WORKING PAPERS

N° 1695

November 2025

“Resale Price Maintenance and Consumer Search”

Andrew Rhodes and Yang Yang

Toulouse School of Economics

Resale Price Maintenance and Consumer Search*

Andrew Rhodes† Yang Yang‡

November 2025

Abstract

We provide a novel pro-competitive rationale for resale price maintenance (RPM). We consider a model where some consumers are fully informed about downstream prices while other consumers are not. When an upstream manufacturer imposes a floor on downstream prices, this qualitatively changes downstream competition—influencing not just the level, but also the dispersion, of prices. The manufacturer optimally imposes a price floor which just eliminates all downstream price dispersion, and this leads to both higher (aggregate) consumer surplus and higher total welfare as compared to the case without RPM.

Keywords: Resale price maintenance, consumer search, price dispersion

*We are grateful to Song Lin, David Ronayne, Chris Wilson, Jidong Zhou, and participants at the 14th Consumer Search Workshop (Hong Kong). We also thank Mingjun Gao for excellent research assistance. Rhodes acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program (grant ANR-17-EURE-0010).

†Toulouse School of Economics. andrew.rhodes@tse-fr.eu

‡School of Economics and Management, Tsinghua University. yangyangedu@gmail.com

1 Introduction

Manufacturers often distribute their products via downstream retailers, and may use resale price maintenance (RPM) to prevent downstream prices from becoming too low. For example, the fashion companies Chloé, Gucci, and Loewe placed restrictions on the maximum discounts that downstream retailers in Europe could apply; the electronics company Philips requested low-price online retailers in France to raise their prices, and used retaliatory measures against those that refused; the leather apparel manufacturer Leegin imposed a minimum resale price in the U.S., and refused to supply retailers that did not adhere to it.[1] Although the treatment of RPM varies across jurisdictions, antitrust authorities usually weigh up lower downstream competition against possible efficiency gains.[2] For instance, it has been argued that RPM could improve efficiency if it leads to higher margins which incentivize retailers to offer better pre-sales services.[3]

In this paper we challenge the view that RPM necessarily lowers downstream competition. Our starting point is that many downstream markets exhibit price dispersion, i.e., different firms charge different prices for the same product. Following the consumer search literature, we rationalize this by the co-existence of consumers with different amounts of information. We show that while RPM forces downstream firms with the very lowest prices to raise them, it also encourages firms with somewhat higher prices to lower them. As a result, RPM affects not just the level but also the dispersion of downstream prices, and so RPM can be *pro-competitive*, meaning it can raise manufacturer profit, consumer surplus, and total welfare. We emphasize that this happens despite our model abstracting away from any possible efficiency benefits due to RPM.

In more detail, in Section 2 we introduce our framework, which builds on the classic model of sales due to Varian (1980). Specifically, two or more downstream firms sell a homogeneous product, and some consumers are well informed about prices ("shoppers") while other consumers are only able to buy from one randomly chosen firm ("non-shoppers"). To this otherwise standard model we introduce an upstream manufacturer, who chooses both a wholesale price at which downstream firms can buy the product, and a floor below which these downstream firms are not allowed to price.

We begin by solving the model in Section 3 for an exogenously given wholesale price.

[1]The European Commission fined Chloé, Gucci, and Loewe a combined 157 million euros, and fined Philips 30 million euros, for these and related practices; see here and here for more details. Leegin was initially required to pay a fine, but this was overturned by the Supreme Court; see here for more details.

[2]In the U.S., following the Leegin case referred to in the previous footnote, "a rule-of-reason" approach was adopted towards RPM. In the European Union, RPM is considered a "hardcore" restriction of competition, but an efficiency defense is possible under Article 101(3); see here for more details.

[3]See paragraph 197 of the European Guidelines on Vertical Restraints here for further examples.

The manufacturer's choice of price floor affects downstream competition in a very natural way. First, when the price floor is sufficiently low, downstream firms randomize over prices using a continuous distribution, exactly as in Varian (1980). The price distribution resolves downstream firms' tension between pricing low to attract shoppers and pricing high to exploit non-shoppers. Second, when the price floor is sufficiently high, downstream firms' competition for shoppers leads to a pure strategy equilibrium where they all price exactly at the floor. Third, when the price floor is intermediate, downstream firms employ a mixed strategy with some novel features: it has a mass point at the bottom, then a gap, and then a continuous part at higher prices. Intuitively, competition for shoppers induces downstream firms to price at the floor with positive probability. This makes it unattractive to charge prices just above the floor, leading to a gap; at sufficiently high prices downstream firms face the same trade-off as in Varian (1980), and hence they draw prices from a continuous distribution. As the price floor increases, the size of the mass point increases while the interval of prices with a continuous distribution shrinks—and so the degree of price dispersion in the downstream market also changes.

Having studied downstream firms' pricing problem, we then solve for the manufacturer's optimal choice of price floor, and examine the welfare impact of RPM. This is a complicated problem, given that downstream prices are drawn from a distribution which depends on the price floor in a complex way. Following the approach of Armstrong and Vickers (2019) (in a different context), we reformulate our problem as one in which downstream firms choose a per-consumer profit level rather than a price. We show that when the price floor is intermediate, the distribution of per-consumer profit improves in the sense of second order stochastic dominance as the price floor increases. We also show that both the manufacturer and consumers are "risk averse" with respect to downstream firms' profit. It then follows that manufacturer profit and total consumer surplus are both quasiconcave and maximized at the *same* price floor, namely the one which just eliminates all downstream price dispersion. Hence, contrary to the prevailing view in the literature, RPM is pro-competitive and raises consumer surplus even absent any efficiency justification. Moreover, we show that RPM also raises total welfare. In addition, these beneficial effects of RPM are larger in markets with a larger number of downstream firms.

Finally, Section 4 extends the above results to the case where the manufacturer chooses both a price floor and a wholesale price. When the demand curve belongs to the "constant curvature class" (e.g., it is linear) we show that the manufacturer optimally chooses exactly the same wholesale price irrespective of whether it uses RPM; as a result, for this class of demands, all the results discussed above apply immediately. When the demand curve does not belong to this class, analytical results are unfortunately not available, but we show via numerical examples that RPM can continue to raise consumer surplus.

Related Literature. The existing literature shows that RPM can be either pro- or anti-competitive, depending on the context (see, e.g., Rey and Vergé, 2008, for a survey). On the one hand, RPM can be anti-competitive because it relaxes interbrand and intrabrand competition (Rey and Vergé, 2010); discourages downstream firms from stocking new products and thus deters upstream entry (Asker and Bar-Isaac, 2014); and facilitates collusion by making it easier to detect price deviations (Jullien and Rey, 2007). On the other hand, RPM can be pro-competitive because it incentivizes retailers to invest in pre-sales services that are valuable to consumers, and prevents free-riding among these retailers (see, e.g., Telser, 1960; Marvel and McCafferty, 1984; Mathewson and Winter, 1984). Different from these papers, we abstract from any investment or free-riding effect; we show that RPM is pro-competitive in a search setting where some consumers have better price information than others.

Our paper contributes to an emerging literature on consumer search and vertical relations. Janssen and Shelegia (2015) introduce an upstream manufacturer into a Stahl (1989) setting where shoppers buy at the lowest price while non-shoppers can search for price quotes at a cost. They show that when non-shoppers do not observe the upstream manufacturer's wholesale price, this allows the manufacturer to squeeze downstream firms, leading consumers to pay higher final prices. Garcia, Honda, and Janssen (2017) consider a related setting, but assume there are competing upstream manufacturers and that both consumers and downstream firms need to search. They show that this double layer of search frictions can lead to a bimodal price distribution at both the upstream and downstream levels.[4] Garcia and Janssen (2018) show that a monopoly manufacturer finds it optimal to offer different downstream firms different wholesale prices, even when they are completely symmetric, so as to encourage consumers to search and thus put downward pressure on downstream margins. More generally, Janssen and Shelegia (2020) stress the importance of whether consumers blame an unexpected price on the downstream firm that charged it or on the manufacturer, since this affects a consumer's belief about other downstream prices and hence influences her decision of whether to search again.

Other papers consider the interaction between consumer search and non-RPM vertical restraints. Asker and Bar-Isaac (2020) show that minimum advertised prices restrict consumer information, allowing an upstream manufacturer to better screen consumers and extract surplus from those with high valuations and high search costs. Lubensky (2017)

[4]Our model also predicts a novel price distribution at the downstream level, with both a mass point at the bottom and a gap. This differs from usual symmetric environments, where firms either mix using the same continuous distribution, or some firms employ a mass point but it is at the top of the price distribution. An exception is Myatt and Ronayne (2025), who consider a model where firms commit to a list price which caps their final price, which leads to an asymmetric pure strategy equilibrium where one firm charges less than the others, even when firms are all symmetric ex ante.

shows that a manufacturer can use non-binding recommended retail prices to inform consumers about its production cost, and steer their search in a way that is beneficial to both itself and to the consumers. However Janssen and Reshidi (2022) show that a policy which forces some sales to occur at a manufacturer's recommended retail price enables the manufacturer to commit to offering different downstream retailers different wholesale prices, to the detriment of final consumers.

