

Free Product Trials*

Dmitry Lubensky[†] and Eric Schmidbauer[‡]

May 13, 2019

Abstract

With a pre-purchase product trial a seller discloses quality but also enables consumers to privately learn match values. This involves a tradeoff between increasing perceived quality, causing an upward shift in demand, and yielding information rents to consumers, causing a rotation in demand. In contrast to classic results, quality is revealed only when it is sufficiently high, and sometimes not at all. Fewer trials are offered when match is relatively more important than quality, and when there are fewer gains-from-trade. “Cooling off” period rules that allow for free returns effectively require the seller to offer a product trial, and we demonstrate that these rules lead to higher consumer surplus, lower producer surplus, and an overall decrease in welfare when there are sufficient gains from trade. When a seller can commit to a trial policy before learning quality, a trial is only offered if quality is sufficiently *low*. (L15) (L5)

Keywords: disclosure, unraveling, demand rotations, product trial.

*For helpful comments we thank Tony Ke, Larry White, and participants at the 2018 Federal Trade Commission’s Consumer Protection Economics Symposium and the 2019 International Industrial Organization Conference. No financial support for this research was received by either author.

[†]Unaffiliated. 500 Aloha St, Seattle, WA 98109. P: 847.858.0328. dmitry.lubensky@gmail.com

[‡]Corresponding author. University of Central Florida, College of Business Administration. 4336 Scorpis St., Orlando, FL 32816. P: 407.823.1204. eschmidb@ucf.edu.

I INTRODUCTION

Firms that hold private information about their product's quality have a variety of means to transmit this information to consumers. Depending on the setting the firm may, for example, signal through price or advertising (Milgrom & Roberts, 1986), subject itself to third-party certification (Lizzeri, 1999), or disclose directly to consumers. Failure to enact one of these measures that is known to be available may cause consumers to draw negative inferences about the product. In particular, if credible disclosure is possible through pre-purchase trial, a free return policy, or other means a firm that does not allow this may be believed to be the lowest type through an unraveling argument (Grossman and Hart 1980; Grossman 1981; Milgrom 1981).

However, when permitting inspection or trial of its product the firm not only transmits its private information but also endows consumers with *their own* private information. For example, when streaming services like Spotify or Netflix offer a free trial, they not only reveal the size of their catalog but in addition allow each consumer to learn how often the content that she is in the mood for is available, how easy this content is to find, how effective the ratings and recommendations are for discovering new content that matches her tastes, etc. Similarly, in offering a test drive a car dealer not only reveals objective attributes like gas mileage and engine performance, but also allows each consumer to experience aspects like seat comfort and positioning, configuration of dash, and the feel of how the car handles. In both examples a product trial unlocks for each consumer information about her value that the seller itself does not know. In general, a consumer's value is composed of the part predicted by the seller's information (quality), and the part orthogonal to the seller's information (match). In offering a product trial which reveals both, a seller can thus

mitigate adverse selection by disclosing quality, but also endows the consumer with information rents by revealing match.

Disclosure of quality or match have each been previously studied in isolation, and in this paper we combine the insights of these two literatures and how they interact into a unified theory of product trials. In our model a firm privately learns its product's quality and decides whether to allow a pre-purchase trial. A trial reveals to each consumer her true value for the product, while without a trial a consumer forms a posterior about her value based on the trial's absence. In equilibrium offering a product trial has two effects: (i) a shift in demand capturing the difference between the quality of this seller and the average seller without a trial and (ii) a rotation in demand due to consumers moving from the same prior about match to a disperse posterior. The second effect can harm profit and thereby provides an endogenous cost for revealing quality, the magnitude of which depends on the relative importance of quality and match in the consumer's utility function.

Our findings are as follows. First, a firm offers a product trial only when its privately observed quality exceeds a threshold at which the benefit from demand shifting is exactly offset by the cost of demand rotation. Thus a full unraveling does not in general occur, and in fact we show it is possible the demand rotation cost is so high that all types pool on not allowing trial. Second, fewer products trials are offered (i.e., the threshold rises) as the importance of match to consumers' utility rises relative to quality. This follows because the endogenous demand rotation cost from allowing a product trial becomes larger as match increases in importance. Third, we consider the effects of a regulation making product trials mandatory, such as the "cooling-off" periods that allow consumers to cancel a sale within a certain number of days of purchase. We demonstrate that such regulation

increases consumer surplus, reduces producer surplus, and has an ambiguous effect on welfare. With a product trial a seller faces a downward sloping demand and its profit-maximizing uniform price causes a welfare loss by excluding some consumers with values above marginal cost. Without a product trial, the seller sets a price equal to the ex-ante value and sells to all consumers, causing a welfare loss by including consumers whose realized values would have been below marginal cost. We demonstrate that this tradeoff is such that the regulation reduces welfare when the production cost is low, but may increase welfare when the production cost is sufficiently high.

Fourth and finally, we find that if the firm has commitment power its optimal policy is to allow a product trial only when its quality is *below* a threshold value. This is because the firm cannot systematically increase beliefs about quality above their true level and so the only consideration is the demand rotation cost. This endogenous “cost” in fact provides a net benefit when quality is low relative to marginal production costs (Johnson & Myatt, 2006) and thus the firm only wants to allow a product trial when the quality is low.¹

Literature review

The question of how much information a firm should give its customers about its product is usually addressed in the literature in one of two settings: either the firm possesses private information or it does not. When the firm has no private information it must decide whether and to what degree to promote consumer learning before purchase, knowing this will lead to better matching but also endow the consumer with an information rent. For example, in Lewis and Sappington (1994) a monopolist can costlessly improve buyers’ knowledge of their taste parameter by allowing buyers

¹Intuitively, when the quality is low the average consumer value can be below cost. The seller needs consumers to learn match values so that it when it prices above cost some consumers are still willing to buy.

to experiment with the product before purchase. More accurate knowledge of taste can improve matching between consumers' preferences and their consumption, thus creating more surplus, some of which can be captured by the firm. However, consumers also receive an information rent. In the model the firm prefers either a fully informative experiment or a completely uninformative one. This effect of dispersion in willingness to pay on firm profits was later systematically explored in Johnson and Myatt's (2006) study of demand rotations. In an application they consider a model based on the setting of Lewis and Sappington (1994) in which advertising helps consumers learn their idiosyncratic taste for a product, and find that maximal or minimal dispersion is desired. In the model of Anderson and Renault (2006) the firm can freely advertise its price, match value, both or neither. The firm itself does not know the match value but can send a signal which the buyer can use to update beliefs. After receiving an ad (or not) the consumer must decide whether to incur a cost to visit the store, upon which it learns price and match and is able to buy. The authors find the firm will use a "minimum match" advertising strategy which guarantees a threshold level of utility to the consumer. Finally, while these earlier papers explore the question of information design, Kamenica and Gentzkow (2011) do so more generally in their analysis of Bayesian persuasion. Here the seller can commit to any distribution of signals for each state provided the signals are "Bayes plausible", i.e. the expected posterior mean equals the prior mean. Kamenica and Gentzkow then explore the optimal design of such signals for the firm. In the present work we make use of many of the insights generated in this literature, while our main point of departure is that the seller has private information about quality. This is crucial because in contrast to the aforementioned literature, the choice to not inform consumers is not benign but rather involves pooling with low

quality seller types in the eyes of consumers.

