Data-enabled learning, network effects and competitive advantage*

Andrei Hagiu† and Julian Wright‡

May 2021

Abstract

We model competition between firms which improve their products through learning from customer data. We explore the implications of the model for competitive dynamics and efficiency. The model allows us to analyze factors affecting an incumbent's competitive advantage such as the shape of firms' learning functions and the extent of data accumulation. We also explore the implications of public policies towards data sharing, user privacy and data acquisitions. Finally, we show conditions under which a consumer coordination problem arises endogenously from data-enabled learning, thus making it possible for consumer beliefs to matter for the incumbent's competitive advantage and the efficiency of the outcome.

Keywords: machine learning, dynamic competition, data sharing, switching costs.

1 Introduction

In recent years, much attention has been focused on the role data can play in providing incumbent firms with a competitive advantage. Digitization, connectivity to cloud-based infrastructures, together with cheaper storage and more effective use of data (i.e. improvements in machine learning algorithms), have made it possible for firms in many industries to translate learning from their customer data into rapid improvements in their products. And with better products, these firms can attract more customers (or more usage from existing

*We would like to thank LMU (and Tobias Kretschmer in particular) for hosting us in May and June of 2019, and the valuable feedback colleagues there gave us as we developed a very early version of this paper. Thanks also to Luis Cabral, Emilio Calvano, Jacques Cremer, Bruno Jullien, and Martin Peitz as well as other participants at the 17th ZEW Conference on ICT, APIOC 2019, IDEI's 2020 Digital Economics Conference, KER 2020, and at online seminars hosted at Centre for Competition Policy, Cornell University, Digital Economy Research Network, MaCCI/EPoS, NUS, Oxera, and TSE, for their comments and suggestions. Gunhaeng Lee and Tat-How Teh provided excellent research assistance.

†Boston University Questrom School of Business, E-mail: ahagiu@bu.edu

‡Department of Economics, National University of Singapore, Singapore, E-mail: jwright@nus.edu.sg

customers), and therefore obtain more data, potentially creating a self-reinforcing cycle that can make it difficult for any new entrant to compete. We call this self-reinforcing cycle "data-enabled learning".[1]

There are a plethora of products and services that leverage data-enabled learning. Some work purely by leveraging data across different users, whereas others personalize the learning at the individual level. As an example of the former, consider Grammarly. This well-known cloud-based service suggests spelling, grammar, tone, style and word choice improvements in exchange for a monthly subscription fee. As users make corrections to the suggestions offered by Grammarly, its language experts and artificial intelligence can use this feedback to continue to improve its future recommendations for all users. Other text processing applications such as chatbots (LivePerson), document management (Luminance) and productivity tools (x.ai) work in a similar way, as do computer vision applications such as driver assistance (Mobileye), autonomous vehicle systems (Cruise), skin cancer detection (SkinVision), farm monitoring (Prospera), home security (Deep Sentinel) and airport screening (SeeTrue), and speech recognition applications such as dictation software (Dragon), translation services (Unbabel) and virtual assistants (Amelia).

In contrast to these applications based on across-user learning, consider Fitbit's premium service, which provides users with personalized advice on health, sleep and fitness based on data collected by their Fitbit device. As customers use their Fitbit device more, the system is able to provide them with better and more helpful advice, which means users tend to prefer it to any new alternative device. Other popular smart connected devices that improve based on individual user data include Ecobee's smart thermostats and cameras (which learn user preferences and habits, and use that to optimize temperature setting and home monitoring), Eight Sleep's smart beds (which adjust the bed's temperature to optimize the user's sleep), Tonal's home gym system (which adjust weights and resistance to optimize workouts), etc. And of course, many applications work by combining both across-user learning and within-user learning, such as paid services that are built on recommendation systems like those offered by Calm, Netflix, Spotify, StitchFix and Tinder.

Data-enabled learning would seem to give incumbent firms a competitive advantage. But how strong is this advantage and how does it differ from that obtained from more traditional mechanisms, such as (i) learning-by-doing, which allows a firm to lower its costs as it produces (and sells) more, thereby being able to offer a lower price and attract more sales, and (ii) standard network effects, which make a firm more valuable to consumers as it attracts more

[1] Some practitioners have called this phenomenon "data network effects". Our use of "data-enabled learning" is intended to be more general since, as discussed below, the phenomenon need not involve any network effect (for example, if the learning is specific to each customer).

consumers, thereby attracting even more consumers and offering even higher value?

To address these questions, and determine the drivers of competitive advantage with data-enabled learning, we develop a theory of dynamic competition in which firms can improve their products via data gleaned from their customers. The theory is based on Bertrand competition between two infinitely lived firms, which can differ in the amount of their prior learning as well as in the shape of their respective learning functions. We allow firms to set prices to compete for consumers every period, taking into account how attracting consumers (and therefore more data) increases the value they can offer to consumers in subsequent periods. Aside from differences in standalone utility and learning functions, the two firms in our model are assumed to be identical. An attractive feature of our particular framework is that, despite allowing for general learning functions, we are able to completely characterize the equilibrium outcomes, determine which firm will win, and each firm's value for any arbitrary starting position of the two firms.

As a baseline, we consider the case with across-user learning only (Section 3), in which the two firms only learn based on the total number of consumers they served in the past. Since consumers' willingness-to-pay can increase over time with more users even if prices remain unchanged, there is a type of dynamic network effect at play. This provides one distinction from traditional learning-by-doing settings, where consumers benefit only via lower prices. Another distinction arises when we allow firms to learn and improve their products while their customers are still consuming them. This is often the case with cloud-based products, but was not possible with traditional products (e.g. physical goods, packaged software). When the product improves while it is still being consumed, consumers need to form expectations over how many other consumers will purchase the same product in order to determine the value they will get from the product, thereby creating a role for consumer beliefs.

We then consider the case of within-user learning only (Section 4), in which any given consumer's willingness-to-pay for a given firm's product only depends on how many times the firm has served that particular consumer in the past. This highlights a further key difference between data-enabled learning and traditional mechanisms, which is the role of customization: firms can often improve their products for each individual customer based on that customer's particular usage experience. Within-user learning implies the more a customer uses a firm's product, the greater is the effective cost of switching to the rival's product. Such data-enabled learning therefore creates an endogenous switching cost that grows over time, but not any type of network effect.

In these baseline models, we show that despite the fact that sometimes both the winning firm and the losing firm price below cost to consumers, the competitive outcome coincides with the socially optimal outcome.[2] This is not at all an obvious result despite our Bertrand competition setting: the determination of which firm wins in a particular period depends on how much each firm is willing to subsidize consumers in the current period, which in turn depends on how much each firm's value increases by winning this period as opposed to losing this period, both of which are only defined recursively. We find the condition determining which firm will win is the same regardless of whether we assume pure across-user learning or pure within-user learning, and whether consumers are forward looking or myopic. Yet unlike the case with across-user learning, if learning is within users, the winning firm's profit is lower when consumers are forward looking than when consumers are myopic. Fundamentally, this is because within-user learning creates endogenous switching costs, so each forward-looking individual consumer takes into account that they would benefit if the losing firm had a stronger competitive position in the future, which in turn makes the winning firm have to price more competitively. We also explain how various factors, such as how far each firm is along its learning curve and the nature of the learning curve, affect the condition for either firm to win.

In Section 5 we combine the two types of learning (across users and within users) and explore when and how this combination by itself can make consumers' beliefs matter for the equilibrium outcome. With myopic consumers we show things remain equivalent to the case of pure across-user learning, in which consumer beliefs don't matter. However, once we allow for forward-looking consumers, beliefs can matter even in the absence of any within-period learning, exogenous switching costs, or strategic pricing by the firms. Specifically, even if the product is available for free and consumers are free to change firms each period, within-user learning still creates an endogenous switching cost. This implies consumers want to choose the product that other consumers are expected to choose in order to also benefit from across-user learning in subsequent periods. For example, this means Google search likely benefits from favorable consumer beliefs even though it is "costless" for users to switch to Bing. We then show a similar coordination problem among consumers arises when firms can set prices, provided across-user learning is sufficiently important relative to within-user learning.

A key implication of our analysis is that even though data can create a type of dynamic network effect, there is no systematic inefficiency in the market outcome, unless the dynamic network effect leads to a consumer coordination problem. On the other hand, even absent any such coordination problem, our analysis implies that increasing the learning of the losing firm generally increases consumer surplus. This might suggest a policy that forces winning firms to share the data obtained with their rivals would be good for consumers. We use our dynamic

[2] The only exception to this occurs with across-user learning when products can improve while they are still being consumed, so there is a role for consumer beliefs to distort the equilibrium outcomes.

framework to explore how data sharing works, and find that it may result in firms competing less aggressively, which in our model means subsidizing consumers less, thus potentially lowering consumer surplus. We also use our model to highlight the unintended consequences of privacy policies. If such policies reduce the rate at which firms can extract useful data from consumers, they will tend to increase the incumbent's competitive advantage, reflecting that the entrant has more scope for new learning and so is affected more by such a policy. Finally, our efficiency result no longer necessarily holds when the firms can acquire data directly (e.g. from a third party) rather than via attracting more consumers. Indeed, we show that if one firm has already reached its maximum learning threshold, that firm will acquire the available data too often relative to what is socially efficient.

2 Related literature

Our paper combines aspects from three large and established literatures. First, it relates to works that provide theoretical models of competition with learning-by-doing (Fudenberg and Tirole, 1983, Dasgupta and Stiglitz, 1988, Salant, 1990, Cabral and Riordan, 1994, Besanko et al., 2010). On the one hand, our treatment of pure across-user learning (in the case without within-period learning), in which learning increases consumers' willingness to pay in subsequent periods, can be viewed as isomorphic to the decreases in firms' costs arising under traditional learning-by-doing models. Yet, our particular model of learning is quite distinct from any in the existing literature. The baseline model with pure across-user learning is closest to that in the classic paper of Cabral and Riordan (1994). A key difference is that they allow for horizontal differentiation between the two firms, whereas we allow for vertical differentiation (both in standalone values and via asymmetric learning functions). Because of the tractability of our particular framework, we are able to obtain several new results, including a general efficiency result, the implications of learning for consumer surplus, and comparative statics on the shape of the learning curves. And more fundamentally, we apply the framework to study questions that do not arise in learning-by-doing settings, specifically policy questions regarding the implications of data sharing, tighter customer privacy, and data acquisitions. Furthermore, the within-user learning version of our model is entirely new, as is the role of endogenous switching costs in that setting. Finally, a further novelty of data-enabled learning relative to learning-by-doing is the possibility for consumer coordination problems to arise, such as when we combine across-user learning with within-period learning or within-user learning.

Second, our paper relates to the substantial literature on network effects. From this literature we borrow the modelling of consumer "beliefs" to select equilibria when consumers play coordination games (Katz and Shapiro, 1986, and Caillaud and Jullien, 2001). At a high level, our model with across-user and within-period learning is related to dynamic models with network effects (e.g. Mitchell and Skrzypacz, 2006, Cabral, 2011, Biglaiser and Cremer 2020, Halaburda et al., 2020), since the combination of these two features works in the same way as exogenous network effects within a period. However, a key difference is that in our setting with learning, willingness-to-pay depends not only on the current stock of consumers, as it would in network effects models, but also on past consumers even if they are no longer consuming from the firm.

Third, our paper relates to IO models of switching costs, such as the classic papers of Klemperer (1987) and Beggs and Klemperer (1992). As noted in the introduction, within-user learning creates an endogenous switching cost that grows the longer a consumer is with a particular firm. To the best of our knowledge, such a feature has not previously been considered in the switching cost literature even though it seems relevant more generally.

Finally, there are some recent works that also study aspects of data-enabled learning, albeit using approaches that are very different from ours. Prufer and Schottmüller (2017) model a dynamic feedback loop in which a firm's current cost of investing in quality is decreasing in the firm's previous period sales, and greater investment in quality leads to higher demand in the current period. This differs from our approach in a few ways. They do not explicitly model consumers' choices or firms' pricing decisions, but focus instead on firms' investment decisions. While they allow for learning across users, they do not consider either within-period learning or within-user learning, and so consumer coordination issues do not arise. Finally, they explore the interesting issue of how learning from data allows the dominant firm to enter connected markets when the existing data the firm has collected can be reused, an issue we do not consider. Farboodi et al. (2019) provide a model of data-enabled learning in which data helps firms choose the best production technique and so higher quality products. They abstract from any strategic firm behavior or consumer decisions, and instead focus on the industry dynamics generated from a continuum of price-taking firms. De Cornière and Taylor (2020) also study the effect of data on competition but focus on showing how different uses of data (including improving a firm's product, but also ad targeting and price discrimination) determine whether data is pro- or anti-competitive. While their framework is much more general in other ways, they don't analyze a fully dynamic setting in which data learning can accumulate. Schaefer et al. (2018) investigate how data drives the quality of internet search results, providing evidence that the quality of search results improve with more data on previous searches, and that personalized information is particularly valuable. Finally, Biglaiser et al. (2019) provides a high-level discussion of the different ways firms can enjoy incumbency advantages, including access to more data: they note that across-user learning and within-user learning are two distinct ways in which data generates a competitive advantage for incumbents.

3 Across-user learning

In this section we focus on across-user learning—each firm improves the product for each consumer based on what it learns from the usage of all its consumers. The learning technology we consider has the general property that the value to consumers of a given firm's product in a period is an increasing function of the measure of consumers who have purchased the product previously (past consumers) from the same firm. We will refer to the model described below as our baseline setting because it turns out that much of the model setup and analysis, as well as many of the results in this section, carry over (with suitable adjustments) to the versions where we allow for within-period learning (Section 3.4), and within-user learning, either by itself (Section 4), or in combination with across-user learning (Section 5).

Two firms, which we refer to as the incumbent I and the entrant E, compete over infinitely many periods. Both firms face a marginal cost of production equal to c and compete in prices. We assume consumers have unit demands in each period, which they can fulfill either from I or E. Although firms charge positive prices in all of our motivating examples, we can allow for the possibility of negative prices by reinterpreting the price as some valuable service that the firm offers consumers in exchange for their data and which is costly for the firm to provide. The (common) one period ahead discount factor is denoted δ, and satisfies $0 < \delta < 1$. We normalize the measure of consumers in each period to one, and so the number of periods that a firm has previously won also represents the total measure of consumers it can learn from. In this formulation of across-user learning, there is no distinction between consumers who are purchasing again (repeat consumers) vs. entirely new consumers. Thus, our setting in this section is compatible with two interpretations:

- Consumers live for one period only. New consumers in a period make a single buying decision and then exit.
- The same consumers are active and make purchasing decisions in every period, i.e. they are infinitely lived, but there is no cost for each of these individual consumers to switch from one firm to the other across periods. In this case, the learning from consumers when they purchase again is assumed to be the same as that from new consumers.

We assume all consumers are atomistic, which is reasonable given the large number of consumers firms are learning from in most of our examples. This means even if consumers are of the second type above, they do not need to consider the impact of their purchase decision on the pricing game or the options they will face in subsequent periods. Later, when we introduce within-user learning (which creates endogenous switching costs), this will no longer be true.

The standalone value of firm i's product (or service) is denoted s_i, where $i \in \{I, E\}$. This value is augmented by learning. In any period, firms are asymmetric Bertrand competitors, so each firm will either sell to all consumers (measure one) or none. As a result, the measure of consumers a firm has learnt from is equal to the number of periods in which it has made positive sales. If firm i has sold in N_i periods in the past, then the value current consumers obtain from firm i's product is $s_i + f_i\left(\min\left\{N_i, \overline{N}_i\right\}\right)$, where f_i is firm i's learning function and $\overline{N}_i$ is the number of periods firm i must sell in to reach its maximum threshold of learning.