Most closely related to our paper is the small literature on consumer search with price restraints. Armstrong, Vickers, and Zhou (2009) show that a price ceiling reduces consumers' incentives to search, potentially raising prices and reducing consumer surplus. Hamilton (1990) considers the search model of Salop and Stiglitz (1977), where consumers observe the price set by each firm but must pay a cost to learn which firm charges which price. He shows that RPM can induce a switch from a two-price equilibrium to a one-price equilibrium, and that under certain conditions this can raise manufacturer profit. However he does not analyze the welfare impact of RPM. Meanwhile, in independent and concurrent work, Baye, Kovenock, and de Vries (2025) study RPM in an asymmetric duopoly Varian model. They provide a condition on demand such that RPM raises manufacturer profit, benefits non-shoppers, but harms shoppers. Different from them, we focus on a symmetric model with a general number of firms, and we show that under weak conditions manufacturer-optimal RPM maximizes both aggregate consumer surplus and total welfare. We also show that the welfare impact of RPM is greater in markets with more downstream firms.[5]

Finally, while there is little empirical research on RPM, there is some recent suggestive evidence in favor of our findings. Williams (2024) studies the impact of Fixed Book Price policies in Europe. These policies aim to prevent discounters from charging very low prices for books, and hence play a similar role to RPM in our model. Interestingly, he finds no evidence that such policies raise average book prices. Moreover, consistent with our model, he finds that they lead to higher book sales. Meanwhile Xia (2024) exploits a high-profile antitrust case in China, which banned a pharmaceutical company from practicing RPM. He finds that RPM leads to lower and less dispersed downstream prices and, as in the previous paper, it also leads to higher sales.

The rest of the paper proceeds as follows. Section 2 introduces our model. Section 3 takes the wholesale price as given, and solves for the manufacturer-optimal price floor and considers its impact on consumers and total welfare. Section 4 extends those results to the case where the manufacturer also chooses the wholesale price. Section 5 concludes.

[5] In a setting where firms compete over both price and quality, Yang (2024) shows that an endogenous price floor may emerge if consumers interpret sufficiently low prices as a bad signal about product quality.

2 Model

An upstream manufacturer supplies a homogeneous product to $n \geq 2$ downstream firms. The manufacturer charges downstream firms a wholesale price w, and also sets a price floor p_F below which downstream firms are not allowed to price. The manufacturer's marginal cost is normalized to zero; downstream firms also incur no additional costs beyond the wholesale price w. There is a unit mass of consumers who are interested in buying the product. A fraction $\lambda \in (0,1)$ of consumers are shoppers, who buy from whichever firm has the lowest price (and choose a firm randomly in case of ties). The remaining $1-\lambda$ of consumers are non-shoppers, who buy from one randomly chosen firm. Conditional on incurring a price p, a consumer buys $q(p)$ units of the product. For all p satisfying $q(p) > 0$, we assume that $q(p)$ is twice differentiable, that $q'(p) < 0$, and also that $q(p)$ is log-concave.

The timing is as follows. First, the manufacturer sets w and p_F. Second, downstream firms simultaneously set prices, ensuring that they weakly exceed p_F. We focus on symmetric equilibria, where all firms draw their price from the same (potentially degenerate) price distribution. Third, shoppers buy from a firm with the lowest price, while non-shoppers buy from a randomly chosen firm.

3 Analysis with an Exogenous Wholesale Price

In this section we assume that the wholesale price w is exogenously fixed, such that the manufacturer only chooses the price floor p_F. (In Section 4 we allow the manufacturer to choose both w and p_F.) Denote by $\pi(p) = q(p)(p-w)$ a firm's per-consumer profit when it sells at price p, and note that it is quasiconcave in p due to the log-concavity of $q(p)$. Let $p^m = \arg\max_p \pi(p)$ be the monopoly price, and let $\pi^m = \pi(p^m)$ be the associated per-consumer profit. To rule out uninteresting cases we assume that $q(w) > 0$, which implies that $\pi^m > 0$.

3.1 Benchmark without RPM

It is instructive to first consider the case where the manufacturer *cannot* use RPM. (Equivalently, the price floor is so low that it is not binding, e.g., $p_F \leq w$.) In this case we have the classic Varian (1980) model of sales. It is well known that there is no (symmetric) pure strategy equilibrium: downstream firms must charge more than w since they have market power over non-shoppers, but if all downstream firms were to hypothetically choose the same price $p > w$, each one could do better by unilaterally slightly undercutting p and winning all the shoppers. The same logic implies that downstream firms' symmetric mixed strategy price distribution also has no mass points. Letting $F(p)$ denote this mixed strategy price distribution, we have:

$$\frac{1-\lambda}{n}\pi^m = \pi(p)\left[\frac{1-\lambda}{n} + \lambda[1-F(p)]^{n-1}\right].$$

In particular, a firm can charge p^m and earn monopoly profits from its share $(1-\lambda)/n$ of the non-shoppers. Alternatively, if it charges $p < p^m$, it earns $\pi(p)$ from its share of the non-shoppers, as well as from the shoppers provided each of its $n-1$ competitors has a higher price—which occurs with probability $[1-F(p)]^{n-1}$. The price distribution $F(p)$ then makes firms indifferent between charging p^m or any lower price in the support of the mixed strategy. Using the above indifference equation we have the standard result:

Lemma 1 (Varian, 1980)**.** *Absent RPM there exists a unique symmetric equilibrium in which firms draw their price from* $[\underline{p}, p^m]$ *according to*

$$F(p) = 1 - \left[\frac{1-\lambda}{\lambda n}\left(\frac{\pi^m}{\pi(p)} - 1\right)\right]^{\frac{1}{n-1}},$$

with $\underline{p}$ *determined by*

$$(1-\lambda)\pi^m = \pi(\underline{p})[1+\lambda(n-1)].$$

Notice that $F(p)$ is continuous and strictly increasing on its support, and so the equilibrium price distribution has no atoms (as explained above) and also no "gaps".

3.2 Equilibrium Downstream Pricing with RPM

Henceforth suppose the manufacturer is able to use RPM. In this subsection we solve for equilibrium downstream prices for any downstream price floor p_F.

It is straightforward to see that if $p_F \leq \underline{p}$ then in equilibrium firms use the standard Varian (1980) price distribution from Lemma 1. In the more interesting case where $p_F > \underline{p}$, equilibrium depends on how p_F compares to p_F^*, where $p_F^* < p^m$ solves

$$\frac{\pi(p_F^*)}{n} = \frac{1-\lambda}{n}\pi^m. \tag{1}$$

To interpret this equation, note that the left-hand side is a firm's profit if it and all other firms charge p_F^* and hence split the market equally, while the right-hand side is a firm's profit if it charges p^m and sells only to its share of the non-shoppers.

Lemma 2. *Suppose the manufacturer imposes a price floor* $p_F \geq p_F^*$*. There is a unique symmetric equilibrium where each downstream firm charges* p_F*.*

Intuitively when $p_F \geq p_F^*$ the price floor is sufficiently high that downstream firms prefer to charge p_F and compete for shoppers, rather than set a higher price to potentially exploit their non-shoppers. Hence in equilibrium downstream firms play a pure strategy equilibrium and all charge p_F.

Lemma 3. *Suppose the manufacturer imposes a price floor $p_F \in (\underline{p}, p_F^*)$. There is a unique symmetric equilibrium where downstream firms draw their price from distribution*

$$G(p) = \begin{cases} 0 & \text{for } p < p_F \\ m & \text{for } p_F \leq p \leq \underline{p}' \\ F(p) & \text{for } p > \underline{p}' \end{cases},$$

where $m \in (0,1)$ and $\underline{p}' \in (p_F, p^m)$ jointly and uniquely solve $m = F(\underline{p}')$ and

$$\pi(p_F)\left[\lambda \frac{1-(1-m)^n}{mn} + \frac{1-\lambda}{n}\right] = \frac{1-\lambda}{n}\pi^m. \tag{2}$$

When the price floor p_F is between $\underline{p}$ and p_F^*, downstream firms use a mixed strategy price distribution $G(p)$. Moreover $G(p)$ has several interesting properties: it has a mass point of size m at the price floor, then a gap up to a critical price $\underline{p}'$, and then for all prices $p \geq \underline{p}'$ it coincides with the Varian (1980) distribution $F(p)$ from Lemma 1.

