

Competing with Big Data*

Jens Prüfer[†] and Christoph Schottmüller[‡]

February 15, 2017

Abstract

This paper studies competition in data-driven markets, that is, markets where the cost of quality production is decreasing in the amount of machine-generated data about user preferences or characteristics, which is an inseparable byproduct of using services offered in such markets. This gives rise to data-driven indirect network effects. We construct a dynamic model of R&D competition, where duopolists repeatedly determine their innovation investments, and show that such markets tip under very mild conditions, moving towards monopoly. In a tipped market, innovation incentives both for the dominant firm and for competitors are small. We also show under which conditions a dominant firm in one market can leverage its position to a *connected market*, thereby initiating a *domino effect*. We show that market tipping can be avoided if competitors share their user information.

JEL classification: D43, D92, L13, L43, L86

Keywords: Big Data, Datafication, Data-driven Indirect Network Effects, Dynamic Competition, Regulation

*We are grateful for comments on an earlier draft to Wieland Müller, Mauricio Rodriguez Acosta, Florian Schütt, Peter N. Sørensen, Xiaoyu Wang and seminar audiences at universities in Copenhagen, Liège, and Tilburg, and the CEPS Workshop on “Competition policy in the digital economy” (Brussels), the SIOE conference (Paris), and the Konkurrenceøkonomisk Forum (Copenhagen). All errors are our own.

[†]CentER, TILEC, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; j.prufer@uvt.nl.

[‡]Department of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, Denmark and TILEC; christoph.schottmueller@econ.ku.dk.

1. Introduction

In recent decades, the rate of technological progress has accelerated and most of it has occurred in fields that draw heavily on machine-generated data about user behavior (Brynjolfsson and McAfee, 2012).¹ This development was coined “the rise of big data” or “datafication” and is explained by two simultaneous, recent technological innovations (Mayer-Schönberger and Cukier, 2013): first, the increasing availability of data, owing to the fact that more and more economic and social transactions take place aided by information and communication technologies which easily and inexpensively store the information such transactions produce or transmit; second, the increasing ability of firms and governments to analyze the novel big data sets. Einav and Levin (2014) ask: “But what exactly is new about [big data]? The short answer is that data is now available faster, has greater coverage and scope, and includes new types of observations and measurements that previously were not available.”

In this paper we attempt to better understand *data-driven markets*: markets where the cost of quality production is decreasing in the amount of machine-generated data about user preferences or characteristics (henceforth: *user information*), which is an inseparable byproduct of using services offered in such markets. Given that it has been documented that some data-driven markets are characterized by imperfect competition and subject to indirect network effects,² we ask under which conditions a duopoly can be a stable market structure in a data-driven market, and when the propensity to market tipping, that is, to monopolization becomes overpowering. We also study under which conditions and how a dominant firm in one data-driven market can leverage its position to another market—including traditional markets that were not data-driven before its entry.

We construct and analyze a dynamic model of R&D competition, where duopolistic competitors repeatedly choose their rates of innovation. The important feature of the model is that it incorporates *data-driven indirect network effects* that arise on the supply side of a market, via decreasing marginal costs of innovation, but are driven by user demand.³ Demand for

¹One example is “ ‘ambient computing’—a future in which robotic assistants are always on hand to answer questions, take notes, take orders or otherwise function as auxiliary brains to whom you might offload many of your chores” (The New York Times, 2015b). A second example are ubiquitous sensors such that all kinds of physical objects can be monitored in real-time, for instance, automobiles collecting information about the car’s and the driver’s performance while driving (which is compared to other drivers), or jet engines (to predict when maintenance is needed) (The New York Times, 2012). See Einav and Levin (2014) for more examples.

²Argenton and Prüfer (2012) produce evidence and theory supporting these characteristics in the search engine industry. Edelman (2015) reports about competitive dynamics in transportation networks. The markets for online search, digital maps, online social networks, or video platforms are all highly concentrated. See <https://www.datanyze.com/market-share/maps/> and <http://www.dreamgrow.com/top-10-social-networking-sites-by-market-share-of-visits-august-2016/>, respectively.

³Data-driven indirect network effects are fundamentally different from *direct network effects*, where con-

the services of one provider generates user information as a costless by-product, which Zuboff (2016) calls “behavioral surplus.” It is private information of the provider who collected it and can be used to adapt the product better to users’ preferences, thereby increasing perceived quality in the future. Thus, higher initial demand reduces the marginal cost of innovation: it makes it cheaper to produce one additional unit of product or service quality, as perceived by users.

We show that, for almost all initial quality differences, the market will eventually tip and one firm will dominate the market. Moreover, we show that such dominance is persistent, in the sense that, once the market has tipped, the weaker firm will never acquire more than a negligible market share in the future. The market is even tipping if it requires continuous, small investments in innovation to keep consumers’ perceived quality constant, which appears to be a reasonable description of dynamic, high-tech markets. Our main result is robust to changes in the time horizon, that is, whether competitors determine innovation investments using a finite time horizon T , for T high, or whether they play a game with an infinite time horizon. We identify a strong *first-mover advantage* in data-driven markets, which leads towards monopolization and is built upon data-driven indirect network effects.

An important feature of a tipped market is that there are very little incentives for both the dominant firm and the ousted firm to further invest in innovation. The reason is that, in the stable steady state where one firm has virtually no demand and the other firm has virtually full demand, the ousted firm knows that the dominant firm offers consumers both a significantly higher quality level and has significantly lower marginal costs of innovation, due to its larger stock of user information. The latter characteristic enables the dominant firm to match any innovative activities of the ousted firm at lower marginal innovation cost and hence keep its quality advantage. As demand follows quality differences in our model, the smaller firm gives up innovating if its quality lags behind the larger firm’s too much. Knowing this, the dominant firm’s best response is to also save on investing in innovation—and still reap the monopoly profit.

Going a step further, we study under which circumstances a dominant position in one data-driven market could be used to gain a dominant position in another market that is (initially) not data-driven. We show that, if market entry costs are not prohibitive, a firm

sumption utility of one consumer increases in the amount of other consumers on the same network and which are, hence, completely demand-driven, for instance, in telecoms (Besen and Farrell (1994), Economides (1996), Shapiro and Varian (1999)). They are also different from *dynamic economies of scale* (or learning-curve effects), which are completely supply-driven, for instance in aircraft manufacturing (Benkard, 2000). In contrast to these mechanisms, data-driven indirect network effects cannot easily be copied by competitors or destroyed by the arrival of a new technology. To complicate matters empirically, these different effects can be overlapping in practice. For instance, online social networks are characterized both by direct and by indirect network effects.

that manages to find a “data-driven” business model, can dominate any market in the long term. We then introduce the concept of *connected markets*, which captures situations where user information gained in one market is a valuable input to improve one’s perceived product quality in another market. We show that user information in connected markets is two-way complementary, such that incentives to acquire user information in one market can justify market entry in another market, and vice versa.

Consequently, if technology firms realize that user information constitutes a key input into the production of quality in data-driven markets, they need to identify other markets where these data can be used as well. In those connected markets, the same results as in the initial markets apply, suggesting a *domino effect*: a first mover in market A can leverage its dominant position, which comes with an advantage on user information, to let connected market B tip, too, even if market B is already served by traditional incumbent firms.

We also study the normative implications of our results. Because a tipped market provides low incentives for firms to innovate further, market tipping may be negative for consumers.⁴ It also deters market entry of new firms, even if they may develop a revolutionary technology. Therefore, we analyze the effects of a specific market intervention that was recently proposed: what if firms with data-driven business models have to share their (anonymized) data about user preferences or characteristics with their competitors?

We show that a dominant firm’s incentives to innovate further do *not* decline after such forced sharing of user information, even in a dynamic model.⁵ Instead, we show that data sharing (voluntary, or not) eliminates the mechanism causing data-driven markets to tip. The intuition is that the key assumption that lead to market tipping in our baseline model—more demand today leads to lower marginal cost of innovation and, hence, to higher equilibrium quality tomorrow—depends on a data collector’s exclusive proprietorship of user information. With mandatory data sharing, both competitors face the same cost function; a firm with initially higher demand does not have a cost advantage in producing quality. As a result, the sharing of user information avoids the negative consequences for innovation that are specific to data-driven markets.⁶ The net welfare effects are ambiguous, though. We show that data sharing is likely to increase welfare if indirect networks effects are sufficiently pronounced or if firms have similar initial market shares or if the market is close to being monopolized already.

⁴From an antitrust perspective, the underprovision of innovation is our *theory of harm* in data-driven markets.

⁵Argenton and Prüfer (2012) suggested a related policy for the search engine market. But their paper only contained a static model, which was hard to interpret over time, a shortcoming addressed by the dynamic model in our paper.

⁶Of course, these markets can still be dominated by one or a few firms, just as any other market. But in that case, we could be more confident that the source of dominance is a continuous, fundamentally superior quality-price ratio and not a windfall innovation-cost reduction from earlier success in the market.

The model offers a rationale why some firms in data-driven markets are highly successful while their competitors fail, and precisely which type of data are crucial to compete in such markets.⁷ Our model can be used to rationalize strategies of firms like Alphabet/Google, which first tipped the search engine market, our most prominent example of a data-driven market. Today, however, Alphabet “has started to look like a conglomerate, with interests in areas such as cars, health care, finance and space” (The Economist, 2016). Our model can also identify the characteristics of industries that may be prone to entry of *data-driven firms*, which has wide-ranging implications for suppliers, buyers, antitrust and regulation authorities in many industries, including some traditional sectors that are not thought of as data-driven today.

In the next section, we present our baseline model and discuss its main assumptions. Section 3 analyzes subgame-perfect Nash equilibria of the model with a finite time horizon T (as well as the limit case $T \rightarrow \infty$). In Section 4, we analyze the incentives and consequences of market entry of a data-driven firm into a traditional market and develop the notions of connected markets and the domino effect. The effects of data sharing among competitors are investigated in Section 5. We analyze robustness and study an extension of the model in Section 6: for the case where the time horizon is infinite, we solve for Markov equilibria and show that the results are also robust if perceived product quality declines exogenously over time. Section 7 connects our model to the literature and discusses applications. Section 8 concludes. All proofs are in the Appendix.

2. The Model

2.1. The Baseline Model: Competing with Big Data

There is a unit mass of consumers each demanding one unit of a good in each period $t \in \{1, 2, \dots, T\}$. Consumers face duopolistic producers $i \in \{1, 2\}$ and value product quality $q_i \geq 0$. The firms’ quality difference is denoted by $\Delta = q_1 - q_2$, such that demand for firm i ’s product in period t is realized as follows:

$$D_1(\Delta) = \begin{cases} \frac{1+\Delta}{2} & \text{if } \Delta \in [-1, 1] \\ 1 & \text{if } \Delta > 1 \\ 0 & \text{if } \Delta < -1 \end{cases} \quad D_2(\Delta) = \begin{cases} \frac{1-\Delta}{2} & \text{if } \Delta \in [-1, 1] \\ 0 & \text{if } \Delta > 1 \\ 1 & \text{if } \Delta < -1. \end{cases} \quad (1)$$

⁷For instance, in February 2016, Alphabet briefly became the world’s most valuable company by market capitalization. Its main competitors in the search engine market, Yahoo and Bing (Microsoft), are reported to have significant troubles, however (The Economist, 2016). Uber was recently named “the world’s most highly valued private startup” (Wall Street Journal, 2015).

We consider goods where consumption or the usage of a service in period t reveals some information about the consumer's preferences or characteristics and where this information can be easily logged, such as in search engines, platform sites for accommodation or car sharing, or digital maps. We call such information *user information*, which grows linearly in D_i and can be stored by the seller automatically and for free. User information is an input into a firm's efforts to improve its perceived product quality and therefore reduces firm i 's cost of innovation. It constitutes *data-driven indirect network effects* in this model.