This specification captures that firms learn more as they obtain more consumer data, and this allows them to increase the value of the product they offer, but only up to a point, i.e. learning is bounded. Later in this section we will show how our results extend to handle unbounded learning functions. The only requirements on f_i are that it is defined for all non-negative real values of its argument, is everywhere weakly increasing, and without loss of generality, strictly increasing in the last step to reach its maximum threshold (i.e. $f_i(\overline{N}_i) > f_i(\overline{N}_i - 1)$). In particular, f_i need not be a continuous function. We normalize $f_i(0) = 0$. Thus, at any point in time there are potentially three sources of asymmetry between firms: (i) differences in standalone value ($s_I \neq s_E$), (ii) differences in the learning curve ($f_I \neq f_E$ and/or $\overline{N}_I \neq \overline{N}_E$), and (iii) differences in how far along each firm is on its learning curve in the current period ($N_I \neq N_E$).

Consumers and firms have perfect and complete information. Our focus is on Markovian strategies and a Markov-perfect equilibrium (MPE). Among MPEs, we also rule out equilibria which are supported by the losing firm pricing in such a way that its value would be negative if it instead won consumers at the prices charged in a particular period and then firms followed their equilibrium pricing strategies in all subsequent periods. This is equivalent to the standard approach of focusing on a Nash equilibrium in which the losing firm does not price below its marginal cost in the one-shot asymmetric Bertrand game. Thus, when we refer to a unique MPE we mean the MPE is unique after eliminating MPEs that involve such weakly dominated strategies. Finally, we use the convention that E wins in case consumers are indifferent when both firms price at the lowest level they would still be willing to sell at.

An implicit assumption in our setup is that a firm will continue to compete even if in equilibrium its value is zero. This is consistent with our assumption that firms do not face any fixed cost to remain in operation and ensures firms always face some form of competition. In reality, obviously firms do face ongoing fixed costs of operating. There are many realistic ways to extend our model such that the losing firm makes positive expected profits in every period and therefore remains active provided fixed operating costs are not too large. In Online Appendix A, we present the details for two such extensions. The most straightforward among them allows for the possibility that in every period, there is some chance there will be a publicly available discovery (e.g. some new public data and/or new type of algorithm) that transforms the firms' existing business by an uncertain augmentation of their standalone values, and makes their past data-enabled learning irrelevant. We show that the resulting conditions for either firm to win are identical to the ones we obtain below after adjusting the discount factor to account for the probability of the public discovery, and all the results obtained in this section go through. Finally, in Online Appendix A we also discuss what happens in our setup when the losing firm exits immediately (given that it makes zero profits) and show that the condition for either firm to win is still preserved.

The relevant state at the start of any period is defined by (N_I, N_E). The corresponding value functions for I and E (i.e. the present discounted values of future profit flows starting with the current period when firms follow their equilibrium pricing strategies in every period) are denoted by $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$. Consider firms starting from the state $(\overline{N}_I, \overline{N}_E)$, so both firms have reached their respective learning thresholds. Then I offers $s_I + f_I(\overline{N}_I)$ to consumers in every period and E offers $s_E + f_E(\overline{N}_E)$ to consumers in every period. Given the resulting asymmetric Bertrand competition, and taking into account that the competitive situation will be the same in every subsequent period, E wins every period starting with the current one if and only if $s_E + f_E(\overline{N}_E) \geq s_I + f_I(\overline{N}_I)$. Furthermore, for each firm $i \in \{I, E\}$ and its rival $j \in \{I, E\}$ with $j \neq i$, we must have

$$V^i(\overline{N}_I, \overline{N}_E) = \frac{1}{1-\delta} \max \left\{ s_i - s_j + f_i(\overline{N}_i) - f_j(\overline{N}_j), 0 \right\}.$$

When at least one of the firms has not yet reached its learning threshold, in order to win the current period, each firm will be willing to offer a subsidy equal to the difference in the discounted value of its future profits if it were to win the current period and its future profits if it were to lose the current period. Taking into account the maximum subsidy each firm is willing to offer in order to win, we can work out which firm will win in each period and its corresponding value function, working backwards from $(N_I, N_E) = (\overline{N}_I, \overline{N}_E)$. Using this approach, we show that for any state (N_I, N_E), there is a cutoff level of $s_E - s_I$ above which E wins and below which I wins. We use two-dimensional backwards induction in (N_I, N_E) to establish the following Proposition (unless stated otherwise, all proofs are in the Appendix).

Proposition 1. *Suppose I has previously sold to $0 \leq N_I \leq \overline{N}_I$ consumers and E has previously sold to $0 \leq N_E \leq \overline{N}_E$ consumers. There exists a unique MPE in which E wins in all periods if and only if $s_E - s_I \geq \Delta(N_I, N_E)$, and I wins in all periods otherwise, where*

$$\begin{aligned}\Delta(N_I, N_E) &= (1-\delta)\left(\sum_{j=0}^{\overline{N}_I - N_I} \delta^j f_I(N_I + j) - \sum_{j=0}^{\overline{N}_E - N_E} \delta^j f_E(N_E + j)\right) \\ &\quad + \delta^{\overline{N}_I - N_I + 1} f_I(\overline{N}_I) - \delta^{\overline{N}_E - N_E + 1} f_E(\overline{N}_E).\end{aligned} \tag{1}$$

Proposition 1 characterizes I's competitive advantage, which is captured by $\Delta(N_I, N_E)$. The higher is $\Delta(N_I, N_E)$, the higher is the level of $s_E - s_I$ required for E to win. This function is proportional to the difference in the present discounted value (PDV) of gross surplus generated from learning across the two firms, comparing the paths where each firm wins in every period from the current period onwards. It is easily verified that $\Delta(N_I, N_E)$ is increasing in N_I and decreasing in N_E. This captures that along the equilibrium path, if a firm wins in the current period, because it benefits from an additional learning period, $\Delta(N_I, N_E)$ shifts such that the same firm will continue to win in each subsequent period, and indeed does so even more easily. This last feature is reminiscent of the "increasing increasing dominance" property explored by Cabral and Riordan (1994). Yet in our setting, E may win and so move ahead of I even though it is currently behind (e.g. $N_E < N_I$). Trivially, this could be because it learns more despite having less data (so $f_E(N_E) > f_I(N_I)$). More interestingly, even if $f_E(N_E) < f_I(N_I)$, this can be true if its learning curve is steeper and/or is increasing for longer, and so ultimately generates more value from getting ahead. For example, inspecting (1), it is easily seen that if $f_E(\overline{N}_E) > f_I(\overline{N}_I)$, then for all δ sufficiently close to 1 (i.e. whenever agents are sufficiently patient), we have $\Delta(N_I, N_E) < 0$, so E has an overall competitive advantage. Thus, with asymmetric learning curves, increasing dominance does not necessarily hold in our setting.

Inspecting (1), it is immediate that I's competitive advantage is increasing in the maximum threshold for I's learning $\overline{N}_I$, decreasing in the maximum threshold for E's learning $\overline{N}_E$, increasing in anything that increases I's learning function f_I, and decreasing in anything that increases E's learning function f_E.

The result in Proposition 1 does not depend on consumers being forward looking. Even when consumers are long-lived, given they face no cost of switching firms in any future period and they are atomistic, each consumer can ignore the implications of their individual choice on the future options they will face. Thus, consumers' optimal choices are made by comparing the surplus offered by each firm in the current period only, which is why the same results would arise if we had assumed consumers were myopic instead.

Proposition 1 was worked out assuming maximum learning thresholds $\overline{N}_I$ and $\overline{N}_E$ exist. The next Proposition extends these results to the case of any increasing learning curves by taking the limit of our results as $\overline{N}_I \to \infty$ and $\overline{N}_E \to \infty$, showing that our framework can also allow for unbounded learning functions.

Proposition 2. *Suppose I has previously sold to $0 \leq N_I < \overline{N}_I$ consumers, E has previously sold to $0 \leq N_E < \overline{N}_E$ consumers. The equilibrium defined in Proposition 1 remains an equilibrium in the limit as $\overline{N}_I \to \infty$ and $\overline{N}_E \to \infty$ (i.e. such that the learning functions continue to increase forever) provided there exists some power function which the learning curves lie below in the limit.*

The condition in Proposition 2 involves a very weak requirement that allows us to capture any reasonable learning function. Moreover, by the same logic underlying this result, all our subsequent characterizations of $\Delta(N_I, N_E)$, $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$ also extend to this unbounded case, by taking the limit as $\overline{N}_I \to \infty$ and $\overline{N}_E \to \infty$.

As part of the proof of Proposition 1 we characterize the value function for each firm, which we summarize here.

Corollary 1. *Suppose I has previously sold to $0 \leq N_I \leq \overline{N}_I$ consumers and E has previously sold to $0 \leq N_E \leq \overline{N}_E$ consumers. In the unique MPE, the firms' value functions are determined as follows:*

- *if $s_E - s_I < \Delta(N_I, N_E + 1)$ then $V^E(N_I, N_E) = 0$ and $V^I(N_I, N_E) = \frac{s_I - s_E}{1-\delta} + \frac{\Delta(N_I,N_E) - \delta\Delta(N_I,N_E+1)}{(1-\delta)^2}$ (no subsidization by E)*
- *if $\Delta(N_I, N_E + 1) \leq s_E - s_I < \Delta(N_I, N_E)$, then $V^E(N_I, N_E) = 0$ and $V^I(N_I, N_E) = \frac{s_I - s_E + \Delta(N_I,N_E)}{(1-\delta)^2}$ (subsidization by E)*
- *if $\Delta(N_I, N_E) \leq s_E - s_I < \Delta(N_I + 1, N_E)$, then $V^E(N_I, N_E) = \frac{s_E - s_I - \Delta(N_I,N_E)}{(1-\delta)^2}$ and $V^I(N_I, N_E) = 0$ (subsidization by I)*
- *if $s_E - s_I \geq \Delta(N_I + 1, N_E)$, then $V^E(N_I, N_E) = \frac{s_E - s_I}{1-\delta} - \frac{\Delta(N_I,N_E) - \delta\Delta(N_I+1,N_E)}{(1-\delta)^2}$ and $V^I(N_I, N_E) = 0$ (no subsidization by I),*

where we define $\Delta(\overline{N}_I + 1, N_E) \equiv \Delta(\overline{N}_I, N_E)$ for any $0 \leq N_E \leq \overline{N}_E$ and $\Delta(N_I, \overline{N}_E + 1) \equiv \Delta(N_I, \overline{N}_E)$ for any $0 \leq N_I \leq \overline{N}_I$.

Corollary 1 characterizes the PDV of the firms' equilibrium profits for different states (N_I, N_E) and different ranges of $s_E - s_I$. An implication of the Corollary is that the losing firm is still willing to offer a subsidy to consumers for some range of parameters when the winning firm's competitive advantage is not too large. Thus, if $\Delta(N_I, N_E + 1) \leq s_E - s_I < \min\left\{\Delta\left(\overline{N}_I, N_E + 1\right), \Delta\left(N_I, N_E\right)\right\}$, then I wins and E subsidizes for a finite number of periods,[3] whereas if $\Delta(\overline{N}_I, N_E + 1) \leq s_E - s_I < \Delta(N_I, N_E)$, then I wins and E subsidizes forever, i.e. even after I reaches its learning threshold. And symmetrically when E wins and I subsidizes. This reflects that in the off-equilibrium situation that the losing firm actually wins the current period, it expects to have a positive value (via some future profits). This implies the winning firm's competitive advantage will be limited by the losing firm's willingness to subsidize consumers in such cases. Moreover, Corollary 1 does not rule out the winning firm also subsidizing consumers. Specifically, if $\Delta(N_I, N_E) \leq s_E - s_I < f_I(N_I) - f_E(N_E)$, then E wins and subsidizes in the current period (and symmetrically for I). The willingness of a firm to offer a subsidy reflects the "prize" that the firm can obtain from winning, which is not only that it obtains an improvement in its product (or moves closer to obtaining this improvement) but that the other firm does not. This underlying logic is similar to that in standard learning-by-doing settings (e.g. Cabral and Riordan, 1994).

As mentioned earlier, in Online Appendix A we consider an extension in which the losing firm makes positive expected profits in every period, because there is a probability λ each period that a publicly available discovery makes their existing data and data-enabled learning redundant. More generally, one can consider any extension where with probability λ each period, the firms' values are set at some levels that do not depend on the current state. With any such extension, the cutoff $\Delta(N_I, N_E)$ for E to win in all periods in which the shock has not yet happened is identical, except δ is replaced by $\delta(1-\lambda)$ to take into account the additional discounting of payoffs by the probability that the current data still remains relevant. After this adjustment to $\Delta(N_I, N_E)$, the firms' value functions (prior to the shock) can be adjusted simply by adding a constant term for each firm everywhere, implying all the results we derive in this section continue to apply, except expressions become more complicated.

3.1 Determinants of competitive advantage

We are interested in the effects on a firm's competitive advantage of changes in technology that increase the amount the firms can learn from data (e.g. improvements in algorithms), the rate at which they can learn from customer data (e.g. restrictions in their ability to use

[3] Specifically, E subsidizes for the first k periods, where $k \leq \overline{N}_I - N_I$ is uniquely defined by the inequality $\Delta(N_I + k, N_E + 1) < s_E - s_I < \Delta(N_I + k + 1, N_E + 1)$.

new consumer data), or shocks that provide firms with more data (e.g. getting access to a new data source that moves them along their learning curves).

To explore these types of questions, in this section we put some more structure on the learning functions. Formally, we assume when firm i has sold to N_i past consumers, the value it offers from learning is $f_i\left(\min\left\{N_i, \overline{N}_i\right\}\right) = \theta_i \max\left\{\min\left\{N_i, \overline{N}_i\right\} - \underline{N}_i, 0\right\}$, where $\underline{N}_i$ and $\overline{N}_i$ are the minimum and maximum thresholds for firm i and $0 \leq \underline{N}_i \leq \overline{N}_i$. This formulation means that the learning curves are S-shaped: the value initially doesn't increases until some minimum amount of data is collected, after which it increases at a constant rate until learning has been exhausted. It includes as special cases linear learning functions up to the threshold ($\underline{N}_i = 0$), and step learning functions ($\overline{N}_i - \underline{N}_i = 1$ and $\underline{N}_i > 0$).

Plugging these functions into (1) leads to the following expression of the cutoff for E to win:

$$\begin{aligned}\Delta(N_I, N_E) \quad &= \quad \theta_I \left(\max\left\{N_I - \underline{N}_I, 0\right\} + \frac{\delta^{\max\{\underline{N}_I - N_I, 0\}+1} - \delta^{\overline{N}_I - N_I + 1}}{1-\delta} \right) \\ &\quad\quad -\theta_E \left(\max\left\{N_E - \underline{N}_E, 0\right\} + \frac{\delta^{\max\{\underline{N}_E - N_E, 0\}+1} - \delta^{\overline{N}_E - N_E + 1}}{1-\delta} \right),\end{aligned} \tag{2}$$

for all $0 \leq N_I \leq \overline{N}_I$ and $0 \leq N_E \leq \overline{N}_E$. Using this cutoff we are able to derive the following comparative statics.