We now explain the construction of $G(p)$ in Lemma 3 in more detail. Intuitively, the price distribution cannot be atomless otherwise firms would use the Varian price distribution from earlier—but some prices in the support of that distribution violate the price floor. At the same time, the atom must be at p_F, since any mass point at a higher price could be profitably undercut. Let m denote the size of the atom at price p_F. Notice that if a downstream firm charges p_F the expected number of shoppers that it sells to is

$$\sum_{k=0}^{n-1} \binom{n-1}{k} \frac{m^k(1-m)^{n-k-1}}{k+1} = \frac{1-(1-m)^n}{mn},$$

because with probability $\binom{n-1}{k} m^k (1-m)^{n-k-1}$ exactly $k = 0, 1, \ldots, n-1$ other firms also charge p_F, in which case each firm gets a share $1/(k+1)$ of the shoppers. It then follows that the left-hand side of (2) is a firm's expected profit from charging p_F. There must be a gap directly above p_F: due to the mass point $m > 0$, slightly increasing price above p_F leads to a discrete drop in the probability of selling to shoppers, and only a small increase in the profit earned on non-shoppers. For prices above the gap, firms face the standard trade-off between competing for shoppers or exploiting non-shoppers, and hence we must have the same distribution as in the Varian (1980) model, i.e., $G(p) = F(p)$. Since $\underline{p}'$ is the supremum of the prices in the gap, by definition the size of the mass point

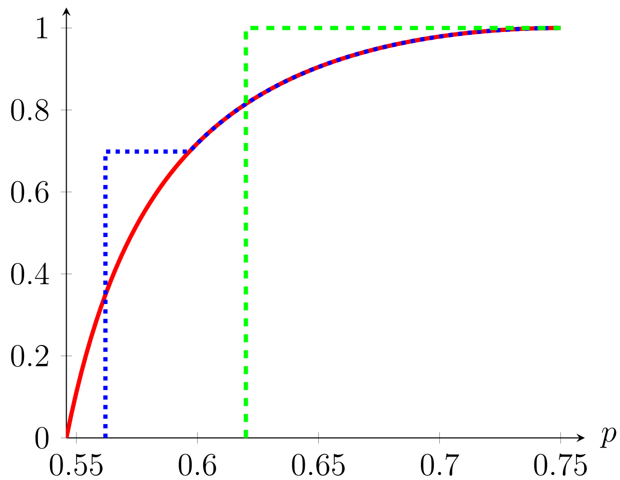

Figure 1: The equilibrium price distribution for a low price floor (red solid curve), intermediate price floor (blue dotted curve), and high price floor (green dashed curve).

must satisfy $m = F(\underline{p}')$. Finally, to ensure that downstream firms are indifferent, their expected profit from charging p_F must equal their expected profit from charging any price on $[\underline{p}', p^m]$, including p^m, and this is ensured by equation (2).

It is straightforward to show that changes in the price floor affect equilibrium pricing in the following way. First, as $p_F \downarrow \underline{p}$, we converge to the Varian price distribution because the size of the mass point tends to zero and $\underline{p}'$ tends to $\underline{p}$. Second, as p_F increases, the mass point m becomes larger and the range of prices $[\underline{p}', p^m]$ over which firms mix continuously shrinks. Third, as $p_F \uparrow p_F^*$, we converge to the pure strategy pricing equilibrium in Lemma 2 because the size of the mass point tends to one and $\underline{p}'$ tends to p^m. Figure 1 illustrates the impact of a price floor on the equilibrium price distribution when $n = 2$, $\lambda = 1/2$, $w = 1/2$, and $q(p) = 1 - p$. The red solid curve depicts the case where $p_F \leq \underline{p}$, i.e., the price floor is sufficiently low that pricing follows Varian (1980). The blue dotted curve depicts the case where $p_F \in (\underline{p}, p_F^*)$, i.e., the price floor is intermediate so the price distribution has a mass point, a gap, and a continuous part. The green dashed curve depicts the case where $p_F \geq p_F^*$, i.e., the price floor is high enough that all firms charge it with probability one.

3.3 Manufacturer-Optimal RPM

We now solve for the manufacturer's optimal choice of price floor p_F. Note that for any fixed wholesale price $w > 0$ the manufacturer seeks to maximize its expected output.

It is straightforward to see from Lemma 2 that expected output is strictly decreasing in $p_F \geq p_F^*$ because downstream retailers all price exactly at the price floor. Now consider the more interesting case where $p_F \in (\underline{p}, p_F^*)$. It is useful to follow Armstrong and Vickers

(2019) and think of downstream firms as choosing a per-consumer profit level π rather than a price p. Denote by $\tilde{G}(\pi)$ a firm's profit distribution, and let $\pi_F \equiv \pi(p_F)$ and $\underline{\pi}' \equiv \pi(\underline{p}')$. It follows from Lemma 3 that $\tilde{G}(\pi)$ has a mass point of size m at π_F, and for $\pi > \underline{\pi}'$ it satisfies

$$\tilde{G}(\pi) = 1 - \left[\frac{1-\lambda}{\lambda n}\left(\frac{\pi^m}{\pi} - 1\right)\right]^{\frac{1}{n-1}}. \tag{3}$$

Now consider the market-wide distribution of per-consumer profit, which we denote by $PD(\pi)$. A shopper buys from the firm with the smallest π, given that $\pi(p)$ is strictly increasing in $p < p^m$. A non-shopper buys from a randomly drawn firm. Hence $PD(\pi)$ has a mass point of size $\lambda[1-(1-m)^n] + (1-\lambda)m$ at π_F, and for $\pi > \underline{\pi}'$ it satisfies

$$PD(\pi) = \lambda\left[1 - [1-\tilde{G}(\pi)]^n\right] + (1-\lambda)\tilde{G}(\pi). \tag{4}$$

We can then prove the following result:

Lemma 4. *Consider a price floor $p_F \in (\underline{p}, p_F^*)$. An increase in p_F:*
(i) Leaves expected per-consumer profit unchanged.
(ii) Improves the distribution of $PD(\pi)$ in the sense of second order stochastic dominance.

An increase in the price floor p_F induces a mean-preserving contraction of the market-wide profit distribution. Intuitively, p^m is always in the support of firms' price distributions, so firms obtain a profit of $(1-\lambda)\pi^m/n$ which is independent of p_F; this explains why changes in p_F do not affect expected profit. Moreover, as p_F increases it becomes more attractive for firms to settle at the price floor, leading to a higher mass point m; this reduces the dispersion of $PD(\pi)$ in the sense of second order stochastic dominance.

Let $Q(\pi)$ denote output as a function of per-consumer profit. We can then write expected manufacturer profit as $w \times \mathbb{E}_{PD}[Q(\pi)]$. It is straightforward to show that log-concavity of demand implies that $Q(\pi)$ is a strictly concave function. Manufacturer profit must then be strictly increasing in $p_F \in (\underline{p}, p_F^*)$, given that higher p_F on this interval induces a mean-preserving contraction in $PD(\pi)$. Since we argued earlier that manufacturer profit is also decreasing in $p_F \geq p_F^*$, the following is immediate:

Proposition 1. *Suppose w is fixed. The manufacturer optimally chooses a price floor $p_F = p_F^*$, where p_F^* is defined by equation* (1)*. All downstream firms then charge p_F^*.*

Because variability in downstream prices reduces the expected quantity that the manufacturer sells, it sets the price floor at p_F^*—just high enough to ensure there is no downstream price dispersion. Recall from equation (1) that $\pi(p_F^*) = (1-\lambda)\pi^m$, and so the manufacturer's optimal price floor is independent of n but is decreasing in λ. Intuitively, when there are more shoppers firms are more willing to price at the floor, and therefore

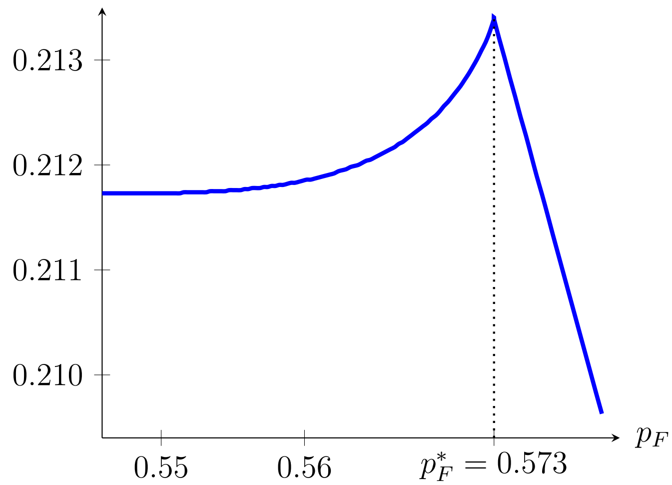

Figure 2: The effect of a price floor on manufacturer profit.

the manufacturer chooses a lower floor. Figure 2 depicts the effect of a higher price floor on manufacturer profit in our running example where $n = 2$, $\lambda = 1/2$, $w = 1/2$, and $q(p) = 1 - p$.

3.4 The Welfare Impact of RPM

We now examine the impact of a price floor on consumer surplus and total welfare, and then link this to the manufacturer's optimal choice of RPM from Proposition 1.

Proposition 2. *Suppose w is fixed.*
(i) Consumer surplus and total welfare are quasiconcave in p_F and maximized at $p_F = p_F^*$.
(ii) The surplus of non-shoppers is quasiconcave in p_F and maximized at $p_F = p_F^*$.
The surplus of shoppers decreases in p_F provided $\pi'(p)$ is log-concave in $p < p^m$.

Proposition 2 shows, perhaps surprisingly, that consumer and total welfare are both maximized at the price floor chosen by the manufacturer. Intuitively, this is because consumers also dislike variation in firms' profit. More precisely, we need to distinguish between whether p_F is above or below p_F^*. The case $p_F \geq p_F^*$ is straightforward: downstream firms charge p_F for sure, and so a higher p_F is worse for both types of consumers and also reduces total welfare.[6] The case $p_F \in (\underline{p}, p_F^*)$ is more interesting. Let $V(\pi)$ denote a consumer's level of surplus given that a firm earns profit π from selling to her. Recalling the definition of $PD(\pi)$, we can write expected consumer surplus as $\mathbb{E}_{PD}[V(\pi)]$. Armstrong and Vickers (2019) prove that log-concavity of $q(p)$ implies that $V(\pi)$ is concave—

[6] In more detail, notice that downstream firms' profit is actually increasing in $p_F \in (p_F^*, p^m)$, then decreasing in $p_F > p^m$. However the sum of consumer surplus, manufacturer profit, and downstream profit for $p_F > p_F^*$ is $\int_p Q(z)dz + pQ(p)$, which is always strictly decreasing in p.