Firms repeatedly set innovation levels $x_{i,t} \geq 0$, such that firm i 's perceived product quality in period t increases by $x_{i,t} = q_{i,t} - q_{i,t-1}$.⁸ A firm that invests in order to increase its quality by x units has to bear the following investment costs in the period of investment: $c(x, D_i) = \gamma x^2/2 + \alpha x(1 - D_i(\Delta))$. $D_i(\Delta)$ is the demand the firm had in the previous period and $\gamma > 0, \alpha \in [0, 1)$ are parameters measuring the difficulty to innovate (γ) and the importance of data-driven indirect network effects (α). We assume $\alpha < 1$ to rule out excessively expensive investment, i.e. the marginal costs of innovation should not be prohibitively high, to make the game interesting.⁹ To avoid messy case distinction, we also assume $\gamma > 1/4$ which limits the size of the investment. In particular, this assumption implies that in a one shot game the optimal investment is less than 2, that is, one period investments will not change the market share from 0% to 100%.

In period 1, we assume some starting value Δ_0 and the respective cost functions of firms 1 and 2. Hence, period 1 should not be thought of as the birth of the industry but the first period of observation. Since we employ subgame-perfect Nash equilibrium and therefore backwards induction as solution concept, actions in prior periods will not change the solution.

The functional form of $c(x, D_i)$ implies that costs are increasing and convex in the rate of innovation and are lower for the firm with the bigger market share in the previous period. Fixed costs of quality do not depend on D_i and are, just as the marginal cost of producing the good or service, apart from inventing it, assumed to be zero.

Now we can define the central concept of this paper, data-driven markets.

Definition 1. (*Data-driven markets*) *A data-driven market is a market characterized by indirect network effects driven by machine-generated data about user preferences or characteristics, s.t. the marginal costs of innovating, $c(x, D_i)$, are decreasing in demand: $c_{x, D_i} < 0$.*

In each period, only one of the two firms can invest in innovation in order to increase its quality, and then demand realizes. In odd periods, firm 1 can invest, whereas in even periods, firm 2 can invest.

⁸For better exposition, we will drop subscripts wherever there is no danger of confusion.

⁹The specific value $\alpha < 1$ implies that in the final period T there are values of Δ_{T-1} (around 0) where both firms would like to invest a positive amount (if they were the one investing in T).

We assume that a firm’s revenue is proportional to its demand and, for notational simplicity, we assume that revenue equals demand. For example, each user can be shown an advertisement and the expected revenue generated by an ad is normalized to 1.

Firm i maximizes its sum of discounted profits, where per-period profits equal demand in periods where firm i cannot invest and equal demand minus costs in periods where firm i can invest. The common discount factor is $\delta \in [0, 1)$. All choices are perfectly observable by the players. We solve this model for subgame-perfect Nash equilibria (SPNEs).

2.2. Discussion of Model Assumptions

Before we analyze the model, we briefly discuss some key assumptions.

Data driven indirect network effects: In the context of search engines, our most prominent example of a data-driven market, the comprehensive computer science literature review in Argenton and Prüfer (2012) signifies the long-documented importance of large amounts of search log/query log data for producing search engine quality. Preston McAfee, the Chief Economist of Microsoft, recently confirmed that the quality of search page results improves in the amount of “rare” queries it receives (where “rare” means the specific search string is unique in the year—but more than 50 % of all queries are rare) (McAfee et al., 2015). Therefore, the search engine with more demand improves its quality faster. He also confirms that a search engine with more demand/scale acquires data on new queries more quickly and that it has more other data to make inferences about users’ queries. He concludes (p.34): “Even at web scale, more data makes search better.”

Demand: Demand, as in equation 1, can be micro founded by a simple Hotelling model. There is a continuum of consumers of mass 1 distributed uniformly between -1 and 1. Firm 1 (firm 2) is located at point -1 (1). A consumer located at z has utility $v + q_1 - (1 + z)/2$ when using firm 1 and a utility of $v + q_2 - (1 - z)/2$ when using firm 2, where $v > 1$ and $q_i \geq 0$ is the quality of firm i . Not consuming gives zero utility and is therefore strictly dominated. Solving for the indifferent consumer yields the demand in (1).

This setup implies that a firm has a monopoly over access to its users, that is, we assume *single homing* of users. This assumption is reasonable because, for instance, a user’s specific query to a search engine or a car sharing platform is only conducted at this search engine/platform. Hence, only this firm learns something about this user’s preferences or characteristics. Over time, the firms generate different data sets about user information.

Revenues: Our assumptions can be interpreted as setting the nominal price to use a firm’s services for consumers to zero but to charge fees to (unmodeled) sellers for access to (targeted)

consumers, for instance, via advertising. Consider the search engine market, online social networks, or other platform markets. Consumers are usually matched “for free” with sellers or advertisers, who pay a fee to the platform (and naturally have to recoup these expenses from consumers).

Notably, some data-driven markets are also *two-sided* implying that players on each market side value the number of players on the other side and constituting cross-market side network effects. But while an extensive literature has studied competition in two-sided markets (e.g. Armstrong (2006), Rochet and Tirole (2006), Hagiu and Jullien (2011)), the importance of data-driven indirect network effects as studied in our paper has gone rather unnoticed. Moreover, many data-driven markets are not two-sided, for instance, maps. Others are semi-two-sided, in the sense that advertising sellers care about the number of consumers but consumers do not care too much about the number of sellers. This holds in particular for markets where ads are only a secondary product, from consumers’ point of view, such as in internet search or social networking.¹⁰

In such a setting, it is straightforward to provide a micro foundation of our revenue assumption: There is a mass of advertisers who can earn profit ν by selling to a consumer. In each period, every consumer is receptive to one ad. That is, for each consumer there is one matching seller such that the consumer will buy the product of the seller if, and only if, the platform he uses displays the matching seller’s ad. A platform can identify the matching seller of a user with probability $\eta > 0$. Assume that platforms charge a price per successfully matched consumer (this pricing scheme is usually called “price per click” in online markets). The platform – being a monopolist over its users on the advertising side of the market – will then charge a price of ν . The sellers will multihome, that is, they are willing to advertise on both platforms, and the revenue of a platform is $\eta * \nu * D_i$ and therefore proportional to its demand, as assumed in the model.¹¹

Alternating moves: The game structure with only one firm being able to act in each period has a long tradition in repeated oligopoly interaction (e.g. in Cyert and de Groot (1970), or Maskin and Tirole (1988)). To appreciate its use, imagine the same model as described above with *both* firms investing every period. It is straightforward to show that even in the static version of this game no pure-strategy equilibria exist.

¹⁰ See the controversial discussion in the literature whether being exposed to (targeted) online advertising brings consumers positive, negative, or zero utility. Edelman et al. (2007), Chen and He (2011), Athey and Ellison (2011), Goldfarb and Tucker (2011), or Taylor (2013) provide more details.

¹¹ Alternatively, if even the price per click, not only the entire bill, were increasing in the precision of the ad, which can be proxied by D_i , the gains from market tipping would be even more pronounced. This would speed up the tipping process shown in section 3. Hence, our assumption of linearly increasing expected revenues in demand is conservative.

The reason is that firms might exit the market: Suppose Δ is close to 1 and $T = 1$. If firm 1 knew that firm 2 was investing zero, firm 1's best response would be to invest just enough to capture the whole market, that is, to set $x = 1 - \Delta$, given that costs are not prohibitive. In this case, however, it is a best response for firm 2 to invest a small amount in x in order to stay in the market, as long as $\alpha < 1/2$, i.e. if marginal costs of investment are not prohibitive.

If firm 2 invests a positive amount, then firm 1 will best respond by investing an even higher amount (note that firm 1's marginal innovation costs are much lower because Δ is close to 1) in order to push firm 2 out of the market anyway. But in this case it is a best response for firm 2 not to invest at all.

Hence, a matching pennies-like situation has emerged: for Δ close to 1, equilibrium strategies have to be mixed. By assuming alternating moves, we can focus on pure-strategy equilibria, which are simpler and more intuitive.

Time horizon: A finite time horizon T can be motivated by managers having either fixed-term contracts or planning to retire at a certain age. More importantly, we think of the finite time horizon more as a technical assumption. We are indeed mainly interested in the case where the time horizon is long, that is, in the limit as $T \rightarrow \infty$.

It is well known that the equilibrium sets in games with an infinite time horizon are large. When we analyze equilibria of games with a finite time horizon T and then take the limit as $T \rightarrow \infty$, we can approximate some equilibria of the game with an infinite time horizon but clearly not all. For instance, Fudenberg and Levine (1983) show that the equilibrium set of the infinitely repeated game is the set of limits of ε -equilibria of finitely repeated games as the number of periods approaches infinity and ε approaches zero. We focus on subgame-perfect equilibria instead of ε -equilibria. Note that, for a finite time horizon, our model has an essentially unique subgame-perfect Nash equilibrium for generic parameter values.¹²

The equilibria that we select have the advantage of (i) satisfying subgame perfection, (ii) being relatively tractable and (iii) retaining intuitive properties of the finitely repeated game. We show some properties that hold for all stationary Markov equilibria of the game with an infinite time horizon in Section 6.2.

3. Analysis

We are interested in the development of market structures over time as a function of initial (exogenous) differences in firms' qualities. Therefore, our central question is, under which

¹²Essential uniqueness means that a firm has a unique optimal investment x in period t for almost all quality differences Δ_{t-1} ; i.e. different equilibria differ only in actions on a negligible set of quality differences.

conditions will a market characterized by data-driven indirect network effects (not) tip?

Definition 2. (Market Tipping): A market is weakly tipping if one firm obtains full demand in some future period (and the other firm does not). A weakly tipping market is strongly tipping if from some period $t' < T - 1$ onwards one firm has full demand in every period in which it invests. A market is absolutely tipping if from some period $t' < T$ onwards one firm has full demand in all following periods.

3.1. Period T

To start with, consider the problem in the final period T and assume that T is *even*, which implies that firm 2 can invest. Firm 2 is faced with a situation where the quality difference after period $T - 1$, Δ_{T-1} , is given and each unit of x it innovates increases q_2 and hence decreases Δ . Firm 2's maximization problem is:

$$\max_{x \geq 0} D_2(\Delta_{T-1} - x) - \gamma x^2/2 - \alpha x(1 - D_2(\Delta_{T-1})). \quad (2)$$

The solution to this maximization problem is:

$$x^T = \begin{cases} 0 & \text{if } \Delta_{T-1} \leq -1 \\ 1 + \Delta_{T-1} & \text{if } -1 < \Delta_{T-1} < -\frac{2\gamma-1+\alpha}{2\gamma+\alpha} \\ \frac{1}{2\gamma} - \frac{\alpha}{\gamma} (1 - D_2(\Delta_{T-1})) & \text{if } \Delta_{T-1} \in \left[-\frac{2\gamma-1+\alpha}{2\gamma+\alpha}, U_\alpha\right] \\ 0 & \text{if } \Delta_{T-1} > U_\alpha \end{cases} \quad (3)$$

where

$$U_\alpha = \begin{cases} 1/\alpha - 1 & \text{if } \alpha \geq 1/2 \\ 1 + (1 - 4\alpha(1 - \alpha))/(4\gamma) & \text{if } \alpha < 1/2. \end{cases}$$

Zero investment and therefore $\Delta_T = \Delta_{T-1}$ emerges if either firm 2 has already grabbed the market ($\Delta_{T-1} \leq -1$) or if investment is prohibitively expensive, which can be the case if both α and Δ_{T-1} are high. The second case leads to $\Delta_T = -1$, that is, firm 2 grabs the complete market in period T . The third case in (3) corresponds to $\Delta_T = \Delta_{T-1} - x$, being interior.¹³

Note that firm 2's investment if Δ_T is interior is decreasing in Δ_{T-1} : A higher Δ_{T-1} implies lower $D_2(\Delta_{T-1})$ (lower market share) and therefore higher (marginal) investment costs due to data-driven indirect network effects.

