacity increasingly helpful in the bad state, but increasingly harmful in the good state, as a correct positive cascade is then likely without a constraint. The first effect dominates for intermediate α, but the second effect dominates when α is sufficiently large, because an incorrect positive cascade is only likely given a very low constraint, which would dramatically reduce profits in the good state. Figure 1(b) depicts the optimal level of capacity for these different parameter values.

In our analysis, we assumed that the seller cannot adjust capacity over time, but this assumption is not crucial. A seller that could adjust its capacity would effectively remove the constraint if an initial sellout triggered a positive cascade. This insight only reinforces our result that restricting capacity can be optimal, as future gains following a cascade would no longer be bounded by the size of the initial constraint. If the seller found it optimal to restrict capacity in our setting, then it would also do so if capacity could be adjusted over time. Moreover, the effective cost of setting a low constraint would be limited to a single period, so the seller would be more aggressive in period 1 in order to make a sellout more likely, and then capture the resulting gains. The following Proposition formally describes this result.

Proposition 2. *Suppose that K_1 is the optimal period-1 capacity constraint when the seller can adjust capacity over time, and that K_0 is the optimal constraint when the seller cannot do so. Suppose furthermore that selling out once at either K_0 or K_1 triggers a positive cascade, i.e.* $\min\{\gamma(1, K_0), \gamma(1, K_1)\} \geq r$*. Then $K_0 \geq K_1$.*

4.2 General Profit Functions

The profit function (10) for a constrained seller applies to situations where a single sell-out triggers a positive cascade. For given prior beliefs, signal accuracy, and capacity, we can always find a value of the consumers' outside option such that one sell-out is indeed sufficient. But more generally, multiple sell-outs may be necessary to trigger a cascade, in particular when capacity is relatively low. Might the seller have an incentive to set such a capacity, which delays the start of a positive cascade, but also increases the probability of selling out?

To address this issue, we write down general profit functions for a seller with capacity K, in the baseline setting. Recall from (6) that $\gamma(l, K)$ denotes the belief of a consumer in cohort $l + 1$ that the state is good, after sell-outs in all l previous periods, and after receiving a bad signal. Let $\mathcal{L} = \{l : \gamma(l, K) > r\}$, which is non-empty by Lemma 3, and denote the smallest element of $\mathcal{L}$ by L: the number of consecutive sell-outs required as of period 1, given capacity K, in order to generate a positive cascade. The value of L is decreasing in K, since selling out at higher capacity presents stronger evidence that the state is good.

Let $\eta_\omega = \sum_{j=K}^{2n} Q_\omega^{base}(j)$ denote the probability of a sell-out, given state $\omega \in \{G, B\}$. Let $S_\omega = \sum_{j=0}^{K-1} jQ_\omega^{base}(j)$ denote expected sales in a given period conditional on not selling out, when the state is ω.

Figure 2: Seller profits and L.

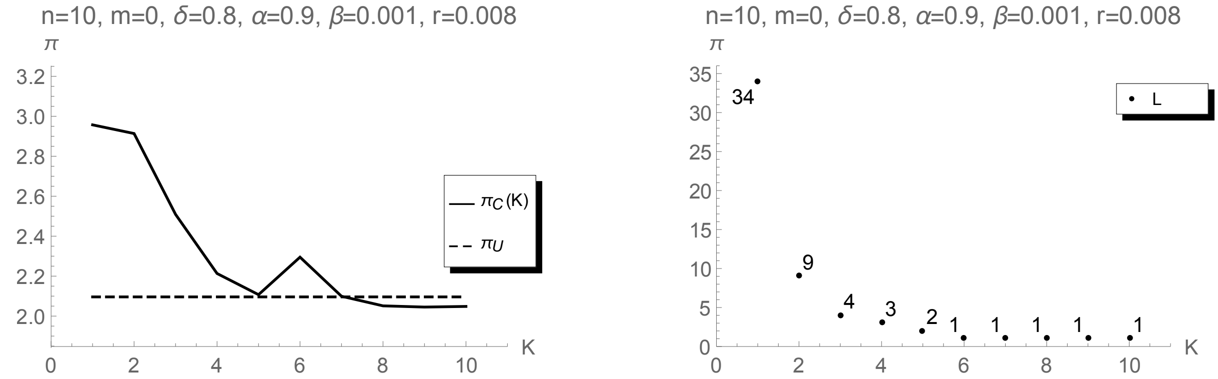

(a) (b)

Expected profits given capacity K and $L \geq 1$ are

$$\pi_c^{base}(K) = \beta \left[\frac{1-(\delta\eta_G)^L}{1-\delta\eta_G}(S_G + \eta_G K) + (\delta\eta_G)^L \frac{K}{1-\delta} \right] + (1-\beta) \left[\frac{1-(\delta\eta_B)^L}{1-\delta\eta_B}(S_B + \eta_B K) + (\delta\eta_B)^L \frac{K}{1-\delta} \right]. \quad (11)$$

When the state is bad, the seller experiences L consecutive sell-outs each with probability η_B, which triggers a positive cascade and yields revenue of K in all periods. Otherwise, the seller earns S_B in the first period where it fails to sell out, and then zero in all later periods. The situation is similar when the state is good, except the probability of a sell-out in each period is η_G rather than η_B, and a seller that fails to sell out earns S_G in that period. As expected, expression (11) reduce to (10) if $L = 1$, so if one sell-out is sufficient to trigger a cascade.

Holding capacity constant at K, profits are strictly decreasing in L, because a large L means that more sell-outs are required to trigger a positive cascade. However, L depends on K, and the seller may in fact want to set a sufficiently low capacity that results in $L > 1$. Such a low capacity makes a sell-out in each period more likely, and thus reduces the probability in each period that a negative cascade begins. Indeed, Figure 2 shows parameter values for which the seller prefers to set $K = 1$, so that $L = 34$ sell-outs are necessary to trigger a cascade.

Figure 2 also shows that profits might not be quasi-concave in K and sudden jumps can occur. An increase in K gives a lower chance of a sell-out in any given period, higher sales conditional on selling out or on being in a positive cascade, and may also mean that fewer sell-outs are necessary for the cascade to begin. The balance between these three forces may sometimes swing sharply. In the figure, an increase in K from 5 to 6 reduces the number of sell-outs necessary to start a cascade from $L = 2$ to $L = 1$, which pushes up profits.

5 Optimal Information Design

We now explore how the informational impact of restricting capacity, in terms of influencing consumer behavior, compares to that of optimal information design. Our seller (sender) chooses capacity without knowing the state, and each consumer (receiver) then receives information that depends on the realized state and on capacity. In this sense, capacity constraints serve as a natural example of a Bayesian persuasion mechanism that can be easily implemented in practice. As capacity is chosen ex ante, and market participants do not observe excess demand, any commitment problem in implementing the desired information structure is avoided. The seller simply cannot serve demand that exceeds capacity. Consumers then directly observe the resulting sell-out; if they did not, the seller could easily disclose that a sell-out occurred.

In what follows, we derive the optimal 'persuasion mechanism' in our setting, and then examine how capacity constraints perform relative to this benchmark. The seller commits to a rule which maps the binary state into a purchase recommendation. The state is realized, and each consumer receives a recommendation according to the chosen rule. Each consumer then makes a purchase decision based on the recommendation and her own private signal, and the seller serves all consumers who want to buy. Thus, the seller's persuasion mechanism substitutes for the social learning process studied in the previous sections.

We consider *information design without elicitation*, where the same rule applies to all consumers, regardless of their private signal. As Kolotilin et al. (2017) and Bergemann and Morris (2017) show, given binary state and action spaces, this forms of information design is equivalent to design with elicitation, in the sense of being able to implement the same consumer behavior.

As such, the rule we derive will constitute the optimal mechanism.[11]

The seller designs the persuasion mechanism to maximize expected sales, which are the same in every period. When later making the comparison with capacity constraints, we will assume that the seller is infinitely patient, and can adjust capacity over time. These assumptions effectively remove the direct cost of restricting capacity, i.e. limiting sales. The seller will therefore set capacity to maximize long-run expected per-period sales, and have the same objective function as in the persuasion setting.

The optimal persuasion mechanism is characterised by a binary message space corresponding to the binary decision set of consumers (see Kamenica and Gentzkow (2011)). However, as buyers are informed, the seller essentially chooses between three options. The first is where all consumers follow the seller's recommendation, which is the 'textbook' form of persuasion. The second is where consumers with bad private signals always follow the seller's recommendation, but consumers with good signals ignore the recommendation and instead always buy. The third is where all consumers ignore the recommendation, and follow their private signals.

We first consider a mechanism where all consumers follow the seller's recommendation. Suppose the seller sends a *buy* recommendation with probability p_G in a good state and with probability p_B in bad state, and otherwise sends a *no-buy* recommendation. The belief of a consumer with a bad signal upon receiving a buy recommendation is

$$\gamma(s=b,buy)=\frac{\beta(1-\alpha)p_G}{\beta(1-\alpha)p_G+(1-\beta)\alpha p_B}\geq r,$$

which implies

$$\beta(1-\alpha)(1-r)p_G\geq(1-\beta)\alpha r p_B. \tag{12}$$

Clearly, if (12) is satisfied, then consumers with $s=g$ prefer to buy. As both p_G and p_B enter opposite sides of (12) with positive signs, setting

[11] We do not allow the seller to directly target different rules at consumers with different private signals. In this sense, the seller is not *omniscient* (see Bergemann and Morris (2017)).

$$p_G = 1, \quad p_B = \frac{\beta(1-\alpha)(1-r)}{(1-\beta)\alpha r}, \tag{13}$$

maximizes expected sales. Our assumption that a consumer would not buy, given only a bad private signal, $P(G|s = b) < r$, implies $p_B < 1$, since

$$p_B < 1 \Leftrightarrow \beta(1-\alpha)(1-r) < (1-\beta)\alpha r \Leftrightarrow \frac{\beta(1-\alpha)}{\beta(1-\alpha) + (1-\beta)\alpha} < r.$$

The seller always sends a buy recommendation in the good state, and only sometimes sends a no-buy recommendation in the bad state, where the latter recommendation is perfectly revealing. Echoing the discussion in Section 4 on restricting capacity, the seller effectively obfuscates when the state is bad, by sometimes sending the same recommendation as when the state is good. This recommendation convinces consumers to buy, even those with bad private signals. The profits from this persuasion mechanism, per consumer and per period, are

$$\pi^* = \beta + (1-\beta)p_B = \beta + \frac{(1-\alpha)(1-r)\beta}{\alpha r}. \tag{14}$$

Now consider a mechanism where only consumers with bad signals follow the seller's recommendation, and let $(\tilde{p}_G, \tilde{p}_B)$ denote the probabilities of buy recommendations. The incentive compatibility constraint for a consumer with a good signal who receives a no-buy recommendation is

$$\gamma(s = g, \textit{not buy}) = \frac{\beta\alpha(1-\tilde{p}_G)}{\beta\alpha(1-\tilde{p}_G) + (1-\beta)(1-\alpha)(1-\tilde{p}_B)} \geq r,$$

or

$$\beta\alpha(1-r)(1-\tilde{p}_G) \geq (1-\beta)(1-\alpha)r(1-\tilde{p}_B). \tag{15}$$

Both (15) and (12) must bind at the optimum, which implies

$$\tilde{p}_G = \frac{\alpha(\alpha\beta - r(\alpha(2\beta-1) - \beta + 1))}{(2\alpha-1)\beta(1-r)}, \quad \tilde{p}_B = \frac{(1-\alpha)(\alpha\beta - r(\alpha(2\beta-1) - \beta + 1))}{(2\alpha-1)(1-\beta)r}. \tag{16}$$

The probability $\tilde{p}_B$ is well-defined as long as $\tilde{p}_G \in [0,1]$. This is the case if

$$\alpha \geq \underline{\alpha}(r,\beta) \equiv \frac{\max\{(1-r)\beta, r(1-\beta)\}}{r+\beta-2r\beta}, \tag{17}$$

which is equivalent to condition (5), that consumers follow their private signals in the absence of other information. The profits from this persuasion mechanism are

$$\tilde{\pi} = \beta[\alpha + (1-\alpha)\tilde{p}_G] + (1-\beta)[\alpha\tilde{p}_B + (1-\alpha)]. \tag{18}$$

As long as $\tilde{p}_G, \tilde{p}_B > 0$, these profits exceed $\beta\alpha + (1-\beta)(1-\alpha)$, which is what the seller would earn if all consumers followed their private signals. Thus, the optimal mechanism either leads all consumers to follow the seller's recommendation, or only those with bad signals to do so.