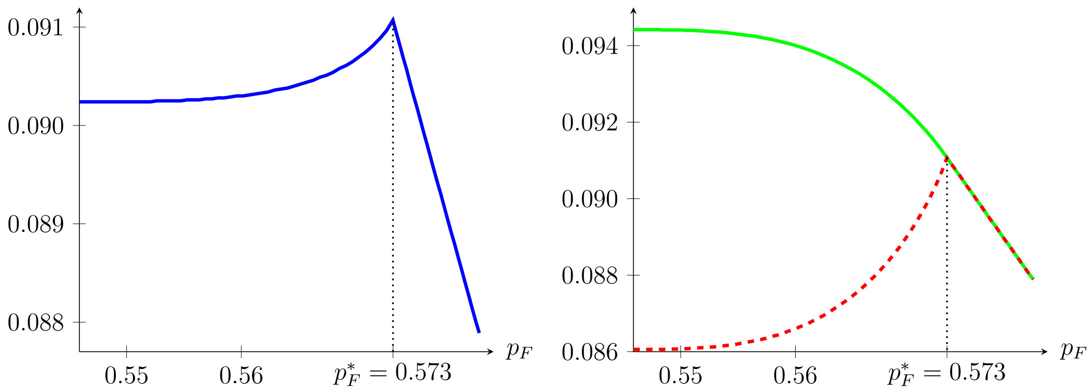

(a) Total consumer surplus (b) Surplus of shoppers and non-shoppers

Figure 3: The effect of a price floor on total consumer surplus (left panel), and on the surpluses of shoppers and non-shoppers (green and red curves respectively, right panel).

consumers are "risk averse" with respect to firm profit. Given the result in Lemma 4 that a higher price floor improves $PD(\pi)$ in the sense of second order stochastic dominance, it follows that expected consumer surplus is increasing in $p_F \in (\underline{p}, p_F^*)$. Since downstream industry profit is constant in $p_F \in (\underline{p}, p_F^*)$ while manufacturer profit is increasing in p_F over this interval, total welfare is also increasing in p_F over this interval.

Proposition 2 also shows that an increase in the price floor affects different consumers differently. In particular, one can show that a higher price floor changes the intensity of competition for each type: as $p_F \leq p_F^*$ increases firms place more mass at the price floor, causing them to earn *more* profit from shoppers but *less* profit from non-shoppers. Non-shoppers then benefit from a higher price floor for two reasons—firms earn less profit from them, and that profit is also less variable. Shoppers, on the other hand, benefit from the reduction in profit variability, but are worse off because firms extract more surplus from them. When $\pi'(p)$ is log-concave the latter effect dominates, and so a higher price floor harms shoppers. It is easy to prove that $\pi'(p)$ is log-concave when $q(p)$ has constant curvature (which is true for, e.g., demands that are linear or exponential).

Figure 3 depicts the effect of a higher price floor on total consumer surplus (left panel), and the surpluses of shoppers and non-shoppers (right panel), in our running example where $n = 2$, $\lambda = 1/2$, $w = 1/2$, and $q(p) = 1 - p$. Note that demand $q(p) = 1 - p$ has constant curvature, so the condition in the proposition is satisfied and therefore shoppers' surplus decreases monotonically in the price floor.

The following is a straightforward corollary of our results so far:

Corollary 1. *Suppose w is fixed. Allowing the manufacturer to use RPM has no effect on downstream firms' profits, but it raises aggregate consumer surplus and total welfare, benefits non-shoppers, and harms shoppers provided $\pi'(p)$ is log-concave.*

Hence a ban on RPM would be detrimental to consumer surplus and total welfare, would harm non-shoppers, but would benefit shoppers under a regularity condition. Moreover, we can also show that these effects are larger in markets with more downstream firms:

Proposition 3. *Suppose w is fixed. The effects of RPM identified in Corollary 1 are larger in markets with larger n.*

This result can be understood as follows. First, as shown earlier, under RPM the optimal price floor is p_F^*, and downstream firms charge this price with probability one. Since p_F^* is independent of n, so is manufacturer profit and the surplus of each consumer type (and hence also aggregate consumer surplus). Second, though, as we show in the proof of Proposition 3, absent RPM outcomes do depend on n. In particular, it is already well-known from the literature that as n increases in a Varian (1980) model, firms adopt more extreme pricing strategies—mainly charging high prices close to p^m to exploit their non-shoppers, but occasionally setting very low prices to compete for shoppers. We demonstrate that greater n leads to a worse per-consumer profit distribution in the sense of second order stochastic dominance. Since we showed earlier that consumer surplus $V(\pi)$ and output $Q(\pi)$ are concave in π, it follows that in the Varian model expected consumer surplus and expected output both decrease in n. We also demonstrate in the proof that the tendency for firms to choose more extreme pricing strategies with higher n is detrimental to non-shoppers (who get a random draw from a distribution with generally high prices) but beneficial to shoppers (because they benefit from the fact that each firm occasionally charges a very low price). Third, it then follows that the beneficial impact of RPM on total consumer surplus, manufacturer profit (and hence also total welfare), and non-shoppers' surplus, as well as the negative impact of RPM on shoppers' surplus, are all larger when there are more firms in the industry.

Remark on asymmetric equilibria. In the above we have focused on symmetric equilibria. However, it is well known that even symmetric models of sales have asymmetric equilibria when $n > 2$. In such equilibria each firm earns the same expected profit $(1-\lambda)\pi^m/n$ as in the symmetric equilibrium in Lemma 1, and prices and profit levels are dispersed. It is straightforward to show that if $p_F = p_F^*$ then the symmetric equilibrium studied above is the *unique* equilibrium. Hence, even allowing for asymmetric equilibria, a price floor at $p_F = p_F^*$ still reduces profit volatility and so leads to higher consumer surplus and manufacturer profit (and thus also total welfare) compared to the case with no RPM.

4 Analysis with an Endogenous Wholesale Price

In this section we allow the manufacturer to choose both a price floor p_F and a wholesale price w. Denote by $\pi(p;w) = (p-w)q(p)$ a firm's per-consumer profit when it sells at price p and faces a wholesale price w, and note that as usual it is quasiconcave in p. Let $p^m(w) = \arg\max_p \pi(p;w)$ be the monopoly price, and let $\pi^m(w) = \pi(p^m(w);w)$ be the associated per-consumer profit.

Using earlier work we can immediately make two observations. First, absent RPM, and for given wholesale price w, downstream firms use the Varian price distribution from Lemma 1 with π^m replaced by $\pi^m(w)$ and $\pi(p)$ replaced by $\pi(p;w)$. Letting $\underline{p}(w)$ denote the lower support, firms mix continuously over the interval $[\underline{p}(w), p^m(w)]$. Secondly, with RPM, and again for a given wholesale cost w, from Proposition 1 the manufacturer chooses a price floor $p_F^*(w)$ that satisfies $\pi(p_F^*;w) = (1-\lambda)\pi^m(w)$, and then downstream firms set price equal to this floor with probability one.

Using the above two observations, we can then write out the manufacturer's problem. When RPM is not possible the manufacturer seeks to

$$\max_w w \int_{\underline{p}(w)}^{p^m(w)} q(p) d[\lambda[1-[1-F(p;w)]^n] + (1-\lambda)F(p;w)], \tag{5}$$

whereas when RPM is possible it seeks to

$$\max_w wq(p_F^*(w)). \tag{6}$$

Solving these optimization problems is challenging, especially the former one due to the complicated way in which a change in w affects the downstream firms' price distribution.[7] Nevertheless we now show that analytical progress is possible for a certain class of demands.

In order to solve (5) and (6), it is useful to view each downstream firm as choosing a *relative profit* β, where β is the fraction of the monopoly profit $\pi^m(w)$ that the firm earns on each sale. Using Lemma 1 we see that absent RPM each retailer chooses β from an interval $[\underline{\beta}, 1]$ where $\underline{\beta} = \frac{1-\lambda}{1+(n-1)\lambda}$ using the distribution function

$$\hat{G}(\beta) = 1 - \left[\frac{1-\lambda}{\lambda n}\left(\frac{1-\beta}{\beta}\right)\right]^{\frac{1}{n-1}} \text{ for } \beta \in \left[\underline{\beta}, 1\right].$$

Proposition 1 implies that with RPM each downstream firm chooses $\beta = 1-\lambda$. Let $Q(\beta;w) \equiv q(p(\beta;w))$ denote demand given relative profit β and wholesale price w, where $p(\beta;w)$ is implicitly determined by $\pi(p;w) = \beta\pi^m(w)$ for $0 \le \beta \le 1$. The following observation will prove very useful:

[7] For an analysis of cost pass-through in the Varian model, see Garrod, Li, Russo, and Wilson (2025).

Lemma 5. *For the class of constant curvature demands* $q(p) = (1 - ap)^b$ *with* $a, b > 0$ *there exists a function* $K(\beta)$ *such that* $Q(\beta; w) = K(\beta)q(w)$.

