Ducking Data Collection

Simon Anderson, Nathan Larson, and Amparo Urbano

October 5, 2019

Abstract

Consumers may forego transactions when they anticipate that information collected from their purchases may be sold on and used to discriminate against them in later markets. We study the sale of consumer data gathered in one market to a monopolist in a second market. As a side effect of purchasing in the first market, consumers reveal to Firm 1 their willingness-to-pay in the second market. For the benchmark case, when consumers do not know their tastes precisely in the (undisclosed) second market, the firm in the first market lowers its price to encourage data collection and fish for information. Although consumers anticipate deleterious consequences in a second market, consumer surplus and total surplus is higher than without the information fishing. Allowing consumers to opt-out worsens market performance. The main driver of this result is the fact that under either policy, consumers who evade having their data sold impose a negative externality on all consumers in the second market (since the second firm anticipates that they will have higher-than-average valuations, and therefore raises its price). We then show there are countervailing effects when consumers know their demand (type) in the second (exploitive) market. If data collection is perfect, it induces a hold-up problem whereby the second market unravels. Data collection raises profit at the expense of consumer surplus; but allowing opt-out has no effect. We model imperfect/partial data collection by assuming that the first firm can only discern the tastes of a fraction of its customers. Then fishing with opt-out is the best policy for firms and worst for consumers (vis-a-vis laissez-faire data collection or no collection). Moreover, laissez-faire data collection benefits consumers (at the expense of the data-harvesting firm) when the data-collection technology is low-yield. When consumer values in the second market are either low or high, pricing on the

second market is mixed (as high value consumers "hide" just up to the point that the second firm begins to prefer to price to them exclusively.) Finally, we draw a distinction between opt-in and opt-out, which can have large welfare consequences: consumers may collectively be substantially better off when it is opting in, rather than opting out, that incurs a small nuisance cost.
Keywords: data collection, opt-in, opt-out, price discrimination

JEL: D43, L12, L13, M37

1 Introduction

While firms have been collecting and exploiting troves of individual-level consumer data for some time, 2018 was a watershed year when these practices captured the public imagination and drew scrutiny from policymakers. A sequence of data breach scandals (including Facebook's contretemps with Cambridge Analytica) laid bare the fact that companies like Google and Facebook have assembled staggeringly detailed information about individual consumers' preferences and that this data can be intentionally marketed or inadvertently leaked in ways that are often adverse to consumer interests. Research indicates that consumers object to the loss of privacy for psychological reasons (the 'creepiness' factors), for fear of having their information used against them in markets, and because of the risks of fraud and identity theft (Tucker 2016, Turow et al. 2009, White et al. 2008, Acquisti et al. 2015). Both the EU (with the adoption of the GDPR) and the US Congress have stressed the urgency of finding appropriate policies to safeguard consumer privacy. However, since well-intentioned policies can often be counterproductive, it is essential to study the incentive issues in markets for consumer data using rigorous economic models.

This project focuses on the market implications of the collection and sale of consumer data, specifically how information can be gathered from one market in order to be used in a second market. However, because customers can be aversely discriminated against in the second market they may wish to hide their information by not participating in the first one. Think for example about opening a bank account and revealing information which the bank can then sell on for use elsewhere: consumer reticence to reveal the information may lead to ducking. In this model we can also address the desire of firms to allow opt out of receiving offers as well as having legislation about giving the consumer the right to opt out.

The model has monopolist firms in two separate markets. As a side effect of purchasing in the first market, consumers reveal to Firm 1 their willingness-to-pay in the second market. We study a data-sharing policy regime, where Firm 1 faces no restrictions on selling this information to the second firm, and an opt-out regime, where Firm 1 can sell the data of its customers but a consumer can opt out by paying a small nuisance cost. Perhaps surprisingly, switching from unrestricted data-sharing to the opt-out policy generally improves the profit of the data-gathering firm and makes consumers worse off. The main driver of this result is the fact that under either policy, consumers who evade having their data sold impose a negative externality on all consumers in the second market (since the second firm anticipates that they will have higher-than-average valuations, and therefore raises its price). Furthermore, because this selection effect can induce unraveling in the second market, default policies matter: consumers may collectively be be substantially better off when it is opting in, rather than opting out, that incurs the small nuisance cost.

Section 2 introduces the model template that we use throughout. Section 3 sketches the data-gatherer's incentive to expand its market share and consumers' incentive to duck data collection in a simple setting with no adverse selection downstream. Section 4 adds the consumer selection effect in a parameterized demand setting that permits clean closed-form results. Section 5 employs a simple example to make the point that default policies on opt-in versus opt-out can have a substantial impact on market outcomes and consumer welfare. Section 6 gives an example (based on analysis not presented here) in which some of our main conclusions can be reversed. Appendices A and B develop the model for alternative assumptions about demand.

1.1 Literature

[To be written.] For this theme we can envisage more elaborate market structures both upstream and downstream. Montes et al. (2018) consider a data-broker selling to downstream Hotelling duopolists and argue that it will sell exclusively to only one. Braulin and Valletti (2015) extend the result to the canonical vertical differentiation model. However, the data is exogenously given and sold en masse (excepting a small extension). Such downstream models (and generalizations, like those of the next theme) can be appended to the collection module, and sales of parsed of information can be addressed too. More intricate is to consider competition in collection markets, and how that plays out in both those markets and in the equilibrium value of information. Belleflamme and Vergote (2016) and Chen et al. (2018) are closest to our opt-out analysis because they permit customers to hide from profiling. The former show (for monopoly) that tracking technology lowers consumer surplus because firms are able to price discriminate, but hiding technology worsens consumer surplus further because the firm raises regular prices to discourage hiding. In Chen et al. (2018), each firm in a Hotelling model can personalize prices for consumers in its target segment and offer a uniform "poaching" price for non-targeted customers. Hiding consumers make it harder to poach, softening competition through higher prices for non-targeted consumers. Both papers suggest, counterintuitively, that privacy regulation empowering consumers may make them worse off.

2 Model Framework

Firms 1 and 2 are monopolists in separate markets, but they share a common pool of consumers indexed by location $x \in [0, 1]$. In some applications a consumer's location could be her actual physical location, but for our purposes it is just a way to identify her. Consumer x has a valuation v_j for Firm j's product and total valuation $v_1 + v_2$ if she acquires both products. She visits Firm 1 first and decides whether to purchase its product at price p_1. If she purchases, Firm 1 learns her type: it observes her location x and value v_2 at Firm 2. If she does not purchase, Firm 1 learns nothing. Depending on the policy regime, Firm 1 may sell its customers' data – that is, reveal their values v_2 – to Firm 2. Then the consumer visits Firm 2 where, if her data has not been sold, she is offered a uniform price p_2. Alternatively, if she bought at Firm 1 and it sold her data to Firm 2, then she is charged a personalized price that extracts her full valuation v_2. Because of the order in which a consumer faces the two firms, we refer to Firm 1's market as upstream and Firm 2's market as downstream. This is just shorthand for the timing of consumer choices; it does not mean that Firm 1 is a supplier to Firm 2. Marginal costs are normalized to zero.

When it is useful, we can write the inverse demand curves in the two markets as $v_1(x)$ and $v_2(x)$. Both firms' marginal costs are normalized to zero. Next we turn to a detailed description of data-sharing. We analyze the following data-sharing policies:

Autarky (A) Selling data is not permitted. The upstream and downstream markets operate independently with each firm setting its monopoly price p_1^A or p_2^A in its own market.

Data-sharing (DS) Selling data is permitted. Firm 1 makes a take-it-or-leave it offer to Firm 2 for the data on all of the upstream purchasers.

Costly opt-out (COO) Selling data is permitted, but a purchaser at Firm 1 can pay a nuisance cost $c > 0$ to opt out of data-sharing. In this case, Firm 1 makes a take-it-or-leave it offer to Firm 2 for the data on all upstream purchasers who did not opt out.

Free opt-out (OO) This is the special case of the costly opt-out regime where the parameter c is taken to the $c \to 0$ limit.

The cost c in the opt-out versions of the model should be thought of as a time and effort cost to the consumer – it could be the time cost of hunting through Firm 1's website to figure out where to click to opt out. This cost should be thought of as small; indeed, we study the $c > 0$ case mainly to provide solid foundations for our main focus, the free opt-out limit.

The equilibrium analysis under autarky is straightforward: each firm chooses its monopoly price, and marginal consumer, to maximize its profit $\pi_i^A = v_i(x)\,x$. Before discussing equilibrium in the other regimes we need a more precise description of the timing and the information each party has when it acts.

Start with the downstream market after Firm 2 has purchased any data on offer. We say that a consumer is "revealed" if Firm 2 knows her valuation v_2; this means she purchased at Firm 1 and did not opt out, if that was a possibility. We say a consumer is "hidden" if Firm 2 does not know her valuation; she may be hidden because she did not buy from Firm 1 or (under the opt-out regimes) because she bought at Firm 1 but opted out of data sharing. As shorthand, under the opt-out policy we will say that a hidden consumer either "waited" (if she did not buy upstream) or opted out. Firm 2 perfectly price discriminates to the revealed consumers; let π_2^{rev} denote its profit from these sales. It offers a uniform price p_2 to any consumer who is hidden; let π_2^{hid} be its profit on these consumers at the optimal p_2. We assume that a revealed consumer cannot buy at the uniform price.

The data sale is one stage earlier. Given the upstream price p_1, Firm 2 forms correct expectations about which consumers will hide and which ones will be revealed if it buys Firm 1's data. It can forecast its gross profit $\pi_2^{rev} + \pi_2^{hid}$ after the data purchase, and it will be prepared to pay up to the difference between this gross profit and its autarky profit, $\pi_2^{rev} + \pi_2^{hid} - \pi_2^A$, for the data. Thus Firm 1 offers the data at price $P = \pi_2^{rev} + \pi_2^{hid} - \pi_2^A$ and fully extracts any improvement in downstream profits relative to autarky.

Next consider, one stage earlier, a consumer's decision about whether to purchase at price p_1 in the upstream market. Let p_2^e be a consumer's expectation about the uniform price p_2 that will be charged to hidden consumers downstream, and let s_2^e be her expected surplus in the downstream market if she is hidden. She expects zero downstream surplus if she is revealed. Under the data-sharing policy regime, she effectively has three choices: buy at Firm 1 and enjoy surplus $v_1 - p_1$, wait and enjoy surplus s_2^e downstream, or don't buy at either firm, if neither surplus is positive. Under the policy regime with opt-out, the consumer has a fourth choice: she can buy at Firm 1, opt out, and then buy at Firm 2's hidden price, enjoying total surplus $v_1 - p_1 + s_2^e - c$.

Finally, at the first stage, Firm 0 sets its price p_1. We study subgame perfect Nash equilibria of the model. For any subgame following a choice p_1, this implies that Firm 2's hidden price p_2 maximizes π_2^{hid}, given correct beliefs about which consumers are hidden, and each consumer's decision at Firm 1 maximizes her total surplus, given correct beliefs about Firm 2's downstream pricing. Correct consumer beliefs imply that $p_2^e = p_2$ and $s_2^e = v_2 - p_2$.[1]

3 Warm-up case: "Rawlsian" consumers

We begin with a special case that is useful mainly as a stepping stone to the more complex settings that come later. Consumer valuations are distributed independently according to $F_1(v_1)$ and $F_2(v_2)$, with autarky (monopoly) prices p_1^A and p_2^A and profits π_1^A and π_2^A respectively. The key assumption is that a consumer does not learn v_2 until *after* participating in Market 1. As described above, Firm 1 learns the v_2 of any consumer who purchases in Market 1. We call these consumers "Rawlsian" because they make decisions in Market 1 under a veil of ignorance about whether they will be high or low value consumers downstream. While this assumption is mainly for expositional convenience, one could motivate it with the idea that "Market 1" stands

[1] These expressions will need to be revised slightly in cases where Firm 1 mixes over two prices.

in for a broad swathe of future markets that the consumer might participate in.

Data-sharing

Under the data-sharing regime DS, a consumer can hope to earn positive surplus in at most one of the two markets: she will be fully extracted downstream if she purchases from Firm 1, or she can decline to buy from Firm 1, remain hidden, and buy at the expected pool price p_2^e if she draws a large enough v_2. All consumers expect the same surplus $s_2 = E(\max(v_2 - p_2^e, 0))$ from the latter option, so they will buy from Firm 1 if and only if $v_1 \geq p_1 + s_2$. Thus a rational consumer does not buy from Firm 1 unless its price is low enough to compensate for the expected surplus she will give up downstream.

The pool of hidden consumers downstream has the same composition – that is, the same distribution of values v_2 – as the full population of consumers, since consumers select into the hidden pool on the basis of their upstream value v_1 but not on the basis of v_2. Consequently, Firm 2's optimal price to the hidden consumers remains $p_2 = p_2^A$. Because consumers have correct beliefs in equilibrium, we have $s_2 = E\left(\max\left(v_2 - p_2^A, 0\right)\right)$; this is just the expected consumer surplus in Market 2 at the monopoly price.

Firm 1 then sells to $q_1 = 1 - F_1(p_1 + s_2)$ consumers in Market 1, earning upstream profit $\pi_1 = p_1 q_1$. Moreover, each of these upstream buyers can be fully extracted downstream, generating profit $\pi_2^{rev} = q_1 E(v_2)$, while the waiters supply downstream profit $\pi_2^{hid} = (1 - q_1)\pi_2^A$. We will write $\Pi_1^{DS} = \pi_1 + \pi_2^{hid} + \pi_2^{rev} - \pi_2^A$ for Firm 1's total profit when it offers Firm 2 a take-it-or-leave-it data-sharing deal that holds Firm 2 to its autarky profit. Then we can write

$$\Pi_1^{DS}(p_1) = (1 - F_1(p_1 + s_2))\left(p_1 + \left(E(v_2) - \pi_2^A\right)\right)$$

Thus data-sharing has two key effects on Firm 1's profit function: its profit per customer rises by $E(v_2) - \pi_2^A$, the downstream profit difference between fully extracting a typical consumer vs. charging her the monopoly price. However, consumers anticipate this, so if it wishes to sell to the same quantity of consumers as under autarky, it must "discount" its price by s_2.

When Firm 1 prices optimally, data-sharing improves total welfare (relative to autarky), and these welfare gains are shared between Firm 1 and consumers. To see why welfare improves, note that Firm 1 can replicate its autarky market share by pricing at $p_1 = p_1^A - s_2$, thereby inducing $q_1^A = 1 - F_1\left(p_1^A\right)$ consumers to purchase upstream (thereby revealing themselves to be fully extracted in Market 2). This improves welfare in Market 2 – the revealed consumers are now served efficiently, and welfare on the rest is the same as under autarky – while upstream welfare is left unchanged. Thus total welfare rises. However, this is not Firm 1's optimal price – it would do better sell to more customers by reducing p_1 even further, and this improves welfare in both Market 1 and (because more consumers are revealed, and thus served efficiently) also in Market 2. To show this last point, we confirm that Firm 1's marginal profit at $p_1 = p_1^A - s_2$ is negative:

$$\left.\frac{d\Pi_1^{DS}}{dp_1}\right|_{p_1=p_1^A-s_2} = \left(1 - F\left(p_1^A\right) - f\left(p_1^A\right)p_1^A\right) - f\left(p_1^A\right)\left(E\left(v_2\right) - \pi_2^A - s_2\right) < 0$$

The first term in parentheses vanishes by definition. (It gives the first-order condition defining Firm 1's optimal autarky price p_1^A.) And the second term is positive: note that $E\left(v_2\right)$ is the total area under the Market 2 demand curve, while π_2^A and s_2 are the monopoly profit and consumer surplus respectively, so the difference $E\left(v_2\right) - \pi_2^A - s_2$ is the monopoly deadweight loss.

To see that both Firm 1 and consumers are better off, we go back to the hypothetical where Firm 1 charges $p_1 = p_1^A - s_2$ and earns $\Pi_1^{DS} = \pi_1^A + q_1^A DWL_2 > \pi_1^A$, where $DWL_2 = E\left(v_2\right) - \pi_2^A - s_2$ is the Market 2 monopoly deadweight loss again. That is, total gross profits rise by $s_2 + DWL_2$ on a consumer who is extracted downstream, reflecting the capture of consumer surplus s_2 and the 'market-expansion' gains DWL_2 from serving the market efficiently. Firm 1 does not keep the former (since it must compensate consumers up front for that extraction), but it benefits from the latter. In this hypothetical, consumers are no worse off. But then, because Firm 1's optimal price is actually lower than $p_1^A - s_2$, consumers will be strictly better off.

The main takeaway so far is that data-sharing can improve welfare *via* market-expansion downstream, and the incentive to profit from data-sharing can lead to lower prices and welfare-improving market expansion upstream as well.

Opt-out

Next we consider how equilibrium outcomes change if consumers can opt-out of having their data shared by paying nuisance cost $c > 0$. As our interest is in the case where c is small or negligible, we immediately assume $c < s_2$, and we will focus on the limiting case where $c \to 0$. Our claim is that under these assumptions the equilibrium reverts back to the autarky outcomes. To sketch the logic, we start by noting that the arguments supporting $p_2 = p_2^A$ remain valid. This means a consumer with upstream value v_1 faces three options: (i) earn surplus $s_1 = v_1 - p_1$ by purchasing at Firm 1 and letting her data be sold; (ii) wait to earn expected surplus s_2 in Market 2; or (iii) purchase upstream, opt out, and participate as a hidden consumer in Market 2, for total surplus $s_1 + s_2 - c$. By our assumption (iii) dominates (i), so she will buy from Firm 1 and opt out if $v_1 \geq p_1 + c$ and wait for Market 2 otherwise. At this point we take the $c \to 0$ limit. Then because Firm makes no data sales and sells only to consumers with $v_1 \geq p_1$, its incentives are exactly as under autarky, and it sets price p_1^A and earns total profit $\Pi_1^{OO} = \pi_1^A$. All consumers enter Market 2 as hidden (either because they didn't purchase at Firm 1, or because they did but opted out), and Firm 2 sets price p_2^A and earns profit π_2^A.

Here, and henceforth, we write CS^A, CS^{DS}, and CS^{OO} for total consumer surplus under the respective regimes. To summarize the points above, we have $\Pi_1^{DS} > \Pi_1^{OO} = \pi_1^A$ and $CS^{DS} > CS^{OO} = CS^A$. Perhaps the surprising result here is that (compared to unrestricted data sales), giving consumers the right to opt out makes them worse off. However, as described above, the logic is straightforward. Under data-sharing, rational consumers cannot be exploited *ex post* without compensating them up front, and they share in the welfare gains from market expansion. However, as we shall see in the next section, there are countervailing effects when consumers know their downstream valuations when they are making decisions at Firm 1.

4 Box-linear demand with imperfect data-gathering

This section is named for the specific functional form assumptions that will be imposed (in the interest of obtaining closed-form solutions). All consumers have common valuation v_1 in Market 1, while valuations in Market 2 are uniformly distributed: $v_2 \sim U[0, 1]$. Thus demand in the two markets is rectangular (box-shaped) or linear respectively. As shall be made clear below, the Market 1 assumption is substantively restrictive, in that it excludes certain effects that could tilt some results in other directions. The Market 2 assumption is less critical; linear demand is helpful for getting closed-forms, but much of the logic would go through with a more general demand function.

Henceforth a consumer knows both v_1 and v_2 from the start of the game. In this section, we assume that data-gathering by Firm 1 is imperfect. For each upstream customer, Firm 1 learns the consumer's true downstream value v_2 with probability $\lambda \in [0, 1]$ and learns nothing about the consumer with probability $1 - \lambda$. Data sales to Firm 2 are as above, but of course Firm 1 can only sell the data of the λ fraction who are revealed; the remainder of its customers are offered Firm 2's uniform price. Consumers understand all of this but do not know in advance whether they will be revealed or hidden if they buy from Firm 1.