In the other main strand of relevant literature the firm possesses private information that it may choose to disclose (see Dranove and Jin 2010 for a literature review on disclosure and third-party certification). It is often assumed in these models that firms can credibly reveal quality at no cost, and a common prediction of such “games of persuasion” is that all quality types will disclose due to an unraveling argument (see Grossman 1981 and Milgrom 1981). That is, in any conjectured equilibrium in which some types pool on non-disclosure, the highest among such types always has incentive to deviate by disclosing. This argument iterates until all types disclose. However, the unraveling prediction is at odds with casual observation that many firms do not disclose and also contrasts with much of the empirical literature demonstrating mandatory disclosure laws seem to have an effect (e.g., Bollinger, Leslie, and Sorensen 2011; Feng Lu 2012; Andrabi, Das, and Khwaja 2017; see Hotz and Xiao 2013 for further citations). This strand of the literature has therefore sought a theoretical basis for at least some non-disclosure. Most obviously, if disclosure is physically costly some firms may not disclose (Grossman & Hart, 1980). Various reasons for non-disclosure related to consumer-naiveté have also been postulated, such as consumers not paying attention to available information, misunderstanding disclosed information, or making incorrect inferences about non-disclosure (Brown, Camerer, and Lovallo 2012; Dranove and Jin 2010; Jin, Luca, and Martin 2015). Our main departure from this strand of literature is that a seller’s disclosure must also generate private information for buyers. In the paper we focus on product trials, in which by disclosing its information the seller induces buyers to physically experience the product, and thereby obtain private information of their own. However, the concept applies more broadly: even

a purely informational disclosure of improved quality (e.g. Spotify announcing it purchased the right to albums by the Beatles, Tesla publishing improved charging times, etc.) can differentially increase consumers' values, thereby involving both a vertical demand shift and a demand rotation.

Our work also relates to the literature dealing with the intersection of disclosure and competition (Hotz and Xiao 2013; Levin, Peck, and Ye 2009; Board 2009; Guo and Zhao 2009). In models such as these, in which match values are public, firms decide in the first stage whether or not to disclose quality and then set prices in the second stage. It is found that non-disclosure can mitigate second stage price competition. Thus in these models the endogenous cost of disclosure is that it makes a competitor's product more substitutable and thus consumers' demand more elastic, and for this reason some non-disclosure may occur. In contrast, if our model were augmented with competition some non-disclosure would occur even without a second stage. That is, if firms simultaneously made disclosure and pricing decisions the demand rotation cost we identify could still imply some non-disclosure. The effect of competition in our model would be to strengthen the consumer's outside option and thus, *ceteris paribus*, shift down a firm's residual demand curve. This would increase the cost of disclosure while not affecting its benefit, and thus induces a lower disclosure threshold. That is, competition would encourage more disclosure.

Finally, we note that other strands of literature have examined the consumer welfare implications of "cooling off" periods and product return policies. These include contributions from the fields of law and economics (Ben-Shahar and Posner 2011; Fishman and Hagerty 2003; Mahoney 1995; Polinsky and Shavell 2010; Rekaiti and Van den Bergh 2000), and industrial organization (Chesnokova 2007; Inderst and Ottaviani 2013; Krämer and Strausz 2015; Petrikaitė 2018), among

others.

The rest of the paper is organized as follows. In Section II we present the model and then in Section III characterize the equilibrium, establish existence, and explore equilibrium properties. In Section IV we consider the effect of mandatory “cooling-off” periods that allow a consumer to return a product within a certain period after purchase, effectively mandating a product trial. In Section V we examine the model in which the seller can commit to a trial policy, and then in Section VI conclude with a discussion of possible extensions.

II MODEL

There is a single seller with constant marginal cost $c \geq 0$ who privately observes quality $\theta \sim F$ on $[\underline{\theta}, \bar{\theta}]$, with $\underline{\theta} \geq 0$ and $\bar{\theta} \leq \infty$. A single consumer with unit demand obtains value

$$v = \theta + \alpha\varepsilon,$$

from purchasing, where $\varepsilon \sim G$ on $[\underline{\varepsilon}, \bar{\varepsilon}]$ with density g and $E[\varepsilon] = \eta$. We interpret ε as the consumer’s match value for the product which is independent of θ , with $\underline{\varepsilon} \geq 0$, while $\alpha \geq 0$ is an exogenous parameter that is common knowledge and measures the importance of match relative to quality.^{2,3} We assume that $1 - G$ is log-concave to ensure that when consumers know their match

²A mathematically equivalent interpretation is that $\theta + \eta$ is the quality and $\varepsilon - \eta$ is the match, where quality is the value to the average consumer and match is the deviation from average. We select the interpretation in the main text only to simplify exposition.

³Although we consider a single-period static setting, little would change if instead a consumer with fixed ε could purchase over multiple periods. In her first period the consumer either receives a free product trial and learns ε , or purchases the product without a trial and learns her ε from experience. Having the same information, her behavior in subsequent periods is therefore independent of whether a trial was offered in the initial one.

value the seller's first order condition characterizes the optimal price. We assume $\underline{\theta} + \alpha\eta \geq c$ so that any type not allowing a product trial is able to sell. The seller observes θ but not ε , decides whether to costlessly allow a product trial, then charges a uniform price p . If the seller allows trial the consumer costlessly observes v and p and decides whether to purchase, where utility is $v - p$ from purchasing and 0 from not purchasing. If the seller does not allow trial the consumer observes p only, holds posterior beliefs μ about quality θ and decides whether to purchase. Thus expected quantity sold equals 1 when there is no trial and $p \in [0, E_\mu(\theta) + \alpha\eta]$, and otherwise equals $q(p, \theta) \equiv 1 - G\left(\frac{1}{\alpha}(p - \theta)\right)$ if there is trial and $p \in [\theta + \alpha\underline{\varepsilon}, \theta + \alpha\bar{\varepsilon}]$. Our solution concept is perfect Bayesian equilibrium, which consists of the firm's product trial and pricing strategy, the consumer's purchasing strategy, and the consumer's beliefs about quality.

III EQUILIBRIUM

We first observe that a seller's payoff depends on the consumer's perception of quality, but not on the actual quality. There is then little scope for price signaling – if price p_1 induces a belief that results in a higher profit than price p_2 for one type of seller, then p_1 is more profitable than p_2 for every type of seller. We therefore restrict attention to equilibria in which beliefs do not depend on the seller's price when a seller does not allow a product trial,⁴ and focus on the firm's decision to allow a product trial or not.

Let μ denote the average quality of the set of types that do not allow trial and $\pi(\theta) \equiv \max_p (p -$

⁴There may exist equilibria in which types that do not offer a product trial set multiple prices. In such an equilibrium the payoff from any of these prices must give the same profit to all types, and therefore consumers must accept higher prices with lower probabilities. Assuming consumers do not use mixed strategies would rule out such equilibria.

$c)(1 - G(\frac{1}{\alpha}(p - \theta)))$ denote the maximized profit of type θ when allowing a product trial. The net benefit of allowing a product trial for type θ is then

$$\Delta(\theta, \mu) = \pi(\theta) - (\mu + \alpha\eta - c). \quad (1)$$

If $\alpha = 0$ so that the match value is irrelevant, the maximized profit simplifies to $\pi(\theta) = \theta - c$, the net benefit of offering a product trial becomes $\Delta(\theta, \mu) = \theta - \mu$, and therefore every type not offering a trial that is better than the average type not offering a trial would deviate to offering a product trial. By this logic the only equilibrium is one in which every seller offers a trial, as demonstrated in Milgrom (1981).