The class of demands with constant curvature includes, amongst others, those derived from the uniform and exponential distributions. Lemma 5 shows that for demands in this class $Q(\beta; w)$ is multiplicatively separable in β and w. This allows us to then rewrite the manufacturer's problem in (5) as

$$\max_w \; w \int_{\underline{\beta}}^1 Q(\beta; w)\, d\widehat{PD}(\beta) \overset{CC}{=} \max_w \; wq(w)\mathbb{E}_{\widehat{PD}}[K(\beta)], \tag{7}$$

where "CC" means that demand has constant curvature, and where $\widehat{PD}(\beta) = \lambda\{1 - [1 - \hat{G}(\beta)]^n\} + (1 - \lambda)\hat{G}(\beta)$ is the market-wide distribution of relative profit. Similarly the manufacturer's problem in (6) can be rewritten as

$$\max_w \; wQ(1 - \lambda; w) \overset{CC}{=} \max_w \; wq(w)K(1 - \lambda). \tag{8}$$

It follows by inspection that (7) and (8) have the same solution, namely $\arg\max \; wq(w)$. Hence, in the case of demands with constant curvature, the ability to practice RPM has *no* effect on the manufacturer's optimal wholesale price. Indeed, this optimal wholesale price is also independent of both n and λ.[8] Moreover, as we pointed out earlier, $\pi'(p)$ is log-concave for demands that have constant curvature. Our results in Section 3 then immediately imply the following:

Proposition 4. *Suppose demand has constant curvature. Allowing the manufacturer to use RPM has no effect on downstream firms' profits, but it raises aggregate consumer surplus and total welfare, benefits non-shoppers, and harms shoppers.*

Recall our running example with $n = 2$, $\lambda = 1/2$, $w = 1/2$, and $q(p) = 1 - p$. Notice that in this example $\arg\max \; wq(w) = 1/2$ and so, given $n = 2$, $\lambda = 1/2$, and $q(p) = 1 - p$, Figures 1-3 illustrate the impact of RPM even when w is endogenous.

Unfortunately, for demands that do not have constant curvature, it is not possible to obtain analytical results on the impact of RPM. This is primarily due to the complexity of the optimization problem in equation (5), which determines the optimal w absent RPM. Nevertheless, numerical analysis confirms that RPM can benefit consumers and increase total welfare even for demand forms that do not have constant curvature. To illustrate this, Figure 4 depicts the impact of RPM when $n = 5$ and $q(p) = e^{-p^4}$, for different levels of λ. (We use this demand function because its curvature varies strongly with p.[9])

[8] This generalizes a result from Janssen and Shelegia (2015) who show that in the Varian model with $n = 2$ and linear demand the optimal wholesale price is independent of λ.

[9] Let $\sigma(p) = \frac{q(p)q''(p)}{[q'(p)]^2}$ denote demand curvature. When $q(p) = e^{-p^4}$ we have that $\sigma(p) = 1 - 3/(4p^2)$, which increases in p and satisfies, e.g., $\lim_{p \to 0} \sigma(p) = -\infty$, $\sigma(1) = 1/4$, and $\lim_{p \to \infty} \sigma(p) = 1$.

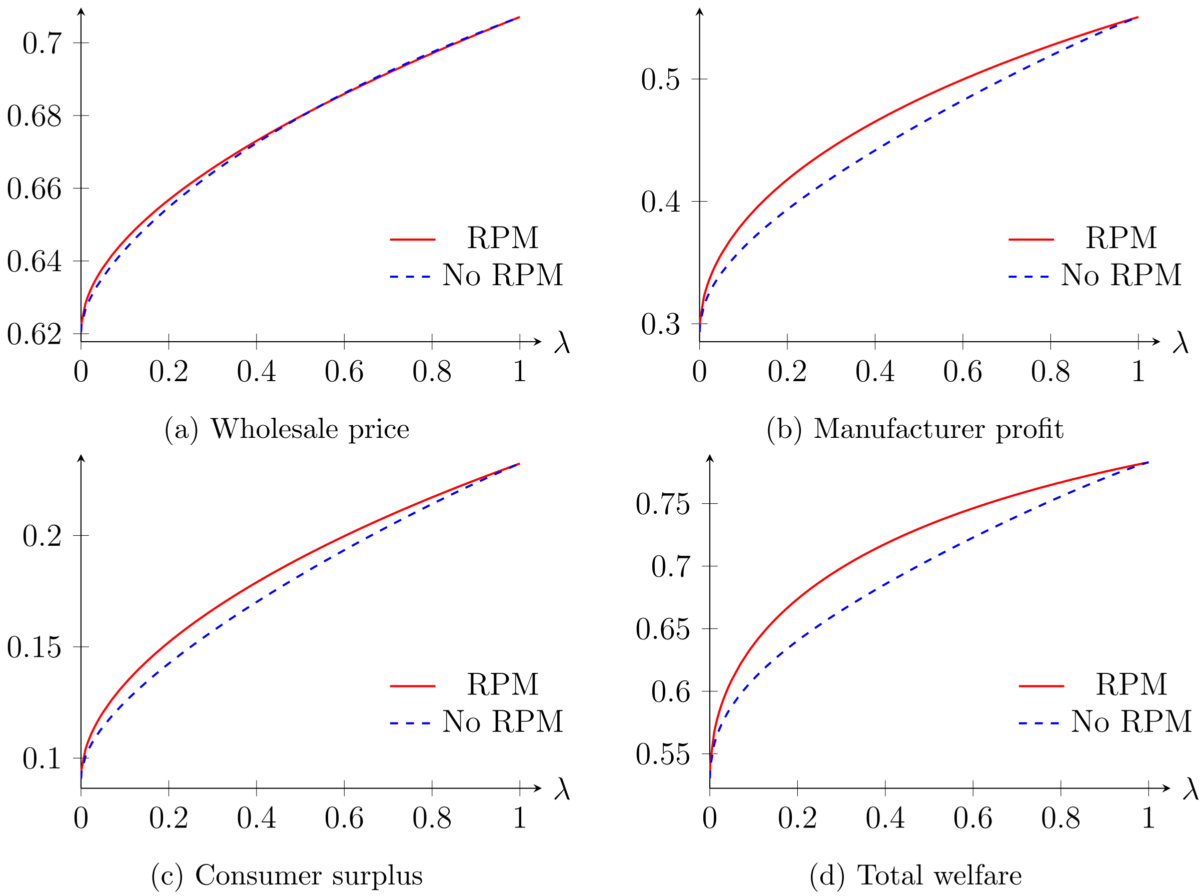

(a) Wholesale price

(b) Manufacturer profit

(c) Consumer surplus

(d) Total welfare

Figure 4: Outcomes with and without RPM when demand has non-constant curvature.

Panel (a) of Figure 4 plots the optimal wholesale price under the two regimes. We observe that the optimal wholesale price is slightly higher under RPM when λ is below around 0.53, but that otherwise RPM leads to a slightly lower wholesale price. Contrary to our earlier results, this implies that RPM is generically *not* neutral for downstream firms. In particular, since downstream industry profit is equal to $\pi^m(w)(1-\lambda)$ (both with and without RPM), and since $\pi^m(w)$ is decreasing in w, the figure implies that RPM harms downstream firms when λ is low but benefits them when λ is higher. Panel (b) of the figure plots manufacturer profit. As expected, RPM is unambiguously good for the manufacturer—starting from whichever w is optimal without RPM, our analysis in Section 3 implies that a switch to RPM improves manufacturer profit, and then the manufacturer can do even better by re-optimizing its wholesale price. Panel (c) shows that RPM always increases consumer surplus. Intuitively, even though RPM raises the optimal w for low values of λ, the increase is small. Hence any resulting upward shift in the price distribution is small, and is dominated by the reduction in profit volatility induced by RPM. Finally, panel (d) shows that RPM also increases total welfare. Intuitively, even though for small λ values RPM reduces downstream industry profit, this is dominated by the increases in both manufacturer profit and consumer surplus.

5 Conclusion

We have examined the effect of RPM in a model of sales where some consumers are well-informed about prices while other consumers are not. The imposition of a price floor qualitatively changes the nature of price competition. In particular, when the price floor is sufficiently high, downstream firms employ a pure strategy and set price exactly equal to the floor. When instead the price floor is more moderate, downstream firms employ a mixed strategy and set price equal to the floor with positive probability, have a gap in their price distribution, and then mix continuously over higher prices. Within this region of moderate price floors, an increase in the price floor induces a "less risky" distribution of market-wide profit, leading to both higher total output and higher consumer surplus. We showed that this implies that the manufacturer's optimal choice of price floor just eliminates all price dispersion in the downstream market, and moreover it is the best price floor from both a consumer and a welfare viewpoint. Nevertheless, RPM has distributional consequences, benefiting non-shoppers but harming shoppers under a regularity condition. Finally, the welfare impact of RPM is larger in markets with more downstream firms.

Our analysis has focused on a price floor that is chosen optimally by an upstream manufacturer. However, in some markets policymakers may impose price floors—such as minimum prices for alcohol and tobacco—in order to try and reduce consumption of those goods.[10] Our analysis suggests that if consumers face informational frictions, and some are better informed than others, then price floors that are not aggressive enough could actually backfire and lead to higher consumption of the affected products.