¹³Note that by $\gamma > 1/4$ and $\alpha \in [0, 1)$, we have $-1 < -(2\gamma - 1 + \alpha)/(2\gamma + \alpha) < U_\alpha$, that is, the case distinction in (3) is well defined.

The optimal investment x^T leads to the following profits for both competitors:

$$V_1^T(\Delta_{T-1}) = \begin{cases} \frac{2\gamma+\alpha-1+(2\gamma+\alpha)\Delta_{T-1}}{4\gamma} & \text{if } \Delta_{T-1} \in \left[-\frac{2\gamma-1+\alpha}{2\gamma+\alpha}, U_\alpha\right] \\ 0 & \text{if } \Delta_{T-1} < -\frac{2\gamma-1+\alpha}{2\gamma+\alpha} \\ D_1(\Delta_{T-1}) & \text{else} \end{cases} \quad (4)$$

$$V_2^T(\Delta_{T-1}) = \begin{cases} \frac{4\gamma+1-2\alpha+\alpha^2}{8\gamma} - \frac{2\gamma+\alpha-\alpha^2}{4\gamma}\Delta_{T-1} + \frac{\alpha^2}{8\gamma}\Delta_{T-1}^2 & \text{if } \Delta_{T-1} \in \left[-\frac{2\gamma-1+\alpha}{2\gamma+\alpha}, U_\alpha\right] \\ \frac{2-\alpha-\gamma}{2} - (\alpha+\gamma)\Delta_{T-1} - \frac{\alpha+\gamma}{2}\Delta_{T-1}^2 & \text{if } -1 < \Delta_{T-1} < -\frac{2\gamma-1+\alpha}{2\gamma+\alpha} \\ D_2(\Delta_{T-1}) & \text{else.} \end{cases} \quad (5)$$

These value functions have two noteworthy characteristics. First, V_1 is increasing while V_2 is decreasing in Δ_{T-1} . This is straightforward: a producer benefits from having higher prior quality. Second, the value functions are piecewise quadratic (or linear). This is an implication of the linear quadratic setup we chose and will simplify the following analysis.

3.2. Period $t < T$

Consider a generic period $t < T$. If t is *odd*, firm 1 can invest and solves the following maximization problem:

$$\max_{x \geq 0} D_1(\Delta_{t-1} + x) - \gamma x^2/2 - \alpha x(1 - D_1(\Delta_{t-1})) + \delta V_1^{t+1}(\Delta_{t-1} + x). \quad (6)$$

The first-order condition (for interior Δ_t and Δ_{t-1} at points of differentiability of V_1^{t+1}) is:

$$\frac{1}{2} - \gamma x - \alpha \frac{1 - \Delta_{t-1}}{2} + \delta V_1^{t+1}'(\Delta_{t-1} + x) = 0. \quad (7)$$

By contrast, if t is *even*, firm 2 invests. The resulting first-order condition is:

$$\frac{1}{2} - \gamma x - \alpha \frac{1 + \Delta_{t-1}}{2} - \delta V_2^{t+1}'(\Delta_{t-1} - x) = 0. \quad (8)$$

We first show an intuitive monotonicity result.

Lemma 1. (Quality monotonicity) (i) $V_1^t(\Delta_{t-1})$ is increasing and $V_2^t(\Delta_{t-1})$ is decreasing in Δ_{t-1} . (ii) Δ_t is increasing in Δ_{t-1} .

In all periods t , firm 1 benefits from higher Δ_t and firm 2 benefits from lower Δ_t . Furthermore, a higher Δ_t leads to a higher Δ_{t+1} . This second result is rather powerful as it implies that an increase in the initial quality difference will lead to a higher quality difference in all following periods.

Let I^t be the set of quality differences Δ_t for which in *all* following periods (up to T) the equilibrium quality difference is in $(-1, 1)$, that is, no firm has full demand in any period; the quality difference is “interior.” Put differently, as long as the quality difference is in I^t , the market remains competitive in all following periods. Hence, the market does not tip.

Lemma 1 implies the following.

Lemma 2. I^t is an interval.

We will now show an important technical property of value functions and quality investment in the interval I^t .

Lemma 3. (Monotonic dynamic quality investment incentives) Assume that the equilibrium investment is strictly positive in all periods if $\Delta_t \in I^t$. Restricted to the interior of I^t , (i) V_i^t is quadratic and convex; (ii) in odd periods, firm 1's investment is linearly and strictly increasing in Δ_{t-1} ; (iii) in even periods, firm 2's investment is linearly and strictly decreasing in Δ_{t-1} .

Lemma 3 states that, for quality differences in I^t , firm i will invest more if it had more demand in the previous period. This means that equilibrium forces do not destroy the basic cost advantage generated by indirect network effects. Instead, having relatively higher quality than a competitor incentivizes a firm to invest even more heavily in the future.

Knowing the incentives for firms to innovate over time, as long as competition is persistent, we now need to study the effects on dynamic equilibrium quality differences and demand.

Proposition 1. (Market tipping for $T \rightarrow \infty$) The length of I^t is less than $\frac{2}{(1+\alpha/(2\gamma))^{t^{T-t}/2}}$. Consequently, the length of I^0 shrinks to zero at exponential speed for $T \rightarrow \infty$.

Proposition 1 shows that the market will weakly tip if the time horizon T is sufficiently long. One firm will have full demand in some periods for almost any initial quality difference if only T is large enough. The idea behind the result is to show that firm 1's (2's) investment is increasing (decreasing) in Δ_{t-1} on I^t . Intuitively, this is not surprising as a higher (lower) Δ_{t-1} implies more user data and lower marginal costs of investment for firm 1 (2).

Now take some $\tilde{\Delta} \in I^{T-2}$. If $\Delta_{T-2} = \tilde{\Delta}$, then investment by firm 1 in $T-1$ will lead to some $\tilde{\Delta}_{T-1}$ and investment by firm 2 in T will lead to some $\tilde{\Delta}_T$. As $\tilde{\Delta} \in I^{T-2}$, no firm will have full demand in later periods and therefore $\tilde{\Delta}_{T-1} \in (-1, 1)$ and $\tilde{\Delta}_T \in (-1, 1)$. Now suppose $\tilde{\Delta} + \varepsilon$ is also in I^{T-2} for some $\varepsilon > 0$. Firm 1's investment in $T-1$ is increasing in Δ_{T-2} , say with slope $s_1 > 0$. Hence, Δ_{T-1} will now be $\tilde{\Delta}_{T-1} + \varepsilon(1 + s_1)$. Firm 2's investment is decreasing in Δ_{T-1} , say with slope $-s_2 < 0$, and therefore Δ_T will be $\tilde{\Delta}_T + \varepsilon(1 + s_1)(1 + s_2)$. Since $\tilde{\Delta} + \varepsilon \in I^{T-2}$, we must have $\tilde{\Delta}_T + \varepsilon(1 + s_1)(1 + s_2) \in (-1, 1)$. How big could ε possibly be? If $\varepsilon > 2/((1 + s_1)(1 + s_2))$, then Δ_T would be more than 2 higher when $\Delta_{T-2} = \tilde{\Delta} + \varepsilon$ compared to when $\Delta_{T-2} = \tilde{\Delta}$. But this is impossible as in both cases $\Delta_T \in (-1, 1)$. Hence, we can conclude that any two points in I^{T-2} are at most $2/((1 + s_1)(1 + s_2))$ apart. That is, the length of I^{T-2} is at most $2/((1 + s_1)(1 + s_2))$. If we could find a lower bound $\underline{s} > 0$ for $s_1 + s_2$ which holds at all points in I^t (for any t), we could argue by backwards induction that the length of I^{T-2n} is less than $2/(1 + \underline{s})^n$ and consequently the length of I^0 is arbitrarily

small when T is sufficiently large. Our linear quadratic setup allows us to find such an \underline{s} easily. However, the same argument applies for any setup where the slope of the investment function – and therefore the effect of indirect network effects on investment – is bounded away from zero on all I^t .

The result of Proposition 1, however, is not entirely satisfactory. We know that some producer will acquire full demand in some period—but what will happen thereafter? Will this firm remain dominant in the following periods or could its competitor turn the market around and have full demand itself in some later period?

Lemma 4. (*Persistent dominance*) (i) *If firm i has full demand in period $t < T - 1$, then firm i will have again full demand in a later period and firm $j \neq i$ will not have full demand in any following period.* (ii) *Take a stationary equilibrium of the game with an infinite time horizon that is the limit of a subgame-perfect equilibrium of the finite-length game, as $T \rightarrow \infty$.¹⁴ If firm i has full demand in period t , then firm i will have full demand in all periods $t + 2n$ for $n \in \mathbb{N}$. Furthermore, firm j will have less demand in all consecutive periods than in $t - 1$.*

Lemma 4 shows that a firm will remain dominant once it has become dominant in the following sense: If firm i has full demand in one period, then firm i will have full demand in more periods afterwards while firm j will never have full demand. Combining Proposition 1 with Lemma 4 implies that in the game with an infinite time horizon one firm will eventually dominate the market by having full demand (at least) every second period while the other firm does not have full demand, i.e. the market is strongly tipping.

Corollary 1. (*Market tipping for $T = \infty$*) *Take an equilibrium of the game with an infinite time horizon that is the limit of subgame perfect equilibria of the game with finite time horizon T as $T \rightarrow \infty$. In this equilibrium, the market is strongly tipping for almost all initial quality differences Δ_0 .*

The situation where firm 2 still has some demand in every second period while firm 1 has full demand in every other period is depicted in the *left* panel of Figure 1. Let us note that innovation stabilizes at a very low level. The intuition is that firm 2 is only motivated to innovate by its profits in the period in which it invests (and not by effects on future profits) because $\delta = 0$ in the left panel. Furthermore, firm 2 does not want to invest a lot because the marginal costs of innovation are rather high due to its low demand. Firm 1 simply undoes firm 2's (low) investment each period and obtains monopoly profits.

If firm 1 cares sufficiently about the future, that is, if δ is not too low, then it will usually be more profitable to push firm 2 completely out of the market, i.e. the market tips absolutely;

¹⁴We show in the appendix that such an equilibrium exists, see section A.2.




Figure 1: Strong tipping (left) and absolute tipping (right). Parameters: $T = 30$, $\alpha = 0.4$, $\gamma = 1$, $\Delta_0 = 0$, $\delta = 0$ (left), $\delta = 0.25$ (right)

see the right panel of Figure 1, where we assume $\delta = 0.25$. Firm 1 will increase q_1 so far that Δ grows above 1 and firm 2 finds it unprofitable to fight back. As soon as this is achieved, however, firm 1 can stop investing forever and enjoy monopoly profits in all remaining periods. The user data that firm 1 then gathers as a monopolist (for free) is not used for innovation but as a barrier to entry. Firm 2 will not try to get back into the market because in this case firm 1 would start using its superior data to immediately push firm 2 out again.

Which of the two scenarios in figure 1 occurs in equilibrium depends on parameter values. The following Lemma gives a clear cut answer in case the parameter α – which represents data-driven indirect network effects – is sufficiently large.

Lemma 5. (*Absolute tipping for high α*) *Let $\alpha \geq 1/2$ and T finite. Then every weakly tipping market is absolutely tipping.*

The reason is that a high α implies high marginal costs of investment (even at zero investment) for a firm with zero market share. Consequently, it is no longer profitable to “invest back” after one’s market share has dropped to zero. In particular, for $\alpha \geq 1/2$, marginal costs are higher than marginal revenue in the current period (which is $1/2$ for our demand function). This implies that investment is not profitable this period and – as the other firm will invest enough next period to again grab the whole market – the investment also does not pay off in the future.