Proposition 3. *Suppose both firms have S-shaped learning functions. (i) I's competitive advantage is increasing in an equal additive increase in the rates of learning θ_I and θ_E, provided the increasing part of I's learning curve lasts (weakly) longer and provided I is (weakly) closer to reaching its maximum threshold (i.e. $\overline{N}_I - \underline{N}_I \geq \overline{N}_E - \underline{N}_E$ and $\overline{N}_I - N_I \leq \overline{N}_E - N_E$, with one inequality strict). (ii) I's competitive advantage is increasing in an equal proportional decrease in the firms' rates of learning θ_I and θ_E, provided I is (weakly) closer to reaching its threshold and E's learning algorithm is at least as good (i.e. $\overline{N}_I - N_I \leq \overline{N}_E - N_E$ and $\theta_E \geq \theta_I$, with one inequality strict) and provided both firms have already obtained some positive value from learning (i.e. $N_I \geq \underline{N}_I$ and $N_E \geq \underline{N}_E$). (iii) If the learning functions are identical (i.e. $\theta_I = \theta_E$, $\underline{N}_I = \underline{N}_E = \underline{N}$, $\overline{N}_I = \overline{N}_E = \overline{N}$) and $N_I > N_E$, then I's competitive advantage is increasing in how close both firms are to obtaining some positive value from learning when neither firm has learnt anything yet (i.e. is larger when both N_I and N_E increase by $k \geq 1$, provided $N_I + k \leq \underline{N}$ and $N_E + k \leq \underline{N}$) and decreasing in how close both firms are to obtaining maximum learning when both have already obtained some positive value from learning (i.e. is smaller when both N_I and N_E increase by $k \geq 1$, provided $N_I > \underline{N}$ and $N_E > \underline{N}$).*

The case in (i) captures what happens if firms gain access to a better algorithm that shifts up what they can learn from their customer data. The logic is that if I has more scope for learning and is closer to reaching its learning threshold, then an equal increase in the customer value that can be provided from learning provides an advantage for I.

The case in (ii) captures firms only being able to use a fraction of any new customer data collected (e.g. due to some new privacy regulation). We model this as the linear portion of firm i's learning curve pivoting to the right from its current position, which represents a slowing down of the learning rate from customer data. This implies, starting from firm i's current position, the value provided by customer data increases at the rate $\rho\theta_i$ instead of θ_i, where $0 < \rho < 1$, but is such that firm i still eventually reaches the same maximum value from learning of $f_i\left(\overline{N}_i\right)$.[4] The result shows that, when I is closer to its threshold and E's learning function is at least as steep as I's, slowing down the rate of learning has a larger effect on E, reflecting that E has more to gain from new customer data, and so loses more when a fraction of that data becomes unusable. The result no longer holds generally when either E or both I and E are yet to move onto the increasing parts of their learning curves. However, one can show it still holds in that case provided firms do not discount the future too much.

The first part of (iii) can be interpreted as considering what happens when both firms discover some new data, so that both firms find themselves closer to where they can start to offer positive value from learning. This discovery increases I's competitive advantage. To understand this result note that if firms are yet to move onto the increasing part of the learning curve, requiring fewer periods of customer data to reach the increasing part brings forward (i.e. discounts less) the advantage that I obtains from being closer to the threshold, thereby increasing its competitive advantage. The effect reverses when both firms are already on the increasing part of the learning curve. Here, a reduction in the number of learning periods both firms need to reach the maximum threshold decreases the number of periods over which I can enjoy its learning advantage, so helps E.

3.2 Welfare analysis

Since the cutoff (1) in Proposition 1 is directly proportional to the difference in the PDV of gross surplus generated from learning if I wins in every period vs. if E wins in every period, we have the following result.

Proposition 4. *The unique MPE outcome characterized in Proposition 1 is socially optimal.*

[4] To avoid integer complications, assume there exist integers $m_I \geq 1$ and $m_E \geq 1$ such that $\rho = \frac{\overline{N}_I - N_I}{\overline{N}_I - N_I + m_I} = \frac{\overline{N}_E - N_E}{\overline{N}_E - N_E + m_E}$. Thus, m_I and m_E are the number of additional periods of learning required for each firm to reach its maximum learning threshold.

Starting from any point, whichever firm creates greater PDV of gross surplus from standalone value and learning assuming it wins every period should be able to offer a higher level of current utility plus subsidy to consumers, and so, given the Bertrand setting, should be able to win in the current period. However, the logic behind the efficiency result in Proposition 4 is more subtle than this. Indeed, the winning firm does not always capture the full PDV of gross surplus that it creates over and above that created by the losing firm because it has to compete with a firm that will sometimes be willing to subsidize in an attempt to win. Nevertheless, the PDV of the subsidy offered by the losing firm is just the PDV of its own future profit assuming it is able to win in the current period. Thus, despite the possibility of subsidies, it turns out that the only way a firm wins (and has a positive value) is if the PDV of gross surplus that it creates by winning every period is greater than that created by the rival firm. This is equivalent to the condition for efficiency to be maximized.

Further emphasizing that the result does not just follow automatically from Bertrand competition, the cutoff for E to win in a finite period version of the model is not generally socially optimal. Indeed, in Online Appendix B we show that the equilibrium cutoff when the firms start with (N_I, N_E) previous sales and the game ends after $T \geq 1$ periods is

$$\Delta(N_I, N_E, T) = (1-\delta)\left(\sum_{j=0}^{T-1} \frac{\delta^j(1-\delta^{T-j})}{1-(T+1)\delta^T + T\delta^{T+1}}(f_I(N_I+j) - f_E(N_E+j))\right),$$

while the socially efficient cutoff in the T-period model is

$$\Delta^S(N_I, N_E, T) = (1-\delta)\left(\sum_{j=0}^{T-1} \frac{\delta^j}{1-\delta^T}(f_I(N_I+j) - f_E(N_E+j))\right).$$

When $T > 1$, we generally have $\Delta(N_I, N_E, T) \neq \Delta^S(N_I, N_E, T)$, however it is easily verified that when $f_i(N_i) = f_i\left(\min\left\{N_i, \overline{N}_i\right\}\right)$ for $i = I, E$ and T tends to infinity, both $\Delta(N_I, N_E, T)$ and $\Delta^S(N_I, N_E, T)$ converge to $\Delta(N_I, N_E)$ defined in (1).

In the finite case with $T > 1$, the distortion $\Delta(N_I, N_E, T) - \Delta^S(N_I, N_E, T)$ can go in either direction depending on the shapes of the learning curves. For example, it is straightforward to verify that with $f_I = f_E \equiv f$ and $N_I > N_E$, this distortion is positive if f is concave, negative if f is convex, and zero if f is linear. On the other hand, if $f_i(x) = \theta_i x$ and $N_I \geq N_E$, the distortion is positive if $\theta_I < \theta_E$, negative if $\theta_I > \theta_E$, and zero if $\theta_I = \theta_E$.

Clearly prices play a key role in the efficiency result of Proposition 4. If prices were constrained to be zero in every period (and there was no other way for firms to invest to provide value to attract users initially, and to monetize them later), there would be no subsidies, meaning consumers would compare $s_I + f_I(N_I)$ and $s_E + f_E(N_E)$ in the current period, and choose E in every period if and only if $s_E - s_I \geq f_I(N_I) - f_E(N_E)$. Future learning benefits would not be taken into account and the equilibrium choice of firm would diverge from the socially optimal one. Competitive subsidies play a key role in ensuring atomistic consumers who make their purchase decisions in every period solely on the basis of current period utility, end up taking into account the future streams of utility created by their current choice of which firm to buy from.

With respect to consumer surplus, note that in equilibrium, the firms' prices in any given period leave consumers indifferent between the two firms. Since along the equilibrium path, once a firm wins, it will win in all future periods, this implies that the PDV of consumer surplus is determined by the surplus offered by the losing firm in each period along the equilibrium path. As a result, learning by the winning firm as it sells to additional consumers each period makes consumers weakly worse off. To see this note first that if the losing firm has already stopped offering a subsidy (e.g. the no-subsidization case in Corollary 1), then learning by the winner does not change the PDV of consumer surplus offered by the loser, which in case firm i is the losing firm is fixed at

$$CS\left(N_I, N_E\right) = \frac{s_i - c + f_i(N_i)}{1-\delta}. \tag{3}$$

In case the losing firm currently offers a subsidy (either one of the two subsidization cases in Corollary 1), the PDV of consumer surplus can be determined by the PDV of total surplus less the PDV of the winning firm's profit stream, i.e.

$$CS\left(N_I, N_E\right) = \frac{s_i - c}{1-\delta} + \sum_{k=0}^{\overline{N}_i - N_i} \delta^k f_i(N_i + k) + \frac{\delta^{\overline{N}_i - N_i + 1}}{1-\delta} f_i(\overline{N}_i) - \left(\frac{s_i - s_j + \Delta(N_I, N_E)}{(1-\delta)^2}\right) \tag{4}$$

in case firm i is the winning firm and firm j is the losing firm. It is straightforward to check that in this case $CS\left(N_I, N_E\right)$ is decreasing in N_i, so consumer surplus is strictly decreasing as the winning firm learns more along the equilibrium path. This reflects that the losing firm's subsidy gets smaller over time as the winning firm's learning advantage continues to increase along the equilibrium path.

The fact that additional learning by the winning firm as it sells to more consumers does not help consumers reflects the fact that in this model, the winning firm extracts all the additional learning benefit above the surplus that the losing firm can offer. While with a more general model of competition or in a setting with elastic aggregate demand we may expect some of the winning firm's learning benefit to be shared with consumers, the result that consumers can be made worse off by the winning firm's learning because the losing firm is no longer willing to subsidize as much in an attempt to compete is likely to be much more general. It implies consumers can be better off if they could somehow deviate by coordinating on joining the losing firm instead of the winning firm for a number of periods. However, given each individual consumer is atomless and can freely switch in each period, she has no reason to internalize this effect since her own decision about which firm to join will not affect the consumer surplus she can obtain next period (in the case she is long lived). This suggests a possible role for a data sharing policy, which we consider in the next section.

3.3 Data policy

In this subsection we use our framework to examine the effect on competition and consumer welfare of three possible types of data policies: (i) requiring data sharing, (ii) enforcing stricter consumer data privacy rules, and (iii) preventing an incumbent firm from acquiring a non-competing firm that contains valuable data.

3.3.1 Data sharing

Requiring the incumbent to share some of its data with the entrant would be one way to help the entrant catch up to the learning of the incumbent, and based on the analysis in Section 3.2, this would seem to be good for consumers. However, if the incumbent's learning is always shared with the entrant, then it may dampen both firms' incentives to invest in building up their data in the first place. In our framework, such investments are endogenously determined by the extent of subsidies offered to consumers. Specifically, once the firm that is behind is able to share in the other firm's learning through a data sharing policy, it will compete less aggressively (a form of free-riding), which can result in higher prices and therefore lower consumer surplus.

To show this tradeoff formally, we consider the introduction of a policy which requires any *new* data collected by firms be shared with each other. Such a policy reflects that past data may not be easily shareable, or even if it was, requiring incumbents to share it would raise issues of ex-post expropriation. Thus, in every period subsequent to the implementation of data sharing, both firms advance their learning by one period (if they are not already at the threshold) regardless of who wins. This means that from the period the policy is implemented onwards, which we treat as being in the first period, the lowest price each firm is willing to charge is c. In every period $j \geq 0$, firm I wins if $s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\}) > s_E + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right)$, and firm E wins otherwise. Consequently, the PDV of consumer surplus after the implementation of the data sharing policy is

$$CS\left(N_I, N_E\right)' = \sum_{j=0}^{\infty} \delta^j \min\left\{s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\}), s_E + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right)\right\} - \frac{c}{1-\delta}. \tag{5}$$

Meanwhile, in the absence of the data sharing policy, when $s_E - s_I < \Delta(N_I, N_E + 1)$ or $s_E - s_I \geq \Delta(N_I + 1, N_E)$, the losing firm does not subsidize so the PDV of consumer surplus is just the PDV of the surplus offered by the losing firm, as in (3). When $\Delta(N_I, N_E + 1) \leq s_E - s_I < \Delta(N_I + 1, N_E)$, the losing firm subsidizes, and the PDV of consumer surplus absent data sharing is given in (4).

Comparing the expressions for consumer surplus, we obtain the following proposition (the proof is in the appendix).

Proposition 5. *Suppose data sharing requires whichever firm wins a given period to share the data obtained with its rival. Suppose one firm is closer to its learning threshold than the other, i.e.* $\overline{N}_I - N_I < \overline{N}_E - N_E$ *or vice versa. Then there exist cutoff levels* $\Theta_l\left(N_I, N_E\right) < \Delta(N_I, N_E)$ *and* $\Theta_h\left(N_I, N_E\right) > \Delta(N_I, N_E)$ *such that:*

- *Data sharing increases the PDV of consumer surplus when* $s_E - s_I \leq \Theta_l\left(N_I, N_E\right)$ *or* $s_E - s_I \geq \Theta_h\left(N_I, N_E\right)$
- *Data sharing decreases the PDV of consumer surplus when* $\Theta_l\left(N_I, N_E\right) < s_E - s_I < \Theta_h\left(N_I, N_E\right)$.

Thus, data sharing increases consumer surplus when the losing firm is at a sufficiently large disadvantage and decreases consumer surplus when the two firms are sufficiently evenly matched. This is intuitive. When the losing firm is at a large disadvantage, it does not subsidize, so there is no downside of data sharing. However, when the two firms are more evenly matched, the losing firm would normally subsidize heavily and data sharing removes this incentive to subsidize, so this is where the downside of data sharing is highest and overwhelms the positive effect of keeping the firms evenly matched in all future periods.

While consumers may be better or worse off under data sharing, it is easily seen that the PDV of total welfare is higher with this type of data sharing because in every period $j \geq 0$, the total surplus created is $\max\left\{s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\}), s_E + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right)\right\} - c$. In other words, both firms are learning in every period so society gets the best of both in every period.

3.3.2 Privacy of consumer data

A stricter privacy policy makes it harder for firms to gather customer data either by restricting the type of data they can make use of or by making it easier for consumers to opt out of sharing their data. We model this as firms only being able to use a fraction of any new customer data collected. Case (ii) of Proposition 3 captured exactly this case by studying what happens along the linear part of a firm's S-shaped learning curve when the value provided by customer data increases at the rate $\rho\theta_i$ instead of θ_i, where $0 < \rho < 1$. The result in Proposition 3 means that, assuming I is initially closer to its threshold and/or E has a steeper learning function, a stricter privacy policy has a disproportionately negative effect on E, reflecting that E has more to gain from new customer data, and therefore more to lose from reductions in the amount of data that is usable. The logic of this effect is easiest to see in the extreme case when I is at (or almost at) its threshold, and E is far away from reaching its threshold. A privacy policy that slows down the amount of useful data that both firms can extract from their customers will have no (or almost no) effect on I, since it has nothing (or almost nothing) left to learn, but it can have a big effect on E which has a lot left to learn.

The following proposition also considers the effect of such a privacy policy on consumer surplus.

Proposition 6. *Suppose both firms have S-shaped learning functions. Suppose I is (weakly) closer to reaching its threshold and E's learning algorithm is at least as good (i.e. $\overline{N}_I - N_I \leq \overline{N}_E - N_E$ and $\theta_E \geq \theta_I$, with one inequality strict), both firms have already obtained some positive value from learning (i.e. $N_I \geq \underline{N}_I$ and $N_E \geq \underline{N}_E$), and I would win absent the policy change (i.e. $s_E - s_I < \Delta(N_I, N_E)$). Then a privacy policy that slows down the rate at which firms learn from customers' data increases I's competitive advantage and weakly decreases consumer surplus, strictly decreasing it if E subsidizes consumers before the policy change.*

The decrease in consumer surplus in Proposition 6 arises because E lowers the subsidy it offers, reflecting that it has less to gain by winning after the (stricter) privacy policy is introduced. As a result, I (which continues to win given its competitive advantage increases) doesn't have to leave as much surplus with consumers to prevent them from switching to E. While we haven't modeled any direct consumer benefits due to tighter privacy rules, any such benefits should not affect the result given consumers can internalize them when deciding which firm to join. We are, of course, abstracting from externalities (positive or negative) that the use of a consumer's private data may impose on other consumers or firms. Our point is only that the competitive effects of privacy policies highlighted here are usually overlooked in existing policy discussions.