[10]For example, in Scotland there is a minimum price of 65p per unit of alcohol (link), while in Minneapolis the city council has imposed a minimum price of $15 per pack of cigarettes (link), and in New York City there are minimum prices for a variety of different tobacco products (link).

References

Armstrong, M., and J. Vickers (2019): "Discriminating against captive customers," *American Economic Review: Insights*, 1(3), 257–72.

Armstrong, M., J. Vickers, and J. Zhou (2009): "Consumer protection and the incentive to become informed," *Journal of the European Economic Association*, 7(2-3), 399–410.

Asker, J., and H. Bar-Isaac (2014): "Raising retailers' profits: On vertical practices and the exclusion of rivals," *American Economic Review*, 104(2), 672–686.

——— (2020): "Vertical Information Restraints: Pro- and Anticompetitive Impacts of Minimum-Advertised-Price Restrictions," *Journal of Law and Economics*, 63(1), 111–148.

Baye, M. R., D. Kovenock, and C. de Vries (2025): "Minimum Resale Price Maintenance in the Absence of Point-of-Sale Service, Agency or Free-Rider Problems," *working paper*.

Garcia, D., J. Honda, and M. Janssen (2017): "The double diamond paradox," *American Economic Journal: Microeconomics*, 9(3), 63–99.

Garcia, D., and M. Janssen (2018): "Retail Channel Management in Consumer Search Markets," *International Journal of Industrial Organization*, 58, 162–182.

Garrod, L., R. Li, A. Russo, and C. M. Wilson (2025): "Understanding Cost Pass-Through When Prices Are Dispersed," Cesifo working paper series, CESifo.

Hamilton, J. H. (1990): "Resale price maintenance in a model of consumer search," *Managerial and Decision Economics*, 11(2), 87–98.

Janssen, M., and E. Reshidi (2022): "Regulating recommended retail prices," *International Journal of Industrial Organization*, 85, 102872.

Janssen, M., and S. Shelegia (2015): "Consumer search and double marginalization," *American Economic Review*, 105(6), 1683–1710.

——— (2020): "Beliefs and Consumer Search in a Vertical Industry," *Journal of the European Economic Association*, 18(5), 2359–2393.

Jullien, B., and P. Rey (2007): "Resale price maintenance and collusion," *RAND Journal of Economics*, 38(4), 983–1001.

Lubensky, D. (2017): "A Model of Recommended Retail Prices," *RAND Journal of Economics*, 48(2), 358–386.

Marvel, H. P., and S. McCafferty (1984): "Resale Price Maintenance and Quality Certification," *RAND Journal of Economics*, 15(3), 346–359.

Mathewson, G. F., and R. A. Winter (1984): "An Economic Theory of Vertical Restraints," *RAND Journal of Economics*, 15(1), 27–38.

Myatt, D. P., and D. Ronayne (2025): "Stable Price Dispersion under Heterogeneous Buyer Consideration," *The RAND Journal of Economics (forthcoming).*

Rey, P., and T. Vergé (2008): "Economics of vertical restraints," *Handbook of antitrust economics*, 353, 390.

——— (2010): "Resale price maintenance and interlocking relationships," *The Journal of Industrial Economics*, 58(4), 928–961.

Salop, S., and J. Stiglitz (1977): "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *The Review of Economic Studies*, 44(3), 493–510.

Stahl, D. O. (1989): "Oligopolistic pricing with sequential consumer search," *American Economic Review*, pp. 700–712.

Telser, L. G. (1960): "Why Should Manufacturers Want Fair Trade?," *Journal of Law and Economics*, 3, 86–105.

Varian, H. R. (1980): "A model of sales," *American Economic Review*, 70(4), 651–659.

Williams, R. J. (2024): "Empirical Effects of Resale Price Maintenance: Evidence from Fixed Book Price Policies in Europe," *Journal of Competition Law & Economics*, 20(1-2), 108–136.

Xia, T. (2024): "Welfare Effects of Resale Price Maintenance: Evidence from the Chinese Pharmaceutical Industry," *Available at SSRN 5118723.*

Yang, Y. (2024): "Price transparency in online markets," *working paper.*

A Omitted Proofs

Proof of Lemma 1. The proof is standard and so omitted. □

Proof of Lemma 2. Clearly if $p_F \geq p^m$ there is a unique equilibrium where all firms charge p_F. In the remainder of the proof we therefore focus on $p_F \in [p_F^*, p^m)$.

Any candidate pure strategy equilibrium must involve all firms charging p_F. To verify that all firms charging p_F is an equilibrium, note that in this putative equilibrium each firm earns $\pi(p_F)/n \geq \pi(p_F^*)/n$. If a firm deviates to a price $p > p_F$ it loses all the shoppers and so earns $(1-\lambda)\pi(p)/n \leq (1-\lambda)\pi^m/n = \pi(p_F^*)/n$, where the equality uses (1). Hence no firm has an incentive to deviate.

Any candidate symmetric mixed strategy equilibrium must be atomless except possibly at p_F, and must have p^m in its support. This implies that expected profit must be $(1-\lambda)\pi^m/n$. However, if a firm charges p_F it gets strictly more than a $1/n$ share of consumers, leading to profit strictly larger than $\pi(p_F)/n \geq \pi(p_F^*)/n = (1-\lambda)\pi^m/n$, which is a contradiction. Hence there is no mixed strategy equilibrium. □

Proof of Lemma 3. The construction of the mixed strategy was explained in the text and so we omit further details here. It is straightforward to show that the left-hand side of (2) is strictly increasing in p_F and decreasing in m, approaches $\pi(p_F)[\,\lambda + (1-\lambda)/n\,]$ as m tends to 0, and approaches $\pi(p_F)/n$ as m tends to 1. Existence and uniqueness of $m \in (0,1)$ and $\underline{p}' \in (\underline{p}, p^m)$ then follow. □

Proof of Lemma 4. To prove part (i), note that for a given "profit floor" π_F expected per-consumer profit is

$$
\begin{aligned}
&[\lambda[1-(1-m)^n] + (1-\lambda)m]\,\pi_F + \int_{\underline{\pi}'}^{\pi^m} \pi \; dPD(\pi) \\
&= m\pi^m(1-\lambda) + \int_{\underline{\pi}'}^{\pi^m} \left[\lambda n[1-\tilde{G}(\pi)]^{n-1} + 1 - \lambda\right] \pi \; d\tilde{G}(\pi) \\
&= \pi^m(1-\lambda),
\end{aligned}
$$

where the second line uses (2) to substitute out π_F and (4) to substitute out for $dPD(\pi)$, and the third line uses (3) to substitute out for $\tilde{G}(\pi)$ and uses the fact that $\tilde{G}(\pi^m) - \tilde{G}(\underline{\pi}') = 1 - m$. Hence expected per-consumer profit is indeed independent of π_F.

To prove part (ii), consider two price floors p_F^0 and p_F^1 satisfying $\underline{p} < p_F^0 < p_F^1 < p_F^*$, and let $\pi_F^0 = \pi(p_F^0)$ and $\pi_F^1 = \pi(p_F^1)$. Let $PD^0(\pi)$ and $PD^1(\pi)$ be the associated per-consumer profit distributions, with associated mass points χ^0 and χ^1, and associated lower bounds $\underline{\pi}'^0$ and $\underline{\pi}'^1$ on the continuous parts of the distributions. Let $\phi(\pi)$ be an arbitrary concave function. We need to prove that $\mathbb{E}_{PD^0}[\phi(\pi)] < \mathbb{E}_{PD^1}[\phi(\pi)]$, or equivalently that

$$\chi^0\phi(\pi_F^0) + \int_{\underline{\pi}'^0}^{p^m} \phi(\pi)dPD^0(\pi) < \chi^1\phi(\pi_F^1) + \int_{\underline{\pi}'^1}^{p^m} \phi(\pi)dPD^1(\pi). \tag{9}$$

As a first step, note that $\underline{\pi}'^0 < \underline{\pi}'^1$ and $PD^0(\pi) = PD^1(\pi)$ for all $\pi \geq \underline{\pi}'^1$. Note also that

$$\frac{\chi^0\phi(\pi_F^0) + \int_{\underline{\pi}'^0}^{\underline{\pi}'^1} \phi(\pi)dPD^0(\pi)}{\chi^1} < \phi\left(\frac{\chi^0\pi_F^0 + \int_{\underline{\pi}'^0}^{\underline{\pi}'^1} \pi\ dPD^0(\pi)}{\chi^1}\right) = \phi(\pi_F^1)$$

where the inequality uses Jensen's inequality, and the equality uses the fact that the mean per-consumer profits are the same.[11] Hence

$$\chi^0\phi(\pi_F^0) + \int_{\underline{\pi}'^0}^{\underline{\pi}'^1} \phi(\pi)dPD^0(\pi) < \chi^1\phi(\pi_F^1).$$

Adding $\int_{\underline{\pi}'^1}^{p^m} \phi(\pi)dPD^0(\pi)$ to the left-hand side and $\int_{\underline{\pi}'^1}^{p^m} \phi(\pi)dPD^1(\pi)$ to the right-hand side, and recalling that $PD^0(\pi) = PD^1(\pi)$ for all $\pi \geq \underline{\pi}'^1$, this simplifies to (9). □