Define

$$\overline{\alpha}(r) \equiv \frac{1}{2}\left(2r - 1 + \sqrt{4r^2 - 8r + 5}\right), \quad \underline{\beta}(r) \equiv r\frac{2r - \overline{\alpha}(r)}{3r-1}, \quad \overline{\beta}(r) \equiv \frac{r}{1-r}[\overline{\alpha}(r) + 1 - 2r],$$

and note that $0 < \underline{\beta}(r) < r < \overline{\beta}(r)$ for all $r \in (0,1)$. The following Proposition establishes when each of the two mechanisms yield higher profits, while taking into account constraint (17).

Proposition 3. *For any $r \in [0,1]$, if $\beta \in [\underline{\beta}(r), \overline{\beta}(r)]$ and $\alpha \in [\underline{\alpha}(r,\beta), \overline{\alpha}(r)]$ then the optimal persuasion mechanism is given by* (13). *Otherwise, the optimal persuasion mechanism is given by* (16) *as long as $\alpha > \underline{\alpha}(r,\beta)$.*

Proposition 3 says that as long as signal precision is not too high, i.e. $\alpha \leq \overline{\alpha}(r)$, the optimal mechanism leads all consumers to follow the seller's recommendation. The restriction $\alpha \geq \underline{\alpha}(r,\beta)$ corresponds to our initial condition that consumers follow their own signals in the absence of any other information. Finally, the restriction $\beta \in [\underline{\beta}(r), \overline{\beta}(r)]$ guarantees that the interval $[\underline{\alpha}(r,\beta), \overline{\alpha}(r)]$ is non-empty.[12]

[12] Proposition 3 is similar to the result obtained in Kolotilin (2018), the only difference being that he also allows for the case $\overline{\alpha}(r) < \alpha < \underline{\alpha}(r,\beta)$. This case is ruled out in our setting by the assumption that consumers follow their private signals in the absence of any other information.

Intuitively, for consumers with good private signals to ignore a no-buy recommendation, the seller must sometimes send this recommendation in the good state, so the recommendation is not perfectly revealing. The effective cost to the seller is that doing so reduces expected sales in the good state, compared to a mechanism where consumers always follow the seller's recommendation. The size of this cost is decreasing in signal precision, because a consumer with a very precise good signal may be willing to ignore a no-buy recommendation that suggests the bad state is quite likely. In general, the scope for persuasion decreases as signals become more precise: if $\alpha \to 1$, then the expected profit of both mechanisms approaches β.

We now look into the optimal choice of capacity as a tool for information design. We again assume that consumers observe sales from previous cohorts, and in particular whether a sellout occurred. To ensure comparability with the optimal persuasion mechanism, we assume that the seller is infinitely patient, and can costlessly adjust its capacity over time. The seller sets its initial capacity to maximize expected per-period sales in the long run, which is equivalent to maximizing the probability of a positive cascade.[13] The question is how this probability compares to that under the optimal persuasion mechanism, either (13) or (16).

First we focus on values of capacity K such that a single sellout triggers a positive cascade. Recall that $\eta_\omega(K, 2n) = \sum_{j=K}^{2n} Q_\omega^{base}(j)$ denotes the probability of sales greater than or equal to K. A consumer with a bad private signal will buy after a single sellout if

$$\frac{\beta(1-\alpha)\eta_G(K, 2n)}{\beta(1-\alpha)\eta_G(K, 2n) + (1-\beta)\alpha\eta_B(K, 2n)} \geq r,$$

or equivalently

$$\beta(1-\alpha)(1-r)\eta_G(K, 2n) \geq (1-\beta)\alpha r\eta_B(K, 2n). \tag{19}$$

Constraint (19) resembles (12) but with one important difference: in the persuasion problem, the seller can choose p_G and p_B independently, but now the choice of capacity jointly determines both probabilities. The firm is interested in maximizing the probability of sell-out, so it will choose the lowest K such that (19) is satisfied. Denote this value of K by $K^*(n)$.[14]

[13] If a positive cascade is triggered, the seller will immediately increase capacity and serve full demand.

[14] Note that for any n, we have $\eta_G(0, 2n) = \eta_B(0, 2n) = 1$. Moreover, $\lim_{n\to\infty} \eta_G(n, 2n) = 1$ and

Now we consider large cohorts. Note that $\lim_{n\to\infty} \eta_G(n, 2n) = 1$ and $\lim_{n\to\infty} \eta_B(n, 2n) = 0$. Thus, for sufficiently large n, it must be the case that $K^*(n) < n$, so that $\lim_{n\to\infty} \eta_G(K^*(n), 2n) = 1$; $\lim_{n\to\infty} \eta_B(K^*(n), 2n) = \frac{\beta(1-\alpha)(1-r)}{(1-\beta)\alpha r} < 1$, which is the same as the recommendation probabilities in persuasion mechanism (13).

It follows that the limit result of restricting capacity is equivalent to the persuasion mechanism where a buy recommendation is always sent in the good state and where consumers always follow all recommendations. Here, a sellout always occurs in the good state and sometimes occurs in the bad state, and it provides just enough good news to convince consumers to buy regardless of their signal (i.e. a positive cascade). A failure to sell out perfectly reveals that the state is bad, so that no consumer buys (i.e. a negative cascade).

We conclude that in large markets, the informational impact of the optimal capacity constraint serves as a good approximation to the above-mentioned persuasion mechanism. As this mechanism is optimal when private signals are relatively imprecise, a seller can achieve optimal information design in large markets with weak private information by restricting capacity to an appropriate level.

One interpretation of the seller's problem, as considered in this section, is in terms of trial sales: the seller can first conduct a trial by making a limited number of units available of its product, and then broadly release the product to all consumers if the trial goes well. Our results suggest that in situations where private information is weak, the seller can do no better than carrying out a single trial, then releasing the product broadly in the case of a sell-out, and otherwise withdrawing the product from the market.

For completeness, we note that the optimal capacity constraint can never induce consumers with good private signals to buy when the seller fails to sell out. Recall, that $K^*(n)$ is the smallest capacity constraint to satisfy (19). Moreover, for $K \in [K^*(n), n]$, we have that $\eta_G(K, 2n)$ approaches 1 and $\eta_B(K, 2n)$ remains bounded away from 1 in a large market.[15] Incentive com-

$\lim_{n\to\infty} \eta_B(n, 2n) = 0$, so $K^*(n)$ exists for sufficiently large n. Finally, due to Lemma 1 part (i), $K^*(n)$ is uniquely defined.

[15]Setting $K > n$ is clearly suboptimal as sell-outs then only occur in the good state.

patibility for a consumer with a good private signal and a no-buy recommendation is

$$\beta\alpha(1-r)[1-\eta_G(K,2n)] \geq (1-\beta)(1-\alpha)r[1-\eta_B(K,2n)],$$

which clearly cannot be satisfied as n approaches infinity, since the left-hand-side approaches zero while the right-hand-side does not. Thus, in situations where (16) constitutes the optimal persuasion mechanism, restricting capacity will result in inferior information design.[16]

We now use the above framework to say more about when the seller will restrict capacity and the size of the optimal constraint. Clearly, a patient seller that can adjust its capacity would never want to remain unconstrained in a large market. Period-1 sales would perfectly reveal the true state, so an unconstrained seller would experience a correct positive cascade with probability β,[17] and an incorrect positive cascade with probability 0. In contrast, a seller that restricts capacity will experience a correct positive cascade with probability β, and an incorrect positive cascade with probability $(1-\beta)\eta_B(K^*(n), 2n)) > 0$.

For the size of the optimal constraint, recall that $\lim_{n\to\infty} \eta_B(K^*(n), 2n) = \frac{\beta(1-\alpha)(1-r)}{(1-\beta)\alpha r}$. The left-hand side is directly decreasing in K^*, since a higher capacity reduces the probability of selling out. Looking at the right-hand side, it follows that the size of the initial constraint set by a patient seller that can adjust its capacity over time, in a large market, is decreasing in the prior, β, and increasing in both signal precision, α, and the value of the outside option, r.

The idea is that a sellout at capacity $K^*(n)$ hides just enough information about the bad state to make a consumer with a bad private signal willing to buy. For such a consumer, a drop say in signal precision makes buying more attractive. As a result the seller can hide on average more unfavourable information by setting a lower capacity. Doing so makes it more likely to sell out, and a single sellout will still generate a positive cascade.

The fact that the optimal capacity is increasing in signal precision stands in contrast with

[16]The key issue is that designing a persuasion mechanism involves two "degrees of freedom", while setting a capacity constraint involves only one. As a mechanism with $p_G = 1$ is essentially a corner solution, any capacity sufficiently far below mean demand yields a sellout for sure in the good state, and the seller can set its exact capacity so as to best approximate p_B. However, as $\tilde{p}_G < 1$, this adds an extra binding constraint to the problem, which makes an approximation of $(\tilde{p}_G, \tilde{p}_B)$ impossible.

[17]This can be seen formally when we present (24), by letting n approach infinity in this expression.

Figure 1(b). What is the reason for this difference? In the current section, the outside option r is fixed, and $K^*(n)$ is the minimal capacity sufficient for triggering a cascade given the outside option. Since the seller is patient, and can adjust capacity over time, its initial capacity choice is essentially costless: it will set $K = K^*(n)$ to maximize the probability of a positive cascade. When plotting Figure 1(b), we allowed the outside option r to vary, so that the minimal capacity to trigger a cascade was always $K = 1$. We also assumed the seller could not adjust capacity over time, which made setting $K = 1$ costly in terms of limiting future sales. As a result, the seller often found it optimal to restrict capacity to a higher level, which was decreasing in signal precision α.

If the seller sets low capacity $K < K^*(n)$, then condition (19) will be violated, but Lemma 3 ensures a cascade will occur after a sufficiently long sequence of sellouts. For (α, β, r, n), and any K, there exists l such that

$$\beta(1-\alpha)(1-r)[\eta_G(K, 2n)]^l \geq (1-\beta)\alpha r[\eta_B(K, 2n)]^l.$$

Following our earlier notation, let L denote the smallest integer l satisfying this inequality, so that L sellouts are needed to trigger a positive cascade.

We now derive the long-run probability of a positive cascade in a large market, given low capacity. Define $\Gamma(l, K)$ as the public belief that the state is good, conditional on sell-outs at capacity K in the first $l \leq L$ periods. That is

$$\Gamma(l, K) = \frac{1}{1 + \frac{1-\beta}{\beta}\frac{[\eta_B(K)]^l}{[\eta_G(K)]^l}}, \tag{20}$$

where $r < \gamma(L, K) < \Gamma(L, K)$ follows from (6) and the definition of L. Moreover, we can write

$$\left[\beta(\eta_G)^L + (1-\beta)(\eta_B)^L\right]\Gamma(L, K) + \left(1 - \left[\beta(\eta_G)^L + (1-\beta)(\eta_B)^L\right]\right)\mu = \beta. \tag{21}$$

The public belief conditional on an eventual positive cascade triggered by L sellouts, multiplied by the probability of such a cascade; plus the average public belief conditional on an eventual negative cascade triggered by a failure to sell out in some $l \leq L$ period (denoted by μ), multiplied by the probability of such a cascade; must equal the prior β.