Note that autarky has the standard monopoly prices $p_1^A = v_1$ and $p_2^A = \frac{1}{2}$, with profits $\pi_1^A = v_1$ and $\pi_2^A = \frac{1}{4}$ respectively. Consumers are fully extracted in Market 1 and receive surplus $\frac{1}{8}$ in Market 2, for a total surplus $CS^A = \frac{1}{8}$. We will lead off with analysis of data-sharing and opt-out when data-gathering is perfect before turning to the general case.

4.1 Perfect data-gathering

For now we assume $\lambda = 1$. For Firm 2 the new implication involves advantageous selection in the pool of hidden consumers in Market 2. Under data-sharing the consumers who reveal themselves tend to be those with low v_2, since these are the consumers with little surplus to protect in Market 2. Consumers with higher than average v_2 tend to remain hidden, hoping to buy at the pool price p_2, but this encourages Firm 1 to raise its price. In this sense, hiding imposes a negative externality on other consumers, since it sweetens the pool of hidden consumers which leads all of them to face a higher price. (Thus what we call advantageous selection for Firm 2 could be considered adverse selection for consumers.) This effect works to undo the earlier conclusion that data-sharing is unambiguously good for both consumers.

4.1.1 Data-sharing (DS)

The key insight is that when data-gathering is perfect, Firm 1 can leverage the selection effect to induce a complete upward unraveling of Market 2, up to a hidden price $p_2 = 1$. To see the logic, suppose Firm 1 has set a price $p_1 < v_1$ that gives consumers positive surplus $s_1 = v_1 - p_1$. Consider a candidate continuation equilibrium with expected downstream price p_2^e. Because a consumer is certain to be fully extracted downstream if she buys from Firm 1, she should buy and reveal herself if $s_2 = v_2 - p_2^e < s_1$, or wait for Market 2 if $s_2 > s_1$. Consequently all of Firm 2's hidden consumers have values $v_2 \in [v_2^*, 1]$, where $v_2^* = p_2^e + s_1$ is the threshold valuation type; that is, they all earn strictly positive surplus at price p_2^e. But this cannot be an equilibrium outcome, as Firm 2 would gain by holding the consumers up with a higher than expected price $p_2 = p_2^e + s_1$. This argument applies whenever $p_2^e + s_1 < 1$, so any equilibrium price must satisfy $p_2 \in [1 - s_1, 1]$. For $s_1 \to 0$, we have $p_2 \to 1$. (Some readers will recognize this as a version of the Diamond paradox from search theory.)

Our description of the full equilibrium of the game will be somewhat informal; we wish to make the case that Firm 1 will set $s_1 \to 0$ by choosing a price p_1 arbitrarily close to its autarky price $p_1^A = v_1$. Under this strategy, Firm 1 induces $p_2 \to v_2^* \to 1$, so the market for hidden consumers disappears, and all consumers buy at Firm 1, are revealed, and are fully extracted downstream. This permits Firm 1 to capture the full, welfare-maximizing social surplus $\pi_1 = v_1$ upstream and also to extract the full social surplus downstream, $\pi_2^{rev} = \frac{1}{2}$. After leaving Firm 2 with its outside option, Firm 1 earns $\Pi_1^{DS} = \pi_1 + \pi_2^{rev} + \pi_2^{hid} - \pi_2^A = v_1 + \frac{1}{4}$. Since Firm 1 has extracted the full social surplus from both markets, minus the outside option it is obliged to pay to Firm 2, this strategy is unambiguously profit-maximizing. As consumers are fully extracted in both markets, we have $CS^{DS} = 0$. We summarize the welfare conclusions below.

Result 1 *When data-gathering is perfect, Firm 1 is better-off and consumers are worse-off with data-sharing, compared to autarky:* $\Pi_1^{DS} > \pi_1^A$ *and* $CS^{DS} < CS^A$.

4.1.2 Opt-out (OO)

As earlier, we start with opt-out at cost $c > 0$ and focus on the $c \to 0$ limit. Here the key insight is that a strictly positive opt-out cost creates the same full unraveling of Market 2 described above, even if c is quite small. Since we are interested in the $c = 0$ limit, consider a continuation after Firm 1 has offered $s_1 = v_1 - p_1 > c$. This ensures that every consumer prefers buying at Firm 1 and opting out (surplus $s_1 - c + s_2$) over waiting for Market 2 (surplus s_2). Thus all consumer buy at Firm 1 and either opt out or let themselves be revealed (surplus s_1); the indifferent consumer satisfies $s_2 = c$, and so the threshold valuation is $v_2 = p_2^e + c$.

From here, the argument for unraveling of Market 2 is just as above. Every consumer choosing to remain in the hidden pool downstream by opting out expects surplus of at least c (to compensate for the opt-out cost), but this leaves them vulnerable to being held up with a higher-than-expected downstream price of $p_2 = p_2^e + c$. As earlier, this logic applies to any $p_2^e < 1 - c$, so any equilibrium price must satisfy $p_2 \in [1 - c, 1]$. As $c \to 0$, the equilibrium price must satisfy $p_2 \to v_2^* \to 1$. Note that this limiting result is true regardless of Firm 1's price, as long as $s_1 > 0$. Note also that in this limit, all consumers purchase at Firm 1 and permit themselves to be extracted downstream. No consumers choose to opt out, even though the cost of doing so is vanishingly low!

Now consider Firm 1's pricing decision in the $c \to 0$ limit. By the argument above, any $s_1 > 0$ suffices to induce all consumers to purchase at Firm 1 and permit themselves to be revealed. Thus, just as in the DS case, Firm 1 does best to offer $s_1 \to 0$ (corresponding to a price $p_1 \to v_1$) and extract the full surplus in both markets (minus paying Firm 2 its outside option). Thus the equilibrium outcome under free opt-out looks identical to the outcome under data-sharing:

Result 2 *When data-gathering is perfect, profits and consumer surplus are identical under regimes DS and OO. That is,* $\Pi_1^{OO} = \Pi_1^{DS} > \pi_1^A$ *and* $CS^{OO} = CS^{DS} = 0 < CS^A$.

Several factors here make the adverse selection problem for consumers (and the consequent unraveling of Market 1) particularly severe. One of these factors is perfect data-gathering – this means that there are no hidden consumers in Market 2 who are hidden by chance, rather than selecting into the hidden pool because of a high v_2. A second factor is the box-shaped demand in Market 1. This makes it inexpensive for Firm 1 to offer a small positive surplus to 100% of its potential consumers; if demand were downward-sloping, it would have to drop its price substantially to achieve the same result. This is important because Firm 1 can intentionally exacerbate downstream adverse selection by increasing its upstream market share; it has an incentive to do this because a rising p_2 encourages consumers to shift from hiding to revealing. Firm 1 will indulge this incentive to exacerbate downstream adverse selection more when it can do so in a cost-effective way (as it can with box-shaped demand upstream).

4.2 Imperfect data-gathering

Now we turn to the case where $\lambda < 1$, where we will be working toward two headline conclusions. First, the selection effect described above is still present, albeit in a less extreme form, and it continues to drive many of the results. Second, we find that switching from unrestricted data-gathering to an opt-out regime unambiguously helps the data-gathering firm and hurts consumers. The first two subsections involve somewhat dry derivation of equilibrium profits and consumer surplus in cases DS and OO. Most of the intuition is held for the subsequent section, where we make comparisons across these policy regimes.

4.2.1 Data-sharing (DS)

The analysis here will be kept fast and light by relying heavily on textbook results for linear demand curves. Consider a situation where Firm 1 has offered surplus $s_1 > 0$ by setting price $p_1 = v_1 - s_1$. We will refer extensively to Figure 1 to illustrate key features of an equilibrium in the downstream market for hidden consumers. Taking a bit more care than earlier, write an unrevealed consumer's anticipated surplus from participating in Market 2 as $s_2 = \max(v_2 - p_2^e, 0)$. With imperfect data-gathering, a consumer who buys at Firm 1 anticipates total expected surplus $s_1 + (1 - \lambda) s_1$ (as there is a $1 - \lambda$ chance she remains hidden). The threshold for indifference between these waiting for Market 2 or purchasing in Market 1 then becomes $s_1 = \lambda s_2$, or $v_2^* = p_2^e + s_1/\lambda$.

Consequently, Firm 2's hidden demand curve is kinked as in Figure 1, as it includes all of the consumers with $v_2 > v_2^*$, but only a fraction $1 - \lambda$ of those with $v_2 < v_2^*$. Thus at prices below v_2^*, Firm 2's inverse demand curve has slope $\frac{1}{1-\lambda}$ as shown. (The dashed line indicates what demand would be if revealed consumers had not been 'carved out.') As earlier, the consumers who intentionally waited to buy at Firm 2 all have valuations that strictly exceed p_2^e (by at least s_1/λ in this case). But now Firm 2's temptation to hold those consumers up with a higher-than-expected price is tempered by the fact that it would lose sales to the consumers who remain hidden because they got lucky. At an equilibrium, as illustrated on the figure, consumers act based on correct price expectations (so $v_2^* = p_2 + s_1/\lambda$) and p_2 must be profit-maximizing for Firm 2. For a linear demand curve with zero marginal cost, the latter condition implies that Firm 2's sales are half of its quantity intercept; that is, $q_2 = \frac{1}{2}\bar{q}$,

Figure 1: The hidden-consumer demand curve in Market 2

where $\bar{q} = 1 - \lambda v_2^*$. We solve these two conditions to characterize the equilibrium price and threshold valuation induced by s_1:

$$p_2 = \frac{1 - s_1}{2 - \lambda} \quad \text{and} \quad v_2^* = \frac{1 + 2\frac{1-\lambda}{\lambda} s_1}{2 - \lambda} \text{ if } s_1 \in \left[0, \frac{\lambda}{2}\right] \tag{1}$$

We note the upper limit on s_1. The threshold valuation is increasing in s_1, and at $s_1 = \lambda/2$ it reaches the corner $v_2^* = 1$, corresponding to the case where all consumers buy at Firm 1 and accept that their downstream values will be revealed. That is, Firm 1 can induce complete unraveling of the downstream hidden consumer market by offering a surplus of at least $\lambda/2$. There are no circumstances where it will offer a surplus larger than this upper limit, as this would simply reduce its upstream profit without changing its profits from data sales. The analysis of Firm 1's optimal price is complicated somewhat by the fact that there are three cases to consider. For low enough λ (when data-gathering is poor), setting $s_1 = \lambda/2$ will be optimal, while for high enough λ (when data-gathering is accurate), it is optimal to set $s_1 = 0$ and price at $p_1 = v_1$. For intermediate λ, offering an 'interior' surplus $s_1 \in (0, \lambda/2)$ is optimal. Our analysis proceeds by assuming this case, deriving Firm 1's optimal s_1, and then checking whether the conditions $s_1 \geq 0$ or $s_1 \leq \lambda/2$ are violated.

We start by characterizing the components of Firm 1's profit, assuming s_1 is interior. Given the linear demand curve, Firm 2 serves $q_2 = (1-\lambda)\,p_2$ hidden consumers, with profit $\pi_2^{hid} = p_2 q_2 = (1-\lambda)\,p_2^2$. A fraction λ of consumers with valuations below v_2^* are revealed and fully extracted; this implies revealed profits $\pi_2^{rev} = \frac{1}{2}\lambda\,(v_2^*)^2$, as illustrated on the figure. In Market 1, Firm 1 sells to a fraction v_2^* of its consumers (all those with $v_2 < v_2^*$). Thus Firm 1's total profit is

$$\Pi_1^{DS}(s_1) = v_2^*(v_1 - s_1) + (1-\lambda)\,p_2^2 + \frac{1}{2}\lambda\,(v_2^*)^2 - \pi_2^A \tag{2}$$

where $\pi_2^A = \frac{1}{4}$. Using (1), this profit function is quadratic in s_1, so it is straightforward to show it is maximized at

$$s_1^* = v_1 - \frac{1}{2}\frac{\lambda}{1-\lambda}. \tag{3}$$

as long as the solution satisfies $s_1^* \in [0, \lambda/2]$. To delineate the three cases, define λ_L and λ_H implicitly by $v_1 = \frac{1}{2}\frac{\lambda_L}{1-\lambda_L}$ and $\frac{\lambda_H}{2} = v_1 - \frac{1}{2}\frac{\lambda_H}{1-\lambda_H}$. Both thresholds can be solved for explicitly (we have $\lambda_L = v_1 + 1 - \sqrt{v_1^2+1}$ and $\lambda_H = \frac{2v_1}{2v_1+1}$), and it is straightforward to confirm that they satisfy $0 < \lambda_L < \lambda_H < 1$. The full description of Firm 1's optimal strategy is then $s_1^* = \lambda/2$ if $\lambda < \lambda_L$, $s_1^* = 0$ if $\lambda > \lambda_H$, and (3) if $\lambda \in [\lambda_L, \lambda_H]$.

We defer intuition for the moment in order to proceed straightaway to Firm 1's optimized profit. We have:

$$\Pi_1^{DS} = \begin{cases} v_1 - \frac{\lambda}{4} & \text{if } \lambda \in [0, \lambda_L) \\ v_2^*(v_1 - s_1^*) + (1-\lambda)\,p_2^2 + \frac{1}{2}\lambda\,(v_2^*)^2 - \frac{1}{4} & \text{if } \lambda \in [\lambda_L, \lambda_H] \\ p_2 v_1 + \left(1 - \frac{\lambda}{2}\right)p_2^2 - \frac{1}{4}\Big|_{p_2 = \frac{1}{2-\lambda}} & \text{if } \lambda > (\lambda_H, 1] \end{cases} \tag{4}$$

where for the middle case v_2^*, p_2, and s_1^* are given by (1) and (3). Note that in the third case where setting $s_1^* = 0$ is optimal, the price and threshold consumer in the downstream market collapse to $p_2 = v_2^* = \frac{1}{2-\lambda}$.

For consumer surplus, we give a simple (and hopefully intuitive) accounting. (The result can be confirmed with a more methodical approach.) Total surplus in Market 1 is $v_2^* s_1^*$ (because v_2^* is the fraction of consumers who purchase). Market 2 surplus accrues only to hidden consumers, as the revealed are fully extracted. If all consumers were hidden, downstream surplus would be $\frac{1}{2}(1-p_2)^2$. Consulting Figure 1, we see that this overstates the actual surplus by a sliver $\frac{\lambda}{2}(v_2^* - p_2)^2$, corresponding to consumers $v_2 \in [p_2, v_2^*]$ who took a chance by purchasing at Firm 1, and (unluckily) were revealed. In summary, we have

$$CS^{DS} = v_2^* s_1^* + \frac{1}{2}(1 - p_2)^2 - \frac{\lambda}{2}(v_2^* - p_2)^2, \tag{5}$$

where the equilibrium variables are as specified in (4).

4.2.2 Opt-out

As usual, we start with $c > 0$ but quickly move to the limiting case where $c \to 0$. The key takeaway here will be that Firm 1 always finds it optimal to offer an infinitessimal surplus $s_1 \to 0$ to the upstream customers. The intuition is that the opt-out cost c is very effective on its own at creating adverse selection downstream, so Firm 1 finds it unnecessary to supplement this effect by reducing its own price.

Consider a situation where Firm 1 has offered $s_1 > c > 0$. As in the perfect data-gathering section, this ensures that every consumer purchases at Firm 1 and either opts out or takes her chances on being revealed. The threshold for indifference is $\lambda s_2 = c$, which implies the threshold consumer and the equilibrium price are related by $v_2^* - p_2 = c/\lambda$. While consumers with values above v_2^* are hidden in Market 2 for a different reason than in the previous section – they bought and opted-out at Firm 1, rather than waiting – the hidden demand curve faced by Firm 2 is still as depicted by Figure 1. The only difference is that c stands in for the role played by s_1 in the figure and in the equilibrium conditions (1). Proceeding to the $c \to 0$ limit, we have the following equilibrium conditions for the downstream market:

$$v_2^* = p_2 = \frac{1}{2 - \lambda}. \tag{6}$$

To summarize, for c small enough, the extent of 'upward unraveling' in the downstream market for hidden consumers is entirely dictated by the opt-out cost, not by the amount of surplus that Firm 1 offers, and for $c \to 0$ the downstream equilibrium tends to a limiting outcome that depends on how accurate data-gathering is, but not on s_1.

The components of Firm 1 profits look, with appropriate adjustments, as they did in the previous section. We have upstream profit $\pi_1 = v_1 - s_1$, since all consumers purchase in Market 1. Market 2 hidden-consumer profit is $\pi_2^{hid} = (1 - \lambda) p_2^2$, and revealed-consumer profit is $\pi_2^{rev} = \frac{1}{2}\lambda p_2^2$. As neither of the latter two terms depends on s_1, Firm 1 clearly does best to set $s_1 \to 0$. Using $\pi_2^A = \frac{1}{4}$, Firm 1's optimal profit under opt out is therefore

$$\Pi_1^{OO} = v_1 + \left(1 - \frac{\lambda}{2}\right) p_2^2 - \frac{1}{4}\Big|_{p_2 = \frac{1}{2-\lambda}}. \tag{7}$$

The fact that Firm 1 fully extracts its upstream purchasers makes consumers surplus rather straightforward. We need only worry about consumer surplus among hidden consumers in Market 2. Furthermore, because v_2^* and p_2 have collapsed together, the correction term from the previous section vanishes, and we simply have

$$CS^{OO} = \frac{1}{2}\left(1 - p_2\right)^2\Big|_{p_2 = \frac{1}{2-\lambda}} \tag{8}$$

4.2.3 Welfare comparisons: who gains and who loses under an opt-out policy

The headline result here will be that a switch from regime DS to OO unambiguously benefits Firm 1 and hurts consumers. For the sake of comparison in this section, let p_2, v_2^*, and s_1^* always refer to equilibrium outcomes under DS, and write $\hat{p}_2 = 1/\left(2 - \lambda\right)$ for the equilibrium price under OO.

Proposition 1 *If $\lambda \in (0, 1)$, Firm 1's profit is strictly higher under regime OO than regime DS.*

Proof. We consider the three cases for λ one at a time. Start with $\lambda > \lambda_H$. In this case, a comparison of (4) and (7) makes it clear that Firm 1 makes identical profits from data sales in either case, but makes strictly greater profit in Market 1 (v_1 versus $p_2 v_1$) under policy OO. Next consider $\lambda < \lambda_L$. Here we have $\Pi_1^{OO} - \Pi_1^{DS} = \left(1 - \frac{\lambda}{2}\right) p_2^2\Big|_{p_2 = \frac{1}{2-\lambda}} - \frac{1}{4}\left(1 - \lambda\right) = \frac{\lambda}{4}\left(\frac{3-\lambda}{2-\lambda}\right) > 0$. The intermediate case $\lambda \in [\lambda_L, \lambda_H]$ is more interesting. Note that $p_2 < \hat{p}_2$ and write

$$\begin{aligned} \Pi_1^{OO} &= v_1 + \left(1 - \lambda\right) p_2^2 + \frac{1}{2}\lambda p_2^2 - \frac{1}{4} + \left(1 - \frac{\lambda}{2}\right)\left(\hat{p}_2^2 - p_2^2\right) \\ & v_2^*\left(v_1 - s_1^*\right) + \left(1 - \lambda\right) p_2^2 + \frac{1}{2}\lambda\left(v_2^*\right)^2 - \frac{1}{4} \end{aligned}$$

Then the profit difference can be simplified to

$$\begin{aligned}
\Pi_1^{OO} - \Pi_1^{DS} &= Y + Z \\
\text{where } Y &= v_2^* s_1^* - \frac{\lambda}{2}\left((v_2^*)^2 - p_2^2\right) \text{ and } Z = (1 - v_2^*)\, v_1 + \left(1 - \frac{\lambda}{2}\right)\left(\hat{p}_2^2 - p_2^2\right)
\end{aligned}$$

Because Z is strictly positive, it suffices to show $Y > 0$. Using the equilibrium condition $v_2^* - p_2 = s_1^*/\lambda$ twice, we have

$$\begin{aligned}
Y &= v_2^* s_1^* - \frac{\lambda}{2}\left(v_2^* - p_2\right)\left(v_2^* + p_2\right) \\
&= v_2^* s_1^* - \frac{s_1^*}{2}\left(v_2^* + p_2\right) \\
&= \frac{1}{2\lambda}\left(s_1^*\right)^2 > 0
\end{aligned}$$

■

Proposition 2 *Consumers are worse off under regime OO than regime DS:* $CS^{OO} \leq CS^{DS}$. *If data-gathering is sufficiently inaccurate ($\lambda < \lambda_H$), the ranking is strict:* $CS^{OO} < CS^{DS}$.