3.3.3 Data acquisition

Our framework can also be used to explore the interesting question of what happens when firms can acquire an outside firm (or its dataset) in order to bolster their data position. Starting from (N_I, N_E), assume there is an opportunity for E and I to acquire a measure $N_A > 0$ of data exclusively. We wish to know which firm will end up acquiring the data, which is equivalent to figuring out which firm is willing to pay more for it (assuming the owner of the data can commit to sell it exclusively), and whether the outcome will be efficient.

If I acquires N_A, then the cutoff for E to win increases to $\Delta(N_I + N_A, N_E)$, whereas if E acquires N_A, then the cutoff for E to win decreases to $\Delta(N_I, N_E + N_A)$, with the understanding that $\Delta(N_I + N_A, N_E) = \Delta\left(\overline{N}_I, N_E\right)$ whenever $N_I + N_A \geq \overline{N}_I$ and $\Delta(N_I, N_E + N_A) = \Delta\left(N_I, \overline{N}_E\right)$ whenever $N_E + N_A \geq \overline{N}_E$. We will use this slight abuse of notation to state our result in the following proposition in order to keep the expressions compact.

Proposition 7. *Suppose there is an opportunity for E or I to acquire a dataset of measure $N_A > 0$. Assume at least one of the two firms (I or E) has not reached its maximum learning threshold. Then there exists a unique Δ^* satisfying*

$$\Delta(N_I, N_E + N_A) < \Delta^* < \Delta(N_I + N_A, N_E) \tag{6}$$

such that I wins the data if and only if $s_E - s_I < \Delta^*$, and E wins the data otherwise. If $\Delta^* < \Delta(N_I + N_A, N_E + N_A)$, then the cutoff in $s_E - s_I$ above which E acquires the data is too low from a social efficiency perspective (i.e. E acquires the data too often), whereas if $\Delta^* > \Delta(N_I + N_A, N_E + N_A)$, then in equilibrium I acquires the data too often.

As shown in the proof in the Appendix, both scenarios in Proposition 7 are possible. Thus, the equilibrium outcome of the data acquisition game can be socially inefficient in either direction. Contrast this with our general efficiency result for the baseline model from Proposition 4. The reason for the difference is that here the competition is for buying data from a third party, who only cares about price and not about the respective utilities currently offered by the two firms. By contrast, when the two firms compete for consumers, these consumers care about the current utility and the price they are being charged.

To illustrate this result, we can look at a special case. Suppose one of the firms—I without loss of generality—is already at the threshold, i.e. $N_I = \overline{N}_I$. Then we have

$$\Delta(N_I + N_A, N_E + N_A) = \Delta\left(\overline{N}_I, N_E + N_A\right) = \Delta(N_I, N_E + N_A) < \Delta^*,$$

where the last inequality follows directly from (6). Thus, in this case I wins the data too often. And conversely when E is already at the threshold. In other words, in the equilibrium of the acquisition game when one firm has already reached its threshold of learning, that firm wins too often relative to what is socially efficient. The idea is that the amount the firm that is at the threshold gives up by not winning (namely its profit when competing with a rival that is below its learning threshold) is greater than the amount that the rival can gain by acquiring the data and winning, since in this case it competes with a rival that is already at its threshold.

A final property of interest is that even though I may acquire the new data in equilibrium, it doesn't necessarily benefit from it. The existence of the new data N_A makes it easier for E to catch up, and I needs to buy the data to prevent E catching up, which can end up hurting it overall.[5]

3.4 Within-period learning

So far consumers have not had to form beliefs about which firm other consumers will choose in any given period. Although there is a self-reinforcing dynamic across periods in which a firm that attracts consumers this period improves its product and finds it easier to attract consumers in subsequent periods, current consumers still did not face the usual coordination problem that arises under classic network effects. Indeed, as noted earlier, even with long-lived consumers, because each consumer was atomistic and could costlessly switch each period, a consumer's decision in one period had no effect on the opportunities they could enjoy in subsequent periods, so there was no reason for any consumer to take into account the decisions of other contemporary consumers.

In this section we show that a consumer coordination problem arises when across-user learning is combined with the assumption that the product continues to improve during the period(s) over which consumers derive utility from the product after paying for it. To do so, we modify the baseline model by allowing firms to learn immediately from consumers purchasing in the current period, rather than learning next period only, as we had assumed previously. As argued in the introduction, within-period learning is a key distinguishing feature of data-enabled learning in the context of cloud-based products, and contrasts with learning in the context of traditional products, where the new and improved features were only embedded in the next releases of the product, which would be sold to future consumers. Because of this feature, other things equal, consumers prefer to buy from the firm they expect others to also buy from, given they benefit from the resulting improvement in the firm's product during the consumption period.

[5] We show this result formally in Online Appendix C.

We consider consumers resolving the resulting coordination problem via two different types of beliefs. First, we assume that in every period, consumers believe all other consumers will coordinate on the equilibrium outcome that is best for them in that period, taking into account the prices they face. This selects the equilibrium which maximizes the consumers' surplus (i.e. net utility) within every period. When consumers live for one period only, which recall is one interpretation of our across-user learning setup, this corresponds to the Pareto optimal equilibrium for consumers for given prices, and with this in mind, we refer to these beliefs as "Pareto beliefs". We then contrast Pareto beliefs with beliefs that always favor the incumbent. The latter means that in each period, consumers believe all other consumers will buy from I whenever this is an equilibrium outcome given the prices charged by the two firms.

Relative to Proposition 1, the case of within-period learning and Pareto beliefs is as if both firms were one period closer to their respective maximum learning thresholds, whereas the case of within-period learning and beliefs favoring I is as if only I were one period closer to its maximum learning threshold. Then following an almost identical proof to that of Proposition 1 we obtain the following result.

Proposition 8. *Suppose I has previously sold to* $0 \leq N_I \leq \overline{N}_I$ *consumers and E has previously sold to* $0 \leq N_E \leq \overline{N}_E$ *consumers. With Pareto beliefs and within-period learning, E wins in all periods if and only if* $s_E - s_I \geq \Delta(N_I + 1, N_E + 1)$*, and I wins in all periods otherwise, where* $\Delta(N_I, N_E)$ *is defined by (1). With favorable beliefs for I and within-period learning, E wins in all periods if and only if* $s_E - s_I \geq \Delta(N_I + 1, N_E)$*, and I wins in all periods otherwise. The equilibrium outcome is socially efficient with Pareto beliefs and involves I winning for a larger range of* $s_E - s_I$ *than is socially efficient under favorable beliefs for I.*

The intuition for these results is the same as before, adjusted for the fact that the winning firm now learns from its current period's consumers. Note the socially optimal outcome takes into account within-period learning, just as the equilibrium conditions do. Moreover, Corollary 1 continues to apply, given that the cutoffs that apply here are defined in the same way as (1), but with the terms in N_I and N_E increased by zero or one depending on the type of beliefs. By the same logic, all other results similarly carry over.

Taking the difference in cutoffs when beliefs favor I vs. when beliefs are Pareto, we obtain

$$\Delta(N_I+1, N_E) - \Delta(N_I+1, N_E+1) = (1-\delta)\left(\sum_{j=0}^{\overline{N}_E - N_E - 1} \delta^j (f_E(N_E + j + 1) - f_E(N_E + j))\right), \tag{7}$$

which is strictly positive whenever $N_E < \overline{N}_E$.[6] The expression in (7) is a measure of the distortion that arises in the cutoff for E to win due to beliefs favoring I, relative to the socially optimal outcome. It shows I's position relative to its learning threshold is irrelevant to the distortion implied by I enjoying favorable beliefs. This is because under both types of beliefs, consumers evaluate the PDV of future product improvements that can be offered by I assuming I will win in every period. Meanwhile, E is effectively viewed by consumers as being one period of learning behind in every period under beliefs that favor I compared to under Pareto beliefs.

Taking the difference between (7) evaluated at $N_E - k$ and (7) evaluated at N_E, the following result follows almost immediately.

Proposition 9. *Suppose there is across-user learning and within-period learning. If E's learning function f_E is weakly concave, the distortion in the cutoff resulting from beliefs favoring I is higher the further away E is from its maximum learning threshold (i.e. the lower N_E is).*

This result reflects that shifting E back by one period of learning has a bigger effect when E is on the steeper (i.e. earlier) part of its learning function, and that the corresponding disadvantage for E arises over additional periods when E is further away from its maximum learning threshold. Put more directly, under the conditions of Proposition 9, the concern that beliefs favoring I give I an inefficient advantage is greater the further away E is from its maximum learning threshold.

4 Within-user learning

In this section we focus on within-user learning. We assume the same measure one of consumers are present in every period. They are infinitely lived and can switch firms in any period. The value to a given consumer of firm i's product in the current period is $s_i + f_i\left(\min\left\{N_i, \overline{N}_i\right\}\right)$ for $i \in \{I, E\}$. The main difference from the case with across-user learning is that here N_i refers to the number of times the consumer has purchased from the same firm before, rather than the total number of consumers that the firm has sold to. Another key difference is that in the presence of within-user learning, it is natural to allow firms to price discriminate across consumers that have different histories. Unlike the case with across-user learning, here the individual consumer's history affects the value offered by each firm's product. Consistent with our perfect information setup, we assume that both

[6] When $N_E = \overline{N}_E$, beliefs are irrelevant, and (7) is equal to zero.

firms can observe how many periods a customer has consumed from each firm previously, and can price discriminate accordingly.[7] Since there is no across-user learning and firms can price discriminate with respect to individual consumers, the analysis can be done independently for each individual consumer. Viewed in this light, the model is consistent with I and E being in different competitive positions with respect to each different individual consumer (or consumer segment), and our results pertain to any such consumer (or consumer segment).

The following proposition characterizes when E wins, when I wins, the social efficiency of the outcome, and the value obtained by firms.

Proposition 10. *Suppose I has sold N_I times in the past to a given consumer and E has sold N_E times to that consumer, where $0 \leq N_I \leq \overline{N}_I$ and $0 \leq N_E \leq \overline{N}_E$. Then a unique MPE exists in which E wins in all periods if and only if $s_E - s_I \geq \Delta(N_I, N_E)$, and I wins in all periods otherwise, where $\Delta(N_I, N_E)$ is defined by (1). The outcome is socially optimal. Moreover, the firms' value functions are determined as follows:*

- *if $s_E - s_I < \Delta(N_I, N_E)$, then $V^E(N_I, N_E) = 0$ and $V^I(N_I, N_E) = \frac{s_I - s_E + \Delta(N_I, N_E)}{1-\delta}$ (with subsidization by E if and only if $\Delta(N_I, N_E + 1) \leq s_E - s_I < \Delta(N_I, N_E)$)*
- *if $s_E - s_I \geq \Delta(N_I, N_E)$, then $V^E(N_I, N_E) = \frac{s_E - s_I - \Delta(N_I, N_E)}{1-\delta}$ and $V^I(N_I, N_E) = 0$ (with subsidization by I if and only if $\Delta(N_I, N_E) \leq s_E - s_I < \Delta(N_I + 1, N_E)$).*

Proposition 10 shows that the cutoff in $s_E - s_I$ that determines which firm wins in each period coincides with that obtained in the across-user learning case, and remains socially optimal.[8] This means the comparative static results in Section 3.1 carry over to the case of within-user learning. However, despite this equivalence, the logic behind the result and its proof are quite different. With across-user learning, we never had to solve the consumers' forward-looking problem given their choice in any given period didn't affect their options in subsequent periods. In contrast, here consumers have to take into account the effect of their current choice of firm on their future options (within-user learning acts as an endogenous switching cost), and so we need to determine the value functions for consumers as well as for firms. In the proof of Proposition 10, we construct the consumers' value functions, and show that consumers never expect to switch firms in equilibrium.

[7] Obviously, each firm will know how many times a consumer has purchased from it before, so can condition its prices on this information. Given the model setup in which the same consumers are present in every period, each firm would therefore also be able to work out how many times a consumer has purchased from the rival firm in the past. In practice, firms may elicit this information from consumers by requiring a consumer supply their past purchase records so as to provide them with a better deal if they switch.

[8] It is easily confirmed, the results in Proposition 10 remain valid when $\overline{N}_I \to \infty$ and $\overline{N}_E \to \infty$, provided learning functions are bounded by some power function.

In contrast to the case of across-user learning, the logic for why the equilibrium outcome under within-user learning coincides with the socially optimal outcome is quite straightforward. Here consumers are forward-looking, so they take into account all future subsidies and utility they will receive from their chosen firm. Starting from any point, whichever firm creates greater PDV of gross surplus from learning, assuming it wins every period, will be able to offer a higher level of utility plus subsidy to consumers, and so will win in the current period given the Bertrand setting. And if a particular firm wins in the current period, this just further increases the value it can offer in future periods, thereby ensuring the same firm wins in all future periods. Consistent with this, in Online Appendix D we confirm that the outcome under within-user learning and a finite time horizon is also always socially optimal, in contrast to what we found with across-user learning.

To further highlight the fact that although the cutoff is the same as with across-user learning, the rationale for the result is different, note that prices no longer drive the equilibrium result. Indeed, if prices were set to zero in all periods (which obviously rules out any subsidies), consumers would simply compare the respective PDVs of their willingness-to-pay for the offerings of the two firms, so the condition for E to win would remain the same. Finally, another key difference compared to the previous case with across-user learning is that the competitiveness of the outcomes (e.g. firms' value functions) now depends on consumers being forward-looking.

Proposition 11. *Suppose there is within-user learning. Then provided either $N_I < \overline{N}_I$ or $N_E < \overline{N}_E$ (so we rule out both firms being at their learning threshold already), the PDV of the winning firm's profit is strictly lower with forward-looking consumers than with myopic consumers, even though the cutoff for either firm to win is the same in both cases.*

The main reason behind the result in Proposition 11 is that the winning firm needs to price more aggressively when facing forward-looking consumers, who take into account future utility when making their choice of firm. The logic is that within-user learning acts as an endogenous switching cost: consumers take into account that the surplus they will be left with in any period is the surplus offered by the losing firm (due to Bertrand competition), and that this surplus will be lower if they buy from the winning firm this period versus if they buy from the losing firm this period. Thus, the winning firm must compensate consumers by pricing lower in the current period compared to the case with myopic consumers. By the same logic, Proposition 11 implies that firms compete more intensely and consumers are better off under within-user learning than under across-user learning, if we keep everything else constant across the two cases. This reflects that with across-user learning, consumers had no reason to consider more than their current surplus when deciding which firm to buy from, and that as a result the firms' value functions with across-user learning are the same as when there is within-user learning and consumers are myopic.

We now turn to the data policy issues discussed in Section 3.3. First, the effect of unanticipated data sharing on consumer surplus works in a very similar way to the case with across-user learning. The PDV of consumer surplus after the implementation of the same type of policy we considered in Section 3.3.1 is the same here, equal to the expression in (5):

$$CS\left(N_I, N_E\right)' = \sum_{j=0}^{\infty} \delta^j \min\left\{s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\}), s_E + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right)\right\} - \frac{c}{1-\delta}.$$

Meanwhile, the PDV of consumer surplus without the data sharing policy is

$$CS\left(N_I, N_E\right) = \begin{cases} \frac{s_E - c}{1-\delta} + \sum_{j=0}^{\overline{N}_E - N_E} \delta^j f_E(N_E + j) + \frac{\delta^{\overline{N}_E - N_E + 1}}{1-\delta} f_E(\overline{N}_E) & \text{if} \quad s_E - s_I < \Delta(N_I, N_E) \\ \frac{s_I - c}{1-\delta} + \sum_{j=0}^{\overline{N}_I - N_I} \delta^j f_I(N_I + j) + \frac{\delta^{\overline{N}_I - N_I + 1}}{1-\delta} f_I(\overline{N}_I) & \text{if} \quad s_E - s_I \geq \Delta(N_I, N_E) \end{cases}.$$

And it is easily seen that $CS\left(N_I, N_E\right)' \leq CS\left(N_I, N_E\right)$ for all (N_I, N_E), which means that data sharing always reduces consumer surplus here, in contrast to the case with across-user learning, where data sharing only reduced consumer surplus when $s_E - s_I$ was in a certain interval. The fact that data sharing is more likely to reduce consumer surplus under within-user learning is consistent with the result above that within-user learning makes the two firms compete more aggressively in the absence of data sharing, which diminishes the potential benefits of data sharing in the first place.