However when $\alpha > 0$ this argument no longer holds. By showing that its quality is better than the average type not offering trial, the seller now also endows consumers with private information about match value, and yields information rents to these consumers. Consequently we will demonstrate that some seller types may not offer trial in equilibrium. We first establish that the product trial policy follows a threshold rule.

Lemma 1 *In any equilibrium the seller follows a threshold policy t and offers a product trial if and only if $\theta \geq t$.*

This follows immediately from the fact that the net benefit of offering trial $\Delta(\theta, \mu)$ from (1) increases monotonically in θ . This fact also implies there is not an equilibrium in mixed strategies, as every type other than the threshold type strictly prefers to either offer trial or not. Letting

$\mu(t) \equiv E[\theta < t]$, the equilibrium value of t is then characterized as follows:

$$\Delta(t) \equiv \Delta(t, \mu(t)) = 0 = \pi(t) - (\mu(t) + \alpha\eta - c)$$

$$0 = \underbrace{(t - \mu(t))}_{B(t)} - \underbrace{(t + \alpha\eta - c - \pi(t))}_{C(t)}.$$

We decompose the net benefit from offering a product trial into two terms. The first, $B(t) \equiv t - \mu(t)$, captures the benefit of increasing the belief from $\mu(t)$ to t conditional on consumers not learning ε . That is, if the threshold firm could disclose only its quality but not the match value, it would increase consumers' willingness to pay by $t - \mu(t)$, raise its price by this amount, and thereby continue both to sell to all consumers and extract all consumer surplus. However, by offering a product trial the firm also endows consumers with private information about ε , rotating demand from a horizontal line at height $\theta + \alpha\eta$ to a downward-sloping demand function $q(p) = 1 - G\left(\frac{1}{\alpha}(p - \theta)\right)$. This effect is captured by the second term in the decomposition, $C(t) \equiv t + \alpha\eta - c - \pi(t)$, which is the change in profits when consumers with belief t learn ε . A demand rotation clearly results in a loss of revenue because the firm can no longer extract all consumer surplus, although the effect on profit will depend on the level of costs c .^{5,6} See Figure 1 for an illustration of this decomposition of pre-purchase product trial's effect on profits of the threshold type t .

We now establish some useful properties of $B(t)$ and $C(t)$ in the following lemma, the main

⁵Thus in the neighborhood of $c = 0$ a demand rotation will always result in a decrease in profit, so that $C(t) > 0$ for any t . To see that profits may increase, suppose a firm's quality is known to be $\theta = 1$ and that match values $\varepsilon \sim U[0, 1]$. Not offering trial earns the firm profits of $1 + \frac{\alpha}{2} - c$ while with trial it earns $\max_p (p - c) \left(1 - \frac{p-1}{\alpha}\right) = \frac{(1+\alpha-c)^2}{4\alpha}$. Then when $c = 1 + \frac{\alpha}{2}$ the firm earns more profits from offering trial than not, so that $C(1) < 0$ in this case.

⁶If the firm could price discriminate the information rent ceded to the consumer by allowing pre-purchase product trial would be lower and thus the demand rotation cost would also be lower. In the limit, perfect price discrimination would lead to the classic unraveling result.

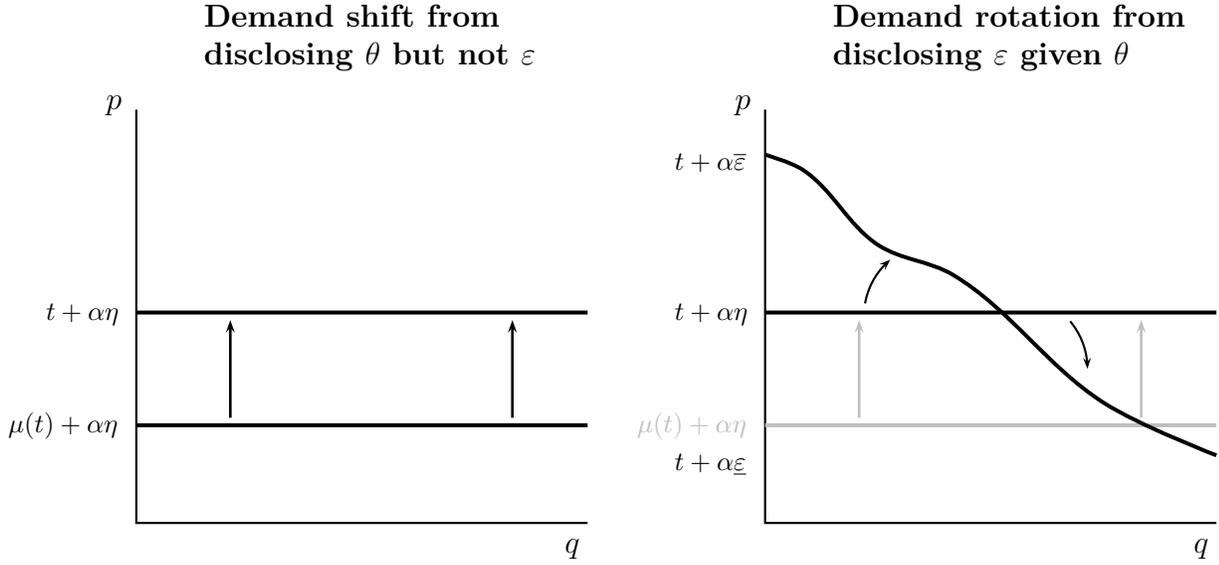


Figure 1: The effect of a free product trial on the demand of the threshold type, decomposed into an upward shift from disclosing quality θ (left panel) and a rotation from disclosing match ε (right panel).

results of which are that the benefit of the vertical demand shift from offering a product trial for the threshold type is increasing in the threshold while the demand rotation cost is also increasing but is bounded above.

Lemma 2

i $B(\underline{\theta}) = 0$ and $B(t)$ increases without bound in t if $\bar{\theta} = \infty$.

ii Given t , $C(t) > 0$ if and only if c is below a threshold value. Given c , $C'(t) > 0$, $C''(t) < 0$, and $C(t)$ achieves its upper bound of $\alpha(\eta - \underline{\varepsilon}) > 0$ when $\bar{\theta} = \infty$.

Proof Part (i). $B(\underline{\theta}) = 0$ is obvious. Also, if $\bar{\theta} = \infty$, we have $\lim_{t \rightarrow \infty} \mu(t) = E[\theta]$ and thus $B(t)$ converges to $t - E[\theta]$ which grows without bound. Part (ii). Define p^* as the profit maximizing

price used in calculating $\pi(\theta)$. Observe that

$$\begin{aligned}
C'(t) &= 1 - \frac{\partial \pi}{\partial t} \\
&= 1 - \left(\frac{p^* - c}{\alpha} \right) g \left(\frac{p^* - t}{\alpha} \right) \\
&= 1 - \left(1 - G \left(\frac{p^* - t}{\alpha} \right) \right) \\
&= G \left(\frac{p^* - t}{\alpha} \right) \\
&> 0.
\end{aligned}$$

In the first line we use $\frac{d\pi}{dt} = \frac{\partial \pi}{\partial t}$, since $\frac{\partial \pi}{\partial p} = 0$ from the envelope theorem. The second line substitutes $\frac{\partial \pi}{\partial t}$, the third uses a substitution from the first order condition for π evaluated at $\theta = t$, namely $\frac{d}{dp} \left((p - c)(1 - G(\frac{1}{\alpha}(p - t))) \right) = 1 - G(\frac{p-t}{\alpha}) - (\frac{p-c}{\alpha}) g(\frac{p-t}{\alpha}) = 0$, while the fourth simplifies.