Proof of Proposition 1. First, we prove that manufacturer profit is strictly increasing in $p_F \in (\underline{p}, p_F^*)$. Note that manufacturer profit is proportional to $\mathbb{E}_{PD}[Q(\pi)]$. From Lemma 4 we know that higher π_F improves $PD(\pi)$ in the sense of second order stochastic dominance, and so it suffices to prove that $Q(\pi)$ is strictly concave. We can write

$$Q'(\pi) = \frac{q'(p)}{\pi'(p)} = \frac{q'(p)/q(p)}{1 + q'(p)(p-w)/q(p)}.$$

Given that for $\pi < \pi^m$ the denominator is positive and $\pi'(p) > 0$, and given that log-concavity of $q(p)$ implies that $q'(p)/q(p)$ is negative and decreasing, it follows that $Q(\pi)$ is indeed strictly concave as required. Second, as explained in the text, manufacturer profit is decreasing in $p_F \geq p_F^*$, and hence it is maximized at $p_F = p_F^*$. □

Proof of Proposition 2. The proof of part (i) follows from arguments in the text, especially the concavity of $V(\pi)$ as proved by Armstrong and Vickers (2019), and so the proof is omitted. Now consider part (ii). Differentiating (2) with respect to π_F $(= \pi(p_F))$ we obtain

$$\frac{\partial m}{\partial \pi_F} = \frac{\frac{1-\lambda}{n} + \lambda\frac{1-(1-m)^n}{mn}}{\frac{\lambda\pi_F}{m^2n}\left[1-(1-m)^n - nm(1-m)^{n-1}\right]} = \frac{m\underline{\pi}'}{\pi_F(\underline{\pi}' - \pi_F)},$$

[11]As mean per-consumer profits are the same we have $\chi^0\pi_F^0 + \int_{\underline{\pi}'^0}^{p^m} \pi\ \ dPD^0(\pi) = \chi^1\pi_F^1 + \int_{\underline{\pi}'^1}^{p^m} \pi\ dPD^0(\pi)$, which simplifies to $\chi^0\pi_F^0 + \int_{\underline{\pi}'^0}^{\underline{\pi}'^1} \pi\ dPD^0(\pi) = \chi^1\pi_F^1$ because $PD^0(\pi) = PD^1(\pi)$ for all $\pi \geq \underline{\pi}'^1$.

where the second equality uses (2) to substitute out π_F, and also uses the fact that $(1-\lambda)\pi^m = \underline{\pi}'\left[\lambda n(1-m)^{n-1}+1-\lambda\right]$. (Note that the latter is derived from the expression for $F(p)$ in Lemma 1 given that $\underline{\pi}' = \pi(\underline{p}')$ and $F(\underline{p}') = m$.) In addition, $\tilde{G}(\underline{\pi}') = m$ implies

$$\frac{\partial \underline{\pi}'}{\partial \pi_F} = \frac{\partial m}{\partial \pi_F}\frac{1}{\tilde{g}(\underline{\pi}')}.$$

Next, the surplus of a non-shopper is $mV(\pi_F) + \int_{\underline{\pi}'}^{\pi^m} V(\pi)\, d\tilde{G}(\pi)$. Its derivative with respect to π_F is

$$\begin{aligned}\frac{\partial m}{\partial \pi_F}V(\pi_F) + mV'(\pi_F) - V(\underline{\pi}')\,\tilde{g}(\underline{\pi}')\,\frac{\partial \underline{\pi}'}{\partial \pi_F} =& [V(\pi_F) - V(\underline{\pi}')]\frac{m\underline{\pi}'}{\pi_F(\underline{\pi}'-\pi_F)} + mV'(\pi_F) \\ >& -V'(\pi_F)\, m\left(\frac{\underline{\pi}'}{\pi_F} - 1\right) > 0,\end{aligned}$$

where the first expression uses the fact that $\tilde{G}(\pi)$ does not depend on π_F for $\pi > \underline{\pi}'$, the first inequality follows from concavity of $V(\pi)$, and the second inequality follows from $V'(\cdot) < 0$ and $\pi_F < \underline{\pi}'$.

Finally, the surplus of a shopper is $[1-(1-m)^n]\, V(\pi_F) + \int_{\underline{\pi}'}^{\pi^m} V(\pi)\, d[1-(1-\tilde{G}(\pi))^n]$. Its derivative with respect to π_F is

$$\begin{aligned}& n(1-m)^{n-1}V(\pi_F)\frac{\partial m}{\partial \pi_F} + \left[1-(1-m)^n\right]V'(\pi_F) - n\,\tilde{g}(\underline{\pi}')\left[1-\tilde{G}(\underline{\pi}')\right]^{n-1}V(\underline{\pi}')\frac{\partial \underline{\pi}'}{\partial \pi_F} \\ &= n(1-m)^{n-1}\left[V(\pi_F) - V(\underline{\pi}')\right]\frac{m\underline{\pi}'}{\pi_F(\underline{\pi}'-\pi_F)} + \left[1-(1-m)^n\right]V'(\pi_F) \\ &= \frac{m(1-\lambda)}{\lambda\pi_F}\left[(\pi^m - \underline{\pi}')\frac{V(\pi_F)-V(\underline{\pi}')}{\underline{\pi}'-\pi_F} + (\pi^m - \pi_F)V'(\pi_F)\right],\end{aligned}$$

where the third line uses (2) and $(1-\lambda)\pi^m = \underline{\pi}'\left[\lambda n(1-m)^{n-1}+1-\lambda\right]$ (which we explained how to derive earlier in the proof). Since concavity of $V(\pi)$ implies that $\frac{V(\pi_F)-V(\underline{\pi}')}{\underline{\pi}'-\pi_F} < -V'(\underline{\pi}')$, and since $-V'(\underline{\pi}') = \frac{q(\underline{p}')}{\pi'(\underline{p}')}$ and $-V'(\pi_F) = \frac{q(p_F)}{\pi'(p_F)}$, a sufficient condition for the square-bracketed term to be negative is that

$$\frac{q\left(\underline{p}'\right)\int_{\underline{p}'}^{p^m}\pi'(z)\,dz}{\pi'\left(\underline{p}'\right)} < \frac{q\left(p_F\right)\int_{p_F}^{p^m}\pi'(z)\,dz}{\pi'\left(p_F\right)}.$$

This inequality holds because $\underline{p}' > p_F$ implies that $q\left(\underline{p}'\right) < q\left(p_F\right)$, and because log-concavity of $\pi'(p)$ implies that $\int_x^{p^m}\pi'(z)dz/\pi'(x)$ is decreasing in x. (Specifically, note that $\frac{\pi''(x)}{\pi'(x)}\int_x^{p^m}\pi'(z)dz \le \int_x^{p^m}\pi''(z)dz = -\pi'(x)$, where the inequality uses log-concavity of $\pi'(p)$. This inequality implies that $\int_x^{p^m}\pi'(z)dz/\pi'(x)$ is decreasing in x.) □

Proof of Proposition 3. Let $\tilde{F}(\pi;n)$ be the distribution of a firm's profit in the Varian model from Lemma 1, and let $\widetilde{PD}(\pi;n)$ be the associated distribution of per-consumer profit. Note that they satisfy respectively

$$\tilde{F}(\pi;n) = 1 - \left[\frac{1-\lambda}{\lambda n}\left(\frac{\pi^m}{\pi} - 1\right)\right]^{\frac{1}{n-1}},$$

and

$$\widetilde{PD}(\pi;n) = \lambda\left[1 - [1-\tilde{F}(\pi;n)]^n\right] + (1-\lambda)\tilde{F}(\pi;n).$$

We will prove that under this Varian distribution, higher n reduces total consumer surplus and total output (and hence manufacturer profit), decreases the surplus of non-shoppers, and increases the surplus of shoppers whenever $\pi'(p)$ is log-concave. Since each consumer's surplus is independent of n under RPM, the result in the proposition then follows.