Second, the effect of a stricter privacy policy on I's competitive advantage is the same as with across-user learning because the cutoff $\Delta\left(N_I, N_E\right)$ is the same here. The expression of the change in consumer surplus due to the implementation of the stricter privacy policy is different, but it turns out by following a very similar proof, to also be negative. Thus, under the same conditions as Proposition 6 but applied to within-user learning, a stricter privacy policy again lowers consumer surplus.

Third and finally, the reasoning for and effects of data acquisition are the same here relative to the case with across-user learning. In particular, Proposition 7 remains valid, even though the exact expression for the equilibrium cutoff Δ^* that determines which firm wins the acquisition game is different.

5 Across-user and within-user learning

The main purpose of this section is to explore the implications of combining across-user and within-user learning; in particular how doing so can sometimes lead to an endogenous network effect in which consumer beliefs matter. We combine the models of Sections 3 and

4 in the following way. First, we assume there is just one fixed set of consumers of measure one who are present in all periods, consistent with the model of within-user learning used in Section 4. And second, we assume that the value a given consumer obtains from firm i's product in the current period depends on the two types of learning—the number of times that this measure one of consumers has previously purchased from the same firm (across-user learning), and the number of times the same consumer has previously purchased from the same firm (within-user learning).

A consumer who has previously consumed n_i times from firm $i \in \{I, E\}$ when firm i has previously served the measure one of consumers N_i times in the past derives utility $s_i + f_i\left(\min\left\{N_i, \overline{N}_i\right\}, \min\left\{n_i, \overline{n}_i\right\}\right)$ in the current period if she chooses firm i's product, where $\overline{N}_i \geq 1$ is the threshold level of across-user learning and $\overline{n}_i \geq 1$ is the threshold level of within-user learning for firm i. We assume the learning function $f_i\left(.,.\right)$ is everywhere weakly increasing in both arguments for $i \in \{I, E\}$. Without loss of generality, we assume that $f_i\left(0, 0\right) = 0$. We consider first the case in which consumers are myopic, and then tackle the more difficult case in which consumers are forward looking.

5.1 Myopic consumers

Assume consumers are myopic, so that they only consider their current period payoffs in making their decisions. We also assume that the initial state is such that the number of times each consumer has chosen firm i in the past corresponds to the number of times the measure one of all consumers have chosen firm i in the past, so we start with $n_i = N_i$. We can then apply our baseline analysis with across-user learning by defining the learning function for firm i as

$$\widetilde{f}_i\left(N_i\right) = f_i\left(\min\left\{N_i, \overline{N}_i\right\}, \min\left\{N_i, \overline{n}_i\right\}\right)$$

so it combines both across-user learning and within-user learning, and the maximum learning threshold for firm i as $\max\left\{\overline{N}_i, \overline{n}_i\right\}$. In the Appendix we prove the following result.

Proposition 12. *Suppose I has previously sold to consumers $0 \leq N_I \leq \overline{N}_I$ times and E has previously sold to consumers $0 \leq N_E \leq \overline{N}_E$ times, where $n_I = N_I$ and $n_E = N_E$. There exists a unique MPE in which E wins in all periods if and only if $s_E - s_I \geq \Delta(N_I, N_E)$, and I wins in all periods otherwise, where*

$$\begin{aligned}\Delta(N_I, N_E) &= (1-\delta)\left(\sum_{j=0}^{\max\left\{\overline{N}_I, \overline{n}_I\right\}-N_I} \delta^j \widetilde{f}_I(N_I + j) - \sum_{j=0}^{\max\left\{\overline{N}_E, \overline{n}_E\right\}-N_E} \delta^j \widetilde{f}_E(N_E + j)\right) \\ &\quad + \delta^{\max\left\{\overline{N}_I, \overline{n}_I\right\}-N_I+1}\widetilde{f}_I(\max\left\{\overline{N}_I, \overline{n}_I\right\}) - \delta^{\max\left\{\overline{N}_E, \overline{n}_E\right\}-N_E+1}\widetilde{f}_E(\max\left\{\overline{N}_E, \overline{n}_E\right\}).\end{aligned}$$

The result reflects that since consumers are myopic, within-user learning works in the same way as across-user learning, as we noted in Section 4. The firm that can offer the highest utility in a given period, wins all consumers in that period, which ensures all consumers go to the same firm in every period. The same logic repeats every period. As a result, there is no coordination problem across consumers, and our existing across-user learning analysis applies. This implies not only does Proposition 1 apply here, but so do all the results of Section 3 provided we use the learning function and maximum threshold defined above.

5.2 Forward-looking consumers

Now assume consumers are forward-looking. As we will see, things quickly become much more complicated. We will focus on the simplest possible case for which consumers can face a coordination problem in which their beliefs matter. Specifically, we assume $\overline{n}_I = \overline{n}_E = \overline{N}_I = \overline{N}_E = 1$ (both firms only need to attract the measure one of consumers for one period to achieve maximal learning), $n_I = N_I = 1$ (I has already achieved its maximum learning), and $n_E = N_E = 0$ (E starts with no learning).

We start by assuming the two firms have their prices exogenously fixed to zero, so they do not engage in any strategic pricing. Even in this simplest possible setting with forward-looking consumers, we will show consumers can face a coordination problem. This reflects the endogenous switching cost created by within-user learning, which provides a reason for each consumer to want to coordinate on the firm they think will attract the other consumers, in order to also benefit from across-user learning. For example, if all consumers are expected to choose I in every period, a consumer that chooses E in every period would forgo the benefits of across-user learning, while a consumer that switches between the two firms at least once would not be reaping the maximum benefits of within-user learning.

There is an equilibrium where consumers all buy from I in every period iff

$$s_E - s_I \leq f_I(1,1) - \delta f_E(0,1).$$

This rules out the best possible deviation for an individual consumer, which is to instead buy from E in every period. Similarly, there is an equilibrium where consumers all buy from E in every period iff

$$s_E - s_I \geq f_I(1,1) - \delta\left(f_E(1,1) - f_E(1,0)\right).$$

This rules out the best possible deviation for an individual consumer, which is to buy from I in the first period and get $f_I(1,1)$, and then switch to E one period later than everyone else, i.e. free-riding on the across-user learning they generate. Thus, there is a non-empty range of $s_E - s_I$ for which there are multiple equilibria, and so a coordination problem for consumers arises, if

$$f_E(0,1) < f_E(1,1) - f_E(1,0).$$

This is equivalent to E's learning function being supermodular given that $f_E(0,0) = 0$. Here, supermodularity means across-user and within-user learning complement each other.

Since consumers are jointly better off buying from E in every period whenever $s_E + \delta f_E(1,1) \geq s_I + f_I(1,1)$, under Pareto beliefs E wins whenever it is an equilibrium for E to win. In contrast, under beliefs that favor I, I wins whenever it is an equilibrium for I to win. Thus, we obtain the following result.

Proposition 13. *Suppose I is at its threshold, E needs one more period of across and within-user learning to reach its threshold, and both firms' prices are fixed to zero in all periods. Moreover, suppose E's learning function is supermodular. Under Pareto beliefs, E wins if and only if*

$$s_E - s_I \geq f_I(1,1) - \delta\left(f_E(1,1) - f_E(1,0)\right),$$

so E does not win often enough compared to the socially efficient cutoff $f_I(1,1) - \delta f_E(1,1)$. *Under favorable beliefs for I, E wins if and only if*

$$s_E - s_I \geq f_I(1,1) - \delta f_E(0,1),$$

so E wins even less often.

The result in Proposition 13, that supermodularity in the learning functions helps create a coordination problem for consumers is quite intuitive. When other consumers are expected to buy from E, this not only increases the value an individual consumer gets from also choosing E (via across-user learning), but it also increases the amount E can offer the individual consumer via within-user learning. In richer settings, the assumption of supermodularity is not necessary. In an earlier version of the paper we established that a similar coordination problem exists when firms have additive learning functions and both of them are one period away from maximal learning, and below we will show supermodularity is not necessary once we allow firms to set prices. On the other hand, it is important to emphasize that both types of learning are necessary for this result. If either type of learning is shut down for both firms (i.e. if either $f_i(x,y) = f_i(x)$ or $f_i(x,y) = f_i(y)$), then the two inequalities in Proposition 13 coincide, and there is no consumer coordination problem, consistent with our discussion in previous sections.

A surprising result in Proposition 13 is that I still wins too often under Pareto beliefs compared to the efficient cutoff. This reflects that E's ability to win in equilibrium is constrained by the incentive of forward-looking individual consumers to free-ride on the across-user learning of other consumers, waiting till other consumers have chosen E before doing so themselves, which makes it harder to support an equilibrium where everyone chooses E. The possibility of an inefficient outcome here (even under Pareto beliefs) parallels a similar finding for our baseline model with pure across-user learning in the absence of firm pricing, although the mechanism here is different.

Turning next to the case that firms can set prices in every period, we can conduct the analysis in the same setting. First, suppose consumers hold Pareto beliefs. This means that when consumers are in state $(1,1,0,0)$, i.e. I has benefitted from one period of both types learning while E has not had any learning whatsoever, and the prices charged by the two firms in the current period are p_I and p_E respectively, consumers choose I in the current period if and only if

$$s_I + f_I(1,1) - p_I + \delta u(1,1,0,0) > s_E - p_E + \delta u(1,1,1,1),$$

and they choose E otherwise. In other words, consumers coordinate on choosing the firm that offers them the highest PDV of utility. It is then straightforward to show (the proof is in Online Appendix E) that with Pareto beliefs, E wins all periods starting from state $(1,1,0,0)$ iff

$$s_E - s_I \geq f_I(1,1) - \delta f_E(1,1),$$

which is the socially efficient cutoff for E to win.

Suppose instead consumers hold beliefs favorable to I. This means that when consumers are in state $(1,1,0,0)$, consumers choose I in the current period if and only if

$$s_I + f_I(1,1) - p_I + \delta u(1,1,0,0) \geq s_E - p_E + \delta u(1,1,0,1),$$

and they choose E otherwise. Note the difference with the case of Pareto beliefs. Here, consumers choose E if and only if it is individually rational for a consumer to choose E even when she expects all other consumers to choose I. This means that to determine the equilibrium we need to characterize what happens in the state $(1,1,0,1)$ of a single consumer that has chosen E for one period while everyone else has always chosen I.[9] By doing so, in Online Appendix E, we are able to prove the following result.

[9]The state $(1,1,0,1)$ means I has benefited from both types of learning, while E has only benefited from within-user learning for the focal consumer.

Proposition 14. *Suppose I is at its threshold, E needs one more period of across and within-user learning to reach its threshold, and both firms price strategically. If* $f_E(0,1) \geq \delta f_E(1,1)$*, then the cutoff for E to win under beliefs favorable to I is the same as the cutoff under Pareto beliefs, which is equal to the efficient cutoff* $f_I(1,1) - \delta f_E(1,1)$*. Otherwise, if* $f_E(0,1) < \delta f_E(1,1)$*, then the cutoff for E to win under beliefs favorable to I is higher than the cutoff under Pareto beliefs.*

The fact that Pareto beliefs lead to the efficient cutoff is consistent with earlier results, showing that efficiency arises when firms compete in prices. Here, endogenous pricing undoes the free-riding problem we identified above with prices exogenously fixed to zero. The reason is that an individual consumer that deviates and chooses I when everyone else chooses E ends up obtaining the same utility as the consumers that choose E—in both cases this is determined by the maximum utility that can be offered by the losing firm I.

As in the case without pricing, it is not surprising that E has a harder time winning when beliefs favor I than under Pareto beliefs. The interesting part of Proposition 14 is that for the outcomes under the two types of beliefs to be different, so that a coordination problem arises, across-user learning must be sufficiently important relative to within-user learning. This can happen with supermodular, submodular or additive learning functions.

6 Conclusion

Data-enabled learning is becoming an essential ingredient to the competitive dynamics between providers of an increasing range of products and services. Our analysis provides several key takeaways regarding the determinants of a firm's competitive advantage in the presence of data-enabled learning. First, everything else equal, a firm's advantage (in winning) is greater, the higher is its learning curve, the further along the firm is on its learning curve, and the larger is the remaining learning potential. However, even if an incumbent starts with a data advantage, if the entrant's learning curve is steeper and can ultimately reach a higher threshold of learning value (possibly due to better technology), through dynamic pricing the entrant may be able to seize the market. Second, anything that slows down the rate of learning (e.g. a privacy policy) will tend to enhance the incumbent's competitive advantage, other things equal. And third, the resulting data network effect only involves a consumer coordination problem when across-user learning is combined with either (i) continued product improvement during the period(s) over which consumers derive utility from the product, or (ii) within-user learning. In these cases, user beliefs that favor the incumbent can be an additional source of competitive advantage.

When firms can price strategically, unless data-enabled learning creates a consumer coordination problem and beliefs are biased in favor of one firm, the equilibrium outcome is efficient. Still, an unanticipated improvement in the learning of the losing firm makes it a stronger competitor and thereby improves consumer surplus. This creates a potential role for a data sharing policy. However, a policy which forces firms to share the data they obtain from winning with the losing firm will also tend to make firms compete less aggressively in the first place, so on balance, can be bad for consumers. A key condition for data sharing policy to improve consumer surplus is that the firm that benefits from such data sharing is sufficiently far behind the leader. The framework also has implications for merger policy: it can explain why an incumbent which has already accumulated a lot of data may have an excessive incentive to acquire a firm that contains a lot of data even if it doesn't compete with it or need the data to learn itself. This is driven by the incumbent's interest in stopping rivals from acquiring the firm and learning from its data, which would allow them to become closer competitors.

Our framework can be productively extended in many different directions. We have referred to some of these throughout the paper, and due to space constraints, placed them in the Online Appendix. There are undoubtedly many others that remain to be worked out. One is to add uncertainty about the effectiveness of data-enabled learning, and/or additional market segments for each of the two firms, along the lines we explore in Section A of the Online Appendix. Another is to try to analyze a more general model incorporating both across-user and within-user learning, building on our analysis in Section 5. One could also leverage our model to study the effects of other policy proposals, such as one that forces firms to give consumers the ability to port their data when switching from one firm to another.

7 Appendix

This appendix contains the proofs of results not already proven in the main text.

7.1 Proof of Proposition 1 and Corollary 1

The state $(N_I, N_E) = (\overline{N}_I, \overline{N}_E)$ is already handled in the main text. Recall the value function of profit in this case is given by

$$V^i(\overline{N}_I, \overline{N}_E) = \frac{1}{1-\delta} \max\left\{s_i - s_j + f_i(\overline{N}_i) - f_j(\overline{N}_j), 0\right\}.$$

Consider the case with $N_I = \overline{N}_I$ and $1 \leq N_E \leq \overline{N}_E$. Suppose $\Delta(\overline{N}_I, N_E)$, $V^I(\overline{N}_I, N_E)$ and $V^E(\overline{N}_I, N_E)$ are given by $\Delta(N_I, N_E)$, $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$ as characterized in Proposition

1 and Corollary 1 with $N_I = \overline{N}_I$ (this is the induction hypothesis).