By Lemma 1, high sales provide good news about the state, so for any fixed K we have

$$\mu < \frac{\Gamma(L-1,K)Q_G^{base}(K-1)}{\Gamma(L-1,K)Q_G^{base}(K-1)+[1-\Gamma(L-1,K)]Q_B^{base}(K-1)} = \frac{1}{1+\frac{[1-\Gamma(L-1,K)]Q_B^{base}(K-1)}{\Gamma(L-1,K)Q_G^{base}(K-1)}}. \quad (22)$$

The average public belief, conditional on failing to sell out in one of the first L periods, cannot exceed the public belief after selling out in exactly $L-1$ periods, and then having period-L sales of $K-1$, only 1 below capacity. By the definition of L, a consumer who observes $L-1$ sellouts will not buy if she receives a bad private signal. This implies that

$$\frac{\Gamma(L-1,K)(1-\alpha)}{\Gamma(L-1,K)(1-\alpha)+(1-\Gamma(L-1,K))\alpha} < r \Rightarrow \Gamma(L-1,K) < \frac{\alpha r}{\alpha r+(1-\alpha)(1-r)} < 1,$$

so $\Gamma(L-1,K)$ is bounded away from 1 for all n and K. Moreover, we have

$$\frac{Q_B^{base}(K-1)}{Q_G^{base}(K-1)} = \frac{\binom{2n}{K-1}\alpha^{2n-K+1}(1-\alpha)^{K-1}}{\binom{2n}{K-1}\alpha^{K-1}(1-\alpha)^{2n-K+1}} = \left(\frac{\alpha}{1-\alpha}\right)^{2(n-K+1)}.$$

Now hold K fixed and let cohort size grow large. Then $\alpha > 1/2$ implies $\lim_{n\to\infty}\frac{Q_B^{base}(K-1)}{Q_G^{base}(K-1)} = \infty$, which combined with $\lim_{n\to\infty}\Gamma(L-1,K) < 1$ and (22) yields $\lim_{n\to\infty}\mu = 0$: in a large market, any failure to sell out perfectly reveals the bad state. The good state must therefore eventually result in a positive cascade: $\lim_{n\to\infty}(\eta_G)^L = 1$.

Substituting into (21) shows that the probability of a positive cascade in the bad state $[\eta_B(K,2n)]^L$ satisfies

$$\left(\beta+(1-\beta)[\eta_B(K,2n)]^L\right)\Gamma(L,K) = \beta \quad (23)$$

in a large market. By definition, L is the lowest value of l for which $\gamma(l,K) \geq r$, or equivalently $\Gamma(l,K) \geq \frac{\alpha r}{\alpha r+(1-\alpha)(1-r)}$. Thus, $\lim_{n\to\infty}\Gamma(L,K) = \frac{\alpha r}{\alpha r+(1-\alpha)(1-r)}$, which by (23) implies

$\lim_{n\to\infty}(\eta_B(K,2n))^L = \frac{\beta(1-\alpha)(1-r)}{(1-\beta)\alpha r} = \lim_{n\to\infty}\eta_B(K^*(n),2n)$.

Thus, the long-run probability of a positive cascade given low capacity, in a large market, is asymptotically equivalent to that under capacity $K^*(n)$, and also under persuasion mechanism (13). This means that an infinitely patient seller, that can adjust its capacity over time, will not go wrong by 'undershooting' capacity $K^*(n)$. Setting $K < K^*(n)$ will increase the probability of a sellout in any period, and mean that more sellouts are needed to trigger a positive cascade. But the probability of an eventual cascade, and long-run expected sales, are left unchanged.

We conclude this section by again assuming that the seller's initial choice of capacity is *irreversible*, and that the market is not necessarily large, to say more about the optimal constraint for a patient seller. From (21), we know that $(\beta[\eta_G(K,2n)]^L + (1-\beta)[\eta_B(K,2n)]^L)\Gamma(L,K) \leq \beta$. Combined with $\Gamma(L,K) > r$, this yields

$$\beta[\eta_G(K,2n)]^L + (1-\beta)[\eta_B(K,2n)]^L < \frac{\beta}{r},$$

so the probability of a positive cascade is bounded above by β/r.

The long-run per period expected profits of a constrained seller is bounded by the probability of a positive cascade times its capacity,

$$\pi_c^{base}(K) = \left(\beta[\eta_G(K,2n)]^L + (1-\beta)[\eta_B(K,2n)]^L\right) K < \frac{\beta}{r}K.$$

Now consider the long-run per period expected profits of an unconstrained seller. Such a seller serves $2n$ consumers in the long run with probability

$$\frac{\sum_{j=n+1}^{2n} Q^{base}(j)}{1-Q^{base}(n)} = \beta + (1-2\beta)\frac{\sum_{j=0}^{n-1} Q_G^{base}(j)}{1-Q^{base}(n)}, \tag{24}$$

which exceeds β for all $\beta < 1/2$. Thus, $\pi_u^{base} > 2n\beta$ holds for all $\beta < 1/2$, which by Proposition 1 are the only values of the prior for which restricting capacity can be optimal. This allows us to formulate the following result regarding the optimal choice of capacity.

Proposition 4. *For δ sufficiently close to 1, only capacity constraints $K > 2nr$ can be profitable* *relative to being unconstrained:* $\pi_u^{base} > \pi_c^{base}(K)$ *for all* $K \leq 2nr$.

This result rules out capacity constraints that are too low, so for which any possible benefit is outweighed by the cost of foregone sales. Whereas 'undershooting' by setting low capacity does little harm to a patient seller that can adjust capacity over time, it is highly undesirable for a seller that cannot. This observation corroborates Proposition 2, which suggested a seller that cannot adjust capacity will set a higher constraint. An implication of Proposition 4 is that a patient seller facing consumers with outside option $r > 1/2$, will never restrict capacity.

6 Informed Consumers

6.1 Main Results

Now we extend the analysis of Section 4 to take into account the presence of informed consumers. We start with the **deterministic model.** Profits for an unconstrained seller are

$$\begin{aligned}\pi_u^{det} = \sum_{j=0}^{2n} jQ^{det}(j) + \delta Q^{det}(n)\pi_u^{det} \\ + \beta\delta \left[\sum_{j=n+1}^{2n} Q_G^{det}(j)\frac{2n}{1-\delta} + \sum_{j=0}^{n-1} Q_G^{det}(j)\left(m + \frac{2n\delta}{1-\delta}\right)\right] \\ + (1-\beta)\delta \left[\sum_{j=n+1}^{2n} Q_B^{det}(j)\frac{2n}{1-\delta} - \sum_{j=n+1}^{2n} Q_B^{det}(j)\left(m + \frac{2n\delta}{1-\delta}\right)\right]. \quad (25)\end{aligned}$$

For an unconstrained seller, period-1 sales of strictly more (less) than n immediately trigger a positive (negative) cascade. This cascade continues for all further periods if it is correct, but not if it is incorrect. An incorrect cascade will be reversed in period 2 as informed consumers' purchase decisions reveal the true state. The result is a correct cascade in all future periods.

According to Lemma 4, there are values of r for which a single sell-out generates a cascade.

For such values of r, profits for a seller with capacity constraint $K \leq n$ are

$$\pi_c^{det}(K) = \sum_{j=0}^{2n} \min\{j, K\} Q^{det}(j)$$
$$+ \beta\delta \left[\sum_{j=K}^{2n} Q_G^{det}(j) \frac{K}{1-\delta} + \sum_{j=0}^{K-1} Q_G^{det}(j) \left(\min\{m, K\} + \frac{K\delta}{1-\delta} \right) \right]$$
$$+ (1-\beta)\delta \left[\sum_{j=K}^{2n} Q_B^{det}(j) \frac{K}{1-\delta} \right]. \quad (26)$$

Comparing (26) with (25) shows that period-1 demand of at least K now triggers a positive cascade. Moreover, an incorrect positive cascade now will never be reversed, even though informed consumers refuse to buy in every period. Their decisions are effectively hidden by the capacity constraint, as they are pooled with uninformed consumers who are rationed.

Theorem 2. *For any $n > 1$, $1 \leq m < n$, $K \leq n$ and $\delta > \sqrt{\frac{2n-K}{2n}}$, there are $(\alpha, \beta) \in (1/2, 1) \times (0, 1/2)$ and $r > 0$ for which the seller can increase its profits above the unconstrained level by restricting capacity to K.*

This result is proved in two steps. First, if the seller is unconstrained, then informed consumers' choices will immediately reverse any incorrect cascade, but an incorrect positive cascade is more likely than an incorrect negative one, because the bad state is more likely. This implies $\pi_u^{\text{det}} < \pi_u^{\mathcal{Q}^{\text{det}}}$. If the seller is constrained, then informed consumers' choices will immediately reverse an incorrect negative cascade, but will never reverse an incorrect positive one, which implies $\pi_c^{\text{det}}(K) > \pi_c^{\mathcal{Q}^{\text{det}}}(K)$. The second step is to show that α and β can be chosen in such a way that Lemma 5 applies.

Intuitively, restricting capacity can help the seller through two channels: by increasing the probability of a positive cascade, and by helping such a cascade (once triggered) to be maintained. For the first channel, restricting capacity makes it more likely that all uninformed consumers want to buy immediately after observing period-1 sales, just as in our baseline setting. For the second channel, the presence of informed consumers will immediately reverse any incorrect cascade if the seller is unconstrained, but will not reverse an incorrect positive cascade if the seller is constrained. The decision of uninformed consumers not to buy is effectively unobservable, as sell-outs continue in all periods, and the seller continues to experiences excess demand.

Looking at consumer behavior as of period 3, only a constrained seller can ever experience an incorrect cascade, more specifically an incorrect positive cascade. This stands in contrast to the baseline, where both types of incorrect cascade could occur, regardless of whether the seller restricted capacity. Yet behavior in the deterministic setting nonetheless resembles that in the baseline with large market size. In the baseline, as the market grew large, the probability of an incorrect negative cascade approached zero, but the probability of an incorrect positive cascade did not if the seller was constrained. The intuition is that large market size, and the presence of informed consumers, both tend to make demand more informative. This will often result in consumers learning, unless the state is bad and the seller restricts capacity.

A similar result to Theorem 2 applies in the **stochastic setting**. Profits for an unconstrained seller are

$$\begin{aligned}\pi_u^{sto} = &\sum_{j=0}^{2n} jQ^{sto}(j) + \delta Q^{sto}(n)\pi_u^{sto} \\ &+ \beta\delta\left[\sum_{j=n+1}^{2n} Q_G^{sto}(j)\frac{2n}{1-\delta} + \sum_{j=0}^{n-1} Q_G^{sto}(j)\left(\frac{\sum_{i=1}^{2n}\binom{2n}{i}\varepsilon^i(1-\varepsilon)^{2n-i}(i+\frac{2n\delta}{1-\delta})}{1-(1-\varepsilon)^{2n}\delta}\right)\right] \\ &+ (1-\beta)\delta\left[\sum_{j=n+1}^{2n} Q_B^{sto}(j)\frac{2n}{1-\delta} - \sum_{j=n+1}^{2n} Q_B^{sto}(j)\left(\frac{\sum_{i=1}^{2n}\binom{2n}{i}\varepsilon^i(1-\varepsilon)^{2n-i}(i+\frac{2n\delta}{1-\delta})}{1-(1-\varepsilon)^{2n}\delta}\right)\right]. \quad (27)\end{aligned}$$

Profits for a seller with capacity constraint $K \leq n$ are

$$\begin{aligned}\pi_c^{sto}(K) &= \sum_{j=0}^{2n} \min\{j, K\} Q^{sto}(j) \\ &+ \beta\delta \left[\sum_{j=K}^{2n} Q_G^{sto}(j) \frac{K}{1-\delta} + \sum_{j=0}^{K-1} Q_G^{sto}(j) \left[\frac{\sum_{i=1}^{2n} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} (\min\{i, K\} + \frac{K\delta}{1-\delta})}{1-(1-\varepsilon)^{2n}\delta} \right] \right] \\ +(1-\beta)\delta &\left[\sum_{j=K}^{2n} Q_B^{sto}(j) \frac{K}{1-\delta} - \sum_{j=K}^{2n} Q_B^{sto}(j) \left[\frac{\sum_{i=2n-K+1}^{2n} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} (i - 2n + K + \frac{K\delta}{1-\delta})}{1 - \sum_{i=0}^{2n-K} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} \delta} \right] \right]. \end{aligned} \tag{28}$$

These profit functions resemble those in the deterministic setting, except incorrect cascades are now eventually corrected, regardless of whether the seller is unconstrained. Thus, unlike in our deterministic setting, the logic of Smith and Sørensen (2000) holds here, in the sense that unboundedly informative signals generate efficient learning in the long run. Incorrect negative cascades are always corrected as soon as a single informed consumer arrives, and the same applies for incorrect positive cascades if the seller is unconstrained. For a constrained seller experiencing an incorrect positive cascade, at least $2n - K + 1$ informed consumers must arrive in the same period and prevent a sell-out, in order for the state to be revealed.