Next we show $C(t) > 0$ if and only if c is below a threshold value. First note that if $c = 0$ then $C(t) > 0$ as consumers obtain an information rent when they learn ε . But $C' > 0$ and C is continuous so it suffices to find a c such that $C(t) < 0$. This occurs, for example, when $c = t + \alpha\eta$ since in this case profits are 0 unrotated but positive when rotated. Thus a threshold value c' exists, and is defined implicitly by $C(0; c') = 0$. Regarding the second derivative, $C' = 1 - \pi'$ and thus $C'' = -\pi'' < 0$ since $\pi'' > 0$ by the log-concavity of $1 - G$. Lastly we establish the result on the upper bound of $C(t)$. First note $\lim_{\bar{\theta} \rightarrow \infty} \pi(\bar{\theta}) = \bar{\theta} + \alpha\underline{\varepsilon} - c$ since for t high enough π obtains a

corner solution $q = 1$ and $p = \bar{\theta} + \alpha\underline{\varepsilon}$. Next, $C(\bar{\theta}) = \bar{\theta} + \alpha\eta - c - \pi(\bar{\theta})$, and thus

$$\begin{aligned} \lim_{\bar{\theta} \rightarrow \infty} C(\bar{\theta}) &= \bar{\theta} + \alpha\eta - c - (\bar{\theta} + \alpha\underline{\varepsilon} - c) \\ &= \alpha(\eta - \underline{\varepsilon}) \\ &> 0, \end{aligned}$$

where the second inequality follows since $\underline{\varepsilon} < E[\varepsilon] \equiv \eta$. ■

We now consider equilibrium existence. In fact it can easily be seen an equilibrium always exists. First, Lemma 1 states any equilibrium will use a threshold policy. Since a threshold type $t \in (\underline{\theta}, \bar{\theta})$ must be indifferent to product trial in equilibrium, the necessary and sufficient condition is $B(t) = C(t)$. If instead $B(t)$ exceeds (is exceeded by) $C(t)$ for all t then the unique equilibrium is pooling on all types offering trial (no types offering trial), so $t = \underline{\theta}$ ($t = \bar{\theta}$). Below we given conditions for an equilibrium with some product trial to exist and be unique.

Proposition 1 *All seller types offer a product trial only if $C(\underline{\theta}) \leq 0$ and α is sufficiently small. Otherwise, there exists a $t \in (\underline{\theta}, \bar{\theta}]$ so that the seller offers a product trial if and only if $\theta \geq t$. Furthermore, if $\mu''(t) < 0$ then the equilibrium is unique.*

Proof That the equilibrium is a threshold was established in Lemma 1, and by Lemma 2 we know that $B(\underline{\theta}) = 0$. If $C(\underline{\theta}) > 0$ then either $C(\theta) > B(\theta)$ for all θ , in which case $t = \bar{\theta}$ and no types offer trial, or there is at least one intersection for some $t \in (\underline{\theta}, \bar{\theta})$, in which case types below t do not offer trial and types above t do. If $C(\underline{\theta}) < 0$ then it is possible that $B(\theta) > C(\theta)$ for all θ , in which case everyone offers a product trial. However, because $C(\theta)$ is increasing and unbounded in α (see

proof of Proposition 2), while $B(\theta)$ does not change in α , once α is sufficiently large an interior intersection is guaranteed. For uniqueness, we have established in Lemma 2 that $C' > 0$, $C'' < 0$, and therefore if $B'' \geq 0$ any intersection of $B(\theta)$ and $C(\theta)$ is unique. It is then easily verified that $B''(t) > 0 \Leftrightarrow \mu''(t) < 0$. ■

The proposition gives conditions ensuring a unique equilibrium in which some quality types do not offer a product trial. Intuitively, $C(\underline{\theta}) > 0$ implies that all quality types incur some positive rotation cost and thus the very lowest types will not want to offer trial (since $B(\underline{\theta}) = 0$). We now consider explicit examples that demonstrate pooling on no trial can occur and an equilibrium in which a strict subset of seller types do not offer a trial need not be unique.

Example 1 *Let $c = 0$ and θ and ε be $U[0, 1]$. Then in equilibrium no types offer trial when $\alpha > 1$, types offer trial if and only if they exceed the threshold $t = \alpha$ when $\alpha \in (0, 1]$, and all types offer trial when $\alpha = 0$. This follows from a direct computation of $B(t)$ and $C(t)$:*

$$\begin{aligned} B(t) &= \frac{t}{2} \\ C(t) &= t + \frac{\alpha}{2} - \pi(t; \alpha) \\ &= t + \frac{\alpha}{2} - \left(\frac{\alpha + t}{2} \right) \left(\frac{1}{2} + \frac{t}{2\alpha} \right) \\ &= \frac{t}{2} + \frac{\alpha}{4} - \frac{t^2}{4\alpha}, \end{aligned}$$

and so $C(t) > B(t)$ for all $t \in [0, 1]$ if and only if $\frac{\alpha}{4} > \frac{t^2}{4\alpha}$ for all $t \in [0, 1]$, or equivalently $\alpha > 1$.

When $\alpha < 1$ we have $B(\alpha) = C(\alpha)$ so that the trial threshold is $t = \alpha$, which is pooling on product trial when $\alpha = 0$ and non-pooling otherwise. Note an interior solution for $\pi(t; \alpha)$ obtains when

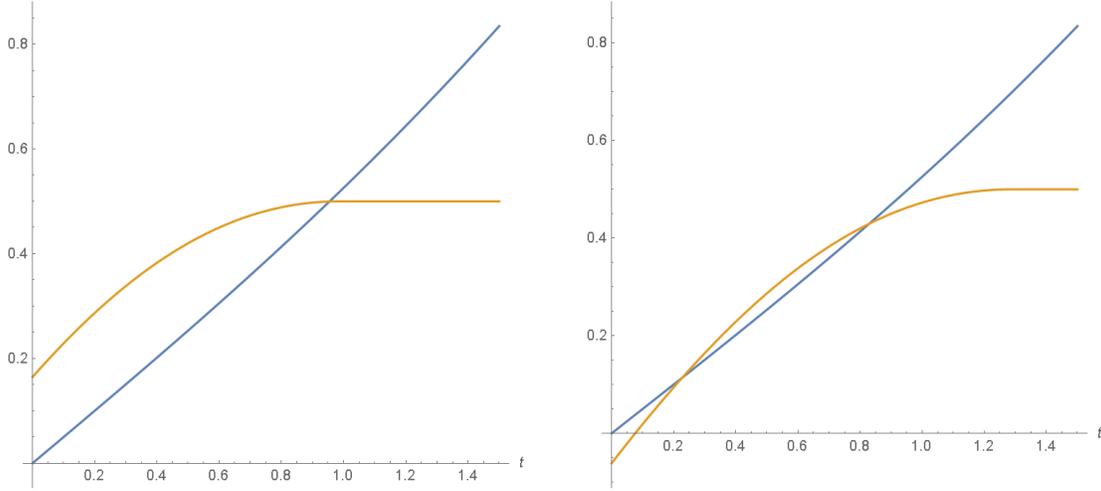


Figure 2: The cost (in orange) and benefit (in blue) of product trial to the threshold type in Example 2. In the left panel $c = \frac{1}{5}$ and a unique equilibrium exists. In the right $c = \frac{1}{2}$ and multiple equilibrium exist.