First, consider total consumer surplus and total output under the Varian distribution. It suffices to show that an increase in n worsens $\widetilde{PD}(\pi;n)$ in the sense of second order stochastic dominance or, equivalently, that $\int_{-\infty}^{\pi} \widetilde{PD}(t;n)dt$ increases in n. Letting $\underline{\pi}(n)$ be the lowest per-consumer profit in the support of the distribution, and noting that $\widetilde{PD}(\underline{\pi}(n);n) = 0$, it is sufficient to prove that $H(\pi;n) \geq 0$ where

$$H(\pi;n) \equiv \int_{\underline{\pi}(n)}^{\pi} \frac{\partial \widetilde{PD}(t;n)}{\partial n} dt.$$

Clearly $H(\underline{\pi}(n);n) = 0$. Moreover $H(\pi^m;n) = 0$. To see the latter, note that using the same steps as in the proof of Lemma 4 per-consumer profit is $\pi^m(1-\lambda)$, and so

$$(1-\lambda)\pi^m = \int_{\underline{\pi}(n)}^{\pi^m} \pi\, d\widetilde{PD}(\pi;n) = \pi^m - \int_{\underline{\pi}(n)}^{\pi^m} \widetilde{PD}(\pi;n)\, d\pi,$$

where the second equality uses integration by parts. Differentiating with respect to n, the above implies that $H(\pi^m;n) = 0$. We now demonstrate that $H(\pi;n) > 0$ for all $\pi \in (\underline{\pi}(n), \pi^m)$. To do this, use the change of variables

$$u = \frac{1-\lambda}{\lambda n}\left(\frac{\pi^m}{\pi} - 1\right),$$

and write

$$\tilde{F}(\pi;n) = 1 - u^{\frac{1}{n-1}} \quad \text{and} \quad \frac{\partial \tilde{F}(\pi;n)}{\partial n} = \frac{u^{\frac{1}{n-1}}}{(n-1)^2}\left(\ln u + \frac{n-1}{n}\right).$$

We then find that

$$\begin{aligned}\frac{\partial \widetilde{PD}(\pi;n)}{\partial n} &= -\lambda[1-\tilde{F}(\pi;n)]^n \ln(1-\tilde{F}(\pi;n)) + \left[\lambda n[1-\tilde{F}(\pi;n)]^{n-1} + (1-\lambda)\right]\frac{\partial \tilde{F}(\pi;n)}{\partial n} \\ &= \frac{(\lambda u + 1 - \lambda)u^{\frac{1}{n-1}}}{(n-1)^2}\left[\ln u + \frac{n-1}{n}\frac{\lambda u n + 1 - \lambda}{\lambda u + 1 - \lambda}\right].\end{aligned}$$

Notice that as π increases from $\underline{\pi}(n)$ to π^m, u decreases from 1 to 0. Notice also that the square-bracketed term in the second line of the above expression is increasing in u, is negative as $u \to 0$ but positive as $u \to 1$. Hence $\partial\widetilde{PD}(\pi;n)/\partial n > 0$ if and only if $\pi < \breve{\pi}(n)$ for some $\breve{\pi}(n) \in (\underline{\pi}(n), \pi^m)$. This implies that $H(\pi;n)$ first increases and then decreases in π; given that $H(\underline{\pi}(n);n) = H(\pi^m;n) = 0$ this means that $H(\pi;n) > 0$ for all $\pi \in (\underline{\pi}(n), \pi^m)$.

Second, consider the surplus of non-shoppers for the Varian distribution, which equals

$$\int_{\underline{\pi}(n)}^{\pi^m} V(\pi)\, d\tilde{F}(\pi;n) = V(\pi^m) - \int_{\underline{\pi}(n)}^{\pi^m} V'(\pi)\, \tilde{F}(\pi;n)\, d\pi,$$

where the equality uses integration by parts. Its derivative with respect to n is

$$-\int_{\underline{\pi}(n)}^{\pi^m} V'(\pi) \frac{\partial \tilde{F}(\pi;n)}{\partial n} d\pi = \frac{n\lambda(1-\lambda)\pi^m}{(n-1)^2} \int_0^1 t_{ns}(u)\, w_{ns}(u)\, du, \tag{10}$$

where the equality uses the same change of variables as earlier in the proof, as well as the fact therefore that $d\pi = -du[n\lambda(1-\lambda)\pi^m]/(1-\lambda+\lambda nu)^2$, and where

$$t_{ns}(u) = u^{\frac{1}{n-1}} \left(\ln u + \frac{n-1}{n} \right) \quad \text{and} \quad w_{ns}(u) = -\frac{V'\left(\pi^m \frac{1-\lambda}{1-\lambda+\lambda nu}\right)}{(1-\lambda+\lambda nu)^2}.$$

Define $T_{ns}(u) = \int_0^u t_{ns}(z)dz$, and note that $T_{ns}(0) = \int_0^0 t_{ns}(z)dz = 0$, and also that

$$T_{ns}(1) = \int_0^1 u^{\frac{1}{n-1}} \ln u\, du + \frac{n-1}{n} \int_0^1 u^{\frac{1}{n-1}}\, du = 0,$$

where the second equality follows from integrating $\int_0^1 u^{\frac{1}{n-1}} \ln u du$ by parts. Since $t_{ns}(u) < 0$ if and only if $u < e^{-(n-1)/n}$, we then conclude that $T_{ns}(u) < 0$ for all $u \in (0,1)$. It is also easy to verify that $w'_{ns}(u) < 0$. Therefore integrating the right-hand side of equation (10) by parts gives

$$-\frac{n\lambda(1-\lambda)\pi^m}{(n-1)^2} \int_0^1 w'_{ns}(u) T_{ns}(u) du < 0.$$

Third, consider the surplus of shoppers for the Varian distribution, which equals

$$\int_{\underline{\pi}(n)}^{\pi^m} V(\pi)\, d\left[1 - [1-\tilde{F}(\pi;n)]^n\right] = V(\pi^m) - \int_{\underline{\pi}(n)}^{\pi^m} V'(\pi) \left[1 - [1-\tilde{F}(\pi;n)]^n\right] d\pi,$$

where the equality uses integration by parts. Its derivative with respect to n is

$$\frac{n\pi^m \lambda(1-\lambda)}{(n-1)^2} \int_0^1 t_s(u) w_s(u) du, \tag{11}$$

where we have again used the change of variables from earlier in the proof, and where

$$t_s(u) = u^{\frac{2-n}{n-1}} (\ln u + n - 1) \quad \text{and} \quad w_s(u) = -\frac{u^2 V'\left(\pi^m \frac{1-\lambda}{1-\lambda+\lambda nu}\right)}{(1-\lambda+\lambda nu)^2}.$$

Letting $T_s(u) = \int_0^u t_s(z)dz$, similar steps as above can be used to establish that $T_s(0) = T_s(1) = 0$ and that $T_s(u) < 0$ for $u \in (0,1)$. We now prove that $w_s'(u) > 0$. To do this, rewrite $w_s(u)$ as a function of π:

$$w_s(u) = -\frac{1}{(\lambda n \pi^m)^2}(\pi^m - \pi)^2 V'(\pi) = \left[\frac{\pi^m - \pi}{(\lambda n \pi^m)^2}\right]\left[\frac{q\,(p)\int_p^{p^m} \pi'\,(z)\,dz}{\pi'\,(p)}\right].$$

Clearly the first square-bracketed term is decreasing in π. The second square-bracketed term is also decreasing in π, given that p increases in π over the relevant range, and given that $q'(p) < 0$ and given that (as shown in the proof of Proposition 2) log-concavity of $\pi'(p)$ implies that $\int_p^{p^m} \pi'\,(z)\,dz/\pi'(p)$ is decreasing in p. Hence the product of the two square-bracketed terms is decreasing in π; as π decreases in u, this implies that $w_s'(u) > 0$ as claimed above. Finally, equation (11) implies

$$\frac{n\pi^m\lambda(1-\lambda)}{(n-1)^2}\int_0^1 t_s(u)w_s(u)du = -\frac{n\pi^m\lambda(1-\lambda)}{(n-1)^2}\int_0^1 w_s'(u)T_s(u)du > 0,$$

where the equality uses integration by parts, and the inequality uses $w_s'(u) > 0$ and $T_s(u) < 0$. □

Proof of Lemma 5. Using $q(p) = (1-ap)^b$ for $a, b > 0$, one can check that

$$\pi^m(w) = \frac{(1-aw)^{1+b}\,b^b}{a\,(1+b)^{1+b}} = q\,(w)^{\frac{1+b}{b}}\,\frac{b^b}{a\,(1+b)^{1+b}}.$$

Hence $\pi\,(p(\beta;w);w) = \beta\pi^m(w)$ is equivalent to

$$[p(\beta;w) - w]\,q\,(p(\beta;w)) = \beta q\,(w)^{\frac{1+b}{b}}\,\frac{b^b}{a\,(1+b)^{1+b}}. \tag{12}$$

Moreover, since $1 - ap = q(p)^{1/b}$, we can also write $a\,(p-w) = q\,(w)^{1/b} - q\,(p)^{1/b}$. Substituting this into (12) and simplifying, and then using $Q(\beta;w) = q\,(p(\beta;w))$, we obtain

$$\left[1 - \left(\frac{Q(\beta;w)}{q\,(w)}\right)^{1/b}\right]\left(\frac{Q(\beta;w)}{q\,(w)}\right) = \frac{\beta b^b}{(1+b)^{1+b}}.$$

This in turn implies that $Q(\beta;w) = K(\beta)q(w)$. (Note that the function $K(\beta)$ is uniquely defined. This is because a unique $p(\beta;w)$ solves $\pi\,(p(\beta;w);w) = \beta\pi^m(w)$ given the quasiconcavity of the profit function, and so $Q(\beta;w) = q\,(p(\beta;w))$ must also be unique.) □

Proof of Proposition 4. We need to prove that constant curvature demand implies that $\pi'(p)$ is log-concave. (The result then follows from arguments in the text.) Note that $\pi'(p) \propto 1 - ap - ab(p - w) > 0$ for $p < p^m$. It then follows from direct calculation that

$$\frac{d^2}{dp^2} \ln \pi'(p) = -\frac{a^2}{(1-ap)^2} \left[(b-1) + (b+1)^2 \left(\frac{1-ap}{1-ap-abp+abw} \right)^2 \right]$$
$$\leq -\frac{a^2}{(1-ap)^2} \left[(b-1) + (b+1)^2 \right] < 0,$$

where the first inequality uses the fact that $p \geq w$ (since downstream firms will never price below their cost). □