Theorem 3. *For any $n > 1$, $K \leq n$ and $\delta > \sqrt{\frac{2n-K}{2n}}$, there are $(\alpha, \beta, \varepsilon) \in (1/2, 1) \times (0, 1/2) \times (0, 1)$ and $r > 0$ for which the seller can increase its profits above the unconstrained level by restricting capacity to K.*

Notice that Theorem 3 holds for some ε, while Theorem 2 holds for all $m \leq n$. In the deterministic setting, an appropriate capacity constraint could stop all further learning after a sell-out. In the stochastic setting, an incorrect positive cascade really only pays off if the probability of having many informed consumers is sufficiently low. That being said, the qualitative message is similar across the two settings: incorrect positive cascades are more difficult to reverse if the number of informed consumers is small relative to capacity.

Although incorrect positive cascades are eventually corrected in the stochastic setting, this takes substantially longer if the seller is constrained. Formally, given an incorrect positive cascade, the expected number of periods until the bad state is revealed is $1/P$, where $P = 1 - \sum_{i=0}^{2n-K} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i}$ if the seller sets capacity constraint K, and $P = 1 - (1-\epsilon)^{2n}$ if the seller is unconstrained. For example, if $2n = 6$ and $\varepsilon = 0.1$, and if initial sales trigger an incorrect positive cascade, then it requires an average of two periods to reveal the true state if the seller is unconstrained, but 787 periods if $K = 3$ and about one million periods if $K = 1$.

Figures 4(a) and 4(b) show how the presence of informed consumers affects the seller's incentive to restrict capacity. It does so indirectly, via the way that $m > 0$ or $\epsilon > 0$ change the probabilities Q^i_ω. It also does so directly, via the way that $m > 0$ or $\epsilon > 0$ explicitly enter the profit functions, because informed consumers can potentially reverse incorrect cascades. It is this latter direct effect that constitutes one of our two channels by which a capacity constraint can increase profits: by masking the choice of informed consumers not to buy, which can prevent an incorrect positive cascade from being reversed.

The medium-gray bell-shaped area in figures 4(a) and 4(b) shows when restricting capacity is optimal in the baseline. The indirect effect of having informed consumers shifts the bell shape to the left (lighter-grey area). More purchase decisions in early cohorts are then likely to reflect the true state; this has an ambiguous effect on the incentive to restrict capacity, as it is analogous to an increase in signal precision. In contrast, the direct effect of having informed consumers unambiguously makes restricting capacity more attractive. The relevant region now consists of both the lighter-grey area and the dark area above, reflecting how informed consumers can quickly reverse profitable positive cascades, but only if the seller is unconstrained.[18]

7 Conclusions

In this paper, we show that a firm may benefit from restricting capacity, so as to create scarcity for its product and increase future sales. Limiting capacity results in coarser information, as consumers who observe a sell-out attach positive probability to all levels of demand that exceed

[18] In the light-grey area, the impact of informed consumers on probabilities is sufficient to justify some capacity constraint. That is, $\pi_u^{\mathcal{Q}^i} < \pi_c^{\mathcal{Q}^i}(K)$ holds for some K. In the dark area, this inequality does not hold, but the seller still prefers to be constrained due to the fact that non-rationed informed consumers might stop incorrect cascades, so that $\pi_u^i < \pi_c^i(K)$ holds for some K, for $i \in \{det, sto\}$.

Figure 3: Areas where capacity restriction can be optimal.

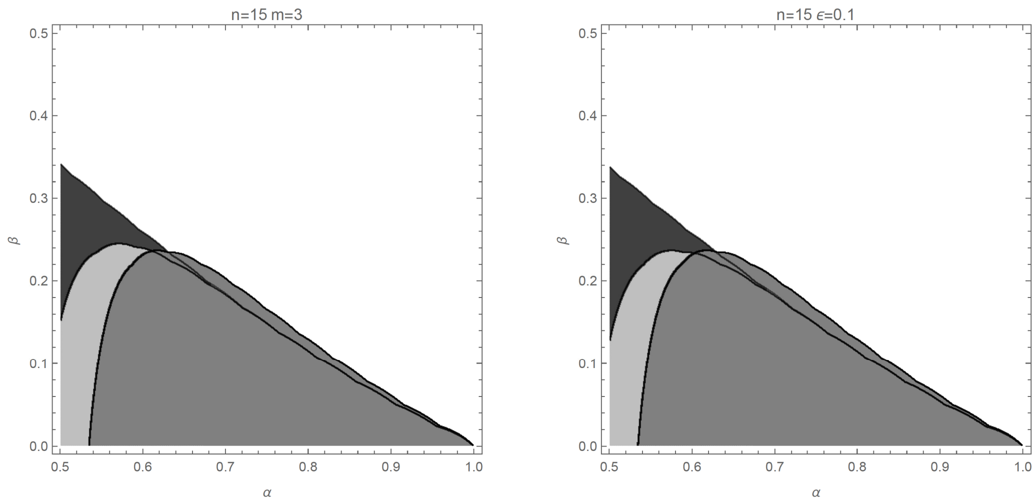

(a) Impact of $m > 0$.

(b) Impact of $\varepsilon > 0$.

capacity. The results show that two main mechanisms the literature suggests may help avoid pathological social learning outcomes, ‘guinea pigs’ and unbounded private signals, can fail to do so, if the firm is able to manipulate the learning environment by a simple instrument such as limiting capacity. We also show that this simple instrument can serve a practical tool for persuading consumers, in the sense of implementing optimal information design in large markets.

Our results rely on the idea that consumers can observe sales and capacity. This is reasonable in many markets, e.g. restaurants, sports and concert tickets, and limited edition products, where sales and capacity are often widely known, but the extent of any excess demand is not. Product scarcity should also affect learning in other settings, but in a way that depends precisely on what consumers can observe. For example, it will matter if concert promoters can ‘paper the house’ by quietly filling seats for free, or if sales figures for certain consumer products become widely reported precisely because shortages occurred. Our mechanism will also continue to apply in situations where capacity is exogenous, in particular if capacity happens to be close to the optimal level. In such situations, our main results will have a slightly different interpretation;

namely, that a firm that must limit production, or use a small venue, may do just as well (or better) than a firm that is not similarly constrained.

Restricting capacity would still help to trigger positive cascades if consumers were partially sophisticated, and believed that others always followed their own private signals (see, e.g., Eyster and Rabin (2010), Gagnon-Bartsch and Rabin (2017), and Dasaratha et al. (2018)). Just as in our analysis, an initial sell-out would result in a cascade if it revealed high enough demand from consumers in the first cohort. A difference is that consumers in later cohorts would become increasingly confident that the state was good after observing more and more sell-outs, and become increasingly willing to buy.

In line with most work on social learning and Bayesian persuasion, our seller is not privately informed, but an alterative would be to assume that the seller receives its own private binary signal about quality. The full analysis of such a model is beyond the scope of this paper. However, it is clear that if the seller's signal precision is low, then there can exist an equilibrium where the seller restricts capacity. Our model can be viewed as a limit case of such a pooling equilibrium where the seller's signal is completely uninformative. A separating equilibrium would require that period-1 consumers always follow their own signal, along with a standard incentive compatibility condition where the seller only wants to restrict capacity after a bad signal. In a similar spirit to our analysis: it is precisely a seller that expects relatively low sales that would profit from restricting capacity.

Our analysis of Bayesian persuasion compares the optimal design of a statistical experiment in the absence of social learning, on the one hand, to the optimal choice of capacity in the presence of social learning, on the other. A related issue is how such experiments might themselves influence the social learning process. For example, suppose that consumers observe each others' purchase decisions and that the seller directly chooses the accuracy of consumers' private signals. Perfectly revealing bad signals, which are optimal in the absence of social learning, would have a clear downside in such a setting: a single consumer's decision not to buy would immediately trigger a negative purchase cascade.

References

Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.

Masaki Aoyagi. Optimal sales schemes against interdependent buyers. *American Economic Journal: Microeconomics*, 2(1):150–82, 2010.

Pak Hung Au. Dynamic information disclosure. *The RAND Journal of Economics*, 46(4): 791–823, 2015.

Abhijit Banerjee and Drew Fudenberg. Word-of-mouth learning. *Games and Economic Behavior*, 46(1):1–22, 2004.

Abhijit V Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, pages 797–817, 1992.

Dirk Bergemann and Stephen Morris. Information design: A unified perspective. 2017.

James Best and Daniel Quigley. Persuasion for the long run. 2017.

Manaswini Bhalla. Waterfall versus sprinkler product launch strategy: Influencing the herd. *The Journal of Industrial Economics*, 61(1):138–165, November 2013.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, pages 992–1026, 1992.

Jacopo Bizzotto, Jesper Rüdiger, and Adrien Vigier. Dynamic persuasion with outside information. 2018.

Subir Bose, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. Dynamic monopoly pricing and herding. *The RAND Journal of Economics*, 37(4):910–928, 2006.

Subir Bose, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. Monopoly pricing in the binary herding model. *Economic Theory*, 37(2):203–241, 2008.

Steven Callander and Johannes Hörner. The wisdom of the minority. *Journal of Economic Theory*, 144(4):1421–1439, 2009.

Boğaçhan Çelen and Shachar Kariv. Observational learning under imperfect information. *Games and Economic Behavior*, 47(1):72–86, 2004.

Yeon-Koo Che and Johannes Hörner. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925, 2017.

Pascal Courty and Mario Pagliero. The impact of price discrimination on revenue: Evidence from the concert industry. *Review of Economics and Statistics*, 94(1):359–369, 2012.

Krishna Dasaratha, Benjamin Golub, and Nir Hak. Social learning in a dynamic environment. 2018.

Laurens G Debo, Christine Parlour, and Uday Rajan. Signaling quality via queues. *Management Science*, 58(5):876–891, 2012.

Patrick DeGraba. Buying frenzies and seller-induced excess demand. *The RAND Journal of Economics*, pages 331–342, 1995.

Piotr Dworczak and Giorgio Martini. The simple economics of optimal persuasion. 2017.

Jeffrey C Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

Erik Eyster and Matthew Rabin. Naive herding in rich-information settings. *American Economic Journal: Microeconomics*, 2(4):221–43, 2010.

Tristan Gagnon-Bartsch and Matthew Rabin. Naive social learning, mislearning, and unlearning. 2017.

Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322, 2016.

Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.

David Gill and Daniel Sgroi. Sequential decisions with tests. *Games and economic Behavior*, 63(2):663–678, 2008.

David Gill and Daniel Sgroi. The optimal choice of pre-launch reviewer. *Journal of Economic Theory*, 147(3):1247–1260, 2012.

Benjamin Golub and Evan D Sadler. Learning in social networks. 2017.

Antonio Guarino, Heike Harmgart, and Steffen Huck. Aggregate information cascades. *Games and Economic Behavior*, 73(1):167–185, 2011.

Helios Herrera and Johannes Hörner. Biased social learning. *Games and Economic Behavior*, 80:131–146, 2013.

Shota Ichihashi. Limiting sender's information in bayesian persuasion. *Available at SSRN 3233072*, 2018.

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

Anton Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018.

Anton Kolotilin and Andriy Zapechelnyuk. Persuasion meets delegation. 2018.

Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964, 2017.

Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012, 2014.

Ting Liu and Pasquale Schiraldi. New product launch: herd seeking or herd preventing? *Economic Theory*, 51(3):627–648, 2012.

Ilan Lobel and Evan Sadler. Information diffusion in networks through social learning. *Theoretical Economics*, 10(3):807–851, 2015.

Jeffrey Mensch. Monotone persuasion. 2018.

Marc Möller and Makoto Watanabe. Advance purchase discounts versus clearance sales. *The Economic Journal*, 120(547):1125–1148, 2010.

Ignacio Monzón. Aggregate uncertainty can lead to incorrect herds. *American Economic Journal: Microeconomics*, 9(2):295–314, 2017.