Now consider the case in which $N_I = \overline{N}_I$ and E has previously sold to $N_E - 1$ consumers. Suppose E charges p^E and I charges p^I in the current period. I wins the current period if $s_I + f_I(\overline{N}_I) - p^I > s_E + f_E(N_E - 1) - p^E$. Otherwise, E wins. If E wins the current period, the PDV of its current and future profits will be $p^E - c + \delta V^E(\overline{N}_I, N_E)$. If it doesn't win the current period, the PDV of its current and future profits will be $\delta V^E(\overline{N}_I, N_E - 1)$. So E is willing to price down to the point where $p^E - c + \delta V^E(\overline{N}_I, N_E) = \delta V^E(\overline{N}_I, N_E - 1)$, or in other words $p^E = c - \delta(V^E(\overline{N}_I, N_E) - V^E(\overline{N}_I, N_E - 1))$. Based on the same logic, I is willing to price down to $p^I = c - \delta(V^I(\overline{N}_I, N_E - 1) - V^I(\overline{N}_I, N_E))$.

Thus, if we denote by $\Omega^i(N_I, N_E)$ the maximum surplus firm i is willing to offer consumers to win in the current period when the current state is (N_I, N_E), then

$$\begin{aligned}\Omega^I\left(\overline{N}_I, N_E - 1\right) &= s_I - c + f_I(\overline{N}_I) + \delta(V^I(\overline{N}_I, N_E - 1) - V^I(\overline{N}_I, N_E))\\ \Omega^E\left(\overline{N}_I, N_E - 1\right) &= s_E - c + f_E(N_E - 1) + \delta(V^E(\overline{N}_I, N_E) - V^E(\overline{N}_I, N_E - 1)).\end{aligned}$$

Given Bertrand competition, I wins the current period iff

$$\Omega^I\left(\overline{N}_I, N_E - 1\right) > \Omega^E\left(\overline{N}_I, N_E - 1\right). \tag{8}$$

Suppose (8) holds. Then in the current period E sets $p^E = c - \delta(V^E(\overline{N}_I, N_E) - V^E(\overline{N}_I, N_E - 1))$ and I wins by setting $p^I = s_I + f_I(\overline{N}_I) - \Omega^E\left(\overline{N}_I, N_E - 1\right)$. Then I's value is

$$V^I(\overline{N}_I, N_E - 1) = \delta V^I(\overline{N}_I, N_E) + \Omega^I\left(\overline{N}_I, N_E - 1\right) - \Omega^E\left(\overline{N}_I, N_E - 1\right),$$

and E's value is $V^E(\overline{N}_I, N_E - 1) = \delta V^E(\overline{N}_I, N_E - 1)$ given it loses and so will be in the same situation next period. Since $\delta < 1$, this implies $V^E(\overline{N}_I, N_E - 1) = 0$.

Conversely, suppose (8) doesn't hold. Then in the current period I sets $p^I = c - \delta(V^I(\overline{N}_I, N_E - 1) - V^I(\overline{N}_I, N_E))$ and E wins by setting $p^E = s_E + f_E(N_E - 1) - \Omega^I\left(\overline{N}_I, N_E - 1\right)$. Then E's value is

$$V^E(\overline{N}_I, N_E - 1) = \delta V^E(\overline{N}_I, N_E - 1) + \Omega^E\left(\overline{N}_I, N_E - 1\right) - \Omega^I\left(\overline{N}_I, N_E - 1\right)$$

and I's value is $V^I(\overline{N}_I, N_E - 1) = \delta V^I(\overline{N}_I, N_E)$ given it loses and will be in the state $(\overline{N}_I, N_E)$ next period.

Combining the above two cases when (8) holds and when it doesn't, we have

$$V^I(\overline{N}_I, N_E - 1) = \delta V^I(\overline{N}_I, N_E) + \max\left\{\Omega^I\left(\overline{N}_I, N_E - 1\right) - \Omega^E\left(\overline{N}_I, N_E - 1\right), 0\right\} \tag{9}$$

$$V^E(\overline{N}_I, N_E - 1) = \delta V^E(\overline{N}_I, N_E - 1) + \max\left\{\Omega^E\left(\overline{N}_I, N_E - 1\right) - \Omega^I\left(\overline{N}_I, N_E - 1\right), 0\right\}. \tag{10}$$

There are two possibilities: $V^E(\overline{N}_I, N_E - 1) = 0$ and $V^E(\overline{N}_I, N_E - 1) > 0$.

Suppose first $V^E(\overline{N}_I, N_E - 1) = 0$. We can solve (9) for $V^I(\overline{N}_I, N_E - 1)$ using the expressions for

$V^I(\overline{N}_I, N_E)$ and $V^E(\overline{N}_I, N_E)$ from the induction hypothesis and the fact that $V^E(\overline{N}_I, N_E-1) = 0$, which gives

$$V^I(\overline{N}_I, N_E-1) = \begin{cases} \frac{s_I - s_E + f_I(\overline{N}_I) - f_E(N_E-1)}{1-\delta} & \text{if} \quad s_E - s_I < \Delta(\overline{N}_I, N_E) \\ \frac{s_I - s_E + (1-\delta)(f_I(\overline{N}_I) - f_E(N_E-1)) + \delta\Delta(\overline{N}_I, N_E)}{(1-\delta)^2} & \text{if} \quad s_E - s_I \geq \Delta(\overline{N}_I, N_E) \end{cases},$$

with $V^I(\overline{N}_I, N_E-1) \geq 0$ iff

$$s_E - s_I \leq (1-\delta)(f_I(\overline{N}_I) - f_E(N_E-1)) + \delta\Delta(\overline{N}_I, N_E) = \Delta(\overline{N}_I, N_E-1).$$

Next, suppose $V^E(\overline{N}_I, N_E-1) > 0$, so $V^I(\overline{N}_I, N_E-1) = \delta V^I(\overline{N}_I, N_E)$. Since $V^I\left(\overline{N}_I, N_E-1\right) \geq V^I\left(\overline{N}_I, N_E\right)$, we must have $V^I(\overline{N}_I, N_E-1) = V^I(\overline{N}_I, N_E) = 0$ in this case. We can solve (10) for $V^E(\overline{N}_I, N_E-1)$ using the expression for $V^E(\overline{N}_I, N_E)$ from the induction hypothesis and the fact that $V^I(\overline{N}_I, N_E-1) = V^I(\overline{N}_I, N_E) = 0$, which gives

$$V^E(\overline{N}_I, N_E-1) = \frac{s_E - s_I + (1-\delta)(f_E(N_E-1) - f_I(\overline{N}_I)) - \delta\Delta(\overline{N}_I, N_E)}{1-\delta},$$

which is positive iff $s_E - s_I > \Delta(\overline{N}_I, N_E-1)$.

Thus, combining the results with $V^E(\overline{N}_I, N_E-1) = 0$ and $V^E(\overline{N}_I, N_E-1) > 0$, we have proven that $\Delta(\overline{N}_I, N_E-1)$, $V^I(\overline{N}_I, N_E-1)$ and $V^E(\overline{N}_I, N_E-1)$ satisfy the characterization in Proposition 1 and Corollary 1, and by induction, this is true for any $0 \leq N_E \leq \overline{N}_E$ when $N_I = \overline{N}_I$. Symmetry can be used to obtain a parallel result for all $0 \leq N_I \leq \overline{N}_I$ when $N_E = \overline{N}_E$.

Now consider any state (N_I, N_E) with $0 \leq N_I \leq \overline{N}_I - 1$ and $0 \leq N_E \leq \overline{N}_E - 1$, and suppose that the results in Proposition 1 and Corollary 1 hold for the states (N_I+1, N_E) and (N_I, N_E+1)—this is the induction hypothesis. Using the same logic as above when the state is (N_I, N_E), we have

$$V^I(N_I, N_E) = \delta V^I(N_I, N_E+1) + \max\{\Omega(N_I, N_E), 0\} \quad (11)$$
$$V^E(N_I, N_E) = \delta V^E(N_I+1, N_E) + \max\{-\Omega(N_I, N_E), 0\} \quad (12)$$

where

$$\begin{aligned}\Omega(N_I, N_E) = \; & s_I + f_I(N_I) + \delta(V^I(N_I+1, N_E) - V^I(N_I, N_E+1)) \\ & - s_E - f_E(N_E) - \delta\left(V^E(N_I, N_E+1) - V^E(N_I+1, N_E)\right).\end{aligned}$$

Again there are two possibilities, depending on whether $\Omega(N_I, N_E)$ is positive or not. If $\Omega(N_I, N_E) \geq 0$, then we must have $V^E(N_I, N_E) = V^E(N_I+1, N_E) = 0$, while if $\Omega(N_I, N_E) \leq 0$, then we must have $V^I(N_I, N_E) = V^I(N_I, N_E+1) = 0$. When $\Omega(N_I, N_E) \geq 0$ (respectively, $\Omega(N_I, N_E) \leq 0$) we can use the expressions of $V^I(N_I+1, N_E)$ and $V^E(N_I, N_E+1)$ (respectively, $V^I(N_I, N_E+1)$ and $V^E(N_I+1, N_E)$) from the induction hypothesis to solve (11) and (12) for $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$. After some careful manipulations, which include rewriting the resulting summation expressions for $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$ in terms of $\Delta(N_I, N_E)$ and $\Delta(N_I, N_E + 1)$, we confirm that the expressions of $\Delta(N_I, N_E)$, $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$ satisfy the characterization in Proposition 1 and Corollary 1 (these additional details are provided in Online Appendix E).

By repeated application of the induction hypothesis, we obtain that the expressions for $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$ in Corollary 1, and the expression for $\Delta(N_I, N_E)$ in Proposition 1 hold for all states (N_I, N_E), where $0 \leq N_I \leq \overline{N}_I - 1$ and $0 \leq N_E \leq \overline{N}_E - 1$.

7.2 Proof of Proposition 2

To show Proposition 1 continues to apply in the limit as $\overline{N}_I \to \infty$ and $\overline{N}_E \to \infty$, we need to prove the expression in (1) continues to be well defined in this limit case. This is equivalent to proving $\lim_{N_i \to \infty} \sum_{j=0}^{N_i} \delta^j f_i(j)$ exists given this also implies $\lim_{N_i \to \infty} \delta^{N_i} f_i(N_i)$ exists since the summation includes the later expression and other positive terms. Consider the power function $f_i(x) = \theta_i x^{\rho_i}$ where $\theta_i > 0$ and $\rho_i > 0$. Using the ratio test, we know that $\lim_{j \to \infty}(\frac{\delta^{j+1}(j+1)^\rho}{\delta^j j^\rho}) = \delta < 1$ which implies the limit exists for any such power function. Then by the limit comparison test, since the relevant terms are all positive, the limits exist for any unbounded learning function f_i provided there exists some $\kappa_i > 0$, such that $\lim_{j \to \infty} \frac{f_i(j)}{j^\rho} \leq \kappa_i$ for $i = \{I, E\}$.

7.3 Proof of Proposition 3

To prove the result in (i) we need to show that the term (in large brackets) multiplied by θ_I in (2) is greater than the corresponding term multiplied by θ_E when $\overline{N}_I - N_I \leq \overline{N}_E - N_E$ and $\overline{N}_I - \underline{N}_I \geq \overline{N}_E - \underline{N}_E$, with one inequality strict.

If $N_E \leq \underline{N}_E$ and $N_I \leq \underline{N}_I$, then the inequality we require simplifies to

$$\delta^{\underline{N}_I - N_I + 1}\left(1 - \delta^{\overline{N}_I - \underline{N}_I}\right) > \delta^{\underline{N}_E - N_E + 1}\left(1 - \delta^{\overline{N}_E - \underline{N}_E}\right),$$

which follows by combining the above assumptions. If $N_E > \underline{N}_E$ and $N_I > \underline{N}_I$, then the inequality we require simplifies to

$$N_I - \underline{N}_I + \delta\frac{1 - \delta^{\overline{N}_I - N_I}}{1 - \delta} > N_E - \underline{N}_E + \delta\frac{1 - \delta^{\overline{N}_E - N_E}}{1 - \delta},$$

which once again can be shown to follow from the above assumptions. Finally, given $\overline{N}_I - N_I \leq \overline{N}_E - N_E$ and $\overline{N}_I - \underline{N}_I \geq \overline{N}_E - \underline{N}_E$, the only remaining case is $N_E \leq \underline{N}_E$ and $N_I > \underline{N}_I$. In this case, the inequality we wish to prove becomes

$$N_I - \underline{N}_I + \delta\left(\frac{1 - \delta^{\overline{N}_I - N_I}}{1 - \delta}\right) > \delta^{\underline{N}_E - N_E + 1}\left(\frac{1 - \delta^{\overline{N}_E - \underline{N}_E}}{1 - \delta}\right).$$

The LHS is lowest when $N_I = \underline{N}_I + 1$ while the RHS is highest when $N_E = \underline{N}_E$, so a sufficient condition for the inequality to hold is that

$$1 + \delta \left(\frac{1 - \delta^{\overline{N}_I - \underline{N}_I - 1}}{1 - \delta} \right) > \delta \left(\frac{1 - \delta^{\overline{N}_E - \underline{N}_E}}{1 - \delta} \right)$$

or equivalently $1 - \delta^{\overline{N}_I - \underline{N}_I} > \delta \left(1 - \delta^{\overline{N}_E - \underline{N}_E} \right)$, which holds because $\overline{N}_I - \underline{N}_I \geq \overline{N}_E - \underline{N}_E$.

For (ii), the new cutoff for E to win becomes

$$\begin{aligned} \Delta(N_I, N_E)' &= \theta_I \left(N_I - \underline{N}_I\right) + \rho\theta_I \left(\frac{\delta - \delta^{\frac{\overline{N}_I - N_I}{\rho} + 1}}{1 - \delta} \right) \\ &\quad - \left(\theta_E \left(N_E - \underline{N}_E\right) + \rho\theta_E \left(\frac{\delta - \delta^{\frac{\overline{N}_E - N_E}{\rho} + 1}}{1 - \delta} \right) \right) \end{aligned} \tag{13}$$

so

$$\Delta(N_I, N_E)' - \Delta(N_I, N_E) = \frac{\delta}{1 - \delta} \left(\theta_I y \left(\overline{N}_I - N_I \right) - \theta_E y \left(\overline{N}_E - N_E \right) \right), \tag{14}$$

where $y(m) = \delta^m - \rho\delta^{\frac{m}{\rho}} - (1 - \rho)$ and $0 < \rho < 1$. The expression in (14) is non-negative when $\theta_E \geq \theta_I$ and $\overline{N}_I - N_I \leq \overline{N}_E - N_E$, and positive when one of the two inequalities is strict. To see this, note $y\left(\overline{N}_E - N_E\right) \leq y\left(\overline{N}_I - N_I\right) < 0$ because $y(0) = 0$ and $y'(m) = \left(\delta^m - \delta^{\frac{m}{\rho}}\right) \ln \delta < 0$.

For (iii), suppose $N_I + k \leq \underline{N}$, which implies $N_E + k < \underline{N}$ because $N_I > N_E$. Then imposing that learning functions are identical on (2), we get

$$\Delta(N_I + k, N_E + k) - \Delta(N_I, N_E) = \frac{(1 - \delta^k)(1 - \delta^{\overline{N} - \underline{N}})(\delta^{\overline{N} - N_I} - \delta^{\overline{N} - N_E})\theta}{(1 - \delta)\delta^{\overline{N} - \underline{N} + k - 1}} > 0.$$

Suppose instead $N_I > \underline{N}$ and $N_E > \underline{N}$. Then

$$\Delta(N_I + k, N_E + k) - \Delta(N_I, N_E) = -\frac{(1 - \delta^k)(\delta^{\overline{N} - N_I + 1} - \delta^{\overline{N} - N_E + 1})\theta}{(1 - \delta)\delta^k},$$

which is negative given $N_I > N_E$.