$t < \alpha$ and so our substitution for $\pi(t; \alpha)$ above is valid when $t \leq \alpha$, which is satisfied.

Example 2 Let $\alpha = \frac{1}{2}$, $c = \frac{1}{5}$, and θ and ε be iid half normal on $[0, \infty)$ with scaled precision parameter 1. A calculation reveals a unique equilibrium threshold $t \approx 0.954$ whereby types offer trial if and only if they exceed t . If instead $c = \frac{1}{2}$ then there are two equilibrium thresholds: $t_1 \approx 0.225$ and $t_2 \approx 0.827$. A graph of $B(t)$ and $C(t)$ for each case is shown in Figure 2.

We now explore equilibrium properties, specifically how the trial threshold changes with production costs c and the importance of horizontal attributes relative to vertical, α .

Proposition 2 When $C(\underline{\theta}) > 0$ and the product trial threshold is unique it falls as the marginal production cost increases and rises as α increases.

Proof Note that both c and α affect $C(t)$ but not $B(t)$. Further, by the construction of the equilibrium in Proposition 1 we know at a unique equilibrium threshold $C(t)$ crosses $B(t)$ from

above, and so the threshold increases in c or α if and only if $C(t)$ increases.

$$\begin{aligned}
\frac{dC}{dc} &= -1 - \frac{d\pi}{dc} \\
&= -1 - \left(\frac{\partial\pi}{\partial c} + \frac{\partial\pi}{\partial p} \frac{\partial p}{\partial c} \right) \\
&= -1 - \frac{\partial\pi}{\partial c} \\
&= -1 - (-q) \\
&< 0.
\end{aligned}$$

The second line performs the total derivative while the third uses the envelope theorem, i.e. $\frac{\partial\pi}{\partial p} = 0$.

The fourth line uses the fact that un-reoptimized profits go down at rate q with respect to c , and the fifth that quantity $q < 1$. Thus $\frac{dC}{dc} < 0$.

We now establish the result on α . Let p^* be the profit maximizing price and let ε^* denote the marginal consumer that purchases at this price; i.e., ε^* is such that $p^* = t + \alpha\varepsilon^*$. This marginal type is derived by taking the seller's first order condition:

$$\begin{aligned}
\pi &= (t + \alpha\varepsilon - c)(1 - G(\varepsilon)) \\
\Rightarrow \pi_{\varepsilon^*} &= 0 = \alpha(1 - G(\varepsilon^*)) - (t + \alpha\varepsilon^* - c)g(\varepsilon^*) \tag{2}
\end{aligned}$$

Next,

$$\begin{aligned}
C(t) &= t + \alpha\eta - c - \pi(t) \\
&= \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} (t + \alpha\varepsilon - c)g(\varepsilon)d\varepsilon - \int_{\varepsilon^*}^{\bar{\varepsilon}} (t + \alpha\varepsilon^* - c)g(\varepsilon)d\varepsilon \\
&= \alpha \int_{\varepsilon^*}^{\bar{\varepsilon}} (\varepsilon - \varepsilon^*)g(\varepsilon)d\varepsilon + \int_{\underline{\varepsilon}}^{\varepsilon^*} (t + \alpha\varepsilon - c)g(\varepsilon)d\varepsilon \tag{3}
\end{aligned}$$

In the second line we multiply $t - c$ by $\int_{\underline{\varepsilon}}^{\bar{\varepsilon}} g(\varepsilon)d\varepsilon = 1$, write explicitly $\eta = \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \varepsilon g(\varepsilon)d\varepsilon$, and substitute for maximized profits π , which integrates the markup $t + \alpha\varepsilon^* - c$ over those types that purchase. The third line splits up the first integral in the second line into an integral over $[\underline{\varepsilon}, \varepsilon^*]$ and over $[\varepsilon^*, \bar{\varepsilon}]$, and then combines with the second integral in the second line. We now differentiate line (3) with respect to α :

$$\begin{aligned}
\frac{dC}{d\alpha} &= \int_{\varepsilon^*}^{\bar{\varepsilon}} (\varepsilon - \varepsilon^*)g(\varepsilon)d\varepsilon + \alpha \left(-\frac{d\varepsilon^*}{d\alpha} \times 0 + \int_{\varepsilon^*}^{\bar{\varepsilon}} -\frac{d\varepsilon^*}{d\alpha} g(\varepsilon)d\varepsilon \right) \\
&\quad + \frac{d\varepsilon^*}{d\alpha} (t + \alpha\varepsilon^* - c)g(\varepsilon^*) + \int_{\underline{\varepsilon}}^{\varepsilon^*} \varepsilon g(\varepsilon)d\varepsilon \\
&= \int_{\varepsilon^*}^{\bar{\varepsilon}} (\varepsilon - \varepsilon^*)g(\varepsilon)d\varepsilon - \frac{d\varepsilon^*}{d\alpha} \left(\alpha(1 - G(\varepsilon^*)) - (t + \alpha\varepsilon^* - c)g(\varepsilon^*) \right) + \int_{\underline{\varepsilon}}^{\varepsilon^*} \varepsilon g(\varepsilon)d\varepsilon \\
&= \int_{\varepsilon^*}^{\bar{\varepsilon}} (\varepsilon - \varepsilon^*)g(\varepsilon)d\varepsilon + \int_{\underline{\varepsilon}}^{\varepsilon^*} \varepsilon g(\varepsilon)d\varepsilon \\
&> 0.
\end{aligned}$$

The first and second lines differentiate the first and second integrals found in line (3) above. The third line substitutes $\int_{\varepsilon^*}^{\bar{\varepsilon}} g(\varepsilon)d\varepsilon = 1 - G(\varepsilon^*)$, reduces and rearranges terms. The fourth line follows from a substitution from the first order condition found in line (2). Finally, the inequality follows

since $\underline{\varepsilon} \geq 0$. ■

The proposition gives comparative statics when the equilibrium is unique. However, it need not be unique as Example 2 demonstrated. In such cases the proposition’s findings still hold provided that at the equilibrium threshold $C(\theta)$ crosses $B(\theta)$ from above, which in particular will occur at the largest threshold whenever $\bar{\theta} = \infty$.⁷

IV “Cooling-Off” Rules

Regulators in many jurisdictions have imposed “cooling-off” periods for certain types of consumer purchases during which consumers are able to cancel a sale.⁸ In effect such a rule mandates a product trial, and we now consider the welfare effects of this regulatory policy (which we call “mandatory trial”).