Ignacio Monzón and Michael Rapp. Observational learning with position uncertainty. *Journal of Economic Theory*, 154:375–402, 2014.

Volker Nocke and Martin Peitz. A theory of clearance sales. *The Economic Journal*, 117(522): 964–990, 2007.

Dmitry Orlov, Andrzej Skrzypacz, and Pavel Zryumov. Persuading the principal to wait. *Journal of Political Economy*, forthcoming, 2018.

HG Parsa, John T Self, David Njite, and Tiffany King. Why restaurants fail. *Cornell Hotel and Restaurant Administration Quarterly*, 46(3):304–322, 2005.

Eduardo Perez-Richet and Vasiliki Skreta. Test design under falsification. 2018.

Jérôme Renault, Eilon Solan, and Nicolas Vieille. Optimal dynamic information provision. *Games and Economic Behavior*, 104:329–349, 2017.

Amin Sayedi. Pricing in a duopoly with observational learning. 2018.

Daniel Sgroi. Optimizing information in the herd: Guinea pigs, profits, and welfare. *Games and Economic Behavior*, 39(1):137–166, 2002.

Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.

Lones Smith and Peter Norman Sorensen. Rational social learning by random sampling. 2013.

Axel Stock and Subramanian Balachander. The making of a "hot product": A signaling explanation of marketers' scarcity strategy. *Management Science*, 51(8):1181–1192, 2005.

Elias Tsakas and Nikolas Tsakas. Noisy persuasion. 2017.

Nick Vikander. Sell-outs, beliefs, and bandwagon behavior. *BE Journal of Theoretical Economics*, forthcoming, 2018.

Ivo Welch. Sequential sales, learning, and cascades. *The Journal of Finance*, 47(2):695–732, 1992.

Appendix A: Proofs

Proof of Lemma 1. **Deterministic model.** Consider $Q_G^{det}(j)$ given by (1) and $Q_B^{det}(j)$ given by (2). Clearly, $\frac{Q_B^{det}(j)}{Q_G^{det}(j)} = 0$ if $2n - j \leq m - 1$. For $2n - j \geq m$ we have

$$\frac{Q_B^{det}(j)}{Q_G^{det}(j)} = \left[\frac{(j-m)!}{(2n-m-j)!}\right]\left[\frac{(2n-j)!}{j!}\right]\alpha^{2(n-j)}(1-\alpha)^{2(j-n)}$$

which is decreasing in j and equals to 1 when $j = n$. Moreover, for $j = n - 1$ we get

$$\frac{Q_B^{det}(n-1)}{Q_G^{det}(n-1)} = \frac{n(n+1)}{(n-m-1)(n-m)}\frac{\alpha^2}{(1-\alpha)^2} > \frac{\alpha^2}{(1-\alpha)^2}$$

and for $j = n + 1$ we get

$$\frac{Q_B^{det}(n+1)}{Q_G^{det}(n+1)} = \frac{(n-m-1)(n-m)}{n(n+1)}\frac{(1-\alpha)^2}{\alpha^2} < \frac{(1-\alpha)^2}{\alpha^2}$$

Now,

$$Q_G^{det}(2n-j) = \binom{2n-m}{2n-j-m}\alpha^{2n-j-m}(1-\alpha)^j = \binom{2n-m}{j}\alpha^{2n-j-m}(1-\alpha)^j = Q_B^{det}(j).$$

Finally,

$$\frac{Q_G^{det}(j)}{Q_G^{det}(2n-j)} = \frac{Q_G^{det}(j)}{Q_B^{det}(j)} > 1$$

for all $j > n + 1$ as the ratio $Q_G^{det}(j)/Q_B^{det}(j)$ is increasing for $j > m$ and $\frac{Q_G^{det}(n)}{Q_B^{det}(n)} = 1$.

Stochastic model. Consider $Q_G^{sto}(j)$ given by (3) and $Q_B^{sto}(j)$ given by (4).

Let $\xi \equiv \alpha + \varepsilon - \alpha\varepsilon$, and note that $\xi \in (1/2, 1)$. Therefore, expressions for $Q_\omega^{sto}(j)$ are exactly as in the deterministic model with $m = 0$ and α replaced with ξ. Moreover, since $\xi > \alpha$ we get $\frac{\xi}{1-\xi} > \frac{\alpha}{1-\alpha}$. Thus, the stochastic model satisfies all four properties.

Baseline model. Follows from the deterministic case with $m = 0$.

□

Proof of Lemma 2. We denote the belief that the state is ω, conditional on past sales $(S_1,\dots S_t)$ and a private signal s, as $P(\omega|S_1,\dots S_t,s)$. Consider $t=2$. Note that

$$P(G|n+1,b)=\frac{1}{1+\frac{1-\beta}{\beta}\frac{\alpha}{1-\alpha}\frac{Q^i_B(n+1)}{Q^i_G(n+1)}}\geq\frac{1}{1+\frac{1-\beta}{\beta}\frac{1-\alpha}{\alpha}}=P(G|g)>r$$

where the first inequality follows from $\frac{Q^i_B(n+1)}{Q^i_G(n+1)}\leq\left(\frac{1-\alpha}{\alpha}\right)^2$, $i\in\{base,det,sto\}$. Thus, the belief of a consumer that quality is good after observing $S_1\geq n+1$ and $s=b$ is better than $P(G|g)$, so the consumer should buy regardless of her private information. In a similar way we get

$$P(G|n-1,g)=\frac{1}{1+\frac{1-\beta}{\beta}\frac{1-\alpha}{\alpha}\frac{Q^i_B(n-1)}{Q^i_G(n-1)}}\leq\frac{1}{1+\frac{1-\beta}{\beta}\frac{\alpha}{1-\alpha}}=P(G|b)<r$$

due to $\frac{Q^i_B(n-1)}{Q^i_G(n-1)}\geq\left(\frac{\alpha}{1-\alpha}\right)^2$. Finally, $P(G|n,b)=P(G|b)$ and $P(G|n,g)=P(G|g)$ due to $\frac{Q^i_B(n)}{Q^i_G(n)}=1$, so if $S_2=n$ then a consumer should follow her own signal. Now, consider $t>2$. Suppose that for all $t'<t-1$, $S_{t'}=n$ holds. Due to $\frac{Q^i_B(n)}{Q^i_G(n)}=1$, this implies that in all cohorts consumers have followed their own signal. Thus, if $S_{t-1}=n$ consumers in cohort t also must follow their signals. Suppose that $S_{t-1}>n$. In this case

$$P(G|S_1,\dots,S_{t-1};b)=\frac{\beta(1-\alpha)Q^i_G(S_{t-1})[Q^i_G(n)]^{t-2}}{\beta(1-\alpha)Q^i_G(S_{t-1})[Q^i_G(n)]^{t-2}+(1-\beta)\alpha Q^i_B(S_{t-1})[Q^i_B(n)]^{t-2}}=\frac{1}{1+\frac{1-\beta}{\beta}\frac{\alpha}{1-\alpha}\frac{Q^i_B(S_{t-1})}{Q^i_G(S_{t-1})}}>r$$

so consumers should buy. Similarly, if $S_{t-1}<n$ we get $P(G|S_1,\dots,S_{t-1};g)<r$ and consumers should not buy.

Now, suppose there exists a first $t'<t-1$ such that $S_{t'}\neq n$. If $S_{t'}<n$ consumers in the next cohort should not buy and a negative cascade starts. If for all $\tau\in[t'+1,t-1]$ $S_\tau=0$, then consumers in cohort t do not gain any additional information, and should also refuse to buy. If for some $\tau\in[t'+1,t-1]$ $S_\tau>0$ (which cannot happen in the baseline case) the purchase must come from an informed consumer and consumers in later cohorts should buy. In a similar vein if $S_{t'} > n$ then the positive cascade starts, it persists if for all $\tau \in [t'+1, t-1]$ $S_\tau = 2n$. It is otherwise reversed, as the decision not to buy comes from an informed consumer, so consumers in later cohorts should not buy. $\square$

Proof of Lemma 3. From (6) it is sufficient to show that $\sum_{j=K}^{2n} Q_B^i(j) < \sum_{j=K}^{2n} Q_G^i(j)$, $i \in \{base, det, sto\}$. For all $j \geq n+1$, we have $Q_G^i(j) > Q_G^i(2n-j)$, which implies $\sum_{j=0}^{K-1} Q_G^i(j) < \sum_{j=2n-K+1}^{2n} Q_G^i(j)$. By adding $\sum_{j=K}^{2n-K} Q_G^i(j)$, we obtain that $\sum_{j=0}^{2n-K} Q_G^i(j) < \sum_{j=K}^{2n} Q_G^i(j)$. Changing the summation order on the left-hand-side gives $\sum_{j=K}^{2n} Q_G^i(2n-j) < \sum_{j=K}^{2n} Q_G^i(j)$. Finally, due to $Q_B^i(j) = Q_G^i(2n-j)$, we obtain $\sum_{j=K}^{2n} Q_B^i(j) < \sum_{j=K}^{2n} Q_G^i(j)$. $\square$

Proof of Lemma 4. As shown in the proof of Lemma 3, $\sum_{j=K}^{2n} Q_B^i(j) < \sum_{j=K}^{2n} Q_G^i(j)$, $i \in \{base, det, sto\}$. Thus,

$$\gamma(1,K) = \frac{1}{1 + \frac{1-\beta}{\beta}\frac{\alpha}{1-\alpha}\frac{\sum_{j=K}^{2n} Q_B^i(j)}{\sum_{j=K}^{2n} Q_G^i(j)}} > \frac{1}{1 + \frac{1-\beta}{\beta}\frac{\alpha}{1-\alpha}} = P(G|b),$$

all of which are independent of r. We can therefore satisfy $\gamma(1,K) > r$ by choosing $r > P(G|b)$ sufficiently close to $P(G|b)$. $\square$

Proof of Lemma 5. For simpler notation we are going to omit the s superscript in our preliminary steps and work with generic $\mathcal{Q}$. Rewrite (7) as

$$[1 - \delta Q(n)]\pi_u^{\mathcal{Q}} = S_1 + \frac{\delta}{1-\delta} 2n S_2, \tag{29}$$

where $S_1 = \sum_{j=1}^{2n} jQ(j)$ and $S_2 = \sum_{j=n+1}^{2n} Q(j)$. Note that as long as $Q(n) < 1$, the term

$[1-\delta Q(n)]$ is bounded away from 0 for $\delta < 1$. In a similar way,

$$\pi_c^Q(K) = S_1 - \sum_{j=K}^{2n}(j-K)Q(j) + \frac{\delta}{1-\delta}K[\sum_{j=K}^{n-1}Q(j) + Q(n) + S_2] >$$
$$S_1 - (2n-K)[\sum_{j=K}^{n-1}Q(j) + Q(n) + S_2] + \frac{\delta}{1-\delta}K[\sum_{j=K}^{n-1}Q(j) + Q(n) + S_2],$$

where the inequality follows from replacing all terms $(j-K)$ with the larger term $2n-K$. Moreover, for $\delta > \frac{2n-K}{2n}$ the penultimate term is smaller than the last one, and thus

$$\pi_c^Q(K) > S_1 + \frac{\delta K - (2n-K)(1-\delta)}{1-\delta}[Q(n) + S_2]$$

This implies that $\pi_c^Q(K) > \pi_u^Q$ if

$$\left\{S_1 + \frac{\delta K - (2n-K)(1-\delta)}{1-\delta}[Q(n) + S_2]\right\}[1-\delta Q(n)] \geq S_1 + \frac{\delta}{1-\delta}2nS_2,$$

which can be rewritten as

$$-\delta Q(n)S_1 + \frac{\delta K - (2n-K)(1-\delta)}{1-\delta}Q(n)[1-\delta Q(n)] \geq$$
$$\frac{S_2}{1-\delta}\left\{2n\delta - [\delta K - (2n-K)(1-\delta)][1-\delta Q(n)]\right\}.$$

As $S_1 \leq 2n$, the above inequality holds as long as

$$\frac{Q(n)}{1-\delta}\left\{[\delta K - (2n-K)(1-\delta)][1-\delta Q(n)] - 2n\delta(1-\delta)\right\} \geq$$
$$\frac{S_2}{1-\delta}\left\{2n\delta - [\delta K - (2n-K)(1-\delta)][1-\delta Q(n)]\right\}$$