Proof. If $\lambda \geq \lambda_H$, then $p_2 = \hat{p}_2$, and the first and third terms of (5) collapse. It is then immediate that $CS^{OO} = CS^{DS}$. Otherwise, if $\lambda < \lambda_H$, we have $s_1^* > 0$ and $p_2 < \hat{p}_2$. The surplus difference may be written

$$\begin{aligned}
CS^{DS} - CS^{OO} &= Y' + Z' \\
\text{where } Y' &= v_2^* s_1^* - \frac{\lambda}{2}\left(v_2^* - p_2\right)^2 \text{ and } Z' = \frac{1}{2}\left(1 - p_2\right)^2 - \frac{1}{2}\left(1 - \hat{p}_2\right)^2
\end{aligned}$$

Because the hidden price is strictly higher under OO, we have $Z' > 0$, and so it suffices to show $Y' \geq 0$, for which we use the equilibrium condition $v_2^* - p_2 = s_1^*/\lambda$ to write

$$Y' = v_2^* s_1^* - \frac{1}{2}\left(v_2^* - p_2\right) s_1^* = \frac{v_2^* + p_2}{2} s_1^* > 0.$$

■

Discussion Figure 2 plots profit and consumer surplus for each regime as a function of the data-gathering accuracy parameter λ. (Consumers' upstream valuation is assumed to be $v_1 = 1$.) At an intuitive level, why is Firm 1 better off when consumers can opt out of data sales than when they cannot? If Firm 1 does not 'discount' to its

Figure 2: Profits and consumer surplus under different data policies. $(v_1 = 1)$

upstream customers (that is, if it offers $s_1 = 0$ rather than $s_1 > 0$), then consumer self-selection leads to the same outcome in the downstream market under either policy DS or OO. However, in the former case, Firm 1 sacrifices upstream sales to customers who wish to remain hidden, while in the latter case it sells to these customers too. This logic seems likely to hold for any distribution of Market 2 valuations, not just the linear demand case. Under policy DS, Firm 1 can try to claw back some of those foregone upstream sales by cutting its price ($s_1 > 0$). Raising s_1 has an effect $d\Pi^{DS}/ds_1$ that can be decomposed into a gain in Market 1 sales, $(v_1 - s_1)\frac{dv_2^*}{ds_1}$, a reduction in inframarginal Market 1 profits, and an increase in downstream profits (from converting additional hidden consumers to revealed). Proposition 1 shows that the sum of the second two effects is negative, so gaining back Market 1 sales by discounting cannot be as profitable for Firm 1 as getting those sales 'for free' under policy OO. Preliminary analysis indicates that this part of the analysis is valid for a broad range of distributions $F_2(v_2)$.[2]

Conversely, why are consumers worse off with the right to opt out than without it? Holding total welfare constant, Firm 1's gain is consumers' loss. However, this is an insufficient explanation because the effect of switching from DS to OO on welfare is ambiguous – there are gains from expanded sales in Market 1, but (it can be shown) there is a countervailing reduction in welfare in Market 2. A slightly better explanation, mirroring the proof of Proposition 2, is that regime DS induces the lower downstream price, $p_2 \leq \hat{p}_2$. It is true that DS-consumers must give up on Market 1 surplus to take advantage of that price, but since they would certainly get zero surplus in Market 1 under OO, regime DS still comes out looking better.

As suggested above, we believe the results would be qualitatively similar with alternative assumptions about demand in Market 2. However, it is worth pointing out two ways in which the Market 1 demand setup is special. First, as noted earlier, when consumers have homogeneous upstream valuations, it costs Firm 1 little to sweeten the pot with a positive s_1 for all of its consumers, thereby inducing strong selection in Market 2. If Market 1 demand were downward sloping, expanding its sales would require increasingly larger discounts, so presumably the equilibrium selection effect in Market 2 would be more limited. Second, the upstream demand setup tends to minimize one unprofitable effect of switching from DS to OO. Holding prices constant,

[2] For example, the logic holds whenever demand $1 - F_2(v_2)$ is convex, but it appears that this condition could be weakened substantially.

some consumers who would have bought at Firm 1 and revealed under DS will buy at Firm 1 and opt out under OO – Firm 1 loses its data profits on these consumers. In the current setup, those losses tend to be small because s_1 tends to be small, and so there are relatively few consumers satisfying $s_2 \in [0, s_1]$ who prefer buying in Market 1 if they face an either-or choice, but who would also like to buy in Market 2.

5 Defaults: Opt-out or Opt-in?

Behavioral economists have advanced the idea that consumers presented with several options have a tendency to stick with whichever one is the default. In the context of consumer data, this idea is in the background of policy debates about whether consumer data should be shareable by default, but with a consumer right to opt out, or whether the default should be no sharing, but with a consumer right to opt in. For many economists, the instinctive response would probably be that there is no effective difference between the two policies as long as any decisions costs are small. Our purpose in this section is to give a simple example in which different default policies lead to drastically different equilibrium outcomes. This is true even though the nuisance cost to a consumer of switching to the non-default choice is vanishingly small, and without resorting to *ad hoc* behavioral preferences. The logic relies on unraveling arguments like those presented earlier.

We retain the demand structure from the last section and specialize to the case where $v_1 = 1$. Data-gathering by Firm 1 is still imperfect, but in a different way than previously. At the time she participates in Market 1, a consumer believes her downstream valuation to be $v_2 \sim U[0, 1]$; if she purchases in Market 1, this is what Firm 1 observes. With probability $\lambda > \frac{1}{2}$, this remains her valuation when she reaches Market 2. However, with probability $1 - \lambda$ she faces a shock after leaving Market 1 that gives her a new, i.i.d. valuation draw $v_2' \sim U[0, 1]$. The main purpose of these assumptions is to create circumstances where a consumer strictly prefers to opt-in (which has not been a possibility heretofore). Our model of opt-out with a nuisance cost $c \to 0$ is as described earlier, and we introduce an opt-in regime (OI) under which Firm 1 can sell Firm 2 what it knows about a customer only if she pays nuisance cost $c \to 0$ to opt in.

5.1 Opt-Out

We claim that despite the change in the information structure, the equilibrium is substantively the same as in Section 4.1: the downstream market for hidden consumers unravels up to $v_2^* = p_2 = 1$, and all consumers have their data sold to Firm 2. To show this, we start with an analysis of Market 2.

Under the new information structure, when Firm 2 buys data 'revealing' that Firm 1 observed a consumer to be type v_2, it knows that her true valuations is v_2 with probability λ and a new, uniform draw from $[0, 1]$ with probability λ. It can be confirmed that Firm 2's most profitable option is to charge this consumer her apparent valuation, $p_2 = v_2$, as long as that apparent valuation is not too low. The consumer purchases if the apparent valuation turns out to be correct, or if she draws a new value above v_2, so this earns Firm 2 a profit $\pi_2^{rev}(v_2) = (1 - (1 - \lambda) v_2) v_2$ on this consumer. Alternatively, if $v_2 < \bar{v} < \frac{1}{2}$, then Firm 2 does better to gamble on a new valuation draw by charging $p_2 = \frac{1}{2}$, thereby earning $\pi_2^{rev}(v_2) = \frac{1}{4}(1 - \lambda)$. Under our assumptions, the threshold $\bar{v}$ satisfies $\bar{v} = \frac{1}{2}\frac{1}{1-\lambda}\left(1 - \sqrt{\lambda(2 - \lambda)}\right) \lesssim 0.14$ for any $\lambda > \frac{1}{2}$. In the case where she is charged her apparent valuation, a revealed consumer still has a chance at a positive surplus, since she might draw a new valuation higher than v_2. In this case, her expected downstream surplus if revealed is $s_r(v_2) = \frac{1}{2}(1 - \lambda)(1 - v_2)^2$. (There is a different expression if $v_2 < \bar{v}$, but we ignore this case for now.)

Next consider sales to hidden consumers in Market 2, under the assumption that Firm 1 has offered surplus $s_1 > c > 0$. As earlier, this ensures that all consumers purchase at Firm 1, with some opting out, and some revealing. At the point when she must decide between these two options, a consumer with apparent downstream value v_2 anticipates expected surplus $s_h(v_2; p_2^e) = \lambda \max(v_2 - p_2^e, 0) + \frac{1}{2}(1 - \lambda)(1 - p_2^e)^2$ if she remains hidden. The first term is her surplus if her apparent surplus turns out to be correct, while the second term reflects expected surplus in the event that she draws a new value. Then this consumer will opt out if $s_h(v_2; p_2^e) - c > s_r(v_2)$ and let herself be revealed otherwise. At this point, it is helpful to note that $s_h(\cdot; p_2^e)$ and $s_r(\cdot)$ satisfy a single-crossing property. The former is increasing in v_2, while the latter is decreasing in v_2, and their intersection is precisely at an apparent valuation equal to the expected downstream price. That is to say, $s_h(p_2^e; p_2^e) = s_r(p_2^e)$, with $s_h(v_2; p_2^e) > s_r(v_2)$ for $v_2 > p_2^e$, and the opposite ranking for $v_2 < p_2^e$. The broad takeaway from this is that consumers with high (apparent) valuations hide and those with low (apparent) valuations reveal, just as earlier. The threshold apparent valuation satisfies $s_h(v_2^*; p_2^e) = s_r(v_2^*) + c$, so we conclude that this threshold type strictly exceeds the expected price: $v_2^* > p_2^e$.

Our next task is to show that in setting its actual hidden-consumer price p_2, Firm 2 confronts the same hold-up issue discussed earlier. We assume $v_2^* > \frac{1}{2}$, as it can be shown (justification deferred) that $v_2^* \leq \frac{1}{2}$ is inconsistent with equilibrium. Then Firm 2 has a pool of $1 - v_2^*$ hidden consumers, of whom a fraction λ have true valuations distributed $v_2 \sim U[v_2^*, 1]$, while the remaining fraction $1 - \lambda$ have new valuations $v_2 \sim U[0, 1]$. Thus, although the rationale is rather different than earlier, the kinked hidden demand curve of Figure 1 remains an accurate depiction. Under the assumption that $v_2^* > \frac{1}{2}$, it is profit-maximizing for Firm 2 to price at the kink: $p_2 = v_2^*$. But since we just showed that $v_2^* > p_2^e$ must hold, this implies holding consumers up with a higher-than-expected price. While we will hand-wave over the fine details, as earlier this rules out an equilibrium at any price below a neighborhood around 1, and as $c \to 1$, the equilibrium threshold v_2^* and price p_2 both converge to 1.[3]

To summarize, by offering any strictly positive surplus $s_1 > c \to 0$, Firm 1 can ensure that the hidden-consumer market unravels upward to a price of $p_2 = 1$, and that all consumers purchase at Firm 1 and permit their data to be sold. Firm 1's optimal strategy is to precipitate this outcome with a vanishingly small surplus offer $s_1 \to 0$. As our main objective in this example is to understand consumer outcomes, we focus on consumer surplus in this equilibrium. Under autarky, consumers earn zero surplus upstream and expected surplus $\frac{1}{8}$ downstream at the monopoly price $p_2^A = \frac{1}{2}$; thus $CS^A = \frac{1}{8}$. In the equilibrium with opt out, upstream consumer surplus is zero as well, and all consumers are revealed downstream. At this point we must tidy up the loose end of revealed consumers with apparent values $v_2 < \bar{v}$. As discussed above, these consumers will face a price of $\frac{1}{2}$, so their expected surplus is $s_r(v_2) = (1 - \lambda)\left(\frac{1}{8}\right)$ which (because $\bar{v} < \frac{1}{2}$) is smaller than the expression $\frac{1}{2}(1 - \lambda)(1 - v_2)^2$ given above. Consequently, consumer surplus can be bounded above:

$$CS^{OO} = \int_0^1 s_r(v_2)\, dv_2 < \frac{1}{6}(1 - \lambda) < CS^A$$

[3]Let us tidy up a few details. If the kink in demand is at $v_2^* \leq \frac{1}{2}$, then the profit-maximing price is $p_2 = \frac{1}{2}$, which is inconsistent with $v_2^* > p_2^e$. And given $v_2^* > \frac{1}{2} > \bar{v}$, we are justified in ignoring for the time being the surplus expression for revealed consumers with $v_2 < \bar{v}$.

5.2 Opt-In

Now consider the situation where consumers must pay a vanishing cost $c \to 0$ to opt in to having their data sold. This may be appealing to a consumer who has a low apparent value and expects to be priced out of Market 2 at the uniform price. By revealing that low apparent value she induces Firm 2 to lower its price, which leaves her with more surplus if she is lucky enough to draw a new, higher value.

As earlier, suppose that Firm 1 offers $s_1 > c$, so once again all consumers will buy at Firm 1. Earlier expressions for $s_h()$ and $s_r()$ remain valid, but now a consumer chooses to opt in if $s_r(v_2) - c > s_h(v_2; p_2^e)$, or stick to the default (remaining hidden) if this inequality is reversed. Just as before, the two surplus expressions cross at $v_2 = p_2^e$, so in this case the condition for the threshold consumer type, $s_r(v_2^*) - c > s_h(v_2; p_2^e)$, implies that we must have $v_2^* < p_2^e$ whenever $c > 0$.

Just as in the previous section, Firm 2's hidden-consumer demand curve is kinked at v_2^*, with profit maximized at $p_2 = v_2^*$, if $v_2^* > \frac{1}{2}$, or $p_2 = \frac{1}{2} = p_2^A$ if $v_2^* \leq \frac{1}{2}$. Just as earlier, this creates an inconsistency, although in the opposite direction: whenever consumers expect a price strictly higher than the autarky price $p_2^A = \frac{1}{2}$, Firm 2 has an incentive to price slightly lower than expected. Consequently, v_2^* and p_2 must fall until $v_2^* < \frac{1}{2}$, at which point Firm 2's optimal price $p_2 = p_2^A$ falls no further. Then the equilibrium condition becomes $s_h\left(v_2^*; p_2^A\right) = s_r(v_2^*) - c$, which is satisfied for v_2^* slightly below $\frac{1}{2}$ for c small. This threshold consumer v_2^* gets a slightly better price by revealing than by staying hidden (v_2^* versus $\frac{1}{2}$), which is just enough to compensate her for the nuisance cost c.

Meanwhile, Firm 1 has an incentive to keep s_1 strictly larger than c (lest it lose sales), but it profits most by reducing s_1 as much as possible: $s_1 \to c$. In the limit, as the nuisance cost vanishes, the condition for the threshold consumer tends to $s_h\left(v_2^*; p_2^A\right) = s_r(v_2^*)$, which is satisfied precisely at $v_2^* = p_2^A$. Thus under the opt-in policy prices tend to their autarky levels, $p_2 = p_2^A$ and $p_1 \to p_1^A = 1$, and 50% of consumers (those with $v_2 < v_2^* = \frac{1}{2}$) opt in. That is to say, the only consumers who opt in are those who do not expect to purchase at p_2^A if their apparent values turn out to be correct.

Clearly the downstream price is quite different under regimes OO and OI, and the fraction of consumers whose data is sold shifts from 100% to 50%. To assess which regime leaves consumers better off overall, we compute CS^{OI}. As consumers are fully

Figure 3: Consumer surplus under different default data policies

extracted in Market 1, we have

$$CS^{OI} = \int_0^{v_2^*} s_r(v_2)\, dv_2 + \int_{v_2^*}^1 s_h\left(v_2; p_2^A\right) dv_2$$

This can be regrouped as

$$CS^{OI} = \lambda CS^A + (1-\lambda) \int_0^1 \max\left(CS^A, \frac{s_r(v_2)}{1-\lambda}\right) dv_2 \geq CS^A$$

The interpretation is that any consumer who faces price p_2^A is just as well off as under autarky, but some consumers ($v_2 \in (\bar{v}, v_2^*)$, for whom $s_r(v_2) = \frac{1}{2}(1-\lambda)(1-v_2)^2 > CS^A = \frac{1}{8}$) face a lower price and are strictly better off.

5.3 Discussion

Based on the results above, we have $CS^{OO} < CS^{OI}$, so consumers are strictly better off when they must active opt in, rather than actively opting out. Figure 3 plots consumer surplus under both policies as a function of the parameter λ. This is just a proof-of-concept example to show that the default – opt-in or opt-out – can matter even if the costs of the non-default choice are negligible and even without new terms (such as a taste for privacy for its own sake) in consumer utility.

The key to the results is that an almost-negligible cost c can tip behavior a lot depending which side of the ledger it falls on. In the example, this "tippiness" is connected to the fact that in any equilibrium, a threshold consumer expects to face the same downstream price whether she opts in or out. This means the threshold consumer is always prone to being tipped either direction by a small cost, and because the new threshold is subject to the same effect, unraveling can happen in either direction. The critical element – the fact that the marginal consumer is close to indifferent indifferent between opting in or out – should be more general than this example. It spills out of the fact that both a profit-maximizing uniform price and a discriminatory price will be set at this marginal consumer's willingness to pay.

Figure 4: Consumer decisions under policies DS and OO

6 Linear demand in both markets

This section gives a very preliminary discussion of data sharing and opt out when the data-gathering firm has downward-sloping demand. Consumers are assumed to draw i.i.d. valuations $v_1 \sim U[0,1]$ and $v_2 \sim U[0,K]$. Figure 4 illustrates the valuation space and gives a concise summary of consumer decisions under policy DS and policy OO with nuisance cost $c > 0$.

References

Anderson, Simon, Alicia Baik, and Nathan Larson, 2019, Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts, CEPR Working Paper 15169.

Belleflamme, P., and Vergote, W., 2016, Monopoly Price Discrimination and Privacy: The Hidden Cost of Hiding, *Economics Letters*, 149, 141 - 144.

Chen, Z., Choe, C., and Matsushima, N., 2018, Competitive Personalized Pricing, mimeo.

Montes, R., Sand-Zantman, W., and Valletti, T., 2015, The value of personal information in markets with endogenous privacy, mimeo.

Taylor, C. and Wagman, L., 2014, Customer Privacy in Oligopolistic Markets: Winners, Losers, and Welfare, *International Journal of Industrial Organization*, 34, 80-84.

Tucker, C.E., 2015. Privacy and the Internet. In Handbook of Media Economics (Vol. 1, pp. 541-562). North-Holland.