4. Connecting Markets: Entering a traditional market with a data-based business model

4.1. Market entry when the time horizon is finite

Consider a market, where a representative incumbent firm operates with traditional investment methods that do not exploit data-driven indirect network effects. We study the strategic situation that arises when a potential entrant, who uses a data-based business model and, hence, is harvesting indirect network effects (henceforth a *data-driven firm*) contemplates to enter the traditional market. As an example, consider the market for road maps in the late 1990s—and what happened to traditional map producers after the entry of digital map providers, most notably Google. We will argue in Section 7 that this case offers a classical application of the theory of *connected markets*, which we present now.

What is important for this model is that exploiting user information creates value for the users but that such exploitation is unique to data-driven firms. Traditional firms do not have the option to create consumption value in this way, for two reasons. First, their product might not provide them with data on usage. Second, personalization might technologically be incompatible with their product. Both reasons are true, amongst many more industries, in the case of traditional paper maps.

Consider the general framework of our T -period model, just as in Section 3. But now assume that firm 2 is an incumbent firm in a market operating a traditional, that is, not data-driven business model. Firm 1 is a potential entrant employing a data-driven business model. Assume that firm 2's costs of investing are $\gamma'x^2/2 + \alpha'x/2$ (i.e. for $\alpha' = \alpha$ firm 2 has the same marginal costs of investment at $x = 0$ as firm 1 has with 50% market share). To rule out excessively expensive investment, we assume $\alpha' < 1$, as before. We also assume that prices are fixed and normalize prices such that demand equals revenue. Hence, the investment x is the only choice variable in all periods after period 1.

Firm 1, by contrast, operates under the same cost function as in the previous sections and has the additional choice in period 1 whether it wants to enter the market (and invest some x of its choosing, which will add to the initial quality difference Δ_0), or not to enter. Entering comes at a fixed cost $F \geq 0$. To simplify notation, we will assume that T is even. Clearly, markets where entry has already taken place emerge as a subgame of this model.

It is straightforward to solve for firm 2's optimal investment in period T which turns out to be:

$$x_2^T(\Delta_{T-1}) = \begin{cases} \frac{1-\alpha'}{2\gamma'} & \text{if } \Delta_{T-1} \in \left[\frac{1-\alpha'-2\gamma'}{2\gamma'}, 1 + \frac{(1-\alpha')^2}{4\gamma'} \right] \\ 1 - \Delta_{T-1} & \text{if } \Delta_{T-1} < \frac{1-\alpha'-2\gamma'}{2\gamma'} \\ 0 & \text{else.} \end{cases}$$

The resulting period T value functions are:

$$V_1^T(\Delta_{T-1}) = \begin{cases} \frac{1}{2} + \frac{1}{2}\Delta_{T-1} - \frac{1-\alpha'}{4\gamma'} & \text{if } \Delta_{T-1} \in \left[\frac{1-\alpha'-2\gamma'}{2\gamma'}, 1 + \frac{(1-\alpha')^2}{4\gamma'} \right] \\ 0 & \text{if } \Delta_{T-1} < \frac{1-\alpha'-2\gamma'}{2\gamma'} \\ 1 & \text{else.} \end{cases}$$

$$V_2^T(\Delta_{T-1}) = \begin{cases} \frac{1}{2} - \frac{1}{2}\Delta_{T-1} + \frac{(1-\alpha')^2}{8\gamma'} & \text{if } \Delta_{T-1} \in \left[\frac{1-\alpha'-2\gamma'}{2\gamma'}, 1 + \frac{(1-\alpha')^2}{4\gamma'} \right] \\ 1 - \frac{\alpha'+\gamma'}{2} + (\gamma' + \alpha'/2)\Delta_{T-1} - \frac{\gamma'}{2}\Delta_{T-1}^2 & \text{if } \Delta_{T-1} < \frac{1-\alpha'-2\gamma'}{2\gamma'} \\ 0 & \text{else.} \end{cases}$$

In this setup, Lemma 1 still holds true in case firm 1 enters because $c_{xD} \leq 0$ for both firms (with equality for firm 2). Consequently, the set of period t quality differences (I^t) for which the equilibrium quality difference in all consecutive periods is in $(-1, 1)$ is an interval, that is, Lemma 2 still holds as well. Lemma 3 is also still true. This leads to the following result:

Proposition 2. (*Tipping tendency in a traditional market*) *Let $T \rightarrow \infty$ and consider the subgame where firm 1 enters. The length of I^0 shrinks to zero and the market tips in a finite number of periods for almost all initial quality differences Δ_0 .*

Corollary 2. (*Entry and tipping in a traditional market*) *Let $T \rightarrow \infty$. (i) For F very high, say $F > \bar{F}$, firm 1 does not enter (regardless of Δ_0). (ii) There exists an $\hat{F} < \bar{F}$ such that for $F \in [\hat{F}, \bar{F}]$ the market tips in favor of firm 1 whenever (Δ_0 is such that) firm 1 enters. (iii) For very low F and not too low Δ_0 , firm 1 might enter and leave the market again in a later period.*

The interesting part of Corollary 2 is (ii), which features an all-or-nothing result (for certain parameter values):¹⁵ either the data-driven firm is deterred from entry, or not. But if it enters the traditional market, it will eventually take it over completely. The mechanism at play is the same as the one studied in the previous section. Conditional on firm 1's market entry, it does not matter in the long-run, anymore, that the product quality of the established firm may be superior (and hence $\Delta_0 < 0$). It is sufficient that firm 1 finds it worthwhile to enter the market and to invest in innovation such that it obtains some positive demand. As soon as this is achieved, the indirect network effects play into firm 1's hands because, from that point onwards, its marginal cost of innovation only decrease. As the traditional firm 2 cannot react by increasing its own quality at the same rate and for the same cost as firm 1, it is bound to lose market share. Because the model contains no other potential disruptions, market tipping cannot be avoided anymore then.

¹⁵Part (iii) includes some cases where the entrant can enter with positive demand but her costs of investment are substantially higher than the incumbent's. It is then profitable to enter and realize profits while waiting to be kicked off the market. The case does not strike us particularly relevant.

In this case, firm 1 managed to transform the traditional market into a *data-driven market*. If another firm with a data-driven business model would show up to compete with firm 1, the model in section 2 would apply henceforth. It follows that the first data-driven firm entering a traditional market has a strong first-mover advantage.

The inevitability of market tipping after entry of firm 1 shifts our attention to the first part of Corollary 2. The long-term structure of the traditional market is decided at the point of time where firm 1 decides about its entry. The decision depends on the cost that firm 1 has to bear to create a product that consumers would accept as a (potentially imperfect, inferior) substitute to the existing products in market. This cost, F , is a function of the product characteristics expected by consumers. It also depends on legal requirements, for instance, to obtain a public authority’s approval or to acquire a license. Consumers’ expectations regarding the must-have features in this market, as is pointed at by the threshold cost level \bar{F} , depend on the actions and quality investments of the traditional incumbent firm. If firm 2 manages to innovate herself by improving consumption utility to such a degree that the market entry cost are prohibitive for firm 1, there will be no entry by firm 1. In this case, the traditional market will not be transformed into a data-driven market.

Before we think one step further, let us define the notion of connected markets.

Definition 3. (*Connected markets*) *Markets A and B are connected if $c_{x_{1,B},D_{1,A}} < 0$ or $c_{x_{1,A},D_{1,B}} < 0$.*

This definition builds on our definition of data-driven markets. Where we characterize a data-driven market as a market with machine-generated indirect network effects, the notion of connected markets focuses on the impact of user information gained in one market for the cost of innovation in another market.¹⁶ We take the connectedness of two specific markets as given. But firms can be creative in developing new business models and, thereby, exploring the degree of connectedness between two markets. Hence, if the market entry costs are not prohibitive, a firm that manages to find a “data-driven” business model can dominate most traditional markets in the long term.

The connectedness of two markets can be used in two ways. If $c_{x_{1,A},D_{1,B}} < 0$, one can imagine a situation where market A in isolation might not tip in favor of firm 1, for instance, because indirect network effects, measured by α , are not important enough. But after entering market B, the additional data gained in market B will allow firm 1 to innovate much cheaper in market A; hence market A may tip (quicker) in favor of firm 1. The additional profits generated in market A might even make entry into market B profitable where entry into B,

¹⁶For instance, think of the user information a search engine gains. A part of that information is geographic (where to drive? where to eat?). Hence, this part of the information is also a valuable input when developing a product for markets of geographic information, such as maps or route planners.

in isolation, would not have been profitable.

The second way to exploit connectedness is associated with $c_{x_{1,B},D_{1,A}} < 0$. Tipping in market A might then make entry in market B feasible. To see this, suppose that entry into market B is prohibitively expensive at the outset but market A tips in favor of firm 1. Firm 1 has therefore more data from market A, which will reduce innovation costs in market B. This might make entry into market B feasible, which will then also tip in favor of firm 1.

Based on Corollary 2 and the definition of connected markets, we obtain the following Result without proof.

Result 1. (*Domino effect*) *Say firm 1 is active in market A and identifies a connected market B where entry is not prohibitive ($F \leq \bar{F}$). Then firm 1 will enter B when it has become sufficiently dominant in A.*

This Result adds on top of Corollary 2 the idea that two characteristics are complementarily helpful in entering and dominating any traditional market: (i) finding a business model that connects a new market with one's home market, that is, to develop a service or product that makes good use of user information gained in one's original market. (ii) possessing a lot of relevant user information in one's home market. Result 1 then states that firm 1 can *leverage* its dominant position from market A to market B. This process can be repeated in markets C, D, etc., which explains the term domino effect. It also gives rise to an empirical prediction: In traditional markets, our model suggests that we observe races between *data-driven firms* to identify data-driven business models utilizing their existing data stocks and *traditional companies* trying to increase data-independent product quality.

4.2. Market entry when the time horizon is infinite

Here we test the validity of our results on the entry of firms with a data-driven business model in a traditional market when firms' horizons are infinite. Specifically, we study stationary Markov equilibria of the game with infinite time horizon that are limits of equilibria of games with finite time horizon as $T \rightarrow \infty$. Still firm 1 is assumed to use a data-driven business model and, hence, can exploit indirect network effects but firm 2 is a traditional firm, without the chance to benefit from indirect network effects.

Lemma 6. (*Market entry in the infinite game*) *Take a stationary Markov equilibrium of the game with infinite time horizon that is the limit of subgame-perfect equilibria in the game with finite time horizon T as $T \rightarrow \infty$. If firm 1 enters and $\Delta_2 > \Delta_0$, then the market will eventually tip in favor of firm 1.*

Lemma 6 states that early movements in quality or market shares are indicative of whether the market will tip. More precisely, if an entrant with a data-driven business model enters a

market and manages to gain positive market share immediately after entry, then the market will tip in her favor.

5. Policy proposal: Data sharing

Our analysis so far has shown that market tipping is a robust phenomenon on markets with data-driven indirect network effects. This result has important policy consequences.

As shown in Lemma 4 and exemplified in Figure 1, the investment incentives of both firms stabilize at a low level once a market has tipped, significantly lower than in the period where both firms are still competing for the market.¹⁷ Reduced innovation incentives of originally R&D-intensive firms are negative for consumers.

Therefore, we study the effects of a regulatory measure that was recently brought up by Argenton and Prüfer (2012) in the context of search engine markets. They propose that competing search engines should be forced to share their (anonymized) search log data (that is, user information, according to our definition) amongst each others. In their static model, this measure would avoid market tipping. Here we broaden the scope to data-driven markets in general, and we set out to analyze the effects of data sharing in our dynamic framework.¹⁸

In the context of our model, the proposal implies that both firms obtain the data of all consumers when innovating. We will show that the forces that lead to market tipping in the earlier sections of this paper are no longer present after this regulatory measure is introduced.