7.4 Proof of Proposition 4

Consider the socially optimal outcome. If I wins in every period, the value created is

$$\frac{s_I}{1 - \delta} + \frac{\delta^{\overline{N}_I - N_I + 1}}{1 - \delta} f_I(\overline{N}_I) + \sum_{j=0}^{\overline{N}_I - N_I} \delta^j f_I(N_I + j)$$

while if E wins in every period, the value created is

$$\frac{s_E}{1-\delta}+\frac{\delta^{\overline{N}_E-N_E+1}}{1-\delta}f_E(\overline{N}_E)+\sum_{j=0}^{\overline{N}_E-N_E}\delta^j f_E(N_E+j).$$

The socially optimal cutoff is the level of $s_I - s_E$ for which these two expressions are equal, which corresponds exactly to the equilibrium cutoff $\Delta(N_I, N_E)$.

7.5 Proof of Proposition 5

The proof follows in three steps. First, we show that $CS\,(N_I, N_E)' - CS\,(N_I, N_E)$, as defined by (3)-(5), is decreasing in $s_E - s_I$ for $s_E - s_I \leq \Delta(N_I, N_E)$ and increasing in $s_E - s_I$ for $s_E - s_I \geq \Delta(N_I, N_E)$. To see this, note

- If $s_E - s_I < \Delta(N_I, N_E + 1)$, then

  $$CS\,(N_I, N_E)' - CS\,(N_I, N_E) = \sum_{j=0}^{\infty}\delta^j \min\left\{s_I - s_E + f_I(\min\left\{N_I + j, \overline{N}_I\right\}), f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right)\right\},$$

  which is clearly (weakly) decreasing in $s_E - s_I$.

- If $\Delta(N_I, N_E + 1) \leq s_E - s_I < \Delta(N_I, N_E)$, then

  $$\begin{aligned} CS\,(N_I, N_E)' - CS\,(N_I, N_E) &= \sum_{j=0}^{\infty}\delta^j \min\left\{f_I(\min\left\{N_I + j, \overline{N}_I\right\}), s_E - s_I + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right)\right\} \\ &\quad + \frac{s_I - s_E}{(1-\delta)^2} + K, \end{aligned}$$

  where K is a term constant in $s_E - s_I$. And since

  $$\sum_{j=0}^{\infty}\delta^j\left(s_E - s_I\right) = \frac{s_E - s_I}{1-\delta} < \frac{s_E - s_I}{(1-\delta)^2},$$

  we conclude that $CS\,(N_I, N_E)' - CS\,(N_I, N_E)$ is strictly decreasing in $s_E - s_I$ on this region.

- The cases with $s_E - s_I \geq \Delta(N_I + 1, N_E)$ and $\Delta(N_I, N_E) \leq s_E - s_I < \Delta(N_I + 1, N_E)$ are respectively weakly and strictly increasing in $s_E - s_I$, following symmetric arguments to the above two cases.

Second, we show that $CS\,(N_I, N_E)' - CS\,(N_I, N_E) < 0$ when $s_E - s_I = \Delta(N_I, N_E)$. Indeed, in this case we have

$$CS\,(N_I, N_E)' \leq \sum_{j=0}^{\infty}\delta^j\left(s_E + f_E(\min\left\{N_E + j, \overline{N}_E\right\})\right) - \frac{c}{1-\delta} = CS\,(N_I, N_E)\,.$$

Furthermore, $CS\left(N_I, N_E\right)' = CS\left(N_I, N_E\right)$ is possible in this case iff

$$s_E + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right) \leq s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\})$$

for all $j \geq 0$. However, since

$$\begin{aligned} s_E - s_I &= \Delta(N_I, N_E) \\ &= (1-\delta)\left(\sum_{j=0}^{\overline{N}_I - N_I} \delta^j f_I(N_I + j) - \sum_{j=0}^{\overline{N}_E - N_E} \delta^j f_E(N_E + j)\right) + \delta^{\overline{N}_I - N_I + 1} f_I(\overline{N}_I) - \delta^{\overline{N}_E - N_E + 1} f_E(\overline{N}_E), \end{aligned}$$

we conclude that $CS\left(N_I, N_E\right)' = CS\left(N_I, N_E\right)$ is only possible in this case if

$$s_E + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right) = s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\})$$

for all $j \geq 0$. Given $\overline{N}_E - N_E > \overline{N}_I - N_I$, this would require $f_E\left(\overline{N}_E\right) = f_E\left(\overline{N}_E - 1\right)$, which is ruled out by assumption. Thus, we must have $CS\left(N_I, N_E\right)' - CS\left(N_I, N_E\right) < 0$ when $s_E - s_I = \Delta(N_I, N_E)$.

Third, we show that $CS\left(N_I, N_E\right)' - CS\left(N_I, N_E\right) > 0$ when $s_E - s_I \leq \min\left\{\Delta(N_I, N_E + 1), f_I\left(N_I\right) - f_E(N_E)\right\}$ and when $s_E - s_I > \max\left\{\Delta(N_I + 1, N_E), f_I\left(N_I\right) - f_E(N_E)\right\}$. In the first case we are in the region where there is no subsidy without data sharing, so we have

$$\begin{aligned} CS\left(N_I, N_E\right)' &= \sum_{j=0}^{\infty} \delta^j \min\left\{s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\}), s_E + f_E\left(\min\left\{N_E + j, \overline{N}_E\right\}\right)\right\} - \frac{c}{1-\delta} \\ &> \sum_{j=0}^{\infty} \delta^j \min\left\{s_I + f_I(\min\left\{N_I + j, \overline{N}_I\right\}), s_E + f_E\left(N_E\right)\right\} - \frac{c}{1-\delta} \\ &= \sum_{j=0}^{\infty} \delta^j \left(s_E + f_E\left(N_E\right)\right) - \frac{c}{1-\delta} = CS\left(N_I, N_E\right). \end{aligned}$$

And, by a symmetric argument, in the case when $s_E - s_I > \max\left\{\Delta(N_I + 1, N_E), f_I\left(N_I\right) - f_E(N_E)\right\}$, we have $CS\left(N_I, N_E\right)' > CS\left(N_I, N_E\right)$.

Together, these three steps, and the continuity of the two functions $CS\left(N_I, N_E\right)$ and $CS\left(N_I, N_E\right)'$ in $s_E - s_I$ imply the result in the proposition.

7.6 Proof of Proposition 6

The result on competitive advantage follows immediately from case (ii) in Proposition 3. Applying the S-curve to (4), the PDV of consumer surplus when I wins equals

$$CS\left(N_I, N_E\right) = \frac{s_I - c}{1-\delta} + \theta_I \sum_{k=0}^{\overline{N}_I - N_I} \delta^k \left(N_I + k - \underline{N}_I\right) + \frac{\delta^{\overline{N}_I - N_I + 1}}{1-\delta} \theta_I \left(\overline{N}_I - \underline{N}_I\right) - \left(\frac{s_I - s_E + \Delta(N_I, N_E)}{(1-\delta)^2}\right)$$

without the policy. Since I wins without the policy, and the policy increases I's competitive advantage, I must continue to win after the policy is introduced. The PDV of consumer surplus with the policy equals

$$CS\left(N_I, N_E\right)' = \frac{s_I - c}{1-\delta} + \theta_I \sum_{k=0}^{\frac{\overline{N}_I - N_I}{\rho}} \delta^k (N_I - \underline{N}_I + \rho k) + \frac{\delta^{\frac{\overline{N}_I - N_I}{\rho}+1}}{1-\delta} \theta_I \left(\overline{N}_I - \underline{N}_I\right) - \left(\frac{s_I - s_E + \Delta(N_I, N_E)'}{(1-\delta)^2}\right),$$

where $\Delta(N_I, N_E)'$ is defined in (13). The change in consumer surplus is

$$\frac{\Delta(N_I, N_E) - \Delta(N_I, N_E)'}{(1-\delta)^2} - \theta_I \sum_{k=0}^{\overline{N}_I - N_I} k\delta^k \left(1-\rho\right) \tag{15}$$

$$-\frac{\theta_I \left(\delta^{\overline{N}_I - N_I + 1} - \delta^{\frac{\overline{N}_I - N_I}{\rho}+1}\right)}{1-\delta} \left(\overline{N}_I - \underline{N}_I\right) + \theta_I \sum_{k=\overline{N}_I - N_I + 1}^{\frac{\overline{N}_I - N_I}{\rho}} \delta^k (N_I - \underline{N}_I + \rho k). \tag{16}$$

Since $\Delta(N_I, N_E)' > \Delta(N_I, N_E)$, the two terms in (15) are negative. Using $\rho = \frac{\overline{N}_I - N_I}{\overline{N}_I - N_I + m_I}$, the expression in (16) can be rewritten as

$$-\delta^{\overline{N}_I - N_I + 1} \frac{\overline{N}_I - N_I}{(1-\delta)^2 \left(\overline{N}_I - N_I + m_I\right)} \left(m_I \left(1-\delta\right) - \left(1 - \delta^{m_I}\right)\right)$$

which is equal to zero if $m_I = 1$, and is negative for $m_I > 1$ given $m_I\left(1-\delta\right) - \left(1-\delta^{m_I}\right) = \left(1-\delta\right)\left(m_I - \sum_{k=1}^{m_I} \delta^{k-1}\right)$.

7.7 Proof of Proposition 7

First, if $s_E - s_I < \Delta\left(N_I, N_E + N_A\right)$, then E is not willing to pay anything for N_A because even if it acquires it, it's not enough to make it win. This also means that E can't have been winning originally without the data, so it makes zero profits regardless of who acquires N_A. In this case, I wins the data since it is willing to pay up to the value it would get if it wins the data less the value it would get if E wins the data—and it is straightforward to verify that this difference is positive (under the assumption that $N_I < \overline{N}_I$ or $N_E < \overline{N}_E$).

Second, if $s_E - s_I \geq \Delta\left(N_I + N_A, N_E\right)$, then we are in the symmetric case where I is willing to pay nothing for the data and E acquires it.

Suppose now we are in the remaining (and interesting) case

$$\Delta\left(N_I, N_E + N_A\right) \leq s_E - s_I < \Delta\left(N_I + N_A, N_E\right),$$

which means that whichever firm acquires the data will win in all periods. Note that we also have

$$\Delta\left(N_I, N_E + N_A\right) \leq \Delta\left(N_I + 1, N_E + N_A\right) \leq \Delta\left(N_I + N_A, N_E\right),$$

with strict inequalities if $N_I < \overline{N}_I$ and $N_E < \overline{N}_E$. Consequently, using the expressions from Corollary 1, E's willingness-to-pay for the data is $W_E\left(s_E - s_I\right)$, where

$$W_E\left(x\right) = \begin{cases} \frac{x - \Delta(N_I, N_E + N_A)}{(1-\delta)^2} & \text{if} \quad \Delta\left(N_I, N_E + N_A\right) \leq x < \Delta\left(N_I + 1, N_E + N_A\right) \\ \frac{x}{1-\delta} - \frac{\Delta(N_I, N_E + N_A) - \delta \Delta(N_I + 1, N_E + N_A)}{(1-\delta)^2} & \text{if} \quad \Delta\left(N_I + 1, N_E + N_A\right) \leq x < \Delta\left(N_I + N_A, N_E\right) \end{cases}.$$

The two cases in this expression correspond to whether I subsidizes after E wins the data or not.

Similarly, I's willingness to pay for the data is $W_I\left(s_E - s_I\right)$, where

$$W_I\left(x\right) = \begin{cases} \frac{-x}{1-\delta} + \frac{\Delta(N_I + N_A, N_E) - \delta \Delta(N_I + N_A, N_E + 1)}{(1-\delta)^2} & \text{if} \quad \Delta\left(N_I, N_E + N_A\right) \leq x < \Delta\left(N_I + N_A, N_E + 1\right) \\ \frac{-x + \Delta(N_I + N_A, N_E)}{(1-\delta)^2} & \text{if} \quad \Delta\left(N_I + N_A, N_E + 1\right) \leq x < \Delta\left(N_I + N_A, N_E\right) \end{cases}$$

Note that $W_E\left(x\right)$ is increasing in x and $W_I\left(x\right)$ is decreasing in x. Furthermore, $W_E\left(\Delta\left(N_I, N_E + N_A\right)\right) = 0$ and $\lim_{x \to \Delta(N_I + N_A, N_E)} W_I\left(x\right) = 0$. Thus, since $W_E\left(x\right)$ and $W_I\left(x\right)$ are continuous, $\Delta\left(N_I, N_E + N_A\right) < \Delta\left(N_I + N_A, N_E\right)$, and combining with the two other cases discussed above, we conclude that there is a unique Δ^* such that I's willingness-to-pay for the data is larger (and therefore I wins) when $s_E - s_I < \Delta^*$ and E's willingness-to-pay is larger (and therefore E wins) for $\Delta^* \geq s_E - s_I$. Since we assumed we are in the situation where at least one firm is not already at its learning threshold, the cutoff Δ^* satisfies (6). Noting that $\Delta\left(N_I + 1, N_E + N_A\right) \leq \Delta\left(N_I + N_A, N_E + 1\right)$ and that the socially efficient condition for E to win the data is $s_E - s_I \geq \Delta\left(N_I + N_A, N_E + N_A\right)$, we obtain the result stated in the Proposition.[10]

7.8 Proof of Proposition 10

As discussed in the text, the analysis can be conducted with a representative consumer. A key difference with the proof of Proposition 1 is that we need to construct the value function for the consumer since the consumer needs to take future utility into account when deciding which firm to buy from. We denote this $u(N_I, N_E)$, which measures the PDV of surplus a consumer expects to obtain from making her optimal choices in every stage when firms play their equilibrium strategies and the consumer has previously bought N_I times from firm I and N_E times from firm E.

[10] The explicit characterization of Δ^* is relegated to Online Appendix C.

It is easily verified that when $(N_I, N_E) = (\overline{N}_I, \overline{N}_E)$, we have

$$\begin{aligned} V^I(\overline{N}_I, \overline{N}_E) &= \frac{\max\left\{s_I + f_I(\overline{N}_I) - s_E - f_E(\overline{N}_E), 0\right\}}{1-\delta} \\ V^E(\overline{N}_I, \overline{N}_E) &= \frac{\max\left\{s_E + f_E(\overline{N}_E) - s_I - f_I(\overline{N}_I), 0\right\}}{1-\delta} \\ u(\overline{N}_I, \overline{N}_E) &= \frac{\min\left\{s_I + f_I(\overline{N}_I), s_E + f_E(\overline{N}_E)\right\} - c}{1-\delta}. \end{aligned}$$

Consider the case with $N_I = \overline{N}_I$ and $1 \leq N_E \leq \overline{N}_E$. Suppose that $\Delta(\overline{N}_I, N_E)$, $V^I(\overline{N}_I, N_E)$ and $V^E(\overline{N}_I, N_E)$ are given by $\Delta(N_I, N_E)$, $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$ as characterized in Proposition 10 with $N_I = \overline{N}_I$, and that the corresponding value function for consumers is

$$u(\overline{N}_I, N_E) = \min\left\{ f_I(\overline{N}_I) + \frac{s_I + \delta f_I(\overline{N}_I) - c}{1-\delta}, \sum_{j=0}^{\overline{N}_E - N_E} \delta^j f_E(N_E + j) + \frac{s_E + \delta^{\overline{N}_E - N_E + 1} f_E(\overline{N}_E) - c}{1-\delta} \right\}.$$

This is the induction hypothesis.

Consider now the state $(\overline{N}_I, N_E - 1)$. The consumer chooses I in the current period iff

$$s_I + f_I(\overline{N}_I) - p^I + \delta u(\overline{N}_I, N_E - 1) > s_E + f_E(N_E - 1) - p^E + \delta u(\overline{N}_I, N_E).$$

I is willing to set p^I down to the point where it is indifferent between winning or losing the current period, i.e. $p^I = c - \delta(V^I(\overline{N}_I, N_E - 1) - V^I(\overline{N}_I, N_E))$. Similarly, E is willing to set p^E down to $p^E = c - \delta(V^E(\overline{N}_I, N_E) - V^E(\overline{N}_I, N_E - 1))$.