A Extension: two-type demand in Market 2 and linear demand in Market 1

Note: For legacy reasons, Markets 1 and 2 are referred to here as Markets 0 and 1. Likewise, the upstream and downstream firms become Firms 0 and 1 respectively.

This section discusses a version of the model in which upstream demand is linear and consumer valuations in the downstream market are either high or low. Each consumer has value $v_0 \tilde{} U[0,1]$ at upstream Firm 0 and $v_1 \in \{v_L, v_H\}$ at downstream Firm 1, with $q = \Pr(v_H)$ the fraction of H-types. Define $\bar{q}$ by $v_L = \bar{q}v_H$. Autarky refers to the situation with no interaction between the two markets. If H-types are rare ($q < \bar{q}$), Firm 1 would set $p_1 = v_L$ in autarky and the downstream market would be efficient. Alternatively, if H-types are common ($q > \bar{q}$), then in autarky Firm 1 would set $p_1 = v_H$ and sell only to them.

When we refer to p_1 below, we mean Firm 1's expected price to "hidden consumers" (anyone who didn't buy from Firm 0 or who opted out). This price is 'expected' in two senses. Consumers have rational expectations about p_1 when making upstream decisions. But also, in this setup, Firm 1 uses only two pure strategies: $p_1 = v_L$ or $p_1 = v_H$. But often, equilibrium will require Firm 1 to mix between these, and we will abuse terminology by referring to p_1 when I really mean $E(p_1)$.

A.1 Example

A.1.1 Assumptions for the example

The diagrams further below are equilibrium results for an example with $v_H = 1$, $v_L = \frac{1}{2}$, thus $\bar{q} = \frac{1}{2}$, and (in the opt-out version) $c \to 0$. I've left q free in order to examine how equilibrium regimes change with the composition of the downstream market. The point of focusing on the $c \to 0$ limit is partly to draw a sharp contrast (relative to no opt-out) and partly to keep the diagrams simpler by eliminating one of the equilibrium regimes.

A.1.2 Summary of equilibrium in the example

In both the versions with and without opt-out, there are essentially three equilibrium regimes, depending on the fraction of the downstream market comprised by high-value consumers (q). Since L-type consumers never get positive downstream surplus, H-type consumers are the only ones who might have an incentive to avoid being revealed, which they can do by not buying at Firm 0 or buying and opting out (if that's an option). Either type of 'evasion' by H-types tilts the composition of the pool of unrevealed, or hidden consumers faced by Firm 1, possibly inducing Firm 1 to prefer a high price ($p_1 = v_H$) rather than a low one.

In the first regime, when q is small, there are too few high-value types for their evasion to affect Firm 1's pricing: it sets $p_1 = v_L$. If consumers can't opt out, Firm 0 cuts its price (relative to $p_0 = \frac{1}{2}$ in autarky) in order to gather info (to be sold on) from more consumers. If consumers have free opt out, then Firm 0 prices no differently than autarky.

In the second regime, when q is moderate, H-type evasion begins to affect how Firm 1 prices to unrevealed consumers. For now, focus on the no opt-out case. Equilibrium requires a gradual rise in Firm 1's price – if p_1 jumped up to $p_1 = v_H$, the H-types would have no incentive to avoid being revealed, and the price hike wouldn't be needed after all. As mentioned above, Firm 1 must price at either v_L or v_H, so a gradual rise in $E(p_1)$ implies Firm 1 must mix. This regime shift involves a downward jump in Firm 0's price. In this second regime, Firm 1's downstream reactions have the effect of making Firm 0's upstream demand from H-types more price sensitive,* and so Firm 0's optimal price jumps down rather than declining smoothly. Over this region, both prices (p_0 and the average p_1 charged to hidden consumers) rise as high-value types become more common. In both of these first two regimes, the value of consumer information is that it permits Firm 1 to fully extract revealed H-types.

In some respects, this second regime doesn't change qualitatively with opt-out. it's still true that H-type evasion shifts the composition of the hidden pool enough that Firm 1 starts to respond with higher prices. But now evasion is easier, and H-types won't stop opting out until their downstream surplus falls to $v_H - E(p_1) = c$. For $c > 0$ this still implies mixing by Firm 1, but in the $c \to 0$ limit that mixture jumps up to $E(p_1) = v_H$ right away. Meanwhile, Firm 0 doesn't cut its price as much as it did without opt-out (since there's less incentive to have more customers if some of them don't permit their information to be sold).

In the third regime, when high-value types are common ($q > \bar{q}$), the value of information changes. In autarky, Firm 1 would set $p_1 = v_H$ and sell only to high-value consumers, so information-sharing expands the downstream market by allowing Firm 1 to also reach revealed L-types with a fully-extracting price. A costly opt-out option is never exercised by consumers of either type, since neither has any downstream surplus to protect. Thus the equilibria with and without opt-out are identical. (The picture is more ambiguous right at $c = 0$, where consumers are indifferent about opting in or out, but there's no ambiguity in the $c \to 0$ limit.) Firm 0's price rises with q: now its info-profits come on the L-type consumers, and as they become rarer, its incentive to cut p_0 below its autarky level diminishes.

A.1.3 Diagrams

In most of these diagrams, equilibrium outcomes without opt-out are purple and outcomes with opt-out are blue.

Prices

Firm 0

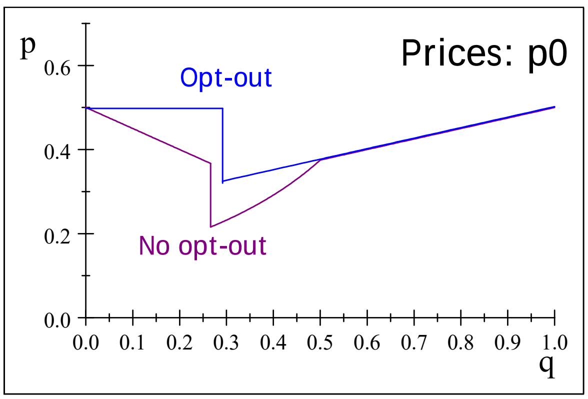

Firm 1

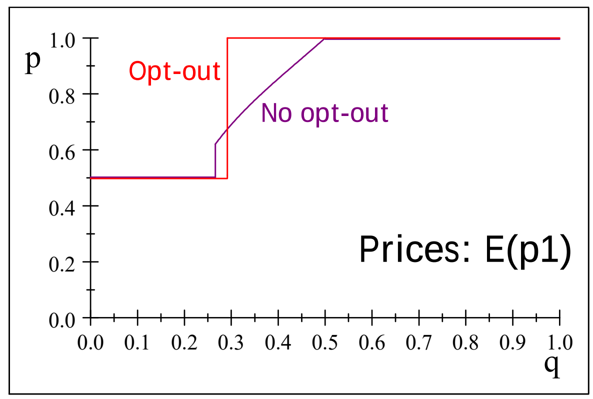

Profits No diagram for Firm 1, as it always just earns its autarky profit: $\pi_1 = v_L$ if $q \leq \bar{q}$, or $\pi_1 = qv_H$ if $q > \bar{q}$. So comparisons of Firm 0's profits (autarky vs info-sharing with and without opt-out) also tell us what's happening with total profits.

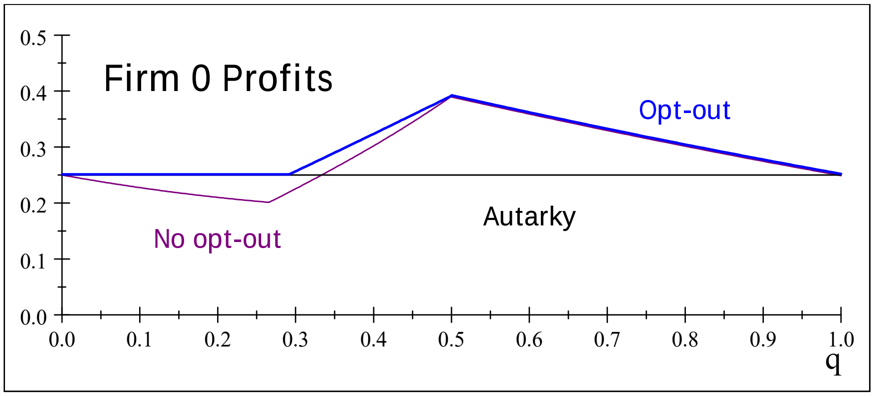

Total Welfare

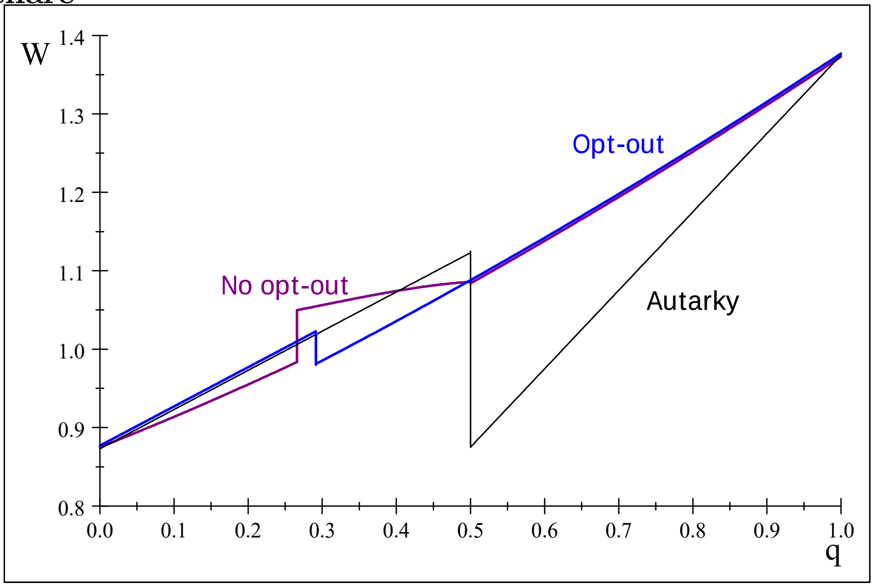

Consumer Surplus Consumer surplus for each downstream type:

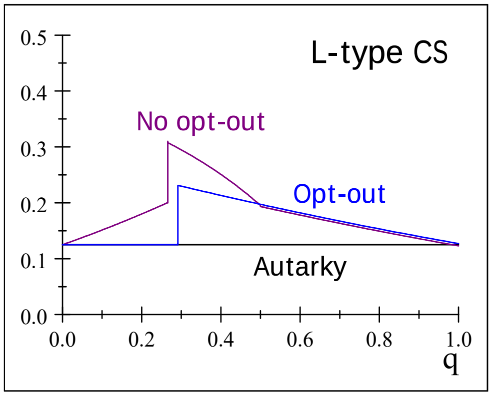

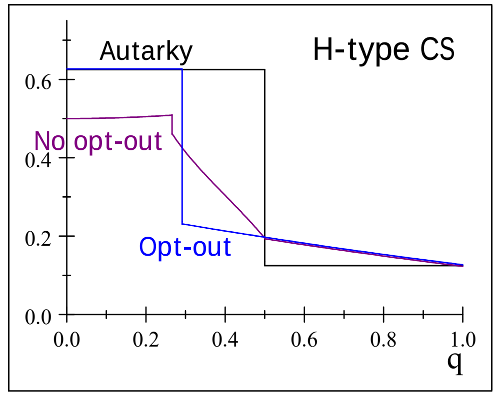

And total consumer surplus:

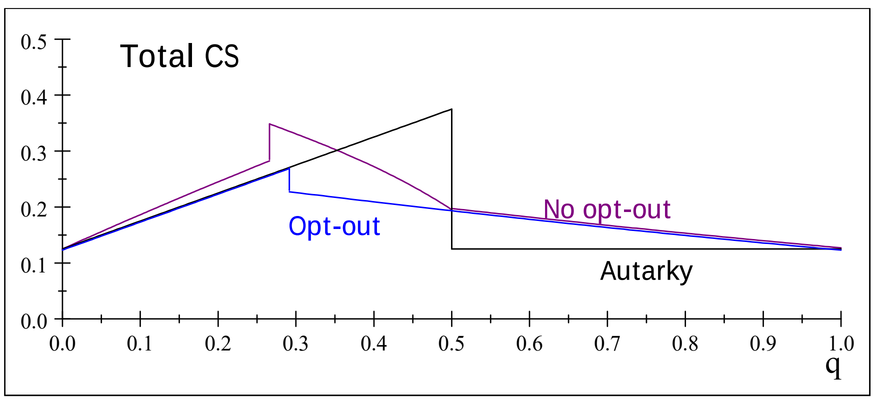

A.2 Conclusions

The headline result is that information sharing is good for *everyone* when it expands the downstream market. In our two-type model, this looks a bit stylized: if high-value types are common enough ($q > \bar{q}$) that the downstream firm would sell only to them, then information sharing permits price discrimination that brings new consumers into the downstream market. By itself, this is good for welfare. And the incentive to profit on this information induces Firm 0 to cut its price, expanding the upstream market, which is also good for welfare. A rising tide lifts both profits (Firm 1 is no worse off) and consumer surplus for both types. (There may be more subtle effects if we distinguish consumers with higher and lower upstream values, but let's set this aside.) None of this depends on whether opt-out is permitted; indeed, in this case consumers are being fully extracted downstream and have no incentive to opt out.

The second headline result is that information sharing, with or without opt-out, is a mixed bag when the downstream market already operates efficiently at autarky. (In our model, for $q < \bar{q}$, Firm 1 would price low and sell to everyone.) Downstream welfare can only fall; upstream welfare may rise if (for the same incentive reason discussed just above) Firm 0 cuts its price, expanding the upstream market. One effect or the other can dominate, but in our example the differences were small relative to the big welfare gain discussed in the previous paragraph.

Let's turn to a few sub-headlines. Some of these are subtle or counter-intuitive, but they are probably also more sensitive to modeling details.

- Selling consumer info may make the upstream firm worse-off. Or to state this more precisely, Firm 0 might wish it could commit to not selling consumer information. When this is true, Firm 0 benefits from free opt-out (since this amounts to a type of commitment device).

- The upstream firm does better when its customers can opt out for free than it does when they can't opt out. (Sometimes strictly, sometimes weakly.)

- if the downstream market would otherwise be efficient, information sharing with no opt-out acts a bit like a transfer of consumer surplus from consumers with high downstream values to consumers with low downstream values. Loosely, the channel is that upstream prices fall and average downstream prices rise. The latter hurts downstream high-types relatively more, whether they face the "pool" price or a fully-extracting discriminatory price. (The low-types have little downstream surplus to lose – they just get priced out.) And lower upstream prices help the low downstream types relatively more. (The high types can't take as much advantage because of the consequences of being revealed.)

- Set information sharing with no opt-out as a benchmark. Changing the rules to permit free opt-out never improves aggregate consumer surplus. Indeed, free opt-out *reduces* aggregate CS if the downstream market would be efficient under autarky.

- Expanding on the last point, it is possible that the option to opt-out may hurt *all* consumer types, including the high-value downstream consumers who exercise the option most eagerly. (In our example this was true whenever high-value types comprised between 29% and 50% of the market.)

A.3 Robustness and caveats

The two-type structure is special, so which of these results are likely to be more general?

- I haven't fully explored this model. While I have a handle on the equilibrium characterization, I don't know if qualitative conclusions change if the downstream market grows more or less valuable relative to the upstream one. (That is, increasing or decreasing both v_L and v_H.) Likewise, for changing the gap between high and low-value consumers.

A.3.1 Likely to be general

- The idea that information sharing can improve welfare when the downstream market would otherwise be inefficient. This is standard "price-discrimination expands markets" stuff.

- The idea that opt-out can be individually rational for consumers but hurt them collectively. The mechanism is that when a higher-value downstream consumer hides in the unrevealed pool, she tends to push up the pool price for everyone. This negative externality is going to be general.

- The idea that the upstream firm might be worse-off when it can sell information (and can't commit to not doing so). It's well-established that there can be value in committing *ex ante* to not holding up consumers *ex post*. And in such cases we would the conclusion that permitting opt-out is good for profits to generalize too.

- The distributional consequences for consumers, at least in broad strokes. Consumers with higher downstream values will be hurt by being revealed, or the threat of it; consumers with lower downstream values will tend to gain from lower upstream prices.

- Qualitative comparisons of upstream and downstream prices, relative to autarky.

A.3.2 Unlikely to be general

- Mixed strategy pricing. This is driven by the two discrete types.
- Sharply demarcated equilibrium regimes. This is an indirect consequence of having discrete types.
- Sharp conclusions about some consumers having no incentive to hide, or no incentive to opt out. With two types, some or all of the consumers always have zero surplus downstream. This won't be true for more general value distributions.

B Perfectly correlated valuations

Note: For legacy reasons, Markets 1 and 2 are referred to here as Markets 0 and 1. Likewise, the upstream and downstream firms become Firms 0 and 1 respectively.

In this section we assume that consumers' valuations in the upstream and downstream market are perfectly correlated. Correlated values might arise because the products share attributes, like fine wine and fine dining, for example, that attract similar consumers. Or correlated values could arise because willingness to pay for both goods rises with income, or with some other consumer attribute. Taking this correlation to be perfect is a simplification that allows us to state clear, sharp policy conclusions.

To preview those conclusions, compared to autarky, profits always rise when data-sharing is permitted. More surprisingly, consumer surplus rises too, unless the downstream market is substantially larger than the upstream market. Starting from a regime with data-sharing, giving consumers the right to free opt-out boosts profits further and unambiguously hurts consumer surplus. Typically consumers are worse-off under data-sharing with free opt-out than they would have been under autarky.

For a regulator interested in protecting consumers, there are two main takeaways. One is that outlawing data sharing does not necessarily make consumers better off.

The second is that data sharing with opt-out is not an innocuous substitute for outlawing data sharing entirely – consumers are worse-off when they can opt out than they are under either of the other two policies.

B.1 Setup and equilibrium under autarky

The valuation of consumer x at Firm 0 is linear in her location: $v_0 = 1 - x$. Her valuation at Firm 1 is proportional to her value at Firm 0: $v_1 = kv_0 = k(1-x)$, for $k > 0$. We say that the downstream market is small or large if $k \in (0,1)$ or $k > 1$ respectively. This abuses language slightly, since it is downstream values, not the size of the market *per se*, that are relatively small or large. The special case of $k = 1$ is complicated by multiple equilibria (arising out of indeterminacies in consumer behavior). Rather than catalog equilibrium in this case exhaustively, we focus on limiting equilibria as k approaches 1 from above or below. Throughout, we will interchangeably identify a consumer by her location x or her upstream valuation v_0.

B.1.1 Autarky

We will state the following standard results without additional comment. Each firm solves a linear-demand monopoly problem. Prices are $p_0 = \frac{1}{2}$ and $p_1 = \frac{k}{2}$, each firm has the same marginal consumer $x = \frac{1}{2}$, and profits are $\pi_0^A = \frac{1}{4}$ and $\pi_1^A = \frac{k}{4}$. Consumer surplus is $CS_0 = \frac{1}{8}$ upstream, $CS_1 = \frac{k}{8}$ downstream and $CS = \frac{1+k}{8}$ in total. Welfare is $W_0 = \frac{3}{8}$ upstream, $W_1 = \frac{3k}{8}$ down, and $W = \frac{3}{8}(1+k)$ in total.

B.1.2 Preliminary comments about the data-sharing (DS) regime

When data is sold, the qualitative features of equilibrium may depend on whether the downstream market is small or large. To illustrate why, we start by coining a phrase "proportionate pricing" to refer to a situation where $p_1 = kp_0$. For example, the equilibrium under autarky has proportionate pricing.