As both producers now have access to the data of all users, the cost function of each firm is $c(x) = \gamma x^2/2$: for innovation purposes, the cost function is specified as if the firm had had full demand in the previous period. Note that this is effectively the same as our baseline model with $\alpha = 0$. In particular, Lemma 1 and Lemma 2 still hold. However, Proposition 1 no longer implies that I^0 shrinks to zero, as $T \rightarrow \infty$, because the bound on its length no longer depends on T if $\alpha = 0$. The following Lemma is similar to Lemma 3 but actually tighter.

Lemma 7. (*Interior innovation with data sharing*) *Assume data sharing. Restricted to the interior of I^t ,*

¹⁷This result is reflected by Edelman (2015), who cites the oral testimony of Yelp’s CEO before the Senate Judiciary Subcommittee on Antitrust, Competition Policy and Consumer Rights on September 21, 2011, and writes: “Google dulls the incentive to enter affected sectors. Leaders of TripAdvisor and Yelp, among others, report that they would not have started their companies had Google engaged in behaviors that later became commonplace.” The problems of TripAdvisor and Yelp can be explained by the theory of connected markets described in Section 4.

¹⁸Argenton and Prüfer (2012) also discuss the technical feasibility and potential legal avenues for implementing their proposal. They tentatively conclude that in both dimensions unsolved issues remain but that, in principle, the proposal is feasible.

- V_i^t is linear;
- firm i 's investment in period t is

$$x_i^t = \frac{1 - \delta^{T-t+1}}{2\gamma(1 - \delta)}. \quad (9)$$

Equation (9) already shows that I^t can only be non-empty if γ is not too small and δ is not too close to 1. This should not be surprising: For $\gamma \rightarrow 0$ investment is costless and, therefore, the firms will invest huge amounts in innovation that lead outside the interior range $(-1, 1)$ for Δ . The same is true for $\delta \rightarrow 1$: If every investment bears fruit forever and there is no discounting, the firms will want to invest arbitrarily large amounts. Equation (9) allows us to compute how much Δ changes over two periods (when each firm can invest once) if $\Delta_{t-1} \in I^t$:

$$x_1^t - x_2^{t+1} = \frac{\delta^{T-t} - \delta^{T-t+1}}{2\gamma(1 - \delta)} = \frac{\delta^{T-t}}{2\gamma} \quad \text{on } I^t.$$

This implies that I^t is shifting upwards over time. The reason for the shift is the alternating move structure of the model. When investing in period 1, firm 1 takes the revenues effect of its investment for all T periods into account. Firm 2 invests a period later. Therefore, its investment has a revenue effect for all but the first period (which is already over). Hence, firm 2's marginal revenue of investing is lower. In contrast to the result with $\alpha > 0$ (see Proposition 1), however, the length of I^0 does not need to shrink for $T \rightarrow \infty$. The just described timing effect is independent of the specific Δ_t in I^t . Hence, the length of I^t does not change while I^t shifts upwards over time. The reason for I^t shrinking to zero in Proposition 1 was the existence of indirect network effects, the competitive effects of which are eliminated by data sharing.

One might wonder by how much I^t shifts. Let \bar{i}^t be the upper bound of I^t . Equation (9) states that \bar{i}^t increases over two periods by $\delta^{T-t}/(2\gamma)$. Summing over $t = 1, 3, \dots, \infty$, this implies that $\bar{i}^\infty - \bar{i}^0 = \delta/(2\gamma(1 - \delta^2))$. If δ is not too high and γ not too low, then this will be a sufficiently low number that allows a non-zero length of I^0 even if $T \rightarrow \infty$.

To give a drastic example, consider $\delta = 0$: Then x^t is the same in all periods and, as long as the first-period investment does not lead to a monopoly in the first period, no firm will dominate in any period.¹⁹ In this sense, mandating data sharing will prevent tipping if (i) δ is not too high and γ not too low and (ii) the market is reasonably symmetric at the time of mandating, that is, if $|\Delta_0|$ is not too high.

While data sharing can prevent market tipping, its welfare consequences are ambiguous in our model. Welfare in a given period consists of the two firms' revenues, which sum to 1, the investing firm's cost, and consumer surplus. Here we adopt the Hotelling interpretation

¹⁹This means unless $\Delta_0 \geq 1 - 1/(2\gamma)$, both firms will be active in all periods.

of our demand structure, where consumers are uniformly distributed between -1 and 1 (see Section 2.2). Total welfare equals the discounted sum of welfare in periods 1 to T .

Welfare effects of data sharing are ambiguous because several effects interact: First, sharing data directly reduces innovation costs. In our model, costs are reduced by $\alpha x(1 - D_i(\Delta_{t-1}))$. Second, quality is higher due to lower marginal costs of innovation and – in particular in later periods – because market tipping may be prevented: Recall that the remaining firm stops innovating after the other firm exited the market. The additional innovation is beneficial for consumers. Third, more consumers may be able to buy from their preferred firm. In the Hotelling interpretation of our demand function, transportation costs of some consumers will be lower if both firms stay in the market. Fourth, higher investments (especially in later periods) imply higher costs. Note that investment costs are duplicated in case both firms stay active in the market. This last effect reduces welfare. Depending on parameter values the overall effect of data sharing on welfare can be positive or negative.




Figure 2: *Data sharing* (compared to *no data sharing*) reduces welfare in shaded regions and improves welfare in white regions. Parameters: $T = 30$, $\gamma = 2$, $\delta = 0.3$, α -grid: $0.1, 0.11, \dots, 1.0$, Δ_0 -grid: $-0.9, -0.91, \dots, 0.9$.

Figure 2 shows some numerical examples. For each (α, Δ_0) combination on the grid we computed the subgame perfect equilibrium with and without data sharing where investments have to be chosen on a fine grid (grid point distance 0.001). The shaded area are those (α, Δ_0) combinations where equilibrium welfare without data sharing is higher than with data sharing. As one would expect, data sharing can reduce welfare only if α is low (the shaded region) as this limits the beneficial first effect. For most parameter values (the white area), data sharing improved welfare. Data sharing is likely to increase welfare if firms have similar initial market share or if the market is close to being monopolized already (in the latter case the first effect

simply reduces the cost of tipping the market completely).

6. Robustness and Extensions

6.1. Decaying quality

An important critique of the notion of data-driven markets (Definition 1) is that user preferences are unstable and subject to fashion trends over time. One could therefore argue that the data about user preferences generated on a data-driven market at one point of time become more and more obsolete. In this case, the product quality perceived by users decays over time if the provider does not constantly invest in quality improvements (or innovation).

In terms of our baseline model, this means that perceived quality q decays if a firm does not innovate. Assume, for concreteness, that $q_{i,t+1} = \mu q_{i,t} + x_{i,t+1}$ where $\mu \in [0, 1]$. Hence, quality decays at rate $1 - \mu$. This implies that, without any investments, $\Delta_{t+1} = \mu \Delta_t$: the quality difference shrinks due to quality decay. Quality decay is therefore a force working *against* market tipping. While it is still true that firms with higher market share have lower marginal costs of investing, firms with higher quality are also more affected by quality decay. Whether firms with a certain quality advantage will be able to tip the market depends on the relative strength of the data-driven indirect network effects compared to the decay effect.

In this section, we show that Proposition 1 still holds as long as μ is not too small. In particular, we will show that the length of I^0 still converges to zero at exponential speed if the sufficient condition $\mu(\mu + \alpha/(2\gamma)) > 1$ holds. That is, our results hold if μ is not too small. The derivation mirrors the one in the main text and we will therefore only quickly describe those steps where the analysis changes.

In the final period T , equations (3), (4) and (5) remain valid if we write $\mu \Delta_{T-1}$ instead of Δ_{T-1} . In particular, V_i^T is linear-quadratic in Δ_{T-1} on I^T .

For $t < T$, the first-order condition on the interior of I^t (for firm 1) changes from (7) to:

$$\frac{1}{2} - \gamma x - \alpha \frac{1 - \Delta_{t-1}}{2} + \delta V_1^{t+1'}(\mu \Delta_{t-1} + x) = 0.$$

Because Lemma 1 is crucial for the remainder of the analysis, we replicate its proof in some more detail in the Appendix. The result is that Lemma 1 remains valid. Lemma 3 also still holds as including μ does not really change the linear structure of the first order condition. Using these intermediate results we obtain the following proposition.

Proposition 3. (*Market tipping with quality decay*) *Assume $\mu(\mu + \alpha/(2\gamma)) > 1$. The length of I^0 shrinks at exponential speed to zero in T .*

Quality decay affects the firm with higher quality more than the firm with lower quality

as decay is modeled proportionally to existing quality. This force counteracts the main mechanism of our paper – namely that higher quality firms have more demand and therefore more user data leading to lower investment costs. We expect to find similar results as in 3 also in other extensions that introduce a way in which higher quality leads to a disadvantage: As long as these effects are not too strong, similar results as in the main section of our paper will hold.

6.2. Infinite Time Horizon

6.2.1. Markov equilibria

To understand whether the tipping result shown in the previous section is an artifact of our equilibrium concept, looking at limits of subgame-perfect Nash equilibria of finite-time horizon games, we now analyze the model with an infinite time horizon.²⁰

Games with an infinite time horizon usually have many equilibria. In particular, there can be equilibria which are not limits of equilibria in T -times repeated games, as $T \rightarrow \infty$. A commonly used restriction, which we also apply here, is to look at *Markov equilibria*.²¹ These are equilibria in which the equilibrium strategy depends only on a “state variable” and not on the full history of the game or the specific time period. The state variable in our setting is the quality difference Δ . In this section, we derive some properties that hold *for all* Markov equilibria. The main purpose is to show that our results on market tipping in the previous section are not a special feature of the equilibrium of the T -times repeated game (as $T \rightarrow \infty$) but that market tipping, in some form, is a robust phenomenon across different equilibria of games with an infinite time horizon when there are data-driven indirect network effects.

To express our results clearly, we define the notions of *steady state* and *stability*.

Definition 4. ((Stable) Steady State): Steady State denotes a quality difference Δ such that, in a given equilibrium, $\Delta_t = \Delta$ implies $\Delta_{t'} = \Delta$, for all $t' = t + 2n$, for $n = 1, 2, \dots$. A Steady State Δ is (strictly) stable if, for some $\varepsilon > 0$, $|\Delta_{t+2} - \Delta| \leq (<)|\Delta_t - \Delta|$ for all $\Delta_t \in (\Delta - \varepsilon, \Delta + \varepsilon)$.

Proposition 4. (Tipping after threshold quality difference) Let $\alpha \geq 1/2$. In every Markov equilibrium, $\underline{\Delta} = -1$ and $\bar{\Delta} = 1$ are strictly stable steady states.

²⁰Such games are usually called “stochastic games”, see (Fudenberg and Tirole, 1991, ch. 13), which might however sound a bit odd in our deterministic setup.

²¹Strictly speaking, we focus in this section on *stationary* Markov equilibria, where strategies do not depend on the time period t . We call these equilibria “Markov equilibria” for short. Note, however, that value functions will still depend on whether a period is odd or even due to the alternating move assumption. A formal definition of Markov equilibrium can be found in Mailath and Samuelson (2006) (ch. 5.5) or in Fudenberg and Tirole (1991) (see ch. 13.1.2 for a discussion of Markov strategies in separable sequential games of perfect information and 13.2.1 for a formal definition of Markov equilibrium).

Proposition 4 considers cases where $\alpha \geq 1/2$, that is, where indirect network effects are sufficiently important in the innovation process. There, a firm with zero demand in the previous period would find zero investments optimal in a one-shot game. This parameter restriction rules out equilibria where, say, firm 1 has full demand in all odd periods but both firms take turns in investing a small amount such that firm 2 has a small, positive market share in even periods and zero market share in odd periods. (With $\alpha \geq 1/2$ firm 2 would make losses in such a situation; see the discussion after lemma 5.) The Proposition implies that a market tips whenever the quality difference between the competitors is sufficiently large.