Now note that for all $\delta \in (0,1)$, the expression

$$f_R \equiv 2n\delta - [\delta K - (2n-K)(1-\delta)][1-\delta Q(n)] = 2\delta^2 nQ(n) + (2n-K)[1-\delta Q(n)]$$

is strictly positive. Moreover, the expression

$$f_L \equiv [\delta K - (2n-K)(1-\delta)][1-\delta Q(n)] - 2n\delta(1-\delta)$$

is also strictly positive if $\delta > \frac{\sqrt{[(2n-K)Q(n)]^2+8n(2n-K)[1-Q(n)]}-(2n-K)Q(n)}{4n-4nQ(n)} \equiv \delta^*(Q(n);n,K)$, where the critical value $\delta^*(Q(n);n,K)$ is increasing in $Q(n)$, equals $\sqrt{\frac{2n-K}{2n}}$ at $Q(n)=0$, and approaches 1 as $Q(n) \to 1$. Note that $\sqrt{\frac{2n-K}{2n}} > \frac{2n-K}{2n}$, and therefore the condition we used for the approximation of $\pi_c^Q(K)$ is automatically satisfied as long as $\delta > \sqrt{\frac{2n-K}{2n}}$. Thus, for any $Q(n) < 1$, there exists $\delta \in \left(\sqrt{(2n-K)/2n}, 1\right)$ such that the right-hand-side of

$$\frac{Q(n)}{S_2} = \frac{Q(n)}{\sum_{j=n+1}^{2n} Q(j)} \geq \frac{f_R}{f_L} \tag{30}$$

is positive and finite. Thus, if $\lim_{s\to\infty} \frac{Q^s(n)}{\sum_{j=n+1}^{2n} Q^s(j)} = \infty$, there exists T_0 such that for all $s > T_0$ the left-hand-side of (30) is larger than the right-hand-side and therefore $\pi_u^{Q^s} < \pi_c^{Q^s}(K)$. Moreover, δ can be chosen arbitrarily close to $\sqrt{\frac{2n-K}{2n}}$, as long as $Q(n)$ is sufficiently close to zero, which proves the second statement of the lemma. □

Proof of Theorem 1. Note that the profit expressions (9) and (10) coincide with equations (7) and (8) for $Q(j) = Q^{base}(j)$. Thus, due to Lemma 5, we can get $\pi_c^{base}(K) \geq \pi_u^{base}$ for any K if there exists a sequence $\{\alpha_s, \beta_s\}_{s=0}^{\infty}$ such that $\{Q^{base,s}(j)\}_{s=0}^{\infty}$ satisfies $\lim_{s\to\infty} \frac{Q^{base,s}(n)}{\sum_{j=n+1}^{2n} Q^{base,s}(j)} = \infty$ and $\lim_{s\to\infty} Q^{base,s}(n) = 0$. To establish the first part we need to show that there exists T_0 such that for all M, we get $Q^{base,s}(n)/\left[\sum_{j=n+1}^{2n} Q^{base,s}(j)\right] > M$ for all $s > T_0$. Equivalently for all M there are (α_s, β_s) such that

$$\frac{A_{\alpha_s} - C_{\alpha_s} M}{M(B_{\alpha_s} - C_{\alpha_s})} > \beta_s$$

where

$$A_{\alpha_s} = \binom{2n}{n} \alpha_s^n (1-\alpha_s)^n, \; B_{\alpha_s} = \sum_{i=n+1}^{2n} \binom{2n}{i} \alpha_s^i (1-\alpha_s)^{2n-i}, \; C_{\alpha_s} = \sum_{i=n+1}^{2n} \binom{2n}{i} \alpha_s^{2n-i} (1-\alpha_s)^i.$$

Now take some sequence $\alpha_s \to 1$. Then, $\lim_{s\to} \frac{A_{\alpha_s}}{C_{\alpha_s}} = \infty$ and therefore there exists T_0 such that for all $s > T_0$ we get $A_{\alpha_s} > MC_{\alpha_s}$. Choose $\beta_s = \frac{1}{2}\frac{A_{\alpha_s} - C_{\alpha_s} M}{M(B_{\alpha_s} - C_{\alpha_s})}$ (note that $B_{\alpha_s} > C_{\alpha_s}$ as $\alpha_s > 1/2$). Then the sequence $\{\alpha_s, \beta_s\}_{s=0}^{\infty}$ yields $\lim_{s\to\infty} \frac{Q^{base,s}(n)}{\sum_{j=n+1}^{2n} Q^{base,s}(j)} = \infty$. Moreover, $\lim_{s\to\infty} Q^{base,s}(n) = 0$, which completes the proof. □

Proof of Proposition 1. We first show that $\beta \geq \frac{1}{2}$ implies $\sum_{j=1}^{2n} jQ^{base}(j) \geq n$ and $\sum_{j=n+1}^{2n} 2nQ^{base}(j) \geq n(1 - Q^{base}(n))$. By $Q^{base}(j) = \beta Q_G^{base}(j) + (1-\beta)Q_B^{base}(j)$ and Lemma 1 (iii), write

$$
\begin{aligned}
Q^{base}(j) - Q^{base}(2n-j) = \\
&\beta Q_G^{base}(j) + (1-\beta)Q_B^{base}(j) - \beta Q_G^{base}(2n-j) - (1-\beta)Q_B^{base}(2n-j) \\
&= \beta Q_G^{base}(j) + (1-\beta)Q_G^{base}(2n-j) - \beta Q_G^{base}(2n-j) - (1-\beta)Q_G^{base}(j) \\
&= (2\beta - 1)[Q_G^{base}(j) - Q_G^{base}(2n-j)], \quad (31)
\end{aligned}
$$

which is positive for all $j \geq n+1$, by $\beta \geq 1/2$ and Lemma 1 (iv). Thus,

$$
\begin{aligned}
\sum_{j=1}^{2n} jQ^{base}(j) &= \sum_{j=0}^{n-1} [jQ^{base}(j) + (2n-j)Q^{base}(2n-j)] + nQ^{base}(n) \\
&\geq \sum_{j=0}^{n-1} \left[\frac{(j + (2n-j))}{2} (Q^{base}(j) + Q^{base}(2n-j)) \right] + nQ^{base}(n) = n, \quad (32)
\end{aligned}
$$

and

$$
\sum_{j=n+1}^{2n} 2nQ^{base}(j) \geq n \left[\sum_{j=0}^{n-1} Q^{base}(j) + \sum_{j=n+1}^{2n} Q^{base}(j) \right] = n[1 - Q^{base}(n)].
$$

We now combine with (7) to obtain

$$
\pi_u^{base} = \frac{1}{1 - \delta Q^{base}(n)} \left(\sum_{j=0}^{2n} jQ^{base}(j) + \frac{2n\delta}{1-\delta} \sum_{j=n+1}^{2n} Q^{base}(j) \right) \geq \frac{1}{1 - \delta Q^{base}(n)} \left(n + \frac{\delta}{1-\delta} n[1 - Q^{base}(n)] \right),
$$

which simplifies to $\pi_u^{base} \geq \frac{n}{1-\delta}$ and therefore $\pi_u^{base} > \frac{n}{1-\delta}$.

Now note that as $K \leq n$, we have $\pi_c^{base}(K) \leq \frac{K}{1-\delta} \leq \frac{n}{1-\delta} < \pi_u^{base}$, which completes the proof. □

Proof of Proposition 2. The optimality of K_0 and K_1 implies that

$$\sum_{j=K_0+1}^{2n}(K_0-j)Q^{base}(j)+\frac{\delta}{1-\delta}K_0\sum_{j=K_0}^{2n}Q^{base}(j) \geq \sum_{j=K_1+1}^{2n}(K_1-j)Q^{base}(j)+\frac{\delta}{1-\delta}K_1\sum_{i=K_1}^{2n}Q^{base}(j)$$

and

$$\sum_{j=K_1+1}^{2n}(K_1-j)Q^{base}(j)+\frac{\delta}{1-\delta}2n\sum_{j=K_1}^{2n}Q^{base}(j) \geq \sum_{j=K_0+1}^{2n}(K_0-i)Q^{base}(j)+\frac{\delta}{1-\delta}2n\sum_{j=K_0}^{2n}Q^{base}(j)$$

Summing it up yields

$$\frac{\delta}{1-\delta}K_0\sum_{j=K_0}^{2n}Q^{base}(j)+\frac{\delta}{1-\delta}2n\sum_{j=K_1}^{2n}Q^{base}(j) \geq \frac{\delta}{1-\delta}K_1\sum_{j=K_1}^{2n}Q^{base}(j)+\frac{\delta}{1-\delta}2n\sum_{j=K_0}^{2n}Q^{base}(j)$$

or

$$\sum_{i=K_1}^{2n}(2n-K_1)Q^{base}(j) \geq \sum_{j=K_0}^{2n}(2n-K_0)Q^{base}(j)$$

which implies that $K_1 \leq K_0$ as $\sum_{j=K}^{2n}(2n-K)Q^{base}(j)$ is a decreasing function of K. □

Proof of Proposition 3. We need to compare $\tilde{\pi}$ and π^*. Define

$$\Delta \equiv \pi^* - \tilde{\pi} = \frac{(1-\alpha)\left(\alpha^2+\alpha-2\alpha r+r-1\right)\left(\beta-\alpha\beta-r(\alpha(1-2\beta)+\beta)\right)}{\alpha(2\alpha-1)(1-r)r}$$

Equation $\Delta = 0$ has four roots:

$$\alpha_1 = \frac{(1-r)\beta}{r+\beta-2\beta r}, \quad \alpha_{2,3} = \frac{1}{2}\left(2r-1 \pm \sqrt{4r^2-8r+5}\right), \quad \alpha_4 = 1.$$

Note that $\alpha_2 = \frac{1}{2}\left(2r-1-\sqrt{4r^2-8r+5}\right) < 0$, $\alpha_3 = \overline{\alpha}(r) \leq 1 = \alpha_4$ and $\Delta < 0$ in a left neighbourhood of $\alpha = 1$. Moreover, $\alpha_3 \geq 1/2$ for all r.

Suppose that $\beta \geq r$. Then we have $\underline{\alpha}(\beta, r) = \alpha_1 \geq 1/2$. It is straightforward to verify that $\alpha_1 > \overline{\alpha}(r)$ if and only if $\beta > \frac{r}{1-r}[\overline{\alpha}(r)+1-2r] \equiv \overline{\beta}(r)$, where $\overline{\beta}(r) > r$. Since in the absence of any recommendation consumers follow their private signals, we require that $\alpha \geq \alpha_1$. Thus, we have $\Delta > 0$ if $r \leq \beta < \overline{\beta}(r)$ and $\alpha_1 < \alpha < \overline{\alpha}(r)$; we have $\Delta < 0$ if $r \leq \beta < \overline{\beta}(r)$ and $\alpha > \overline{\alpha}(r)$, or if $\beta > \overline{\beta}(r)$ and $\alpha > \alpha_1$.

Now suppose that $\beta < r$. Then we have $\underline{\alpha}(\beta, r) = \frac{r(1-\beta)}{r+\beta-2r\beta} > \frac{1}{2} > \alpha_1$. Moreover, $\underline{\alpha}(\beta, r) < \overline{\alpha}(r)$ if and only if $\beta > r\left[\frac{2r-\overline{\alpha}(r)}{3r-1}\right] \equiv \underline{\beta}(r)$, where $\underline{\beta}(r) < r$. Thus, we have $\Delta > 0$ if $\underline{\beta}(r) < \beta < r$ and $\underline{\alpha}(\beta, r) \leq \alpha < \overline{\alpha}(r)$; we have $\Delta < 0$ if $\underline{\beta}(r) < \beta < r$ and $\alpha > \overline{\alpha}(r)$, or if $\beta < \underline{\beta}(r)$ and $\alpha \geq \underline{\alpha}(\beta, r)$.