Equilibrium prices generally will not be proportionate under data-sharing, but the concept is still useful as a benchmark. Consumer v_0 enjoys surplus $s_0 = v_0 - p_0$ from a purchase at Firm 0 and surplus $s_1 = kv_0 - p_1$ from a hidden purchase at Firm 1. If prices are proportionate, these surpluses satisfy $s_1 = ks_0$. This means the sets of consumers with weakly positive surplus at each firm are identical. Furthermore, if forced to choose between s_0 and s_1, all consumers would prefer purchasing at Firm 0 if $k < 1$, or at Firm 1 if $k > 1$. In effect, choosing s_0 or s_1 is the decision that consumers face under regime DS, but proportionate pricing cannot be an equilibrium outcome. If the downstream market is small and consumers expect $p_1 = kp_0$, then all consumers with values $v_1 \geq p_1$ will buy at Firm 0 and be revealed. But then Firm 1 faces no hidden consumers who are prepared to pay p_1, and it would be better off with a lower price. Alternatively, if the downstream market is large and consumers expect $p_1 = kp_0$, then all consumers will wait to buy at Firm 1. But this leaves Firm 0 with zero profit; it would be better off cutting p_0 until it makes some sales. Based on this logic, in the DS model we expect to find $p_1 < kp_0$ in equilibrium when $k < 1$, and $p_1 > kp_0$ when $k > 1$.

Those conditions on relative prices pin down which consumers buy at Firm 0 and reveal themselves, and which consumers wait and buy at Firm 1's uniform price (if at all). If the downstream market is small and $p_1 < kp_0 < p_0$, then $s_0 \geq s_1$ is equivalent to $(1-k)\,v_0 \geq p_0 - p_1$. Thus higher-value consumers will buy at Firm 0 and lower-value consumers wait. On the other hand, if the downstream market is large and $p_1 > kp_0 > p_0$, then the same $s_0 \geq s_1$ condition may be expressed as $(k-1)\,v_0 \leq p_1 - p_0$. In this case, higher-value consumers wait, and it is lower-value consumers who buy at Firm 0 (if at all). In both cases the underlying logic is the same: the more valuable product commands a relative price premium, and this sorts consumers into the highest-value types (who are prepared to pay the premium for the more valuable product) and lower-value types (who settle for the lower-value product at a relatively low price).

B.2 Equilibrium with data-sharing

We shall keep the analysis light and, hopefully, intuitive. Many of the technical details are relegated to the appendix. Based on the discussion above, we deal separately with the cases where the downstream market is small or large.

B.2.1 Small downstream market ($k < 1$)

The figure illustrates an equilibrium of the subgame after Firm 0 sets its price. The threshold consumer x^* is defined by her indifference $s_0^* = s_1^*$ between buying at Firm 0 (thereby revealing herself and getting zero surplus downstream) and waiting to buy

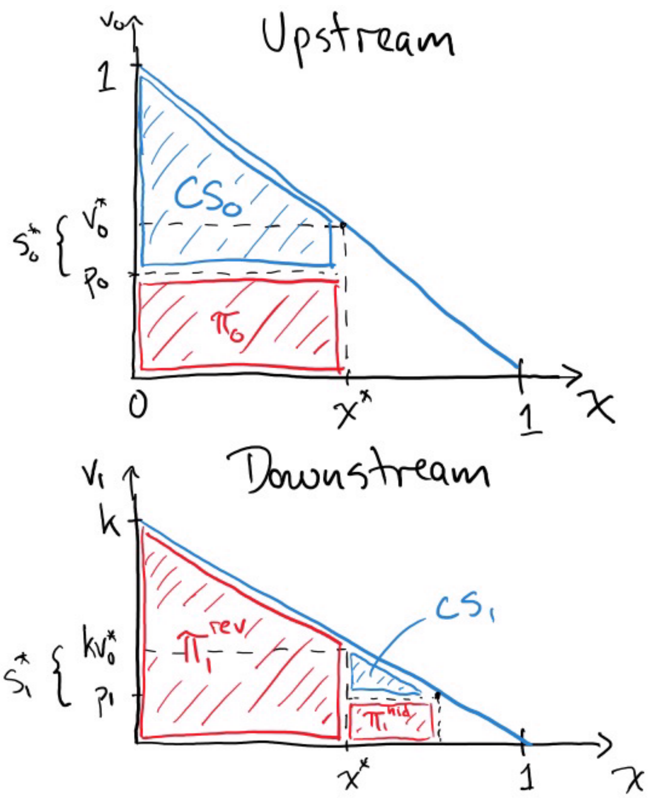

Figure 5:

at Firm 1. Consumers with higher values than hers buy at Firm 0 and enjoy total consumer surplus CS_0, while consumers with lower values wait, buy at Firm 1 if their values are not too low, and enjoy total consumer surplus CS_1. Firm 1's uniform price p_1 is set to maximize its profit π_1^{hid} on the hidden consumers, all of whom are to the right of x^*. Consumers to the left of x^* bought at Firm 0, had their values revealed to Firm 1, and are fully extracted downstream, with total profit π_1^{rev}.

We will write π_0 for Firm 0's profit on sales in its own market, and $\Pi_0 = \pi_0 + P$ for its overall profit, including the sale of data to Firm 1 at price P. Recall that Firm 0 optimally sells its customer data at a price that fully extracts the gains to Firm 1 over its autarky profit: $P = \pi_1^{rev} + \pi_1^{hid} - \pi_1^A$. Thus we have $\Pi_0 = \pi_0 + \pi_1^{rev} + \pi_1^{hid} - \pi_1^A$. There is a one-to-one relationship between Firm 0's price p_0 and the marginal consumer (x^* or v_0^*) induced by that price, so as a matter of convenience we may treat either p_0, x^*, or v_0^* as Firm 0's choice variable. For the sake of intuition, we shall write Firm

0's first-order condition in terms of the marginal consumer x^*.

$$\frac{d\Pi_0}{dx^*} = \underbrace{p_0 - x^* \left|\frac{dp_0}{dx^*}\right|}_{\text{upstream}} + \underbrace{kv_0^* - p_1}_{\text{downstream}} = 0$$

The first term – the change in upstream profit π_0 – is almost a standard monopoly marginal profit expression, with gains on the marginal consumer and losses on the inframarginal ones. The difference is that the price Firm 0 can get from its marginal consumer is not her full valuation. The second term reflects the gain in downstream profit from converting consumer x^* from a hidden buyer paying the uniform price to a revealed and fully-extracted buyer. This second term is just s_1^*, and since $s_0^* = s_1^*$ for this marginal buyer, we can also write $d\Pi_0/dx^* = v_0^* - x^* \left|dp_0/dx^*\right|$.

At this point, it is useful to invoke standard results for linear demand curves to conclude that Firm 1's price must satisfy $p_1 = \frac{1}{2}kv_0^*$. (Firm 1 faces linear demand from hidden consumers with price intercept kv_0^*; its optimal price is half of this intercept.) From this we can establish the marginal consumer's surplus ($s_1^* = s_0^* = \frac{1}{2}kv_0^*$), and therefore the amount by which p_0 must be "discounted" below v_0^* to make consumer x^* indifferent to buying upstream: $p_0 = \left(1 - \frac{k}{2}\right) v_0^*$. With this in hand, we can write the first-order condition as

$$\frac{d\Pi_0}{dx^*} = v_0^* - \left(1 - \frac{k}{2}\right) x^* = 0$$

Compare this to Firm 0's first-order condition under autarky: $d\Pi_0/dx^* = v_0^* - x^* = 0$, which leads to $x^* = \frac{1}{2}$. Under data-sharing, Firm 0 has an unambiguously stronger incentive to increase sales x^*, even at the cost of a lower price p_0, because a larger customer base means more data to sell. It is clear that Firm 0's optimal choice must satisfy $x^* > \frac{1}{2}$ and $p_0 < v_0^* < p_0^A = \frac{1}{2}$.

B.2.2 Large downstream market ($k > 1$)

The hold-up problem at Firm 1 When the downstream market is large ($k > 1$), hidden consumers face a type of hold-up problem at Firm 1, and as a consequence Firm 1 will end up mixing between a low price and a high price in equilibrium. The diagram below will help to illustrate the hold-up problem.

The diagram represents a hypothetical situation (not necessarily an equilibrium

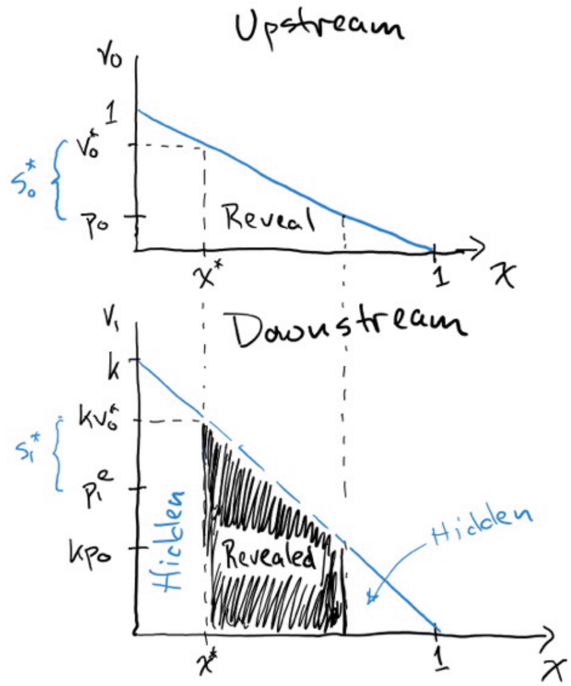

Figure 6:

outcome) where hidden consumers expect the pure strategy downstream price p_1^e. As discussed earlier, that price must satisfy $p_1^e > kp_0$ (since Firm 0 would make no sales otherwise). This means there will be a threshold consumer x^* who expects equal surplus at both firms: $s_0^* = s_1^*$, with $s_0^* = v_0^* - p_0$ and $s_1^* = kv_0^* - p_1^e$. Consumers with higher values than consumer x^* will wait, while consumers with lower values will buy at Firm 0 and reveal themselves as long as v_0 exceeds p_0. In the lower panel, the revealed consumers are blacked out to focus attention on the demand Firm 1 faces from hidden consumers. Its hidden demand has endogenously partitioned into high-value consumers (with $v_1 \geq kv_0^*$) and substantially lower-value consumers (with $v_1 \leq kp_0$).

The key point is that consumers expect a price p_1^e that it is not in Firm 1's interest to deliver. At price p_1^e, Firm 1 sells only to the high-value hidden consumers. But all of these consumers would be prepared to pay at least kv_0^*, so Firm 1 would be strictly better-off charging $p_1 = kv_0^*$ instead. Any candidate equilibrium with a pure strategy p_1 fails on these grounds, so an equilibrium will require Firm 1 to mix. The fundamental issue is a hold-up problem reminiscent of the Diamond paradox.

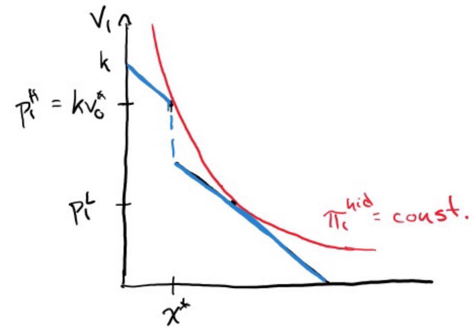

Figure 7:

The marginal high-value "hider" x^* turns down a positive surplus at Firm 0 because she expects an equally good surplus at Firm 1. But Firm 1 knows this and has an incentive to hold her up for this surplus when she arrives in the downstream market. This hold-up problem does not depend critically on the perfect correlation of values. Later we argue that by setting its price low enough the upstream firm may be able to induce this sort of hold-up quite generally, and it may have an incentive to do so since the result (rational consumers tip toward buying upstream) is good for data sales.

Equilibrium with mixed high/low pricing downstream If we excise the revealed consumers from Firm 1's demand curve, its demand from hidden consumers looks like the diagram below. Because of the "step" in this demand curve, Firm 1's pricing decision involves comparing the profit at a high price $p_1^H = kv_0^*$ aimed at the high-value hidden consumers with the most profitable low price p_1^L on the lower-value portion of the demand curve. As discussed above, there is no equilibrium where Firm 1 exclusively charges p_1^H or p_1^L, so it must mix between them, and for this both prices must be equally profitable, as depicted on the diagram.

Suppose Firm 1 charges p_1^L with probability α and p_1^H with probability $1-\alpha$. The following technical result helps to characterize an equilibrium.

Lemma 1 *In any equilibrium,* $\alpha = \frac{1}{k}$ *and* $p_1^L = kp_0$.

The logic of the lemma concerns the threshold consumer x^*, with upstream value v_0^*, who is indifferent between waiting for a chance at Firm 1's low price versus buying at Firm 0 and revealing herself. More precisely, it concerns this threshold consumer and how she relates to consumers with lower values. If she is indifferent because she expects a high chance at a moderately low price ($\alpha > \frac{1}{k}$ and $p_1^L > kp_0$), then all lower-value consumers will either reveal at Firm 0 or not buy at either firm. But in this case, Firm 1 has no reason to charge a low price – it would be better-off always charging p_1^H. Alternatively, if the threshold consumer is indifferent because she expects a small chance of a very low price ($\alpha < \frac{1}{k}$ and $p_1^L < kp_0$), then all lower-value consumers strictly prefer to wait and take their chances at Firm 1. But if there are no revealed consumers, then Firm 1 faces its original demand curve and should simply charge its autarky price instead. The only way that consumer expectations about Firm 1's pricing can be compatible with Firm 1's *ex post* incentives is if $\alpha = \frac{1}{k}$ and $p_1^L = kp_0$.

One implication of the lemma is that consumers with values below that of the threshold consumer expect downstream surplus $s_1 = \alpha\left(kv_0 - p_1^L\right) = v_0 - p_0 = s_0$. That is, *all* consumers with values $v_0 \in [p_0, v_0^*]$ are indifferent between revealing themselves with a purchase at Firm 0 versus waiting for the chance of a low price at Firm 1. We focus on equilibria where the indifferent consumers with lower values wait, $v_0 \in [p_0, p_0 + w]$, and those with higher values reveal, $v_0 \in [p_0 + w, v_0^*]$. The measure w who wait will be pinned down by equilibrium conditions; later we will discuss the possibility of breaking consumer indifference differently.

The figure below illustrates the conditions for an equilibrium of the subgame after Firm 0 sets a price p_0. We summarize the main conditions in the following proposition.

Proposition 3 *In the data-sharing model with a large downstream market, the subgame equilibrium after Firm 0 sets its price satisfies the following relationships between p_0, the highest-value consumer (x^*, v_0^*) to be indifferent between waiting and revealing, and the mass w of indifferent consumers who wait:*

(I) $p_0 = \sqrt{x^*(1-x^*)} = \sqrt{v_0^*(1-v_0^*)}$

(II) $w = p_0 - x^*$

Note that the consumers who purchase at neither firm are those with upstream values $v_0 < p_0$; there are p_0 of these non-purchasers. Relationship (II) is a consequence of the fact that pricing at $p_1^L = kp_0$ and selling to $x^* + w$ hidden consumers must be the most profitable option on the lower part of the hidden demand curve. This means that Firm 1 earns $\pi_1^L = kp_0^2$ by pricing low to its hidden consumers. Its indifference between pricing low and high implies $\pi_1^{hid} = \pi_1^L = \pi_1^H = p_1^H x^*$, and this indifference leads to relationship (I).

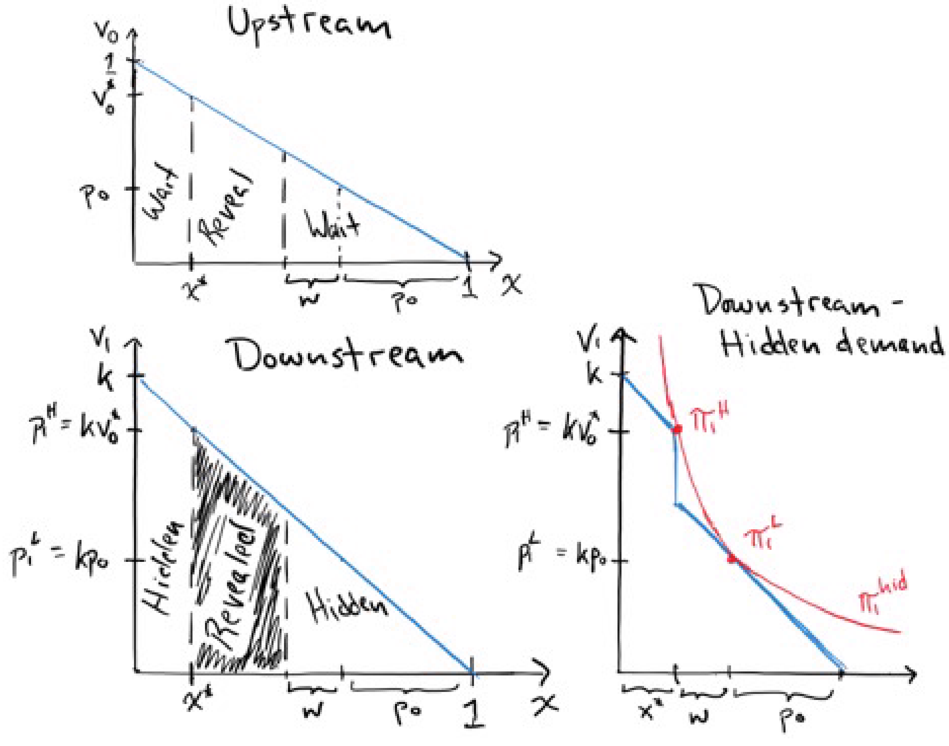

Upstream
Wait
Reveal
Wait
w
p₀
1
x
Downstream
Hidden
Revealed
Hidden
Downstream -
Hidden demand

Figure 8:

The upstream firm's profit and first-order condition

Just as in the case of a small downstream market, Firm 0's total profit may be written $\Pi_0 = \pi_0 + \pi_1^{rev} + \pi_1^{hid} - \pi_1^A$. In its own market Firm 0 sells only to consumers with values $v_0 \in [p_0 + w, v_0^*]$. Using (II), there are $q_0 = 1 - 2p_0$ of these revealed consumers. In order to highlight the parallels with the $k < 1$ case, we write the first-order condition below using q_0 as the upstream firm's strategic variable.

$$\frac{d\Pi_0}{dq_0} = \frac{d\pi_0}{dq_0} + \left(v_1^{rev} - p_1^L\right) = 0$$

The first term is the marginal profit on upstream sales. And as in the $k < 1$ case, the second term represents the downstream profit from converting an additional consumer from paying the uniform price p_1^L to being fully extracted.[4] However, the valuation v_1^{rev} of a marginal revealed consumer is a more complicated object in this case since Firm 0 gains consumers at two margins (the left and right borders of the "revealed" region in the figure). The value v_1^{rev} is an appropriately weighted average of consumer valuations at these two boundaries. Using the expression for upstream profit, this first-order condition can also be expressed as:

$$q_0 = q_0^A + \left(v_1^{rev} - p_1^L\right)$$

where $q_0^A = \frac{1}{2}$ is the monopoly quantity under autarky. As in the $k < 1$ case, the potential profits from selling its customers' data induce Firm 0 to make more sales than it would have done if data-sharing were forbidden.

B.3 Data-sharing with costly opt-out

Now we turn to the model where Firm 0 can sells its purchasers' data to Firm 1, but any consumer who purchases at Firm 0 can pay the nuisance cost c to opt out of having her data sold. For c small but positive, this strikes us as a reasonable depiction of how a policy that permits both data sales and opt out, but requires consumers to actively opt out, might look. Our ultimate goal with this model is to study equilibrium outcomes in the $c \to 0$ limit, as a tractable approximation of outcomes when c is positive but small. In order to get to that limit, we first must

[4] Rather than write the profit on the marginal hidden consumer as p_1^L, we could also write it as the expected profit $\frac{x^*}{x^*+w} p_1^H$ on a sale at the high price – they are the same.

establish some facts about equilibrium when $c > 0$.