Going a step further, we study how large the set of initial quality differences is that finally leads to market tipping. Especially, we are interested in the role of firms' discount factor, δ .

Lemma 8. *For every $\varepsilon > 0$, there exists a $\bar{\delta} > 0$ such that the market tips for all initial quality levels apart from a set of measure less than ε if $\delta < \bar{\delta}$.*

Lemma 8 states that the market tips for almost all initial quality differences if the discount factor δ is sufficiently low: If firms do not value the future too much, the market will tip. To understand the intuition of this result, say the market should tip for all initial quality differences but an interval of ε length. For any $\varepsilon > 0$, there is a discount factor $\bar{\delta}$ such that the Lemma is true whenever $\delta < \bar{\delta}$. This can be understood as a continuity property: In the one shot game, firm 1's (firm 2's) investment is increasing (decreasing) in the initial quality difference due to the indirect network effects.²² Hence, there is only a single quality difference at which the two firms' investments would be equal. For quality differences above (below) this level, firm 1 (firm 2) invests more than its rival. Consequently, the market would tip for all but this one initial quality level if myopic players repeatedly play the game. While full myopia corresponds to $\delta = 0$, Lemma 8 shows that this idea still holds approximately for low but positive discount factors.

6.2.2. Multiplicity of Markov Equilibria: Numerical Analysis with Finite State Space

The results in the previous subsection do not shed light on the market outcome if the discount factor is high and initial market shares are approximately equal. This subsection contains a numerical analysis that addresses these issues and illustrates why stronger analytical results are hard to come by: We show that there is a multiplicity of Markov equilibria and in some of these equilibria the market does not tip but reach an interior steady state. However, our tipping result is robust in the sense that the market tips if the data-driven indirect network effects are sufficiently strong. The multiplicity of Markov equilibria also shows why there

²²This monotonicity does not hold in the regions where the investing firm grabs the entire market, but then the market will obviously tip, even within a single period!

is a need for equilibrium selection among the Markov equilibria. This is a more technical motivation for our focus in earlier sections on the limit of subgame-perfect equilibria of finite-length games as the time horizon T tends to infinity, which is a relatively intuitive selection rule.

We will use a slight variation of our model. In particular, we will assume in this subsection that firms cannot invest any arbitrary amount but that only qualities on a finite grid are feasible. The reason for this change in modeling is the following: Solving for Markov equilibria in models with an infinite state space is subject to several technical problems. The usual way is to solve either via value-function or policy-function iteration. However, these methods may not converge to an equilibrium because the value function operator is not a contraction mapping in games (in contrast to single agent decision problems where it usually is). Furthermore, these methods will determine only one equilibrium (if they converge) while there might be many equilibria – possibly with very different properties and outcomes. See Iskhakov et al. (2015) and their references for a more thorough discussion of these problems.

Assuming a finite grid of qualities allows us to use the methods described in Iskhakov et al. (2015) to solve for *all* pure-strategy Markov equilibria. This is done through an algorithm similar to backwards induction. However, backward induction is carried out on the state space and not in time, as in the previous sections. If both firms have the maximally feasible quality on the finite grid, then no one can invest now or in any future period and therefore the value functions are determined by the parameters. Next, states in which one firm has the maximal quality and the other has the second-highest feasible quality are analyzed. The latter firm has two possible decisions: invest to the maximum quality, or don't investment at all. The value of the first possible decision was derived in the previous step and the latter decision leads to a steady state where again the value is determined by the parameters. Iterating further, one obtains all Markov equilibria of the game; see Iskhakov et al. (2015) for the details.

To see why there can be multiple equilibria in our setting, consider the third step, in which we analyze the state where both firms have the second-but-highest feasible quality. Assume for now that firm 1 invests if it is its turn to invest. If it is firm 2's turn to invest, firm 2 has to choose between not investing (and falling behind next period as firm 1 will invest) and investing to the maximal quality. Say, it is optimal to invest in this case. If, however, firm 1 does not invest in this state (when it is its turn), firm 2's decision problem will be different: If firm 2 does not invest, it will not fall behind and instead a steady state in the second-but-highest quality results. This is more attractive than falling behind and it might well be optimal for firm 2 not to invest in this case.

Hence, there can be two equilibria here. One in which both firms invest when it is their turn and one where neither does. At every interior point of the quality grid, these situations

might repeat and the number of Markov-perfect equilibria can grow exponentially. Table 1 shows that multiplicity is indeed a prevalent phenomenon in our setting.²³

n	2	3	4	5	6	7	8	9	10
# eq.	2	4	8	16	38	96	113	113	113

Table 1: Number of pure-strategy Markov equilibria for different grid sizes (n is number of grid points). The distance between grid points is fixed to 0.25. Parameter values: $\alpha = 0.75$, $\gamma = 1$, $\delta = 0.75$, lowest feasible quality is 0.

As the number of equilibria is rather large in any reasonably fine grid, equilibrium selection becomes necessary for any meaningful analysis or prediction. In the previous sections, we selected equilibria that were limits of subgame-perfect Nash equilibria of finitely repeated games. Here, we use two different selection methods as a robustness check, which differ in how they deal with the multiplicity described above.

In the *Invest Selection*, we resolve the multiplicity by always selecting the possibility where both firms invest. In the *Steady State Selection*, we resolve the multiplicity by always selecting the equilibrium where neither invests.

The steady-state selection arguably tries to prevent tipping by creating (interior) steady states wherever possible. In table 2, we report the long-run outcomes for each equilibrium for different strengths α of the data-driven indirect network effects and different initial quality differences. We distinguish between outcomes where one firm obtains 100% market share in the long run, where both firms have positive market share and less than maximum quality, and situations where both firms have maximum quality.²⁴

Unsurprisingly, tipping occurs more often under the invest selection. Notably, however, Table 2 displays that tipping in favor of firm 1 occurs even in some cases where the initial quality level of firm 1 is smaller than firm 2's (where $q_{1,0} < q_{2,0} = 2.5$). The reason for this result is that, in our alternating-move game, firm 1 has a first-mover advantage because it can invest in quality already in period 1, whereas firm 2 has to wait until period 2. Market tipping in favor of firm 1 in this situation implies that, on the equilibrium path, firm 1 uses its first-mover advantage and heavily invests in $q_{1,1}$ before firm 2 can react. This increases Δ and, consequently, increases firm 2's innovation cost, $c(x, D_2(\Delta))$. Note, however, that this advantage cannot explain all tipping in favor of firm 1: as the left columns of Table 2 show, if $q_{1,0}$ is sufficiently low as compared to $q_{2,0}$, the market tips in favor of firm 2.

²³The code calculating these and the following equilibria is available on the authors' websites.

²⁴The latter can occur if investment costs are very low and one could argue that the grid should be larger for these parameter values.

More importantly, *market tipping occurs under both selection rules if α is sufficiently high*. That is, our main result, that data-driven indirect network effects lead to market tipping, appears to be robust to different equilibrium selection methods. Just as above, we can also show that, in this numerical example, equilibrium innovation investments are zero once a market has tipped. Absent other changes, such as quality decay (see Section 6.1), the dominant firm permanently keeps a huge quality lead without ever having to innovate again.

	$q_{1,0} = 0.5$	$q_{1,0} = 1.0$	$q_{1,0} = 1.5$	$q_{1,0} = 2.0$	$q_{1,0} = 2.5$	$q_{1,0} = 3.0$	$q_{1,0} = 3.5$
Invest Selection							
$\alpha = 0.0$	tip2	tip2	MaxQ	MaxQ	MaxQ	MaxQ	MaxQ
$\alpha = 0.2$	tip2	tip2	MaxQ	tip1	tip1	tip1	tip1
$\alpha = 0.4$	tip2	tip2	tip1	tip1	tip1	tip1	tip1
$\alpha = 0.6$	tip2	tip2	tip1	tip1	tip1	tip1	tip1
$\alpha = 0.8$	tip2	tip2	tip1	tip1	tip1	tip1	tip1
$\alpha = 1.0$	tip2	tip2	tip1	tip1	tip1	tip1	tip1
Steady State Selection							
$\alpha = 0.0$	tip2	IntQ	IntQ	IntQ	MaxQ	MaxQ	MaxQ
$\alpha = 0.2$	tip2	IntQ	IntQ	IntQ	tip1	tip1	tip1
$\alpha = 0.4$	tip2	tip2	IntQ	IntQ	tip1	tip1	tip1
$\alpha = 0.6$	tip2	tip2	IntQ	tip1	tip1	tip1	tip1
$\alpha = 0.8$	tip2	tip2	tip2	tip1	tip1	tip1	tip1
$\alpha = 1.0$	tip2	tip2	tip2	tip1	tip1	tip1	tip1

Table 2: Long run outcomes: Given $q_{2,0} = 2.5$, the table shows the steady state outcome to which the market converges for different values of firm 1’s starting quality $q_{1,0}$ and different strengths of the data-driven indirect network effects α . “tip i ”: i has 100% market share; “MaxQ”: both firms have the maximal quality; “IntQ”: interior steady states in which both firms have positive market share. Parameters: $\gamma = 1$, $\delta = 0.75$, quality grid $\{0.0, 0.01, \dots, 5.0\}$

7. Data-driven Markets in Literature and Practice

In this section, we connect to the literature and use our model to explain some recent developments (and sketch potential future ones). For the purpose of this exercise, it is irrelevant whether the user information a data-driven firm has access to consists of aggregate or individual-level data about user preferences or characteristics. For instance, if a firm has access to the geodata locating many users’ mobile phones, in real time, it can “predict” traffic

jams or the amount of people attending a certain event. Here it is not necessary to know the identity of the individual mobile phone owners. In turn, if a firm can track an individual user over time, for instance because the user is logged in to some service or exhibits other individual characteristics (such as a specific combination of software versions on her mobile phone), then the firm can personalize its services. Both types of user information can be combined, such that an individual user can receive suggestions about where to go, whom to meet, or what to dine, which are based on the preferences of other users with similar characteristics.

7.1. Search engines

The seminal example to understand the dynamics of data-driven markets better is the market for search engines and Alphabet/Google in it, whose “mission is to organize the world’s information and make it universally accessible and useful.”²⁵ Zuboff (2016) explains: “Most people credit Google’s success to its advertising model. But the discoveries that led to Google’s rapid rise in revenue and market capitalization are only incidentally related to advertising. Google’s success derives from its ability to predict the future—specifically the future of behavior.” According to her, in the firm’s early years (since 1998), user information inherent in search logs were not structurally stored. It took until the year 2001 that Google’s management realized that this information could be used by sellers of other goods or services to identify consumers who have a high probability of buying and to make consumers tailored contract offers, which fit their individual preferences or consumption patterns.²⁶

That Google only started to systematically store and exploit user information in 2001, combined with our model of a data-driven entrant in traditional markets (Section 4), suggests an explanation for why the firm’s search engine could overtake the hitherto incumbents, AOL and Yahoo, in 2003 (measured by US market share, see Argenton and Prüfer (2012), p. 90). AOL and Yahoo had offered their users curated web entry points, where individual staff members would try to catalogue and rank websites—a “traditional” business model. By contrast, Google’s reliance on its algorithm, together with the automatic growth of its user information database over time qualified it as a data-driven firm, as defined above. In such a situation, Proposition 2 predicts the market tipping that actually occurred, not only in the US market but also in Europe and beyond.

Argenton and Prüfer (2012) introduced the idea of search log data-based indirect network effects as a crucial dimension of search engine competition in the (law and) economics lit-

²⁵<https://www.google.com/about/company/>.