As for $\beta \geq r$, we have $\alpha_1 = \underline{\alpha}(\beta, r)$, so combining cases $\beta \geq r$ and $\beta < r$ gives the result in the proposition. □

Proof of Theorem 2. The expression for the unconstrained profits (25) can be rewritten as

$$\pi_u^{det} = \sum_{j=0}^{2n} jQ^{det}(j) + \delta Q^{det}(n)\pi_u^{det} + \delta \sum_{j=n+1}^{2n} Q^{det}(j)\frac{2n}{1-\delta} + \\ (2\beta-1)\delta \sum_{j=0}^{n-1} Q_G^{det}(j)\left(m + \frac{2n\delta}{1-\delta}\right) \quad (33)$$

and therefore for $\beta < 1/2$ we have

$$\pi_u^{det} \leq \pi_u^{\mathcal{Q}^{det}}$$

with $\pi_u^{\mathcal{Q}^{det}}$ defined by (7) and probabilities $\mathcal{Q}^{det}$ defined by (1) and (2). Moreover, equation (26) can be rewritten as

$$\pi_c^{det} = \sum_{j=0}^{2n} \min\{j, K\} Q^{det}(j) + \delta \sum_{j=K}^{2n} Q^{det}(j) \frac{K}{1-\delta} + \beta\delta \left[\sum_{j=0}^{K-1} Q_G^{det}(j) \left(\min\{m, K\} + \frac{K\delta}{1-\delta} \right) \right] \quad (34)$$

and therefore

$$\pi_c^{det}(K) \geq \pi_c^{\mathcal{Q}^{det}}(K)$$

where $\pi_c^{\mathcal{Q}^{det}}(K)$ is defined by (8). Therefore, $\pi_c^{det}(K) \geq \pi_u^{det}$ if for all m there is a sequence $\{\alpha_s, \beta_s\}$ such that $\{Q_\omega^{det,s}(j)\}_{s=0}^{\infty}$ satisfies the requirements of Lemma 5. That is, we have to show that there is a sequence such that, for any number M, there is T_0 such that

$$Q^{det,s}(n) \Big/ \left(\sum_{j=n+1}^{2n} Q^{det,s}(j) \right) = Q^{det,s}(n) \Big/ \left(\beta_s \sum_{j=n+1}^{2n} Q_G^{det,s}(j) + (1-\beta_s) \sum_{j=n+1}^{2n} Q_B^{det,s}(j) \right) > M$$

or

$$\beta_s < \frac{Q^{det,s}(n) - M \sum_{j=n+1}^{2n} Q_B^{det,s}(j)}{M \left(\sum_{j=n+1}^{2n} Q_G^{det,s}(j) - \sum_{j=n+1}^{2n} Q_B^{det,s}(j) \right)}$$

Because of the properties of $Q_\omega^{det}(j)$ established in Lemma 1, we have that for all (α_s, β_s), $\sum_{j=n+1}^{2n} Q_G^{det,s}(j) - \sum_{j=n+1}^{2n} Q_B^{det,s}(j) > 0$ holds. Now, take a sequence $\alpha_s \to 1$. Then,

$$\lim_{s\to\infty} \frac{Q^{det,s}(n)}{\sum_{j=n+1}^{2n} Q_B^{det,s}(j)} = \lim_{s\to\infty} \frac{\binom{2n-m}{n}}{\sum_{j=n+1}^{2n} \binom{2n-m}{2} \alpha_s^{n-j} (1-\alpha_s)^{j-n}} = \infty$$

Then choose $\beta_s = \frac{1}{2} \frac{Q^{det,s}(n) - M \sum_{j=n+1}^{2n} Q_B^{det,s}(j)}{M \left(\sum_{j=n+1}^{2n} Q_G^{det,s}(j) - \sum_{j=n+1}^{2n} Q_B^{det,s}(j) \right)}$, which is positive for s large enough. Therefore, for such a sequence (α_s, β_s), we have $\lim_{s\to\infty} \frac{Q^{det,s}(n)}{\sum_{j=n+1}^{2n} Q^{det,s}(j)} = \infty$. Moreover, $\lim_{s\to\infty} Q^{det,s}(n) =$

0. Thus, all conditions of Lemma 5 are satisfied, which completes the proof. □

Proof of Theorem 3. The profit function of the unconstrained seller (27) can be rewritten as:

$$\pi_u^{sto} = \sum_{j=0}^{2n} jQ^{sto}(j) + \delta Q^{sto}(n)\pi_u^{sto} + \delta \sum_{j=n+1}^{2n} Q^{sto}(j)\frac{2n}{1-\delta} + (2\beta-1)\delta \sum_{j=0}^{n-1} Q_G^{sto}(j) \left(\frac{\sum_{i=1}^{2n} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} (i + \frac{2n\delta}{1-\delta})}{1-(1-\varepsilon)^{2n}\delta} \right) \quad (35)$$

Thus, for $\beta < 1/2$ we have

$$\pi_u^{sto} < \pi_u^{Q^{sto}}$$

Now our aim is to show that for all α, β, K, n, there is a (small enough) value of ε such that the constrained seller's profits (28) are smaller than the profits defined by (8). If the seller sets a capacity constraint K, its profits (28) can be rewritten as

$$\begin{aligned}\pi_c^{sto}(K) = \sum_{j=0}^{2n} \min\{j, K\} Q^{sto}(j) + \sum_{j=K}^{2n} Q^{sto}(j)\frac{K}{1-\delta} \\ + \beta\delta \sum_{j=0}^{K-1} Q_G^{sto}(j) \left[\frac{\sum_{i=1}^{2n} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} (\min\{i, K\} + \frac{K\delta}{1-\delta})}{1-(1-\varepsilon)^{2n}\delta} \right] \\ - (1-\beta)\delta \sum_{j=K}^{2n} Q_B^{sto}(j) \left[\frac{\sum_{i=2n-K+1}^{2n} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} (i - 2n + K + \frac{K\delta}{1-\delta})}{1-\sum_{i=0}^{2n-K} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i}\delta} \right] \end{aligned} \quad (36)$$

Thus, $\pi_c^{sto}(K) < \pi_c^{Q^{sto}}(K)$ if and only if

$$\left(\frac{\beta}{1-\beta}\right) \left(\frac{\sum_{j=0}^{K-1} Q_G^{sto}(j)}{\sum_{j=K}^{2n} Q_B^{sto}(j)} \right) > \frac{\left[\frac{\sum_{j=2n-K+1}^{2n} \binom{2n}{j} \varepsilon^j (1-\varepsilon)^{2n-j} (j-2n+K+\frac{K\delta}{1-\delta})}{1-\sum_{j=0}^{2n-K} \binom{2n}{j} \varepsilon^j (1-\varepsilon)^{2n-j}\delta} \right]}{\left[\frac{\sum_{j=1}^{2n} \binom{2n}{j} \varepsilon^j (1-\varepsilon)^{2n-j} (\min\{j,K\}+\frac{K\delta}{1-\delta})}{1-(1-\varepsilon)^{2n}\delta} \right]} \quad (37)$$

Notice that $\min\{j, K\} \geq 0$ and $j - 2n + K \leq K$. Thus, to establish (37), it will be enough to show that

$$\left(\frac{\delta\beta}{1-\beta}\right)\left(\frac{\sum_{j=0}^{K-1} Q_G^{sto}(j)}{\sum_{j=K}^{2n} Q_B^{sto}(j)}\right) > \left[\frac{\sum_{j=2n-K+1}^{2n}\binom{2n}{j}\varepsilon^j(1-\varepsilon)^{2n-j}}{1-\sum_{j=0}^{2n-K}\binom{2n}{j}\varepsilon^j(1-\varepsilon)^{2n-j}\delta}\right] \Big/ \left[\frac{\sum_{j=1}^{2n}\binom{2n}{j}\varepsilon^j(1-\varepsilon)^{2n-j})}{1-(1-\varepsilon)^{2n}\delta}\right] \quad (38)$$

Note that $\sum_{j=0}^{K-1} Q_G^{sto}(j) > Q_G^{sto}(0) = (1-\xi)^{2n} > (1-\alpha)^{2n}$ where $\xi \equiv \alpha + \varepsilon - \alpha\varepsilon$. Also, $\sum_{j=K}^{2n} Q_B^{sto}(j) < 1$. This implies that

$$\left(\frac{\delta\beta}{1-\beta}\right)\left(\frac{\sum_{j=0}^{K-1} Q_G^{sto}(j)}{\sum_{j=K}^{2n} Q_B^{sto}(j)}\right) > \frac{\delta\beta}{1-\beta}(1-\alpha)^{2n}$$

For all $\alpha < 1$, $\beta \in (0,1)$, the right-hand-side is strictly positive and independent from ε.

For $\delta < 1$, both $1-\sum_{j=0}^{2n-K}\binom{2n}{j}\varepsilon^j(1-\varepsilon)^{2n-j}\delta$ and $1-(1-\varepsilon)^{2n}\delta$ are strictly positive for all $\varepsilon \in [0,1]$. Moreover,

$$\lim_{\varepsilon\to 0}\frac{\sum_{j=2n-K+1}^{2n}\binom{2n}{j}\varepsilon^j(1-\varepsilon)^{2n-j}}{\sum_{j=1}^{2n}\binom{2n}{j}\varepsilon^j(1-\varepsilon)^{2n-j}} = \lim_{\varepsilon\to 0}\varepsilon^{2n-K}\frac{\sum_{j=2n-K+1}^{2n}\binom{2n}{j}\varepsilon^{j-2n+K}(1-\varepsilon)^{2n-j}}{\sum_{j=1}^{2n}\binom{2n}{j}\varepsilon^j(1-\varepsilon)^{2n-j}} = \lim_{\varepsilon\to 0}\varepsilon^{2n-K}\frac{\binom{2n}{2n-K+1}}{2n} = 0$$

Thus, for all α and β there exists $\varepsilon_0(\alpha,\beta)$[19] small enough, such that the inequality (38) is satisfied for all $\varepsilon < \varepsilon_0(\alpha,\beta)$.

Similarly to the proof of Theorem 2, take a sequence $\alpha_s \to 1$ and corresponding sequence β_s such that $\pi_u^{\mathcal{Q}^{sto,s}} < \pi_c^{\mathcal{Q}^{sto,s}}$ starting from some s. Choose $\varepsilon_s = \frac{1}{2}\varepsilon_0(\alpha_s,\beta_s)$. For such ε_s, starting from some s, we have $\pi_c^{sto}(K) > \pi_u^{sto}$. □

[19]We omit other arguments for simpler notation

Appendix B: General Profit Functions with Informed Consumers

Here we derive general profit functions for a seller with capacity K, in both the deterministic and stochastic setting. As in Section 4.2, let $\mathcal{L} = \{l : \gamma(l, K) > r\}$, and denote the smallest element of $\mathcal{L}$ by L: the number of consecutive sell-outs required as of period 1, given capacity K, in order to generate a positive cascade.

Let $\eta^i_\omega = \sum_{j=K}^{2n} Q^i_\omega(j)$, $i \in \{det, sto\}$, denote the probability of a sell-out, given state $\omega \in \{G, B\}$. Let $S^i_B = \sum_{j=0}^{K-1} jQ^i_B(j)$ denote expected sales in a given period conditional on not selling out, when the state is bad, and let similarly, let $\tilde{S}^{det}_G$ denote expected sales as of a period where a sell-out does not occur, and the state is good. For the deterministic setting, write

$$\tilde{S}^{det}_G \equiv \sum_{j=0}^{K-1} Q^{det}_G(j)\left(j + \delta m + \frac{\delta^2 K}{1-\delta}\right),$$

where unlike in the baseline, the presence of informed consumers will immediately reverse an incorrect negative cascade.

Expected profits given capacity K and $L \geq 1$, in the deterministic setting, are

$$\pi^{det}_c(K) = \beta\left[\frac{1-(\delta\eta^{det}_G)^L}{1-\delta\eta^{det}_G}(\tilde{S}^{det}_G + \eta^{det}_G K) + (\delta\eta^{det}_G)^L \frac{K}{1-\delta}\right] + (1-\beta)\left[\frac{1-(\delta\eta^{det}_B)^L}{1-\delta\eta^{det}_B}(S^{det}_B + \eta^{det}_B K) + (\delta\eta^{det}_B)^L \frac{K}{1-\delta}\right], \quad (39)$$

which reduces to (26) if $L = 1$.