B.3.1 The hold-up problem, and mixing by Firm 1

To understand how the opt-out option affects the markets, the key insight is that consumers face a hold-up problem at Firm 1 that is very much like the hold-up problem they faced in model DS when the downstream market was large. As a consequence, there is no pure strategy price p_1 that hidden consumers could expect and that Firm 1 would also stick to. This means that like the $k > 1$ case of model DS, all equilibria, regardless of k, involve mixing by Firm 1.

We sketch the intuition here; proofs are left to the appendix. A consumer facing a price p_0 in the upstream market compares her surplus $s_0 = v_0 - p_0$ from an upstream purchase with her expected surplus s_1^e from buying as a hidden consumer downstream, and also with the total payoff from $s_0 + s_1^e - c$ from preserving her downstream surplus by opting out when she purchases at Firm 0. This consumer's optimal decision can be classified as follows.

- **Opt out** if both upstream and downstream surplus exceed the nuisance cost ($s_0 \geq c$ and $s_1 \geq c$). That is, buy at Firm 0, opt out of data-sharing, and then buy again as a hidden consumer at Firm 1. If there are any opt-outs, they will be the highest-value consumers.

- **Buy once, in the better market** if the condition above is not met and at least one of the markets provides positive surplus. That is, **reveal** (buy at Firm 0 without opting out) if $s_0 \geq s_1^e$ and s_0 is positive. Or **wait** (with the expectation of buying at Firm 1) if $s_1^e > s_0$ and s_1^e is positive.

- **Wait** (expecting to not purchase at Firm 1 either) if both s_0 and s_1^e are negative.

For c small enough, any equilibrium will have consumers with high enough valuations opting out. And for any $c > 0$, an equilibrium must have some consumers who buy at Firm 0 without opting out – that is, there must be some consumers whose values are revealed to Firm 1. To understand why, consider the counterfactual where all consumers are hidden when they reach Firm 1, either because they opted out or because they did not purchase at Firm 0. Then Firm 1 will simply charge its autarky price $p_1^A = \frac{k}{2}$. But then consumers with valuations v_1 just above p_1^A would not find it

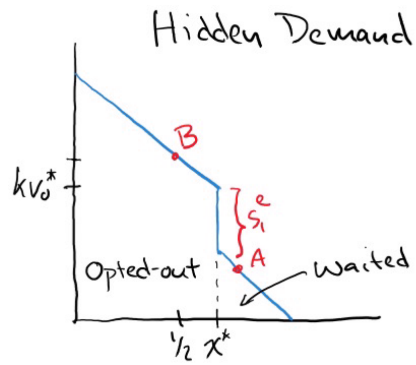

Figure 9:

worthwhile to opt out. If they get positive surplus at Firm 0 (and we can show that Firm 0 will ensure that they do), then they will reveal instead.

This means that Firm 1 will face hidden-consumer demand that looks like the figure below. With some range of mid-value revealed consumers carved out, Firm 1's hidden demand is composed of a high-value segment of consumers who opted out and a low-value segment of consumers who waited, with a valuation gap between them.

In the figure, Firm 0's marginal opt-out was consumer x^* with upstream valuation v_0^*. This consumer would only pay c to opt out if she expected a strictly positive surplus s_1^e downstream, which means that she must expect Firm 1 to price at a point such as A on its hidden demand curve. However, given the cliff in the demand curve, Firm 1 may be tempted to hold up hidden consumers by pricing at $p_1 = kv_0^*$ or higher instead. In this case, there is no equilibrium with a pure strategy p_1. This temptation to hold up consumers is acute if there are many high-value opt-outs relative to the lower-value consumers who waited:

Lemma 2 *If Firm 0 prices at or below its autarky level ($p_0 < p_0^A$) then, for sufficiently small c, Firm 1 does not use a pure strategy in the downstream equilibrium.*

This result holds under much weaker conditions when the downstream market is large, but for simplicity of presentation we state a version that applies to both the $k < 1$ and $k > 1$ cases. The figure gives a general sense of the argument (although the details are more involved). If there were a downstream equilibrium with a pure strategy price, it would have to involve consumers expecting and Firm 1 delivering a price at a point like A. Meanwhile point $B = \left(\frac{1}{2}, \frac{k}{2}\right)$ represents the sales and price combination that Firm 1 chooses under autarky. If consumers expect this low p_1 and p_0 is also low, then enough of them will opt out that B becomes a feasible choice on Firm 1's hidden-consumer demand curve. But then Firm 1 can gain by holding up consumers with the higher than expected price $p_1 = p_1^A = \frac{k}{2}$.

One implication of this result is that the upstream firm can induce a hold-up problem in the downstream market by setting a sufficiently low price. While formal results will come later, it is useful to have an informal sense of why Firm 0 might wish to do this. In the hypothetical of the figure, suppose that Firm 1 does deviate to a higher price like $p_1 = p_1^A$. But then consumers $x \in \left[\frac{1}{2}, x^*\right]$ anticipate no downstream surplus, stop paying c, and become revealed consumers. Furthermore, those consumers just to the left of $x = \frac{1}{2}$ who now expect downstream surplus less than c also become revealed. But in this case, Firm 1 will wish to increase p_1 again until it is fully extracting the marginal opt-out consumer. And so forth. By launching this chain of events with a low price p_0, Firm 0 can precipitate an unraveling process that substantially increases the number of revealed consumers. And this is to Firm 0's advantage, since it profits on the sale of these consumers' data. Although this is informal disequilibrium logic, we shall see that, at least in broad strokes, it is correct.

B.3.2 Downstream equilibria with high-low mixing by Firm 1 when $c > 0$

Based on the arguments above, we turn to strategies in which Firm 1 mixes between a high price and a low price for hidden consumers. As earlier, let p_1^H and p_1^L be these prices, and let α be the probability of the low price. The figure below helps to illustrate conditions for an equilibrium in the downstream market after Firm 0 has set the price p_0.

Much of the notation and intuition transfers from the mixed equilibrium in model DS, but there are some important differences. Consumers with upstream values above v_0^* (to the left of x^*) opt out. Those with upstream values between $\hat{v}_0$ and v_0^* buy at Firm 0 and are revealed – these "missing consumers" are the reason for the vertical segment of the hidden-consumer demand curve. The threshold consumer $\hat{v}_0$ is identified by her indifference between her surplus $\hat{s}_0$ at Firm 0 and her expected surplus $\hat{s}_1 = \alpha \left(k\hat{v}_0 - p_1^L\right)$ at Firm 1. Consumers with upstream values below $\hat{v}_0$ wait and either buy from Firm 1 if they get the low price (if $kv_0 \geq kv_0^L := p_1^L$) or never buy at all (if $v_0 < v_0^L$).

In equilibrium, the marginal opt-out consumer x^* must be indifferent to preserv-

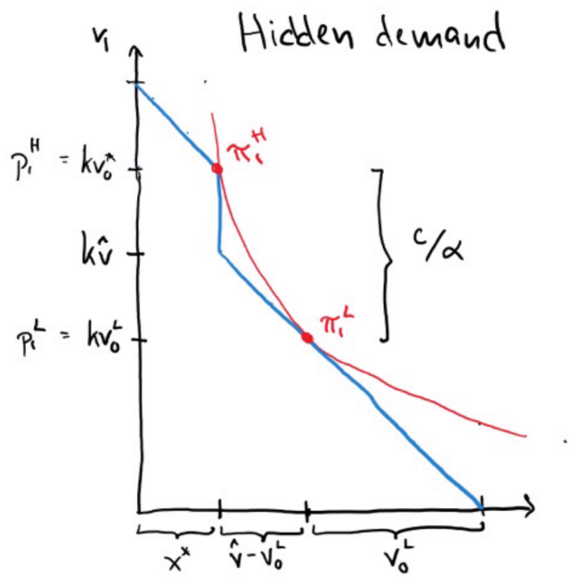

Figure 10:

ing her expected downstream surplus $s_1^* = \alpha k \left(v_0^* - v_0^L\right)$ by paying c; the condition $k\left(v_0^* - v_0^L\right) = \frac{c}{\alpha}$ is illustrated on the diagram. For Firm 1, the low price p_1^L must satisfy a first-order condition for optimality; by standard results for linear demand this means that its sales $x^* + \left(\hat{v}_0 - v_0^L\right)$ are equal to its "non-sales" v_0^L. And Firm 1's profits π_1^H and π_1^L at its high and low prices must lie on the same isoprofit curve, as depicted. With a bit of work, these equilibrium conditions may be summarized as follows.

$$
\begin{aligned}
\text{(I)} \qquad v_0^L &= \sqrt{v_0^* \left(1 - v_0^*\right)} = \sqrt{x^* \left(1 - x^*\right)} \\
\text{(II)} \qquad p_0 &= \left(2 - \frac{c}{v_0^* - v_0^L}\right) v_0^L - \left(1 - \frac{c}{v_0^* - v_0^L}\right)\left(1 - v_0^*\right)
\end{aligned}
$$

Condition (I) follows immediately from $\pi_1^H = \pi_1^L$, and pins down Firm 1's marginal consumer at its low price in terms of the marginal opt-out consumer v_0^*. Condition (II) connects both of these thresholds to Firm 0's price via the indifference of consumer $\hat{v}_0$. Over the economically relevant range, p_0 is monotonically decreasing in v_0^*, so we may treat either p_0, v_0^*, or x^* as Firm 0's choice variable. We now move directly to the $c \to 0$ limit where these equilibrium conditions simplify substantially.

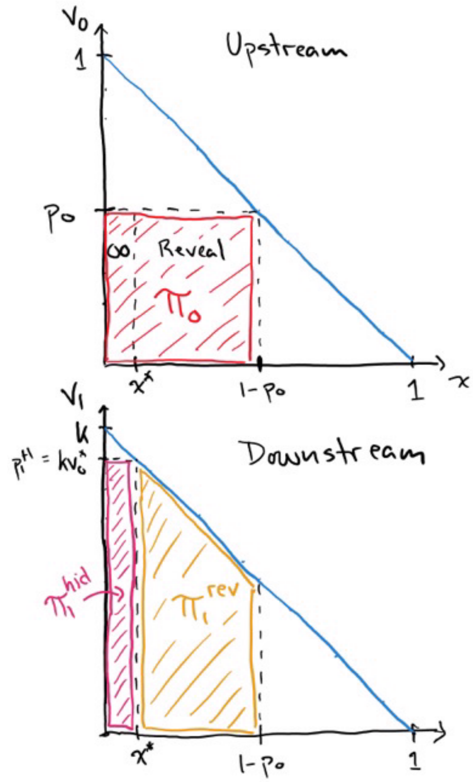

Figure 11:

B.3.3 Equilibrium in the free opt-out ($c \to 0$) limit

We are now in a position to state results about the overall equilibrium of the free opt-out limit of model COO. The downstream equilibrium conditions (I) and (II) carry over – (I) is unchanged, and (II) simplifies to

$$\text{(II')} \qquad p_0 = 2v_0^L - (1 - v_0^*)$$

Firm 0's various sources of profit are illustrated in the figure below.

Firm 0 sells its own product to all consumers with values $v_0 \geq \hat{v}_0$ (of whom some opt-out and some reveal). But notice that α tends to zero with c, so hidden consumers get Firm 1's low price vanishingly often, and consequently the threshold consumer $\hat{v}_0$ satisfies $\hat{v}_0 \to p_0$. So Firm sells to everyone whom it offers positive surplus, earning upstream profit $\pi_0 = p_0 (1 - p_0)$. And downstream, Firm 0 fully extracts any consumer who bought upstream, except for the consumer surplus of the opt-outs. As usual, its total profit is $\Pi_0 = \pi_0 + \pi_1^{rev} + \pi_1^{hid} - \pi_1^A$; using the figure, this may be written

$$\Pi_0 = p_0 (1 - p_0) + \frac{k}{2} \left(1 - p_0^2 - (x^*)^2\right) - \frac{k}{4}$$

To a first approximation, Firm 0 earns standard uniform price profits on its customers upstream, and perfectly discriminatory profits on them downstream. The latter gives Firm 0 an incentive to price below its autarky level and make more sales, and that incentive will be stronger when the downstream market is larger. We can see this in the firm's first-order condition – note that x^* is increasing in p_0, so the optimal price is unambiguously below Firm 0's autarky price $p_0^A = \frac{1}{2}$:

$$\frac{d\Pi_0}{dp_0} = (1 - 2p_0) - k \left(p_0 + x^* \frac{dx^*}{dp_0}\right) = 0$$

C Policy Comparison

In this section we compare profits and consumer surplus under three policy regimes: autarky (A), data-sharing (DS), and data-sharing with free ($c \to 0$) opt-out (OO). For profits, it suffices to look either at the total profit Π or the upstream firm's profit Π_0, since the downstream firm is always held to its autarky profit. The headline result is that we have an unambiguous ranking: $\Pi_{OO} > \Pi_{DS} > \Pi_A$. That is, Firm 0 gains when it can sell data, but it does even better when consumers can opt out of data-sharing. Meanwhile, consumers are always worst off when they can opt out: $CS_{OO} < CS_A$ and $CS_{OO} < CS_{DS}$. If the downstream market is not too much larger than the upstream one, consumers are best-off under data-sharing: $CS_{DS} > CS_A$; when the downstream market is very large, this ranking is reversed.

C.1 Ranking profits under different policies

Let $\Pi = \Pi_0 + \pi_1^A$ be total profit.

C.1.1 Data-sharing vs Autarky

When the downstream market is larger than the upstream one ($k > 1$), the argument is fairly simple. Referring back to equilibrium conditions (I) and (II) for the data-sharing case, it is clear that Firm 0 can induce all of the consumers to reveal by

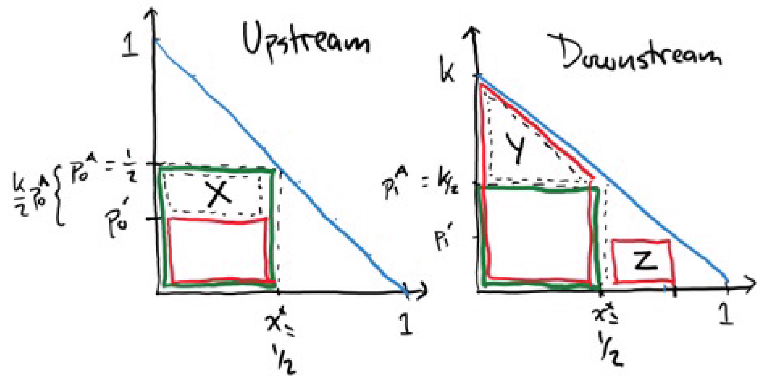

Figure 12:

setting $p_0 = 0$. This strategy earns no profit upstream, but extracts the full surplus from every consumer downstream, with total profit $\Pi = \frac{k}{2}$, which exceeds the total autarky profit $\frac{1+k}{4}$. Since this strategy is feasible but not necesarily optimal, *a fortiori* we have $\Pi_{DS} > \Pi_A$.

For the $k < 1$ case, consider the diagram below where the autarky profit in each market is in green, and autarky sales in each market are $x^* = \frac{1}{2}$.

Under the data-sharing policy, by choosing an appropriate price (labeled p_0') Firm 0 has the ability to induce exactly the same consumers to buy who would have bought under autarky. Profits from this strategy are outlined in red. This strategy is less profitable than autarky upstream, since p_0' must be below p_0^A to compensate consumers for revealing their downstream values. However, this strategy is more profitable than autarky downstream: it extracts the full surplus of all the autarky consumers, and it adds additional uniform-price profits on the hidden consumers to the right of x^*. Let the total profit under this particular data-sharing strategy be Π'. The difference $\Pi' - \Pi_A$ is just the difference in areas $Y + Z - X$ corresponding to these downstream and upstream profit changes. We have $p_0^A - p_0' = \frac{k}{2}$ and so $X = \left(p_0^A - p_0'\right) x^* = \frac{k}{4}$. Downstream, $Y = \frac{1}{2} x^* \left(k - p_1^A\right) = \frac{k}{4}$, and $Z = \frac{k}{16}$. So the upstream "discount" required to entice consumers to reveal their values is exactly balanced by the gains from fully extracting them downstream: $X = Y$. Consequently, $\Pi' - \Pi_A = Z > 0$. Since this particular strategy by Firm 0 is feasible under regime DS but not necessarily optimal, we have $\Pi_{DS} > \Pi_A$.

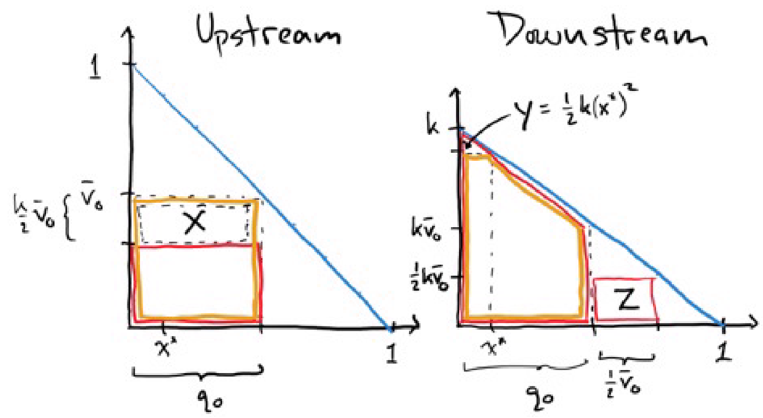

Figure 13:

C.1.2 Free opt-out vs Data-sharing

Once again we need separate arguments for the $k < 1$ and $k > 1$ cases.

Small downstream market $(k < 1)$

In both cases, DS or OO, Firm 0 sells to all consumers with values above some threshold. For the sake of consistent notation, let us say that Firm 0 sells to consumers $x \in [0, q_0]$. Under either policy regime, we may treat q_0 as the firm's choice variable and write $\Pi_{DS} = \max_{q_0} \Pi_{DS}(q_0)$ and $\Pi_{OO} = \max_{q_0} \Pi_{OO}(q_0)$. The objective of the partly graphical argument below is to show that the profit in the opt-out regime dominates profit under DS at *any* relevant level of upstream sales: $\Pi_{OO}(q_0) > \Pi_{DS}(q_0)$ whenever $q_0 \geq \frac{1}{2}$. The optimal sales under policy DS satisfy $q_0^{DS} \geq \frac{1}{2}$, so this establishes that $\Pi_{OO} \geq \Pi_{OO}\left(q_0^{DS}\right) > \Pi_{DS}$.