²⁶Google is not the only data-driven firm to engage in fortune telling. Amazon was recently issued a patent on a novel *Method and System for Anticipatory Package Shipping* (Patent number US008615473 (December 24, 2013), <http://pdfpiw.uspto.gov/.piw?docid=08615473>). This method contains that Amazon ships products it expects customers to buy but have not bought yet, based on previous orders and other user information.

erature. They document the tipping of the search-engine market since 2003 and construct a simple model that explains a strong tendency towards monopolization, based on indirect network effects. However, their static model cannot convincingly capture dynamic incentives such as the incentives of a leading search engine to invest in R&D once it would have to share its search-log data with competitors. The paper at hand improves on this weakness.

Chiou and Tucker (2014) recently reacted to Argenton and Prüfer (2012) and studied empirically how the length of time that search engines retain their server logs affects the apparent accuracy of subsequent searches, which could be interpreted as a measure of search engine quality. For instance, one out of three changes in the data retention policy of a search engine that Chiou and Tucker (2014) study was Yahoo’s decision to anonymize its personalized user information after 90 days, in December 2008. They find no empirical evidence for a negative effect from the reduction of data retention on the accuracy of search results. This is an important finding and should be taken seriously by privacy regulators. However, it is not surprising in the light of the model studied above: such anonymization, if done properly, eliminates a search engine’s potential to identify or re-engineer a user’s identity.²⁷ But the change in Yahoo’s policy did not derogate the search engine’s aggregate amount of data on users’ clicking behavior, which is the driver for indirect network effects. See the note at the beginning of this section.

Burguet et al. (2015) set out to identify the main sources of market failure in the markets that search engines intermediate. Complementary to our approach, they focus on the reliability of the organic search results of a dominant search engine and take search engine quality as given. They show that improvements in an alternative (non-strategic) search engine induce the dominant search engine to improve search reliability, which benefits consumers: just as in our framework, more competition (in our case, via data sharing) leads to more innovation and higher quality of results. It also benefits consumers. Burguet et al. (2015) refrain from studying dynamic effects, which we do in this paper and which explains market tipping.

²⁷Chiou and Tucker (2014) complement empirical findings that “targeted advertising” techniques increase purchases (Luo et al. 2014), prices (Mikians et al. 2012), and sellers’ profits (Shiller 2013). Bringing these results together implies that search engines do not rely on information about individual users to provide high-quality results but that their advertising profits, which drive innovation incentives, depend on personalization. Dengler and Prüfer (2016) provide a microfoundation for anonymization choices of consumers with high willingness-to-pay but limited strategic sophistication to hide their identities from a seller who makes use of data-supported targeting techniques. de Cornière and de Nijs (2016) endogenize the amount of information that users disclose to firms and study the incentives of an ad-supported platform to disclose information about its users to advertisers prior to an auction. They find that disclosure can lead to higher prices even without price discrimination (due to better targeting). Taylor (2004), Acquisti and Varian (2005), Calzolari and Pavan (2006), Casadesus-Masanell and Hervas-Drane (2014), Campbell et al. (2015) and de Cornière (2016) offer more entry points to this literature. A great survey is Acquisti et al. (2016).

In Halaburda et al. (2016), two competing platforms repeatedly set prices. Consumers not only value product quality but also benefit from direct network effects. If those are strong enough, consumers may choose to buy a product with inferior quality from a “focal” platform. Halaburda et al. (2016) complements our paper in several important aspects: the focus on the pricing, not the quality decision; the reliance of direct, not data-driven indirect network effects, and the normalization of production costs to zero, as opposed to positive costs of innovation that are decreasing in a firm’s output.

Edelman (2015) underlines the opportunity of dominant firms on data-driven market to use their market power to speed up monopolization, via tying their main product with other services. He proposes “to open all ties,” that is, to allow competitors to wholly replace Google’s offerings rather than to present consumers with parallel offerings from both Google and its competitors (p.399). As the analysis above indicates, which does not assume any abusive behavior of a dominant firm, ruling out certain conduct, such as tying, is unlikely to prevent a dominant data-driven firm from completely tipping the market. The only proposal we are aware of that may be able to achieve that is data sharing (see Section 5).²⁸

This view is supported by Lianos and Motchenkova (2013), who show in a two-sided market setting that, “similar to Argenton and Prüfer (2012), the desired reduction in the asymmetry in the size of network effects can be achieved through the remedy to require search engines to share their data bases and data on previous searches” (p.451). Moreover, Lianos and Motchenkova (2013) show that a dominant monopoly platform results in higher prices and underinvestment in quality-improving innovations by a search engine relative to the social optimum. They also show that monopoly is sub-optimal in terms of harm to advertisers in the form of excessive prices, harm to users in the form of reduction in quality of search results, as well as harm to society in the form of lower innovation rates in the industry.

7.2. Connecting markets: digital maps and beyond

An epitome of connecting markets is Google’s entry and takeover of the market for road maps.²⁹ Until the 1990s, most road maps were printed on paper and served by a set of oligopolistic incumbents that offered consumers horizontally or vertically differentiated versions of maps. By the nature of printed paper, those maps were static, in the sense that

²⁸Some people may think that sharing of data with user information would allow competitors of a dominant firm to reengineer a dominant firm’s algorithm (its key resource and tool of innovation), potentially aided by machine learning. Others doubt this: “The Great Google Algorithm is not a set of ranking factors; rather, it is a collection of protocols, operating systems, applications, databases, and occasional information retrieval processes. [...] The Great Google Algorithm changes at an exponential rate” (<http://www.seo-theory.com/2011/01/07/why-you-cannot-reverse-engineer-googles-algorithm>).

²⁹As of May 2016, Google Maps had a market share above 90 percent, measured by $\frac{\# \text{ Websites using given technology}}{\text{Total } \# \text{ websites using any of the technologies}}$. See <https://www.datanyze.com/market-share/maps/>.

consumers could not customize them for a specific purpose and recustomize them for another purpose later. The digitization of maps lifted the latter restriction, such that users of online maps or downloaded maps, for instance in cars' GPS navigation systems, could zoom in and out and ask for the "best" route to a certain destination. The unique feature distinguishing Google from its competitors, when introducing Google Maps, was that a share of users' queries that the firm received in its main search engine business was geography-related. As such, Google received a huge stock of geographical information about user preferences and characteristics as a by-product of its main search engine business. Today, many features in Google Maps are fed by such data and not copyable (at the same quality level) by competitors that lack such large amounts of user information.³⁰ Zuboff (2016) comments: "Google recently announced that its maps will not only provide the route you search but will also suggest a destination." Necessarily, such a product feature would be based on the amount of relevant user information.

The case of tipping the road maps market exemplifies further markets that are likely to be entered by data-driven firms with huge stocks of user information. The theory of connected markets and the domino effect also suggest an explanation for the (at first sight) unfocused strategy of Alphabet, Google's parent company. In February 2016, *The Economist* (2016) wondered: "Today Alphabet is a giant advertising company with the potential to become a giant in other sectors as well—although exactly which ones, no one is yet sure. [...] The firm has started to look like a conglomerate, with interests in areas such as cars, health care, finance and space, as it tries to find the next big thing." This may just be the behavior of a firm in search of a suitable connected market, as described in Result 1.³¹ Ironically, it

³⁰For instance, Google Maps contains information about restaurants and bars that lists "popular times" for each day. "To determine popular times, Google uses data from users who have chosen to store their location information on Google servers" (<https://support.google.com/business/answer/6263531?hl=en>). Providing such information only adds value if the share of users "to store their location information" on a firm's servers is sufficiently large to be representative.

³¹Cars, originally a "traditional" industry, may be on the list of data-driven firms: "A high-end car, for instance, has the digital horsepower of 20 personal computers and generates 25 gigabytes of data per hour of driving. [...] Apple and Google are pressing carmakers to install the operating systems they have designed for cars' entertainment systems, which in practice will suck up all sorts of other data about the car and its occupants" (*The Economist*, 2015). Similar to the road map industry, the "trick" in the car market may be not to compete in the dimension of engineering and design quality, where incumbents' experience is decades ahead. Instead, attacking the traditional business model with a new business model, *self-driving* cars, looks like a close application of the theory of connected markets: "The car processes both map and sensor information to determine where it is in the world. [...] The software predicts what all the objects around us might do next. It predicts that the cyclist will ride by and the pedestrian will cross the street" (<https://www.google.com/selfdrivingcar/how/>). Google has the necessary geodata from its GoogleMaps and StreetView applications, among others, which gives it a headstart in the dimension of relevant user information.

is already implied by Google’s privacy policy: “We may combine personal information from one service with information, including personal information, from other Google services [...] [Y]our activity on other sites and apps may be associated with your personal information in order to improve Google’s services and the ads delivered by Google.”³²

8. Conclusion

The process of datafication is around us and progressing with staggering speed. From an economic perspective, a key feature of this process is the growing importance of data-driven indirect network effects, which combine the automatic demand-side creation of information on users’ preferences and characteristics, as a by-product of using goods and services that are connected to the internet, with a reduction in the marginal cost of innovation on the supply side. Due to this combination, unlike direct network effects, two-sided market network effects, or dynamic economies of scale, these data-driven indirect network effects cannot be easily copied by competitors or be made irrelevant by the random arrival of the next revolutionary innovation.³³

The results of our model on market tipping and connected markets suggest a race. On the one hand, technology firms with large stocks of existing data on user preferences and characteristics will be looking to identify data-driven business models utilizing these data stocks in other industries. On the other hand, traditional companies will be trying to increase data-independent product quality in order to make it prohibitively costly for those data-driven firms to enter their markets in the first place. Several firms—most clearly Google, the self-proclaimed “data company”—have apparently understood this mechanism. But other companies, notably some in manufacturing or car making, may have missed the message.³⁴ Therefore, we applied the model and exemplified the domino effect by showing that Google’s strategy to invest in many apparently unrelated markets can be rationalized by our model: these markets are either already connected (by user information driving indirect network effects in each of them) or the firm is trying to identify business models where user information from existing markets can serve as a valuable input into traditional markets.³⁵

³²<https://www.google.com/policies/privacy/>

³³See also footnote 3.

³⁴For instance, during a recent motor show, the head of production at Mercedes said amidst discussions about the future of the automobile industry, “we created the automobile, and we will not be a hardware provider to somebody else,” according to The New York Times (2015a). Our model suggests a more cautious prediction.

³⁵Google’s success suggests that the firm is good in connecting markets. In August 2016, Google was reported to have seven products with more than one billion users each (<http://fortune.com/2016/08/25/facebook-google-tech-companies-billion-users/>). Many of them benefit from access to a common pool of user information.

The policy proposal to require data sharing of anonymized user information among competitors in data-driven markets coincides with major policy initiatives of the European Commission. Beginning in May 2018, the General Data Protection Regulation will impose an obligation on firms to enable individuals to take their personal data with them when they quit using an online service.³⁶ That is, differing from the proposal studied in Section 5, data sharing will be implemented on the user side, not on the producer side, in Europe. A thorough comparison of both types of mandatory data sharing is up for future research.

Similarly, the proposal studied in Section 5 is silent about organizational and institutional issues. What type of data should be shared in which market, and precisely by whom? At which intervals? Should competitors be asked to share data bilaterally, in a network of dyads? Or should there be a third party, for instance a centralized public authority that collects and distributes the data from and among competitors? Or should such an authority be a private industry association that is run by and on behalf competitors? These are just some of the important questions, on top of a battery of legal issues, that have to be answered before policy makers could seriously consider to take action.

In more general terms, this is only the second paper, after Argenton and Prüfer (2012), that has brought the idea of data-driven indirect network effects as common denominator of data-driven markets, to the attention of academic economists. As such, it offers many avenues for future research, both theoretically and empirically. On the empirical side, the fundamental mechanism of treating demand side-generated user information as input into the supply side-run innovation process must be studied and verified in various industries. On the theoretical side, the nexus of innovation and personalization, that is, the use of past user data to improve the service not only in general but in particular for the user whose interaction generated the data, is an important topic for future research.