Remark 1 *A mandatory trial increases consumer surplus and reduces producer surplus.*

The policy’s effect on consumer surplus is readily ascertained. In equilibrium a consumer either faces a seller with quality $\theta \in [\underline{\theta}, t)$ from whom she buys without a trial, or a seller with quality $\theta \in (t, \bar{\theta})$ from whom she receives a trial. In the former case, the consumer pays exactly her expected value (taking the expectation over all $\theta \in [\underline{\theta}, t]$ and $\varepsilon \in [\underline{\varepsilon}, \bar{\varepsilon}]$), thus receiving no surplus. In the latter case, the consumer is charged $p^*(\theta)$ and buys only when her value exceeds this price, thereby

⁷The largest threshold may be of interest since, as we later argue, from an ex-ante perspective the seller prefers less product trial and it results in higher total surplus.

⁸For example, in the European Union consumers have 14 days to withdraw from a sales contract that was executed over the internet, phone, or mail (*Official Journal of the European Union*, L 304, November 22, 2011, p. 64–88). In the U.S. the Federal Trade Commission’s regulation 3084-AB10 gives consumers a 3-day right to cancel certain sales over \$25 made at home or over \$130 made at a seller’s temporary location (80 Fed. Reg. 1329). In addition, many U.S. states provide broader consumer protections, such as in California where the cooling-off period is extended to 7 days for certain purchases (Cal. Civ. Code §1689.6).

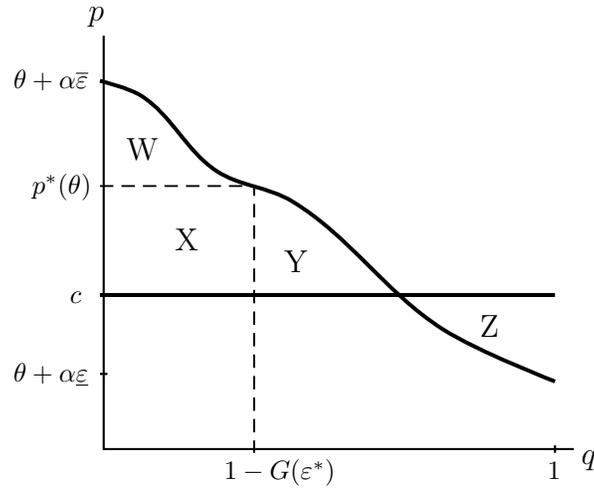


Figure 3: Welfare effect of product trial

receiving surplus. A mandatory trial policy then has no effect for $\theta > t$, but for $\theta < t$ allows consumers to collect surplus by buying only when their value exceeds $p^*(\theta)$.

It is also straightforward to demonstrate that producer surplus is reduced by the policy. In equilibrium every type θ has the option to offer a product trial, and hence receive the payoff they would under the policy. Every type $\theta < t$ that chooses not to offer a trial is therefore made worse off by the policy, and thus producer surplus is lower.

The effect of mandatory product trials on total welfare is more nuanced. To build intuition, consider the demand curve faced by a particular seller θ that offers a product trial, as depicted in Figure 3. As before, let $\varepsilon^*(\theta)$ denote the threshold match value at which a consumer purchases when the seller sets the optimal price, i.e. $p^*(\theta) = \theta + \alpha\varepsilon^*(\theta)$. Without the trial every consumer buys regardless of ε , resulting in a total of surplus of $\theta + \alpha\eta - c$, which in terms of the figure is given by $W + X + Y - Z$, while with the trial only consumers above ε^* buy, resulting in a total surplus

of $W + X$. The welfare impact of the mandatory trial therefore equals $Z - Y$, and is expressed as

$$\Delta W(\theta) = \underbrace{\int_{\theta + \alpha\varepsilon \in [\theta + \alpha\underline{\varepsilon}, c]} (c - (\theta + \alpha\varepsilon)) f(\varepsilon) d\varepsilon}_{\text{Welfare benefit of trial (Z)}} - \underbrace{\int_{\theta + \varepsilon \in (c, \theta + \alpha\varepsilon^*]} (\theta + \alpha\varepsilon - c) f(\varepsilon) d\varepsilon}_{\text{Welfare cost of trial (Y)}}. \quad (4)$$

A product trial thus has the welfare benefit (Z) of avoiding trade with consumers whose willingness to pay is below the production cost and the welfare loss (Y) of forgone gains from trade from those whose values are above the production cost but below the monopoly price. The relative size of these two effects depends on the parameters of the model as we demonstrate below.

Lemma 3 *There exists a smallest value \underline{a} and a largest value \bar{a} so that $Z \geq Y$ whenever $\theta - c \leq \underline{a}$ and $Z \leq Y$ whenever $\theta - c \geq \bar{a}$.*

Proof If $\theta - c = -\alpha\bar{\varepsilon}$ then the area $Y = 0$ and the area $Z > 0$, while when $\theta - c = -\alpha\underline{\varepsilon}$ then $Z = 0$ and $Y > 0$. Then, because Y and Z are both continuous in $\theta - c$ the lemma follows. ■

Intuitively, Z measures welfare loss (without a product trial) from served consumers with negative gains from trade while Y measures the welfare loss (with a product trial) from unserved consumers with positive gains from trade. A high value of $\theta - c$ implies a small Z and a large Y , and a low value of $\theta - c$ implies the reverse. Using Lemma 3 and defining \hat{c} such that $C(\underline{\theta}, \hat{c}) = 0$, i.e. the lowest type $\underline{\theta}$ is indifferent to a demand rotation, we derive sufficient conditions for when a mandatory trial policy helps or harms welfare as follows.

Proposition 3 *Mandatory product trials strictly harm welfare whenever $c \leq \underline{\theta} - \bar{a}$ and strictly improve welfare whenever $c \geq \bar{\theta} - \underline{a}$ and $c < \hat{c}$.*

Proof First, if $c \leq \underline{\theta} - \bar{a}$ then $\theta - c \geq \bar{a}$ for all θ , and therefore $Z \leq Y$ for all θ by Lemma 3. In particular this is true for $\theta \in [\underline{\theta}, t]$ for whom the mandatory trial policy binds, so it then remains to be shown that this set of types is nonempty (i.e. $t > \underline{\theta}$). This in turn follows from the fact that type $\underline{\theta}$ is strictly better off from offering the product trial, because $B(\underline{\theta}) = 0$ and $C(\underline{\theta}) = W + Y - Z \geq W > 0$.

Next, if $c \geq \bar{\theta} - \underline{a}$ then $\theta - c \leq \underline{a}$ for all θ , and therefore the policy improves welfare for every θ it forces to switch to offering a trial. However this is a strict improvement in welfare only if $[\underline{\theta}, t]$ is nonempty. To guarantee this, it must again be the case that type $\underline{\theta}$ strictly prefers not to offer a product trial, that is $C(\underline{\theta}) \geq 0$, which is ensured by $c < \hat{c}$. ■

To interpret Proposition 3, observe that at a very low c it is both the case that many types do not offer a trial in equilibrium and that forcing each of those types to offer a trial harms welfare. As c increases, the set of types not offering a product trial shrinks while the negative welfare impact of a product trial diminishes for all types. The question is whether there exists a $c < \hat{c}$ at which the net welfare effect becomes positive while the set of no-trial types $[\underline{\theta}, t]$ is nonempty. In general this cannot be guaranteed, and we now explore this further with a specific example.