The figure compares total profits at the same q_0 under policy DS (profits outlined in red) and under OO (profits outlined in orange). Write $\bar{v}_0 = 1 - q_0$ for the upstream value of the marginal upstream customer. In the upstream market, Firm 0 can charge this marginal consumer her full value under policy OO, whereas it is only able to charge her $\left(1 - \frac{k}{2}\right)\bar{v}_0$ under policy DS. Thus upstream profits are better by $X = \frac{k}{2}\bar{v}_0 q_0$ under opt-out. In the downstream diagram, x^* represents the marginal consumer who opts out when the policy is OO. Policy DS extracts the full surplus on the downstream customers who bought at Firm 0. Policy OO does almost as well, but it misses the consumer surplus on the first x^* consumers who opt out and only pay the high hidden price p_1^H. Furthermore, under policy DS there are also market-expanding profits $Z = \frac{1}{4}k\bar{v}_0^2$ on lower-value hidden consumers. So the downstream market is more profitable by $Y + Z$ under policy DS. The total difference in profits is $\Pi_{OO}(q_0) - \Pi_{DS}(q_0) = X - Z - Y$. Profit rectangles X and Z have the same height, but for $q_0 \geq \frac{1}{2}$, Z is less than half as wide. Thus $Z \leq \frac{1}{2}X$ for $q_0 \geq \frac{1}{2}$. And one can confirm that the number of opt-outs is small enough (we have both $x^* < \bar{v}_0$ and $2x^* < q_0$) that $Y < \frac{1}{2}X$. It follows that we have the ranking $\Pi_{OO}(q_0) > \Pi_{DS}(q_0)$ for $q_0 \geq \frac{1}{2}$ claimed above, and so it follows (as argued above) that $\Pi_{OO} > \Pi_{DS}$.

Large downstream market ($k > 1$)

Here too, the idea is to match up the DS and OO profit functions in whichever way is most felicitous to a comparison. Recall that under policy DS, if Firm 0 wishes its highest-value sale to be consumer x^*, then it must set $p_0^{DS} = \sqrt{x^*(1 - x^*)}$, and in this case the lowest-value consumer it sells to is $\hat{x} = 1 - p_0^{DS} - w = 1 - \left(2p_0^{DS} - x^*\right)$. Next recall that under policy OO, if Firm 0 wishes its marginal opt-out consumer to be x^*, then it must set $p_0^{OO} = 2\sqrt{x^*(1 - x^*)} - x^*$, and it sell to all consumers who are higher-value than $\hat{x} = 1 - p_0^{OO}$.

Although x^* has a slightly different interpretation under each policy, in both cases it may be treated as Firm 0's choice variable, so we may write profits as $\Pi_{DS}(x^*)$ and $\Pi_{OO}(x^*)$ respectively. Notice that the induced prices will satisfy $p_0^{OO} = 2p_0^{DS} - x^*$, so we have the very convenient feature that the "final sale" $\hat{x}$ will be the same under both policies! The diagram below shows total profits under each policy at an exemplary x^*. Profits under policy DS are outlined in red, and profits under OO are outlined in orange.

By construction, total downstream profits are the same under either policy. In both cases the revealed consumers $x \in [x^*, \hat{x}]$ are fully extracted. And in both cases the profits on hidden consumers are equivalent to selling to consumers $x \in [0, x^*]$ at the reservation value of consumer x^*. However, policy OO is more profitable upstream. Under free opt-out, Firm 0 can induce the same set $[x^*, \hat{x}]$ of revealed consumers without "discounting" its price, and it makes the additional sales on the opt-out consumers to boot. Thus we unambiguously have $\Pi_{OO}(x^*) > \Pi_{DS}(x^*)$ for any x^*. In particular, if $x^* = x_{DS}^*$ is profit-optimal under policy DS, then we have $\Pi_{OO}(x_{DS}^*) > \Pi_{DS}(x_{DS}^*) = \Pi_{DS}$. Of course x_{DS}^* is not necessarily the optimal choice under policy OO, so *a fortiori* we have $\Pi_{OO} > \Pi_{DS}$.

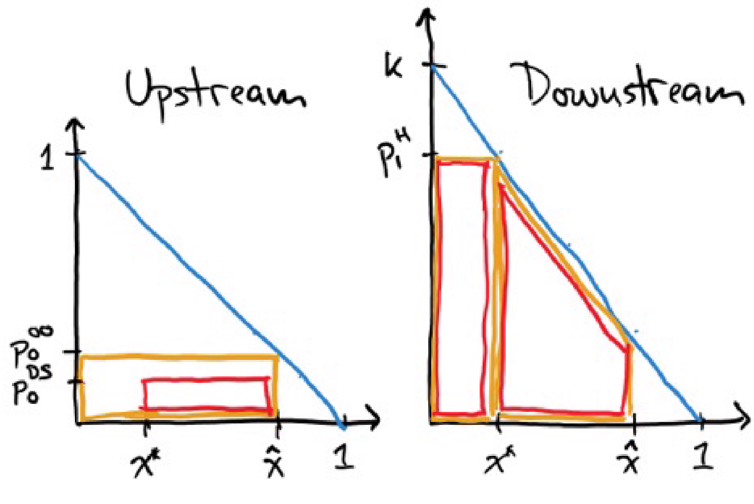

Figure 14:

C.2 Ranking consumer surplus under different policies

C.2.1 Data-sharing vs Autarky

Total consumer surplus is the sum of surplus in each market: $CS = CS_0 + CS_1$. The diagram below illustrates consumer surplus in under autarky and under data-sharing when the downstream market is small ($k < 1$). Switching to data-sharing involves an upstream gain $CS_0^{DS} - CS_0^A$ for consumers (due to the lower price p_0^{DS}) and a loss of surplus $CS_1^A - CS_1^{DS}$ downstream. We will show that the upstream gain outweighs the downstream loss. First, all of Firm 0's autarky customers benefit from the price cut that comes with data-sharing, so the upstream gain is at least $CS_0^{DS} - CS_0^A \geq \frac{1}{2}\Delta p$, where $\Delta p = p_0^A - p_0^{DS}$. And the downstream loss can be trivially bounded by $CS_1^A - CS_1^{DS} < CS_1^A$. So to prove that consumers gain from the switch to data-sharing, it suffices to show that $\frac{1}{2}\Delta p > CS_1^A = \frac{k}{8}$. It is straightforward to show that this condition holds at the equilibrium p_0^{DS}, so we have $CS_{DS} > CS_A$ when the downstream market is small.

When the downstream market is large ($k > 1$) we will be content to show that consumers are better-off under autarky for k sufficiently large, but better-off under data-sharing for k closer to 1. The first point is a consequence of the fact that by setting p_0^{DS} sufficiently close to zero, Firm 0 can induce all consumers to reveal, thereby extracting all of the surplus under the dowstream demand curve. The equilibrium p_0^{DS} trades off this incentive against the sacrifice of upstream profits, but as k grows

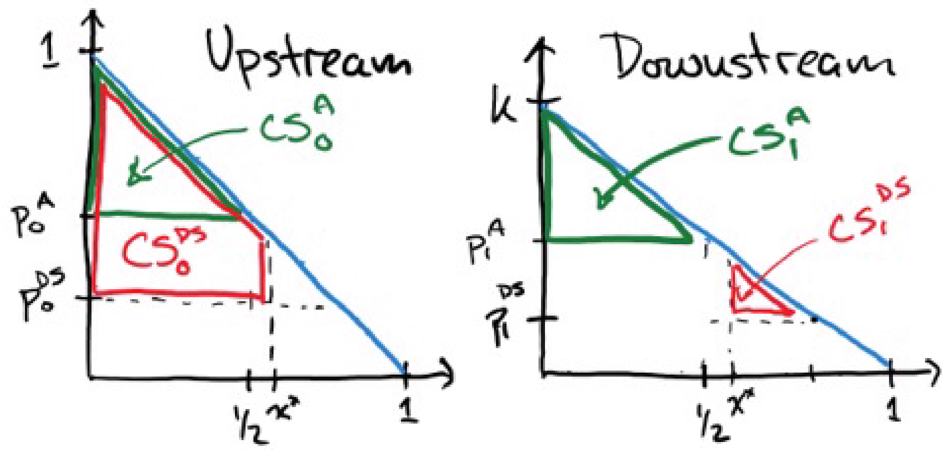

Figure 15:

larger the equilibrium p_0^{DS} and C_1^{DS} both tend to zero, while CS_0^{DS} tends to $\frac{1}{2}$. Recall that autarky consumer surplus is $CS_0^A = \frac{1}{8}$ and $CS_1^A = \frac{k}{8}$. For consumers, the gains in the smaller upstream market cannot compensate for losing (in the limit) all of their surplus in the larger downstream market, so we have $CS_{DS} < CS_A$ for k sufficiently large.

At the other extreme, for k near 1, the equilibrium price under model DS is $p_0^{DS} \approx 0.14$. Because every consumer has the option to buy at this price, including those who choose to wait, total consumer surplus satisfies $CS_{DS} > \frac{1}{2}(1-0.14)^2 > CS_A$ for k close to 1. Numerical computations suggest that consumers are better-off under data-sharing for $k \lesssim 2.4$, and better-off under autarky for $k \gtrsim 2.4$.

C.2.2 Free opt-out vs the worst alternative policy

Because we intend to demonstrate that consumer surplus is lower under the free opt-out regime than under either of the other two policies, we focus on comparisons between C_{OO} and whichever alternative, CS_{DS} or CS_A, is worse.

Opt-out vs. Autarky The figure below illustrates consumer surplus under autarky (green) and free opt-out (orange). Under opt-out, consumers gain from lower prices upstream but lose almost all of their surplus downstream. Recall that the equilibrium has Firm 1 offering hidden consumers its high price p_1^H with probability tending to one, so only the opt-out consumers to the left of x^* retain any surplus

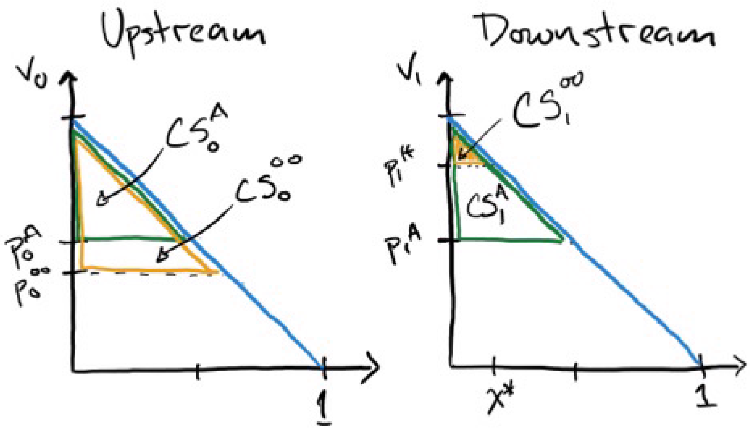

Figure 16:

downstream.

It can be established numerically that $CS_{OO} < CS_A$ whenever the downstream market is sufficiently large; in this case, sufficiently large means $k \gtrsim 0.25$. The logic begins with the fact that there are quite few opt-outs; regardless of k, the OO equilibrium has $x^* < 0.1$ and so $CS_1^{OO} < 0.005k$. Given this, Firm 0's first-order condition implies $p_0^{OO} \approx \frac{1}{2+k}$, so under the opt-out regime, upstream consumers get a "discount" of roughly $p_0^A - p_0^{OO} \approx \frac{1}{2}\frac{k}{2+k}$. One can show that (as long as the downstream market is not too small) this discount is not enough to compensate them for losing almost all of their downstream surplus.

Opt-out vs. Data-sharing when $k > 1$ Here we use the same hypothetical that we did in comparing profits. Fix a range $[x^*, \hat{x}]$ of revealed consumers, with $\hat{x} = 1 - \left(2\sqrt{x^*(1-x^*)} - x^*\right)$. Let $CS_{DS}(x^*)$ be the consumer surplus under policy DS when Firm 0 induces this range of revealed consumers by pricing at $p_0^{DS} = \sqrt{x^*(1-x^*)}$. Then let $CS_{OO}(x^*)$ be the consumer surplus under policy OO when Firm 0 induces the same range of revealed consumers by pricing at $p_0^{OO} = 2\sqrt{x^*(1-x^*)} - x^* = 2p_0^{DS} - x^*$.

The figure depicts $CS_{DS}(x^*)$ (in red) and $CS_{OO}(x^*)$ (in orange). We have already discussed the story behind $CS_{OO}(x^*)$, but the diagram for $CS_{DS}(x^*)$ requires some explanation. Strictly speaking, the upstream consumer surplus is area Y. Areas X and Z involve tricks for displaying some of the downstream surplus on the upstream diagram. Recall that consumers $x \in \left[\hat{x}, 1 - p_0^{OO}\right]$ are indifferent between buying upstream and waiting for a chance at Firm 1's low price. While they wait in

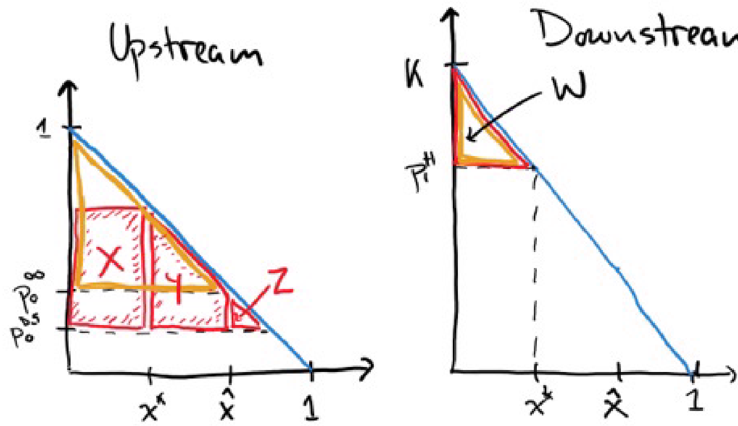

Figure 17:

equilibrium, because of their indifference their consumer surplus may be represented by area Z. Consumer x^* is also indifferent between her surplus s_1^* at Firm 1 and her surplus s_0^* at Firm 0; the latter is the height of rectangle X. Her surplus at Firm 1 is $s_1^* = \alpha\left(kv_0^* - p_1^L\right) + (1-\alpha)\left(kv_0^* - p_1^H\right) = \alpha\left(kv_0^* - p_1^L\right)$ (since $p_1^H = kv_0^*$). A consumer to her left with the higher value $v_0 > v_0^*$ waits for Firm 1 and earns surplus $s_1 = s_1^* + k\left(v_0 - v_0^*\right) = s_0^* + \left(kv_0 - p_1^H\right)$. Area X represents the first term of this surplus, summed over $x \in [0, x^*]$, and area W represents the second term.

Expressions for each consumer surplus function may be read directly off the figure. We have:

$$
\begin{aligned}
CS_{DS}\left(x^*\right) &= \frac{1}{2}\left(1-p_0^{DS}\right)^2 + \frac{1}{2}(k-1)\left(x^*\right)^2 \\
CS_{OO}\left(x^*\right) &= \frac{1}{2}\left(1-p_0^{OO}\right)^2 + \frac{1}{2}k\left(x^*\right)^2
\end{aligned}
$$

As usual, let x_{DS}^* and x_{OO}^* be the equilibrium levels of x^* under each policy, with $CS_{DS} = CS_{DS}\left(x_{DS}^*\right)$ and $CS_{OO} = CS_{OO}\left(x_{OO}^*\right)$ the equilibrium levels of consumer surplus. The strategy for showing $CS_{DS} > CS_{OO}$ has three steps. First we show that $CS_{DS}\left(x^*\right) > CS_{OO}\left(x^*\right)$ holds (under suitable conditions). Then we show that $x_{OO}^* > x_{DS}^*$. And finally we show that $CS_{OO}'\left(x^*\right)$ is negative (again, under suitable conditions). Together these steps imply that $CS_{OO}\left(x_{OO}^*\right) < CS_{OO}\left(x_{DS}^*\right) < CS_{DS}\left(x_{DS}^*\right)$, as claimed.

For the sake of brevity, we will just sketch the steps. For the first step, $CS_{DS}\left(x^*\right) - CS_{OO}\left(x^*\right)$ has a positive term proportional to the price difference $p_0^{OO} - p_0^{DS}$ and a negative term proportional to $(x^*)^2$, so it suffices to show that the former dominates the latter. For the second step, we draw attention to the diagram in the profit comparison of these two cases. At a given x^*, Π_0^{DS} and Π_0^{OO} coincide on their downstream profits; they differ only on upstream profits π_0^{DS} and π_0^{OO}. This makes it relatively simple to evaluate the difference in marginal profits $\frac{d\Pi_0^{OO}}{dx^*} - \frac{d\Pi_0^{DS}}{dx^*}$, and this lets us establish that $\left.\frac{d\Pi_0^{OO}}{dx^*}\right|_{x^*=x^*_{DS}}$ is strictly positive. Finally, $CS'_{OO}(x^*)$ has a positive (downstream) term and negative (upstream) term, but the profit first-order condition can be used to establish that the negative term dominates for $x^* \leq x^*_{OO}$.

D References

Acquisti, A., Taylor, C. and Wagman, L., 2016. The economics of privacy. Journal of Economic Literature, 54(2), pp. 442-92.

Anderson, S., Baik, A. and Larson, N., 2015. Personalized pricing and advertising: An asymmetric equilibrium analysis. Games and Economic Behavior, 92, pp. 53-73.

Anderson, S., Baik, A. and Larson, N., 2019. Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts. CEPR Discussion Paper 13793

Belleflamme, P. and Vergote, W., 2016. Monopoly price discrimination and privacy: The hidden cost of hiding. Economics Letters, 149, pp. 141-144.

Belleflamme, P., Wynnw Lam, W.M. & Vergote, W. (2017), Price discrimination and dispersion under asymmetric profiling of consumers.

Bergemann, D. and Bonatti, A., 2011. Targeting in advertising markets: implications for offline versus online media. The RAND Journal of Economics, 42(3), pp. 417-443.

Bergemann, D. and Bonatti, A., 2015. Selling cookies. American Economic Journal: Microeconomics, 7(3), pp. 259-94.

Bester, H. and Petrakis, E., 1995. Price competition and advertising in oligopoly. European Economic Review, 39(6), pp.1075-1088.

Bester, H. and Petrakis, E., 1996. Coupons and oligopolistic price discrimination. International journal of industrial organization, 14(2), pp.227-242.

Braulin, F.C. and Valletti, T., 2016. Selling customer information to competing firms. Economics Letters, 149, pp.10-14.

Conitzer, V., Taylor, C.R. and Wagman, L., 2012. Hide and seek: Costly consumer privacy in a market with repeat purchases. Marketing Science, 31(2), pp.277-292.

Calzolari, G. and A. Pavan (2006). On the optimality of privacy in sequential contracting. Journal of Economic Theory 130 (1), 168{204.

Fudenberg, D. and Tirole, J., 2000. Customer poaching and brand switching. RAND Journal of Economics, pp.634-657.

Jann, Ole and Christoph Schottmuller (2018) An Informational Theory of Privacy. Forthcoming, The Economic Journal

Iyer, G., Soberman, D. and Villas-Boas, J.M., 2005. The targeting of advertising. Marketing Science, 24(3), pp.461-476.

Montes, R., Sand-Zantman, W. and Valletti, T.M., 2019. The value of personal information in markets with endogenous privacy. Management Science, 65(3) 1342–1362.

Posner, E.A. and Weyl, E.G., 2018. Radical Market. Princeton University.

Taylor, C.R., 2004. Consumer privacy and the market for customer information. RAND Journal of Economics, pp.631-650.

Tucker, C.E., 2015. Privacy and the Internet. In Handbook of Media Economics (Vol. 1, pp. 541-562). North-Holland.

Turow, J., 2017. The aisles have eyes: How retailers track your shopping, strip your privacy, and define your power. Yale University Press.

Turow, J., King, J., Hoofnagle, C.J., Bleakley, A. and Hennessy, M., 2009. Americans reject tailored advertising and three activities that enable it.

Valletti, T. & Wu, J. (2016), Big data versus small data: Consumer profiling with data requirements.

Villas-Boas, J.M., 1999. Dynamic competition with customer recognition. The Rand Journal of Economics, pp.604-631.

White, T.B., Zahay, D.L., Thorbjørnsen, H. and Shavitt, S., 2008. Getting too personal: Reactance to highly personalized email solicitations. Marketing Letters, 19(1), pp. 39-50.