If we denote by $\Omega^i(N_I, N_E)$ the PDV of the maximum surplus firm i is willing to offer consumers to win in the current period when the current state is (N_I, N_E), then

$$\begin{aligned} \Omega^I\left(\overline{N}_I, N_E - 1\right) &= s_I - c + f_I(\overline{N}_I) + \delta(V^I(\overline{N}_I, N_E - 1) - V^I(\overline{N}_I, N_E) + u(\overline{N}_I, N_E - 1)) \\ \Omega^E\left(\overline{N}_I, N_E - 1\right) &= s_E - c + f_E(N_E - 1) + \delta(V^E(\overline{N}_I, N_E) - V^E(\overline{N}_I, N_E - 1) + u(\overline{N}_I, N_E)). \end{aligned}$$

Given Bertrand competition, I wins the current period iff

$$\Omega^I\left(\overline{N}_I, N_E - 1\right) > \Omega^E\left(\overline{N}_I, N_E - 1\right).$$

Using the same logic as in the Proof of Proposition 1 and Corollary 1, we obtain

$$\begin{aligned} u(\overline{N}_I, N_E - 1) &= \min\left\{\Omega^I\left(\overline{N}_I, N_E - 1\right), \Omega^E\left(\overline{N}_I, N_E - 1\right)\right\} && (17) \\ V^I(\overline{N}_I, N_E - 1) &= \delta V^I(\overline{N}_I, N_E) + \max\left\{\Omega^I\left(\overline{N}_I, N_E - 1\right) - \Omega^E\left(\overline{N}_I, N_E - 1\right), 0\right\} && (18) \\ V^E(\overline{N}_I, N_E - 1) &= \delta V^E(\overline{N}_I, N_E - 1) + \max\left\{\Omega^E\left(\overline{N}_I, N_E - 1\right) - \Omega^I\left(\overline{N}_I, N_E - 1\right), 0\right\}. && (19) \end{aligned}$$

There are two possibilities: $\Omega^I\left(\overline{N}_I, N_E - 1\right) \geq \Omega^E\left(\overline{N}_I, N_E - 1\right)$ or $\Omega^I\left(\overline{N}_I, N_E - 1\right) < \Omega^E\left(\overline{N}_I, N_E - 1\right)$.

Suppose first $\Omega^I\left(\overline{N}_I, N_E - 1\right) \geq \Omega^E\left(\overline{N}_I, N_E - 1\right)$, which immediately implies $V^E(\overline{N}_I, N_E - 1) = 0$. Using the expressions of $V^I(\overline{N}_I, N_E)$, $V^E(\overline{N}_I, N_E)$ and $u(\overline{N}_I, N_E)$ from the induction hypothesis, we then solve (17) and (18) for $V^I(\overline{N}_I, N_E - 1)$ and $u(\overline{N}_I, N_E - 1)$ to obtain

$$
\begin{aligned}
V^I(\overline{N}_I, N_E - 1) &= \frac{s_I - s_E + \Delta(\overline{N}_I, N_E - 1)}{1-\delta} \\
u(\overline{N}_I, N_E - 1) &= \frac{s_E + \delta^{\overline{N}_E - N_E + 2} f_E(\overline{N}_E) - c}{1-\delta} + \sum_{j=0}^{\overline{N}_E - N_E + 1} \delta^j f_E(N_E - 1 + j). \qquad (20)
\end{aligned}
$$

Now suppose $\Omega^I\left(\overline{N}_I, N_E - 1\right) < \Omega^E\left(\overline{N}_I, N_E - 1\right)$, which means E wins the current period, so $V^I(\overline{N}_I, N_E - 1) = \delta V^I(\overline{N}_I, N_E)$, which implies $V^I(\overline{N}_I, N_E - 1) = V^I(\overline{N}_I, N_E) = 0$ because we must have $V^I(\overline{N}_I, N_E - 1) \geq V^I(\overline{N}_I, N_E)$. Furthermore, (17) and (19) become

$$
\begin{aligned}
u(\overline{N}_I, N_E - 1) &= \frac{s_I + f_I(\overline{N}_I) - c}{1-\delta} \qquad (21) \\
V^E(\overline{N}_I, N_E - 1) &= \frac{s_E - s_I - \Delta(\overline{N}_I, N_E - 1)}{1-\delta}.
\end{aligned}
$$

Thus, combining the results with $\Omega^I\left(\overline{N}_I, N_E - 1\right) \geq \Omega^E\left(\overline{N}_I, N_E - 1\right)$ and $\Omega^I\left(\overline{N}_I, N_E - 1\right) < \Omega^E\left(\overline{N}_I, N_E - 1\right)$, we have proven that $\Delta(\overline{N}_I, N_E - 1)$, $V^I(\overline{N}_I, N_E - 1)$ and $V^E(\overline{N}_I, N_E - 1)$ satisfy the characterization in Proposition 10, and $u(\overline{N}_I, N_E - 1)$ is given by the minimum of the two expressions in (20) and (21). By induction, this is true for any $0 \leq N_E \leq \overline{N}_E$ when $N_I = \overline{N}_I$. Symmetry can then be used to obtain a parallel result for all $0 \leq N_I \leq \overline{N}_I$ when $N_E = \overline{N}_E$.

Now consider any state (N_I, N_E) with $0 \leq N_I \leq \overline{N}_I - 1$ and $0 \leq N_E \leq \overline{N}_E - 1$. The induction hypothesis is that the result in Proposition 10 holds for the states $(N_I + 1, N_E)$ and $(N_I, N_E + 1)$ and that the value functions for consumers in these two states are

$$
u(N_I + 1, N_E) = \min\left\{ \begin{array}{c} \sum_{j=0}^{\overline{N}_I - N_I - 1} \delta^j f_I(N_I + 1 + j) + \frac{s_I + \delta^{\overline{N}_I - N_I} f_I(\overline{N}_I) - c}{1-\delta}, \\ \sum_{j=0}^{\overline{N}_E - N_E} \delta^j f_E(N_E + j) + \frac{s_E + \delta^{\overline{N}_E - N_E + 1} f_E(\overline{N}_E) - c}{1-\delta} \end{array} \right\}
$$

$$
u(N_I, N_E + 1) = \min\left\{ \begin{array}{c} \sum_{j=0}^{\overline{N}_I - N_I} \delta^j f_I(N_I + j) + \frac{s_I + \delta^{\overline{N}_I - N_I + 1} f_I(\overline{N}_I) - c}{1-\delta}, \\ \sum_{j=0}^{\overline{N}_E - N_E + 1} \delta^j f_E(N_E + 1 + j) + \frac{s_E + \delta^{\overline{N}_E - N_E} f_E(\overline{N}_E) - c}{1-\delta} \end{array} \right\}.
$$

Using the same logic as above, we have

$$
u(N_I, N_E) = \min\left\{\Omega^I(N_I, N_E), \Omega^E(N_I, N_E)\right\}, \qquad (22)
$$

where

$$
\begin{aligned}
\Omega^I(N_I, N_E) &= s_I - c + f_I(N_I) + \delta(V^I(N_I + 1, N_E) - V^I(N_I, N_E + 1) + u(N_I + 1, N_E)) \\
\Omega^E(N_I, N_E) &= s_E - c + f_E(N_E) + \delta(V^E(N_I, N_E + 1) - V^E(N_I + 1, N_E) + u(N_I, N_E + 1))
\end{aligned}
$$

and

$$V^I(N_I, N_E) \quad = \quad \delta V^I(N_I, N_E+1) + \max\{\Omega(N_I, N_E), 0\} \tag{23}$$
$$V^E(N_I, N_E) \quad = \quad \delta V^E(N_I+1, N_E) + \max\{-\Omega(N_I, N_E), 0\}, \tag{24}$$

where $\Omega(N_I, N_E) \equiv \Omega^I(N_I, N_E) - \Omega^E(N_I, N_E)$.

Again there are two possibilities, depending on whether $\Omega(N_I, N_E) > 0$ or not. If $\Omega(N_I, N_E) > 0$, then we must have $V^E(N_I, N_E) = V^E(N_I+1, N_E) = 0$, while if $\Omega(N_I, N_E) \leq 0$, then we must have $V^I(N_I, N_E) = V^I(N_I, N_E+1) = 0$. For each of these cases, we can use the expressions of $u(N_I+1, N_E)$, $u(N_I, N_E+1)$, $V^I(N_I, N_E+1)$ and $V^E(N_I+1, N_E)$ from the induction hypothesis to solve (22), (23) and (24) for $V^I(N_I, N_E)$, $V^E(N_I, N_E)$ and $u(N_I, N_E)$. Then we can rewrite the resulting summation expressions in $V^I(N_I, N_E)$ and $V^E(N_I, N_E)$ in terms of $\Delta(N_I, N_E)$ to obtain

$$\begin{aligned} V^I(N_I, N_E) &= \max\left\{\frac{s_I - s_E + \Delta(N_I, N_E)}{1-\delta}, 0\right\} \\ V^E(N_I, N_E) &= \max\left\{\frac{s_E - s_I - \Delta(N_I, N_E)}{1-\delta}, 0\right\} \end{aligned}$$

and

$$u(N_I, N_E) = \min\left\{ \begin{array}{c} \sum_{j=0}^{\overline{N}_I - N_I} \delta^j f_I(N_I+j) + \frac{s_I + \delta^{\overline{N}_I - N_I + 1} f_I(\overline{N}_I) - c}{1-\delta}, \\ \sum_{j=0}^{\overline{N}_E - N_E} \delta^j f_E(N_E+j) + \frac{s_E + \delta^{\overline{N}_E - N_E + 1} f_E(\overline{N}_E) - c}{1-\delta} \end{array} \right\}.$$

By repeated application of the induction hypothesis, the result in Proposition 10 holds for all states (N_I, N_E), with $0 \leq N_I \leq \overline{N}_I$ and $0 \leq N_E \leq \overline{N}_E$.

Finally, we can confirm that consumers do not want to switch firms along the equilibrium path. Suppose the current state is (N_I, N_E). If $s_E - s_I < \Delta(N_I, N_E)$, then consumers choose I in the current period, so next period the state will be (N_I+1, N_E). Since $\Delta(N_I, N_E)$ is increasing in N_I, this means in the subsequent period consumers will choose I again, and so on until $N_I = \overline{N}_I$. When the state is $(\overline{N}_I, N_E)$, all subsequent periods will have an identical state and outcome, in which consumers choose I. Conversely, if $s_E - s_I \geq \Delta(N_I, N_E)$, then consumers choose E in the current period, so next period the state will be (N_I, N_E+1). Since $\Delta(N_I, N_E)$ is decreasing in N_E, this means in the subsequent period consumers will choose E again, and so on until $N_E = \overline{N}_E$. At the state $(N_I, \overline{N}_E)$, all subsequent periods will have an identical state and outcome, in which consumers choose E.

7.9 Proof of Proposition 11

For any given state (N_I, N_E), even though the interpretation of N_i is different, the firms' value functions with myopic consumers and within-user learning are the same as with across-user learning, given that in both cases consumers have no reason to consider more than their current surplus when deciding which firm to buy from, and along the equilibrium path, the consumer (in the case of within-user learning) or all consumers (in the case of across-user learning) choose the same (winning firm) each period. We thus use the value functions from Corollary 1 and compare them with the value functions from Proposition 10.

There are four different cases to consider, depending on the value of $s_E - s_I$. If $s_E - s_I < \Delta(N_I, N_E + 1)$ then $V^I(N_I, N_E)$ is higher with myopic consumers since

$$\frac{s_I - s_E}{1-\delta} + \frac{\Delta(N_I, N_E) - \delta\Delta(N_I, N_E + 1)}{(1-\delta)^2} > \frac{s_I - s_E + \Delta(N_I, N_E)}{1-\delta} > 0$$

given $\Delta(N_I, N_E) > \Delta(N_I, N_E + 1)$. If $\Delta(N_I, N_E + 1) \leq s_E - s_I < \Delta(N_I, N_E)$, then $V^I(N_I, N_E)$ is higher with myopic consumers since

$$\frac{s_I - s_E + \Delta(N_I, N_E)}{(1-\delta)^2} > \frac{s_I - s_E + \Delta(N_I, N_E)}{1-\delta} > 0.$$

By symmetry, $V^I(N_I, N_E)$ is also higher with myopic consumers when $s_E - s_I \geq \Delta(N_I, N_E)$.

7.10 Proof of Proposition 12

We apply exactly the same two-dimensional backwards induction steps as in the proof of Proposition 1 to the learning function $\widetilde{f}_i(N_i)$ starting from the threshold $n_i = N_i = \max\left\{\overline{N}_i, \overline{n}_i\right\}$, and working backwards. Provided all consumers have the same history in any given period, as they do in equilibrium, each firm must offer a single price to everyone. Note that if some individual consumer deviates and chooses something different from other consumers, it has no effect on the firms' profits and we can ignore it (the deviating consumer is of measure zero, and the measure of consumers attracted by the winning firm remains equal to one). Moreover, such a consumer doesn't consider what happens to them after they deviate in terms of their future utility off the equilibrium path since they are myopic and so we don't need to work out their value in these off-equilibrium states to determine if consumers want to deviate in the first place. As a result, there is no need to specify value functions in states in which $n_i \neq N_i$, which is what allows us to apply the proof of Proposition 1 to the learning function $\widetilde{f}_i(N_i)$ without any modification.

References

Beggs, A. and P. Klemperer (1992) "Multi-period competition with switching costs," *Econometrica*, 60(3), 651-666.

Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010) "Learning-by-doing, organizational forgetting, and industry dynamics," *Econometrica*, 78(2), 453-508.

Biglaiser, G. and J. Crémer (2020) "The value of incumbency when platforms face heterogenous consumers," *American Economic Journal: Microeconomics*, 12(4), 229-269.

Biglaiser, G., E. Calvano, and J. Crémer (2019) "Incumbency advantage and its value," *Journal of Economics and Management Strategy*, 28(1), 41-48.

Cabral, L. M (2011) "Dynamic price competition with network effects," *Review of Economic Studies*, 78, 83-111.

Cabral, L. M. and M. H. Riordan (1994) "The learning curve, market dominance, and predatory pricing," *Econometrica*, 62(5), 1115.

Caillaud, B. and B. Jullien (2001) "Competing cybermediaries," *European Economic Review*, 45(4-6), 797-808.

Dasgupta, P. and J.Stiglitz (1988) "Learning-by-doing, market structure and industrial and trade policies," *Oxford Economic Papers*, 40(2), 246-268.

De Corniere, A. and G. Taylor (2020) "Data and Competition," Working paper.

Farboodi, M., R. Mihet, T. Philippon, and L. Veldkamp (2019) "Big data and firm dynamics," *American Economic Review Papers and Proceedings*, 109, 38-42.

Fudenberg, D. and J. Tirole (1983) "Learning-by-doing and market performance," *Bell Journal of Economics*, 14(2) 522-530.

Halaburda, H. B. Jullien, and Y. Yehezkel (2020) "Dynamic competition with network externalities: Why history matters," *RAND Journal of Economics*, 51(1), 3-31.

Katz, M. L. and C. Shapiro (1986) "Technology adoption in the presence of network externalities," *Journal of Political Economy*, 94(4), 822-841.

Klemperer, P. (1987) "Markets with consumer switching costs," *Quarterly Journal of Economics*, 102(2), 375-394.

Mitchell, M. and A. Skrzypacz (2006) "Network externalities and long-run market shares," *Economic Theory*, 29(3), 621-648.

Prufer, J. and C. Schottmüller (2017) "Competing with big data," *Tilburg Law School Research Paper.*

Salant, D. (1990) "Some Stochastic Oligopoly Races for Experience," *Working paper.*

Schaefer, M., G. Sapi and S. Lorincz (2018) "The effect of big data on recommendation quality: The example of internet search," No 284. Düsseldorf Institute for Competition Economics Discussion Papers.