Let $\alpha = 1$, $G(\varepsilon) = \varepsilon$ on $[0, 1]$, and $F(\theta) = \gamma + (1 - \gamma)\theta$ also on $[0, 1]$. That is, the distribution of match values ε is uniform while the distribution of qualities θ has an atom of size $\gamma \in (0, 1]$ at $\theta = 0$ and is otherwise also uniform. This atom provides an easy way to parameterize the disclosure benefit for the threshold type, given by

$$B(t) = t - \mu(t) = t - \frac{(1 - \gamma)\frac{t^2}{2}}{\gamma + (1 - \gamma)t}. \quad (5)$$

Observe that $\gamma = 0$ corresponds to the standard uniform case with $B(t) = \frac{1}{2}t$, that $B(t)$ increases in γ , and that $\lim_{\gamma \rightarrow 1} B(t) = t$, which is the largest possible benefit.

We first solve for the equilibrium threshold t . When offering a product trial the maximized profit for type θ is $\pi(\theta) = \frac{1}{4}(1 + \theta - c)^2$, and thus

$$C(t) = t + \frac{1}{2} - c - \frac{1}{4}(1 + t - c)^2. \quad (6)$$

The equilibrium threshold is then implicitly defined as $B(t) = C(t)$.

Next, with the equilibrium threshold t in hand we can evaluate the welfare consequences of mandatory product trials. Using $\Delta W(\theta)$ from line (4), the increase in ex-ante welfare from a mandatory trial is

$$\int_0^t \Delta W(\theta) dF(\theta) = \gamma[Z(0) - Y(0)] + (1 - \gamma) \int_0^t [Z(\theta) - Y(\theta)] d\theta, \quad (7)$$

where the first term on the right hand side accounts for the welfare effect on the atom at $\theta = 0$ and the second term integrates the welfare effect over the remaining types $(0, t)$ that do not offer a trial in equilibrium. The expressions for $Z(\theta)$ and $Y(\theta)$ as defined in line 4 are as follows:

$$\begin{aligned} Z(\theta) &= \int_0^{c-\theta} (c - (\theta + \varepsilon)) d\varepsilon = \frac{(c - \theta)^2}{2} \\ Y(\theta) &= \int_{c-\theta}^{\varepsilon^*} (\theta + \varepsilon - c) d\varepsilon = \frac{\varepsilon^{*2}}{2} - (c - \theta)\varepsilon^* + \frac{(c - \theta)^2}{2} = \frac{1}{8}(1 + \theta - c)^2, \end{aligned}$$

where in the line for $Y(\theta)$ the term $\varepsilon^* \equiv 1 - G(p^* - \theta) = 1 - (\frac{1}{2}(1 + \theta + c) - \theta) = \frac{1 + \theta - c}{2}$.

Now for every combination of parameter values (c, γ) we can compute the equilibrium threshold t , and then using this threshold compute the net effect of mandated product trials on welfare, as depicted in Figure 4. We find that for every value of γ the policy reduces welfare when c is low. When c is high, welfare can increase from this policy but only if γ is also sufficiently high. In particular, if there is no atom at $\theta = 0$ then the policy always reduces welfare. The intuition is that when γ is small so too is the benefit of offering a trial, and thus at intermediate qualities the seller does not offer a trial in equilibrium. Inducing these seller types with sizable gains from trade to offer trials harms welfare, and turns out to more than offset the welfare benefit of inducing trials from seller types with lower qualities. However, as the size of the atom γ grows the benefit $B(t)$ grows, the threshold t falls, and the policy begins to affect mostly those types for whom a product trial would increase welfare.

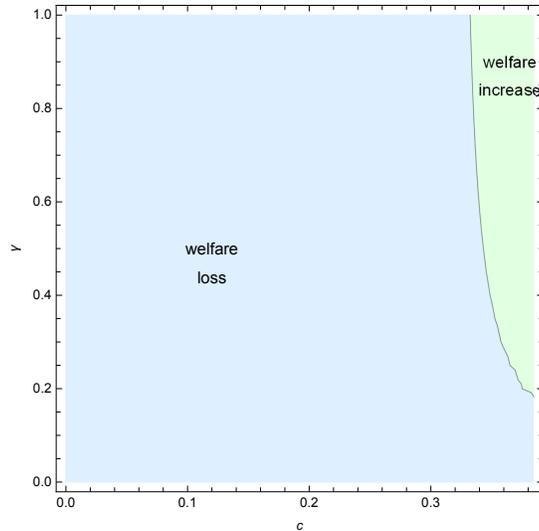


Figure 4: The regions in (c, γ) space where a mandatory trial policy harms and improves welfare (blue and green regions respectively). In this example $G(\varepsilon) = \varepsilon$ on $[0, 1]$ and $F(\theta) = \gamma + (1 - \gamma)\theta$ on $[0, 1]$.

V Commitment power

The question of whether and how much information to provide consumers before purchase has been explored recently in the literature on Bayesian persuasion and information design. In contrast to this literature, in our model the seller cannot commit to the decision of whether to offer a product trial before it learns the quality θ . Although we make this assumption for the sake of realism, we would also like to better understand what happens when this assumption is relaxed, and thereby facilitate comparison with the Bayesian persuasion literature.

Formally, a policy d is a function from the set of states $[\underline{\theta}, \bar{\theta}]$ into a probability of offering a product trial in $[0, 1]$. While we consider all possible commitment policies we conclude below a pure strategy threshold policy is always optimal.

Proposition 4 (Commitment) *There exists a unique threshold t^* such that types offer a product trial if and only if they are **below** t^* . If $C(\underline{\theta}) > 0$ then $t^* = \underline{\theta}$ so the optimal policy is to never offer a trial.*

Proof Let d be *any* trial policy (not necessarily a threshold policy) and denote by μ_d the Bayesian posterior mean of θ given policy d and no product trial. Let $\Delta(\theta, d) = \pi(\theta) - (\mu_d + \alpha\eta - c)$ be the increase in profits to type θ from offering trial when policy d is anticipated by consumers. Then $\Delta(\theta, d) = B(\theta, d) - C(\theta)$, where $B(\theta, d) = \theta - \mu_d$ and $C(\theta) = \theta + \alpha\eta - c - \pi(\theta)$. Let D be the set of θ that do not offer a trial under policy d . Then for any trial policy d ,

$$E_{\theta \in D}[B(\theta, d)] = E_{\theta \in D}[\theta - \mu_d] = 0$$

since μ_d is precisely the Bayesian posterior mean given policy d and no product trial. Simply put, because B is linear in θ , it follows that $E[B(\theta)] = B(E[\theta])$. Thus there is no ex-ante demand shifting benefit from any trial policy.

The seller's optimal policy is thus straightforward – commit to offering a trial for all types θ for which $C(\theta) \leq 0$. Thus, if $C(\underline{\theta}) \geq 0$ it is optimal for the seller to commit to no trial, and if $C(\underline{\theta}) < 0$ then, because C is monotonically increasing there exists a threshold t^* so that it is optimal to commit to offering trial for all types $\theta < t^*$. ■

Our results are in contrast to standard models that do not incorporate a demand rotation effect (i.e., where $\alpha = 0$). In such models when $\underline{\theta} \geq 0$ the firm is indifferent to any trial threshold, whereas in our model the firm may strictly prefer no trial at all (e.g., when $c = 0$). Furthermore, the optimal policy with commitment, in which low types offer a trial and high types do not, is exactly the opposite of what occurs in equilibrium without commitment. Therefore, whether the seller can credibly commit to a product trial policy plays a crucial role in the types of product trial outcomes one ought to expect.