E Appendix

F Opt-in Example

This is a proof-of-concept example where (1) some consumers choose to opt in, and (2) outcomes are different when consumers choose to opt-in, vs choose to opt-out. The idea is that low-value consumers who can't afford the hidden price p_1 may want to opt-in. But they must have some reason to do this, so there must be some reason that Firm 1 can't fully extract them. Here the reason is that Firm 0's information about a customer's v_1 is imperfect, so Firm 1 sometimes leaves this consumer some surplus.

F.1 Setup

Consumer valuations Consumers are the unit interval. All consumers have the *same* value $v_0 = 1$ at Firm 0. This simplifies the analysis and keeps the focus on what's happening downstream. Each consumer has a type $v \in [0, 1]$ that is a noisy indicator of her downstream value at Firm 1 – this is what Firm 0 observes about her, if she buys upstream. The consumer knows her v too, but doesn't observe her true v_1 until she shows up at Firm 1. With probability $q > \frac{1}{2}$, that value is $v_1 = v$, but with probability $1 - q$ her value is a new uniform draw: $v_1 \sim U[0, 1]$. So q may be thought of as the accuracy of Firm 0's information.

Autarky outcome The "remixing" of downsteam values is inessential to Firm 1 in this case – the final distribution of v_1 is uniform and so $p_1^A = \frac{1}{2}$. Firm 0's monopoly price is $p_0^A = 1$. Profits are $\pi_0^A = 1$ and $\pi_1^A = \frac{1}{4}$. Consumers get 0-surplus upstream and $CS_1 = \frac{1}{8}$, so total consumer surplus is $CS_A = \frac{1}{8}$.

Notation Write $s_0 = 1 - p_0$ for each consumer's surplus from a purchase at Firm 1. Let $s_r(v)$ be the expected surplus at Firm 1 of a type-v consumer who is revealed. If p_1 is Firm 1's price to hidden consumers, let $s_h(v\ ;\ p_1)$ be the expected surplus of a type-v consumer who isn't revealed.

F.2 Preliminary analysis

How does Firm 1 price to revealed consumers? Remember $q > \frac{1}{2}$, so more often than not when Firm 1 faces a type-v consumer, her value is really $v_1 = v$. Unless v is quite small, Firm 1's best option is to set $p_1(v) = v$. The consumer buys unless her value has been shuffled (prob $1-q$), and her new draw is below v (prob. v). So the revealed profit on this customer is $\pi_1^{rev}(v) = (1-(1-q)v)v$. [More detail: for v small enough, Firm 1 is better off setting $p_1 = \frac{1}{2}$ (on the chance that the consumer has drawn a new value). That earns $\pi_1^{rev}(v) = \frac{1}{4}(1-q)$.]

How does Firm 1 price to hidden consumers? Suppose consumer types $v \geq v^*$ are hidden and types $v < v^*$ are revealed. So Firm 1 has $1 - v^*$ hidden consumers. It knows that a fraction q of them will have $v_1 = v$, and the others will be new $U[0,1]$ draws, so its hidden demand will be a mix of $v_1 \in [v^*, 1]$ (weight q) and $v_1 \in [0,1]$ (weight $1-q$). This hidden demand looks like the figure below. When the threshold consumer satisfies $v^* \geq \frac{1}{2}$, it is optimal to price at the kink: $p_1 = v^*$. Otherwise, if $v^* < \frac{1}{2}$, the monopoly price $p_1 = \frac{1}{2}$ is optimal.

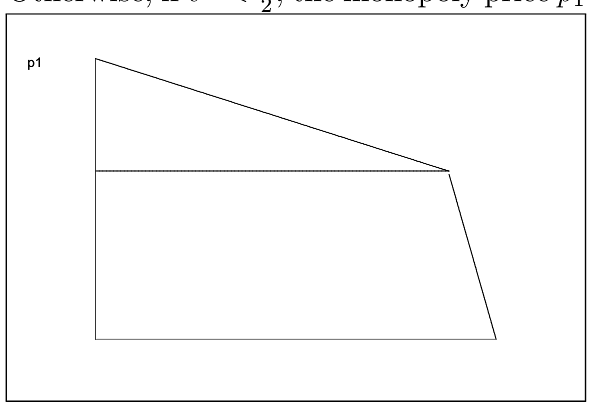

The surplus from hiding vs. being revealed A type-v consumer who is hidden gets surplus $\max(v - p_1^e, 0)$ if her value turns out to be $v_1 = v$. Otherwise, her expected surplus is the CS of a uniformly-drawn consumer facing price p_1^e: $\frac{1}{2}(1-p_1^e)^2$. So her total expected surplus is $s_h(v\ ;\ p_1^e) = q(v - p_1^e) + \frac{1}{2}(1-q)(1-p_1^e)^2$ if $v_1 > p_1^e$, or $s_h(v\ ;\ p_1^e) = \frac{1}{2}(1-q)(1-p_1^e)^2$ if $v_1 \leq p_1^e$. A revealed consumer who faces price $p_1(v) = v$ gets no surplus unless v doesn't reflect her true v_1. In this case (prob $1-q$), she gets the expected CS of a uniformly drawn consumer facing $p_1 = v$. So $s_r(v) = \frac{1}{2}(1-q)(1-v)^2$. [This won't be correct for very low v – see above – but we can ignore that for now.]

Which surplus is larger: $s_h(v\ ;\ p_1^e)$ or $s_r(v)$? While this comparison doesn't correspond exactly to the choice consumers make, it is useful as a benchmark. Holding p_1^e fixed, the hidden surplus $s_h(v\ ;\ p_1^e)$ is increasing in v, while the revealed surplus $s_r(v)$ is decreasing in v. So a consumer forced to choose between these two surpluses will tend to prefer to hide if v is high and reveal if v is low. In fact, the threshold type is $v = p_1^e$: $s_h(v\ ;\ p_1^e) > s_r(v)$ iff $v > p_1^e$. There is some logic to this. A consumer of type $v = p_1^e$ faces the same price whether she hides or reveals; hence her indifference. If $v > p_1$, she faces a higher price when she reveals than when she hides, whereas if $v < p_1^e$ revealing gets her a lower price.

F.3 Data-sharing

Under this policy, Firm 0 can sell data and consumers can't opt out. Let v^* be the threshold consumer, so $v > v^*$ don't buy at Firm 0 and are hidden, while $v < v^*$ do buy at Firm 0 and are revealed. The threshold consumer satisfies $s_0 + s_r(v^*) = s_h(v^*\ ;\ p_1^e)$. Remember, the consumer with $v = p_1^e$ will be indifferent between s_r and s_h. But here we are comparing $s_0 + s_r$ to s_h, so if Firm 0 offers positive surplus, the threshold consumer must be strictly higher: $v^* > p_1^e$. But – see above – Firm 1 will want to price at $p_1 = v^*$, which isn't consistent with $v^* > p_1^e$. This is basically the same hold-up problem we have seen elsewhere. Any interior ($v^* < 1$) equilibrium threshold would have to satisfy $s_0 + s_r(v^*) = s_h(v^*\ ;\ v^*)$, but since $s_r(v^*) = s_h(v^*\ ;\ v^*)$ this is impossible if $s_0 > 0$ – the lefthand side (buying at Firm 0 and revealing) is always better. So if $s_0 > 0$, the equilibrium must have *all* consumers buying at Firm 0 and revealing. Since Firm 0 can induce this outcome with any $s_0 > 0$, the equilibrium must have $s_0 \to 0$, $p_0 \to 1$, $p_1 = 1$, and all consumers reveal.

Technical notes: (1) there is another possibility – $v^* < \frac{1}{2}$ and Firm 1 sets $p_1 = \frac{1}{2}$ – but this makes the contradiction worse. (2) Technically, Firm 0's optimal s_0 – as small as possible, but strictly positive – isn't well defined. But this is just a technicality. (3) Remember, revealed consumers with v low enough get offered $p_1(v) = \frac{1}{2}$ rather than $p_1(v) = v$. But that's is still better than the hidden price $p_1 = 1$, so this doesn't affect their preference for revealing.

Intuition: if there were an interior threshold consumer v^*, this consumer would have to expect to face a downstream price of v^* whether she is hidden or revealed. But in this case, she might as well reveal so that she can also get surplus upstream.

F.3.1 Profits

Upstream profit is $\pi_0 = 1$, and there are no hidden consumer profits. Revealed consumer profits are a little messy: $\pi_1^{rev} = \int_0^1 \max\left(\left(1-\left(1-q\right)v\right)v\ ,\ \frac{1}{4}\left(1-q\right)\right)\ dv$. Total profit is $\Pi_{DS} = \pi_0 + \pi_1^{rev} - \pi_1^A = \frac{3}{4} + \pi_1^{rev}$. The messiness is entirely in the cutoff between the two revealed profit terms in the integral. But as a lower bound, $\pi_1^{rev} > \int_0^1 \left(1-\left(1-q\right)v\right)v\ dv = \frac{1}{3}q + \frac{1}{6}$, so $\Pi_{DS} > \frac{11}{12} + \frac{1}{3}q$. So since $q > \frac{1}{2}$, we unambiguously have $\Pi_{DS} > \Pi_A$.

F.3.2 Consumer surplus

Consumers get zero surplus upstream and (since all are revealed), $CS_1 = \int_0^1 s_r\left(v\right)\ dv$ downstream. So total consumer surplus is $CS_{DS} = CS_1 = \frac{1}{6}\left(1-q\right)$. Since $q > \frac{1}{2}$, this is unambiguously worse than autarky ($CS_A = \frac{1}{8}$).

F.4 Opt-out ($c \to 0$ limit)

Under this policy, Firm 0's customers have their data sold by default, but they can opt out at cost $c > 0$. We focus on the $c \to 0$ limit, so that opting-out is ever so slightly more inconvenient than the default. The punchline is that the equilibrium is essentially identical to the data-sharing case.

Since we plan to take c very small, suppose for now that Firm 0 offers enough surplus to cover that cost: $s_0 > c$. In this case, no consumer will decline to purchase at Firm 0. The only question is whether or not to opt-out. The general form of the downstream surpluses if revealed or hidden ($s_r\left(v\right)$ or $s_h\left(v\ ;\ p_1^e\right)$) doesn't change at all, so it is still true that higher v types will prefer to be hidden (in this case, by opting out) and lower v types will prefer being revealed. The surplus comparison is

$$\text{Opt-out:}\quad s_0 + s_h\left(v\ ;\ p_1^e\right) - c \qquad \text{vs.} \qquad \text{Reveal:}\quad s_0 + s_r\left(v\right)$$

So a consumer opts out iff v is high enough that $s_h\left(v\ ;\ p_1^e\right) - c > s_r\left(v\right)$. But the same hold-up problem crops up. If there is an interior threshold consumer $v^* \geq \frac{1}{2}$, then Firm 1 will set its hidden price at $p_1 = v^*$. In this case, $s_h\left(v^*\ ;\ p_1^e\right) = s_h\left(v^*\ ;\ v^*\right) =$

$s_r(v^*)$. That is, the putative threshold consumer faces the same price whether hidden or not, hence the same dowstream surplus whether hidden or not. But in this case, she's strictly better off saving c by not opting out. So she's not indifferent, and the true threshold must be higher. Since this argument applies at any interior v^*, we are in the same situation as the data-sharing model. For any $c > 0$ and $s_0 > c$, the equilibrium has $v^* = 1$, $p_1 = 1$, all consumers reveal, and no one opts out. Firm 0 has an incentive to induce this outcome, so the limit case has $s_0, c \to 0$, $p_0 \to 1$, and exactly the same outcomes as the data-sharing model.

F.4.1 Profits and consumer surplus

So $\Pi_{OO} = \Pi_{DS} > \Pi_A$ and $CS_{OO} = CS_{DS} = \frac{1}{6}(1-q) < CS_A$.

F.5 Opt-in ($c \to 0$ limit)

Under this policy, Firm 0's customers do not have their data sold by default, but they can opt in at nuisance cost $c > 0$. Again we will take $c \to 0$, since the idea is that consumers who are approximately indifferent about what to do just stick to the default. The punchline is that the outcome here will be different than when consumers must actively opt out.

As earlier, suppose that Firm 0 offers $s_0 > c$, so once again all consumers will buy at Firm 0. Their choice is between the default (which amounts to staying hidden) and opting-in. The surplus from each is

$$\text{Default (stay hidden):} \quad s_0 + s_h(v\ ;\ p_1^e) \qquad \text{vs.} \qquad \text{Opt-in:} \quad s_0 + s_r(v) - c$$

It's the same comparison as before, except that the hassle cost c now tips a nearly indifferent consumer in favor of staying hidden rather than opting in. This is going to lead us to an unraveling argument like the earlier ones, but in the opposite direction.

Remember, Firm 1 will set its hidden price at $p_1 = v^*$ whenever the threshold consumer is $v^* \geq \frac{1}{2}$. Otherwise it just sets $p_1 = \frac{1}{2}$. We can't have an equilibrium with a threshold consumer $v^* > \frac{1}{2}$. If we did, the hidden price would be $p_1 = v^*$, and consequently $s_h(v^*\ ;\ p_1^e) = s_h(v^*\ ;\ v^*) = s_r(v^*)$, just as above. That is, the putative threshold consumer expects the same price whether hidden or revealed, so she strictly prefers sticking to the default (staying hidden), since that saves c. So she's not really the threshold consumer, and v^* must be lower. This unraveling argument drives v^*

down until it is just below $\frac{1}{2}$, at which point the hidden price stops falling and stays at $p_1 = \frac{1}{2}$. So the equilibrium threshold satisfies

$$s_h\left(v^* \ ; \ \frac{1}{2}\right) = s_r\left(v^*\right) - c$$

The threshold consumer $v^* < \frac{1}{2}$ gets a slightly better price by revealing (v^* vs. $\frac{1}{2}$), and this is just enough to compensate for the cost c. Firm 0 has an incentive to keep $s_0 > c$ (otherwise it would lose some sales), but only slightly – i.e., $p_0 \to 1 - c$.

In the $c \to 0$ limit, the condition for the indifferent consumer becomes $s_h\left(v^* \ ; \ \frac{1}{2}\right) = s_r\left(v^*\right)$, which is satisfied at $v^* = \frac{1}{2}$. So under the opt-in policy, the equilibrium has $p_0 \to 1$, $p_1 = \frac{1}{2}$, and exactly half of the consumers opt-in (those with $v \in \left[0, \frac{1}{2}\right]$). So there is less data-sharing under opt-in than under opt-out.

F.5.1 Profits

Firm 0 earns $\pi_0 = 1$ upstream. Of the $v \in \left[\frac{1}{2}, 1\right]$ consumers who don't opt in, all will buy at the hidden price except for those who draw a new value less than $p_1 = \frac{1}{2}$. That means hidden sales are $\left(1 - v^*\right)\left(1 - (1-q)\,p_1\right) = \frac{1}{2}\left(1 - \frac{1}{2}(1-q)\right)$, and hidden profits are $\pi_1^{hid} = \frac{1}{8}(1+q)$. Revealed-consumer profits are $\pi_1^{rev} = \int_0^{\frac{1}{2}} \max\left(\left(1 - (1-q)\,v\right)v \ , \ \frac{1}{4}(1-q)\right) dv$. Total profit is $\Pi_{OI} = \pi_0 + \pi_1^{hid} + \pi_1^{rev} - \pi_1^A = \frac{7}{8} + \frac{1}{8}q + \pi_1^{rev}$. Using the same lower bound as earlier to get $\pi_1^{rev} > \frac{1}{24}q + \frac{1}{12}$, we have $\Pi_{OI} > \frac{23}{24} + \frac{1}{6}q$. Since $q > \frac{1}{2}$, this is unambiguously better than autarky.

F.5.2 Consumer surplus

Consumers get nothing upstream. Total downstream surplus is $CS_1 = \int_0^{\frac{1}{2}} s_r(v)\,dv + \int_{\frac{1}{2}}^1 s_h\left(v \ ; \ \frac{1}{2}\right) dv$, and so $CS_{OI} = CS_1$. With direct calculations, we have $CS_{OI} = \frac{5-2q}{24}$. Since $q < 1$, we have $CS_{OI} > CS_A$. And given the other rankings, we have $CS_{OI} > CS_A > CS_{DS} = CS_{OO}$. In plain English, consumers are best-off when data-sharing is permitted but consumers must actively opt-in. The outcome is not the same if data-sharing is the default but consumers can opt out – in that case, consumers are worse off.

F.6 Discussion

This is just a proof-of-concept example to show that the default – opt-in or opt-out – can matter even if the costs of the non-default choice are negligible and even without new terms (privacy for its own sake?) in consumer utility.

The key to the results is that an almost-negligible cost c can tip behavior a lot depending which side of the ledger it falls on. And I think it's fair to say that this "tippiness" depends on the fact that, in equilibrium, a threshold consumer expects to face the same downstream price whether she opts in or out. This is a feature that is more general than this example. (It shows up in the perfect correlation model that we worked through too.) When there is no uncertainty about types, it is a fairly straightforward property: Firm 1 wants to price at the full willingness-to-pay of its marginal hidden customer, $p_1 = v_1^*$, and if she switches to revealing herself, it wishes to fully extract her in that case too. If observable types are an imperfect indicator of v_1, as here, it depends a bit on the modeling whether that property carries through.

THESE ARE CORRECT CS PLOTS:

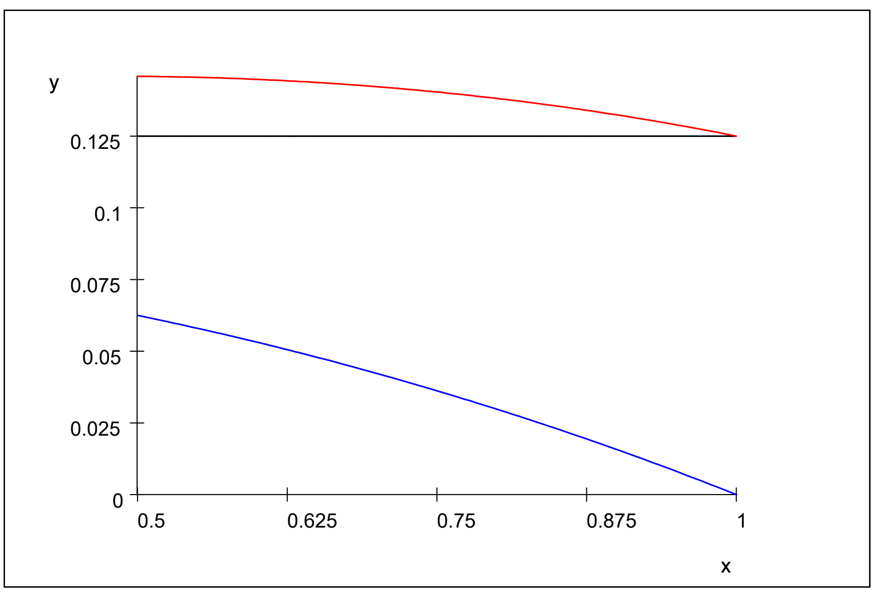

Same, prepare for figure:

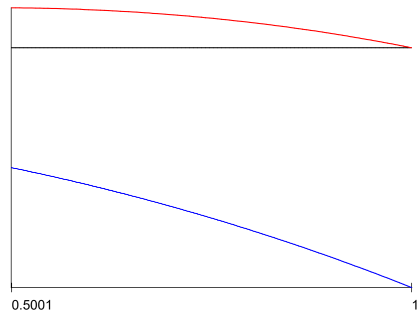
0.5001
1