In the stochastic setting, expected sales as of a period where a sell-out does not occur, and the state is good, are

$$\tilde{S}^{sto}_G \equiv \sum_{j=0}^{K-1} Q^{sto}_G(j)\left(j + \delta\frac{\sum_{i=1}^{2n}\binom{2n}{i}\varepsilon^i(1-\varepsilon)^{2n-i}(\min\{i,K\} + \frac{K\delta}{1-\delta})}{1-(1-\varepsilon)^{2n}\delta}\right),$$

as the incorrect negative cascade is reversed in the first subsequent period that an informed consumer arrives. Moreover, let

$$R_B^{sto} = \frac{K}{1-\delta} - \frac{\sum_{i=2n-K+1}^{2n} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} (i - 2n + K + \frac{K\delta}{1-\delta})}{1 - \sum_{i=0}^{2n-K} \binom{2n}{i} \varepsilon^i (1-\varepsilon)^{2n-i} \delta},$$

where the second term reflects the fact that incorrect positive cascades will eventually be reversed, but only when at least $2n - K + 1$ informed consumers arrive in the same period. Expected profits given capacity K and $L \geq 1$, in the stochastic setting, are

$$\pi_c^{sto}(K) = \beta \left[\frac{1 - (\delta\eta_G^{sto})^L}{1 - \delta\eta_G^{sto}} (\tilde{S}_G^{sto} + \eta_G^{sto} K) + (\delta\eta_G^{sto})^L \frac{K}{1-\delta} \right] + (1-\beta) \left[\frac{1 - (\delta\eta_B^{sto})^L}{1 - \delta\eta_B^{sto}} (S_B^{sto} + \eta_B^{sto} K) + (\delta\eta_B^{sto})^L R_B^{sto} \right], \quad (40)$$

which reduces to (28) if $L = 1$.

Figure 4 illustrates how profits depend on how many consumers are informed. Profit functions (39) and (40) are directly decreasing in m and ε respectively, given $\beta < 1/2$, since informed consumers are more likely to know that the state is bad rather than good. If we held fixed the number of sellouts necessary to trigger a positive cascade, say at $L = 1$, then this would suggest that the presence of informed consumers would unambiguously reduce profits. Nonetheless, Figure 4 shows that constrained profits are saw shaped, because more informed consumers can also reduce the number of sell-outs at a given capacity that are required to trigger a positive cascade. For example, an increase in m can lead the seller to choose lower K, because one sell-out will now trigger a cascade at this lower capacity, and result in higher profits.

Figure 4: Seller profits as function of m and ε.

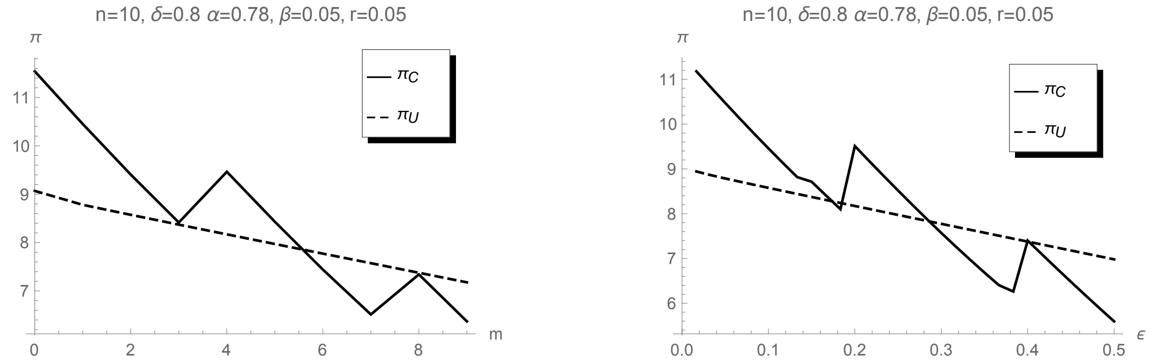

(a) Deterministic Model

(b) Stochastic Model

Appendix C: Pricing

In this appendix we briefly discuss how pricing can affect our results. The seller's optimization problem, seen only as a function of capacity, is complicated by the fact that profits are not necessarily quasi-concave in K. Considering both capacity and price complicates matters further still, since the seller's optimal price is directly linked to its capacity choice. To illustrate the main features of the optimal pricing problem, and maintain tractability, we proceed as follows. First, we assume the seller sets a price at $t = -1$ that is fixed for all periods. Second, we normalize the consumer's outside option to zero and now interpret r as the price, so the consumer problem is exactly as before, and the seller's problem involves the optimal choice of r. Third, we focus on the baseline case, without informed consumers, which makes our framework more similar to Bose et al. (2008), and we compute the optimal price for a constrained seller numerically.[20]

In line with Bose et al. (2008), an unconstrained seller can choose a pooling price such that both consumers with good and bad signals want to buy:

$$r_p = \frac{\beta(1-\alpha)}{\beta(1-\alpha) + \alpha(1-\beta)},$$

[20]That being said, the fact that Bose et al. (2008) consider dynamic rather than static pricing, and do not consider restricted capacity, limits the comparability of our results.

which yields profits

$$\Pi_p = \frac{2n}{1-\delta} r_p.$$

Any lower pooling price $r < r_p$ would not increase demand and would result in lower profits.[21]

Alternatively, an unconstrained seller can charge a separating price such that only consumers with good signals want to buy:

$$r_s = \frac{\alpha\beta}{\alpha\beta + (1-\alpha)(1-\beta)}.$$

Doing so yields profits

$$\Pi_s = \pi_u^{base} r_s,$$

with π_u^{base} given by (9). Any higher price $r > r_s$ would lead to zero sales, whereas any lower separating price $r_p < r \leq r_s$ would not increase demand. The optimal pricing problem is solved by a direct comparison of Π_p and Π_s.

For a constrained seller, we only consider separating prices, $r_p < r \leq r_s$, since a seller that charged a pooling price would prefer to be unconstrained. Given capacity K and separating price r, profits are

$$\Pi = \pi_c^{base}(K) r, \tag{41}$$

with $\pi_c^{base}(K)$ given by (11). The higher the separating price r, the more sell-outs L are required to trigger a positive cascade, so $\pi_c^{base}(K)$ depends on r.

To say more about the optimal price given capacity K, define $L_{\max}$ as the largest value of L for which $\gamma(L-1, K) < r_s$. That is, consumers with good signals in the first cohort will want to buy at price $r = \gamma(L, K)$ if and only if $L \leq L_{\max} - 1$.[22] Demand is independent of price whenever $r \in (\gamma(L, K), \gamma(L+1, K))$, so the optimal price is an element of $\{\gamma(1, K), \ldots, \gamma(L_{\max} - 1, K), r_s\}$.[23] The seller faces the following trade-off: charge a higher price and wait for more sell-outs to

[21] If $m \geq 1$, then in the good state all consumers would still buy at price r_p in all periods, but in the bad state sales would be limited to $(2n - m)$ in the first period and zero thereafter.

[22] Such a value $L_{\max}$ exists by Lemma 3.

[23] Notice that $L_{\max}$ sell-outs are required to trigger a cascade at price $r = r_s$, which is the highest possible separating price.

trigger a cascade, or charge a lower price and hope to trigger a cascade more quickly. Figure 5 illustrates how this trade-off is usually resolved in favour of a higher price and longer waiting for a positive cascade, unless K is very close to n and both α and δ are low.

Panels 5(a) and 5(c) show that the optimal price is often equal to the unconstrained separating price r_s, regardless of which capacity $K \leq n$ the seller finds to be optimal. Differences between these prices, when they do exist, are very small in magnitude. Thus, the seller's problem will often boil down to whether restricting capacity increases expected demand, holding price fixed at $r = r_s$. This suggests that the conclusions from our main analysis, where the seller only optimized over capacity, will tend to carry through to a model with pricing.

Given the relatively high optimal price, multiple sell-outs $L \approx L_{\max}$ are often required to start a cascade, where panel 5(d) shows a saw-like pattern. As signal precision increases, the seller requires fewer sell-outs to trigger a cascade for a given capacity, which pushes L down. But as precision increases, panel 5(b) shows that the seller will sometimes reduce its capacity, so that more sell-outs are required to trigger a cascade, which pushes L sharply up. Panel 5(e) suggests that the overall comparison of profits is qualitatively similar to what we had before: the seller finds it optimal to restrict capacity when the signal accuracy α is neither too high nor too low, similar to Figure 1. The only difference is that for low signal accuracy, an unconstrained seller now prefers the pooling price.

Figure 5: Pricing

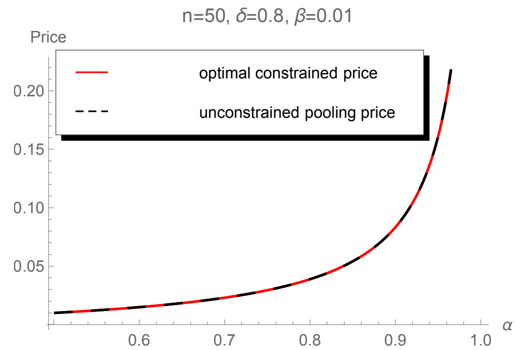

(a) Prices

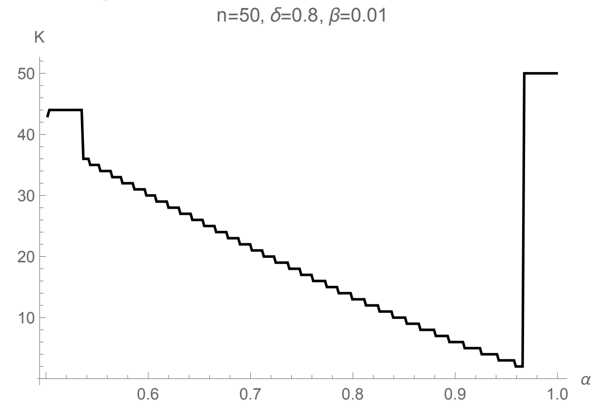

(b) Optimal Capacity

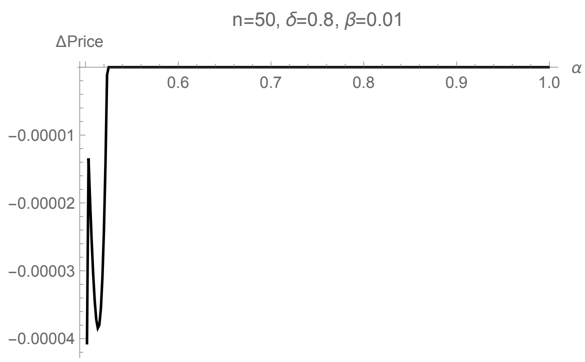

(c) Price difference

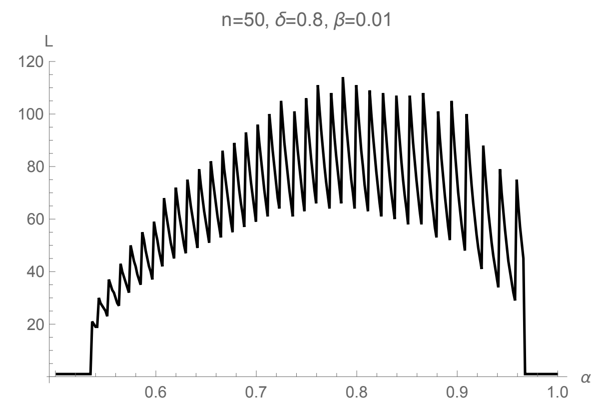

(d) Number of sell-outs

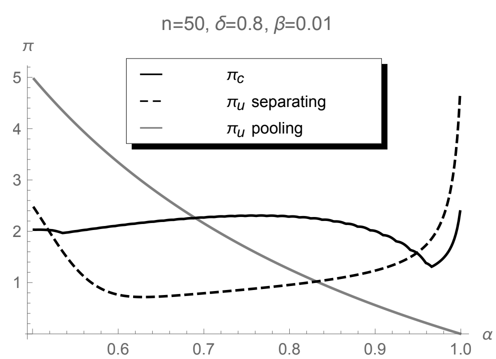

(e) Profit comparison