Example 3 *We now reconsider Example 2 in which $c = \frac{1}{5}$ and the firm lacked commitment power. It was found that types offer a trial if and only if they exceed $t^* \approx 0.955$. However, from the left panel of Figure 2 we see that $C(0) > 0$ and thus a firm with commitment power never offers trial.*

By contrast, when $c = \frac{1}{2}$ we have

$$\begin{aligned} C(0) &= \alpha\eta - c - \pi(0) \\ &= \frac{1}{2} \times 1 - \frac{1}{2} - \pi(0) \\ &= -\pi(0) \\ &< 0, \end{aligned}$$

and so the firm with commitment power will sometimes offer a product trial. Specifically, a calculation shows the firm will commit to offering a trial only for types below $t^* \approx 0.076$.

VI CONCLUSION

When a firm allows consumers to sample its product prior to purchase, it not only discloses the product's quality but also allows each consumer to privately learn their idiosyncratic match value. In equilibrium the seller faces a tradeoff: by allowing a pre-purchase trial the seller yields information rents to consumers, but by not allowing the trial the seller implicitly signals that the product is of low quality. The yielding of information rents can be interpreted as an endogenous disclosure cost, and as such negates the standard unraveling logic leading to full disclosure. Instead, we find that in equilibrium the seller only offers a product trial if their quality is sufficiently high, and sometimes not at all. The model predicts fewer trials for products in which the horizontal component of value is more important than vertical quality, as well as for products with lower gains from trade. Interestingly, a seller that can commit to a product trial policy acts in a manner that is the exact

opposite of equilibrium behavior, as the seller only offers a trial if her quality is sufficiently *low*. Finally, a policy mandating product trials leads to an increase in consumer surplus, a decrease in producer surplus, and has an ambiguous impact on welfare. Because consumers are fully rational, a trial does not help them directly by dispelling confusion but instead by endowing them with private information, for which the seller is forced to compensate in equilibrium. In terms of welfare, a product trial reduces the quantity sold both to consumers with values above and below marginal cost, which may on net increase or decrease total gains from trade.

The model easily admits several natural extensions. For example, suppose that a proportion of consumers are naive in that they do not update their beliefs when a firm fails to offer a product trial (e.g., Brown et al. 2012; Jin et al. 2015). The effect of the presence of these consumers is to lower the benefit $B(t)$ of offering the trial, as the upward demand shift occurs only for a subset of buyers. On the other hand, the demand rotation cost $C(t)$ would not be affected by naive consumers since not offering the trial has no impact on a consumer's belief about ε , as the firm has no private information about it. Thus, using logic similar to that in Proposition 2, the presence of naive consumers would result in fewer product trials.

As another extension consider the impact of third party reviewers in our setting. Suppose the reviewers provide a public signal σ about quality θ , inducing buyers to expect the seller's quality to be distributed according to $F(\theta|\sigma)$. Taking this distribution as given, from this point the seller's decision of whether to offer a trial is characterized exactly by the current model, with some resulting equilibrium threshold $t(\sigma)$. One can then evaluate the effects of an intermediary by studying the distribution of signals σ it provides, and thereby the distribution of thresholds $t(\sigma)$ that it induces.

Alternatively, suppose reviewers provide consumers with private information about match ε . In this case, the seller expects consumers to already have some private information about match, and therefore the cost of offering a product trial and revealing the rest of this private information is reduced. This would imply that third party reviewers would induce more product trials, acting as a complement rather than a substitute.

References

- Anderson, S. P., & Renault, R. (2006). Advertising content. *American Economic Review*, *96*(1), 93–113.
- Andrabi, T., Das, J., & Khwaja, A. I. (2017). Report cards: The impact of providing school and child test scores on educational markets. *American Economic Review*, *107*(6), 1535–63.
- Ben-Shahar, O., & Posner, E. A. (2011). The right to withdraw in contract law. *The Journal of Legal Studies*, *40*(1), 115–148.
- Board, O. (2009). Competition and disclosure. *The Journal of Industrial Economics*, *57*(1), 197–213.
- Bollinger, B., Leslie, P., & Sorensen, A. (2011). Calorie posting in chain restaurants. *American Economic Journal: Economic Policy*, *3*(1), 91–128.
- Brown, A. L., Camerer, C. F., & Lovallo, D. (2012). To review or not to review? limited strategic thinking at the movie box office. *American Economic Journal: Microeconomics*, *4*(2), 1–26.
- Chesnokova, T. (2007). Return policies, market outcomes, and consumer welfare. *Canadian Journal of Economics/Revue canadienne d'économique*, *40*(1), 296–316.

- Dranove, D., & Jin, G. Z. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4), 935–963.
- Feng Lu, S. (2012). Multitasking, information disclosure, and product quality: Evidence from nursing homes. *Journal of Economics & Management Strategy*, 21(3), 673–705.
- Fishman, M. J., & Hagerty, K. M. (2003). Mandatory versus voluntary disclosure in markets with informed and uninformed customers. *Journal of Law, Economics, and organization*, 19(1), 45–63.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3), 461–483.
- Grossman, S. J., & Hart, O. D. (1980). Disclosure laws and takeover bids. *The Journal of Finance*, 35(2), 323–334.
- Guo, L., & Zhao, Y. (2009). Voluntary quality disclosure and market interaction. *Marketing Science*, 28(3), 488–501.
- Hotz, V. J., & Xiao, M. (2013). Strategic information disclosure: The case of multiattribute products with heterogeneous consumers. *Economic Inquiry*, 51(1), 865–881.
- Inderst, R., & Ottaviani, M. (2013). Sales talk, cancellation terms and the role of consumer protection. *Review of Economic Studies*, 80(3), 1002–1026.
- Jin, G. Z., Luca, M., & Martin, D. (2015). *Is no news (perceived as) bad news? an experimental investigation of information disclosure* (Tech. Rep.). National Bureau of Economic Research.
- Johnson, J. P., & Myatt, D. P. (2006). On the simple economics of advertising, marketing, and product design. *American Economic Review*, 96(3), 756–784.

- Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *The American Economic Review*, *101*(6), 2590–2615.
- Krähmer, D., & Strausz, R. (2015). Optimal sales contracts with withdrawal rights. *The Review of Economic Studies*, *82*(2), 762–790.
- Levin, D., Peck, J., & Ye, L. (2009). Quality disclosure and competition. *The Journal of Industrial Economics*, *57*(1), 167–196.
- Lewis, T. R., & Sappington, D. E. (1994). Supplying information to facilitate price discrimination. *International Economic Review*, 309–327.
- Lizzeri, A. (1999). Information revelation and certification intermediaries. *The RAND Journal of Economics*, 214–231.
- Mahoney, P. G. (1995). Mandatory disclosure as a solution to agency problems. *The University of Chicago Law Review*, *62*(3), 1047–1112.
- Milgrom. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 380–391.
- Milgrom, & Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, *94*(4), 796–821.
- Petrikaitė, V. (2018). A search model of costly product returns. *International Journal of Industrial Organization*, *58*, 236–251.
- Polinsky, A. M., & Shavell, S. (2010). Mandatory versus voluntary disclosure of product risks. *The Journal of Law, Economics, & Organization*, *28*(2), 360–379.
- Rekaiti, P., & Van den Bergh, R. (2000). Cooling-off periods in the consumer laws of the ec

member states. a comparative law and economics approach. *Journal of Consumer Policy*,
23(4), 371–408.