³⁶Regulation (EU) 2016/679 of the European Parliament and of the European Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (<http://ec.europa.eu/justice/data-protection/>).

A. Appendix: Proofs

A.1. Derivations for Period T

The optimal interior x^T in (3) follows directly from the first order condition. The main task is to derive the boundaries of the range of interior investments. This range is limited by the following conditions: First, $-1 \leq \Delta_T = \Delta_{T-1} - x^T \leq 1$. Second, $V_2^T \geq 0$ and third $x^T \geq 0$. For $\alpha \geq 1/2$, the marginal costs of investment at $\Delta_{T-1} = 1$ are higher than the marginal revenue and therefore zero investment is optimal already for some $\Delta_{t-1} \leq 1$. Clearly, $V_2^T \geq 0$ and $\Delta_T \leq 1$ in this case and the binding upper bound on Δ_{T-1} is derived from $x^T \geq 0$ which can then be rewritten as $\Delta_{T-1} \leq 1/\alpha - 1$. For $\alpha < 1/2$, there are some $\Delta_{T-1} > 1$ such that investment is still profitable and the optimal investment at these $\Delta_{T-1} > 1$ is $1/2\gamma - \alpha/\gamma > 0$ (as $D_2(\Delta_{T-1}) = 0$). Profits for firm 2 are in this case $1/2 - \Delta_{T-1}/2 + 1/(8\gamma) - \alpha/(2\gamma) + \alpha^2/(2\gamma)$ and therefore $V_2^T \geq 0$ if $\Delta_{T-1} \leq 1 + (1 - 4\alpha(1 - \alpha))/(4\gamma)$. Note that this condition is also sufficient for $\Delta_T \leq 1$ and we obtain the expression for U_α from these considerations.

The condition $\Delta_T \geq -1$ can be rewritten (using the optimal interior investment) as $\Delta_{T-1} \geq -1 + 1/2\gamma - (1 - D_2(\Delta_{T-1})\alpha)/\gamma$. Depending on whether Δ_{T-1} is below or above 1, the condition is $\Delta_{T-1} \geq -1 + 1/(2\gamma + \alpha)$ or $\Delta_{T-1} \geq -1 + (1 - 2\alpha)/(2\gamma)$. As, by the assumption $\gamma > 1/4$, the former bound is below 1, $\Delta_{T-1} \geq -1 + 1/(2\gamma + \alpha)$ is the relevant lower bound. The profits in (4) and (5) follow then simply from plugging (3) into the profit functions.

A.2. Existence of a Stationary Equilibrium that is the Limit of Subgame-perfect Nash Equilibria

We will show that there exists a stationary Markov equilibrium that is the limit of subgame-perfect Nash equilibria in games with finite time horizon as the time horizon T approaches infinity. Note that a Markov equilibrium can essentially be denoted by the value function. We will therefore concentrate on those. For every time horizon T take a subgame perfect Nash equilibrium and denote the first period value function of player i as $V_i^{1,T}$. Note that player i 's value function is bounded from below by 0 and bounded from above by $1/(1 - \delta)$ (i.e. the revenue of capturing the whole market for all times without investing anything). By Lemma 1, the value functions are monotone. Take an increasing sequence of T s and the corresponding sequence of value functions of firm 1 $(V_1^{1,T})_T$.

We will show that this sequence $(V_1^{1,T})_T$ has a pointwise converging subsequence. To do so we consider the metric space of increasing functions mapping into $[0, 1/(1 - \delta)]$. First consider the restrictions of the functions in $(V_1^{1,T})_T$ to the rational domain \mathbb{Q} . As \mathbb{Q} is countable, the diagonal theorem, see for example appendix A14 in Billingsley (2008), establishes that there exists a subsequence of $(V_1^{1,T})_T$ that converges pointwise on \mathbb{Q} . With a slight abuse of notation,

let $(V_1^{1,T})_T$ be this subsequence in the remainder and let \tilde{V} be the pointwise limit. Note that \tilde{V} is monotone (on \mathbb{Q}) because all $V_1^{1,T}$ are monotone. Now consider again the original domain \mathbb{R} and let \mathbb{D} be the set of points on which $(V_1^{1,T})_T$ does not converge. For any $d \in \mathbb{D}$, we claim that $\lim_{q \in \mathbb{Q} \nearrow d} \tilde{V}(q) < \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q)$. By monotonicity of \tilde{V} on \mathbb{Q} , the opposite strict inequality is impossible. If, however, $\lim_{q \in \mathbb{Q} \nearrow d} \tilde{V}(q) = \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q)$, then $(V_1^{1,T}(d))_T$ must converge to $\lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q)$, contradicting that $d \in \mathbb{D}$. To make the latter point clear, suppose to the contrary that there is an $\varepsilon > 0$ such that for any T' there exists a $T > T'$ such that $|V_1^{1,T}(d) - \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q)| > \varepsilon$. For concreteness, let us assume that we can find such a $T > T'$ such that $V_1^{1,T}(d) - \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q) > \varepsilon$ (the opposite case analogous). Take a $q' > d$ such that $\tilde{V}(q') < \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q) + \varepsilon/2$. As $(V_1^{1,T})$ converges point-wise, there exists a T'' such that $\tilde{V}(q') + \varepsilon/2 > V_1^{1,T}(q')$ for all $T > T''$. Hence, $V_1^{1,T}(q') < \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q) + \varepsilon$ for $T > T''$ but this (together with the monotonicity of $V_1^{1,T}$) contradicts that $V_1^{1,T}(d) - \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q) > \varepsilon$ for some arbitrarily large T . This establishes that $\lim_{q \in \mathbb{Q} \nearrow d} \tilde{V}(q) < \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q)$ for all $d \in \mathbb{D}$.

As \tilde{V} is monotone on \mathbb{Q} and \mathbb{Q} is dense in \mathbb{R} , the condition $\lim_{q \in \mathbb{Q} \nearrow d} \tilde{V}(q) < \lim_{q \in \mathbb{Q} \searrow d} \tilde{V}(q)$ can hold at only countably many points d . Hence, \mathbb{D} has a countable number of elements. But then the diagonal theorem can be applied again to show that there exists a subsequence of $(V_1^{1,T})_T$ that converges point-wise on \mathbb{R} . For this subsequence of T , take the corresponding subsequence of $(V_2^{1,T})_T$ and, using the same steps, we can get a subsequence such that also $(V_2^{1,T})_T$ converges point-wise on \mathbb{R} . Let $V_1^{1,*}$ and $V_2^{1,*}$ be these limit value functions. Using the second period value functions corresponding to the elements of the sequence of value functions converging to $(V_1^{1,*}, V_2^{1,*})$ and applying the same steps again gives us a subsequence of value functions (this time for the even periods where firm 2 is investing) converging point-wise. The resulting $(V_1^{1,*}, V_2^{1,*}, V_1^{2,*}, V_2^{2,*})$ is a stationary Markov equilibrium (if the Bellman equation was not satisfied for one player at some Δ , it would also be violated for this player in a subgame-perfect Nash equilibrium for T sufficiently high).

A.3. Proof of Lemma 1

We know from Section 3.1 that Lemma 1 is true in $t = T$. We proceed by induction. Assuming that the statement is true for $t + 1$, we will now show that it is true for t . For concreteness, say t is odd, i.e. firm 1 can invest. We consider the last statement of the lemma first: Take two values of Δ_{t-1} ; a high one, Δ^h , and a low one, Δ^l . Denote firm 1's optimal investment by $x(\Delta_{t-1})$. Now suppose – contrary to the lemma – that $\Delta_t^h = \Delta^h + x(\Delta^h) < \Delta^l + x(\Delta^l) = \Delta_t^l$. We will show that this leads to a contradiction. Optimality of the investment $x^h = x(\Delta^h)$ requires that investing x^h leads to a higher value than investing $\Delta^l + x^l - \Delta^h$ when the quality

difference is Δ^h :

$$\begin{aligned}
& D_1(\Delta^h + x^h) - c(x^h, D_1(\Delta^h)) + \delta V_1^{t+1}(\Delta^h + x^h) \\
& \geq D_1(\Delta^l + x^l) - c(\Delta^l + x^l - \Delta^h, D_1(\Delta^h)) + \delta V_1^{t+1}(\Delta^l + x^l) \\
& \Leftrightarrow D_1(\Delta^h + x^h) - D_1(\Delta^l + x^l) + \delta V_1^{t+1}(\Delta^h + x^h) - \delta V_1^{t+1}(\Delta^l + x^l) \\
& \geq c(x^h, D_1(\Delta^h)) - c(\Delta^l + x^l - \Delta^h, D_1(\Delta^h)).
\end{aligned}$$

Similarly, investing x^l must lead to a higher value than investing $x^h + \Delta^h - \Delta^l$ if the quality difference is Δ^l :

$$\begin{aligned}
& D_1(\Delta^l + x^l) - c(x^l, D_1(\Delta^l)) + \delta V_1^{t+1}(\Delta^l + x^l) \\
& \geq D_1(\Delta^h + x^h) - c(\Delta^h + x^h - \Delta^l, D_1(\Delta^l)) + \delta V_1^{t+1}(\Delta^h + x^h) \\
& \Leftrightarrow D_1(\Delta^h + x^h) - D_1(\Delta^l + x^l) + \delta V_1^{t+1}(\Delta^h + x^h) - \delta V_1^{t+1}(\Delta^l + x^l) \\
& \leq c(\Delta^h + x^h - \Delta^l, D_1(\Delta^l)) - c(x^l, D_1(\Delta^l)).
\end{aligned}$$

Taking these two optimality conditions together we obtain

$$c(x^l, D_1(\Delta^l)) - c(\Delta^h + x^h - \Delta^l, D_1(\Delta^l)) \leq c(\Delta^l + x^l - \Delta^h, D_1(\Delta^h)) - c(x^h, D_1(\Delta^h)). \quad (\text{A.1})$$

We will show that this last inequality cannot hold. Note that $\Delta^h > \Delta^l$ implies that $x^h < \Delta^h + x^h - \Delta^l$. Therefore, the strict convexity of c in x implies that

$$c(x^l, D_1(\Delta^l)) - c(\Delta^h + x^h - \Delta^l, D_1(\Delta^l)) > c(\Delta^l + x^l - \Delta^h, D_1(\Delta^l)) - c(x^h, D_1(\Delta^l))$$

as the difference in x is the same on both sides of the inequality but the cost difference is evaluated at a lower x on the right hand side. As D_1 is strictly increasing in Δ and $\Delta^h > \Delta^l$, the assumption $c_{xD_1} < 0$ implies that the right hand side of the previous inequality is lower when evaluated at $D_1(\Delta^h)$ instead of $D_1(\Delta^l)$ (this is the point where we use $\Delta^l + x^l > \Delta^h - x^h$ which implies $\Delta^l + x^l - \Delta^h > x^h$), i.e.

$$\begin{aligned}
c(x^l, D_1(\Delta^l)) - c(\Delta^h + x^h - \Delta^l, D_1(\Delta^l)) & > c(\Delta^l + x^l - \Delta^h, D_1(\Delta^l)) - c(x^h, D_1(\Delta^l)) \\
& > c(\Delta^l + x^l - \Delta^h, D_1(\Delta^h)) - c(x^h, D_1(\Delta^h)).
\end{aligned}$$

But this contradicts (A.1). We can therefore conclude that Δ_t is increasing in Δ_{t-1} .

To show that $V_1^t(\Delta_{t-1})$ is increasing in Δ_{t-1} consider again $\Delta^h > \Delta^l$ and let x^l be the optimal choice under Δ^l :

$$\begin{aligned}
V_1^t(\Delta^h) & = \max_x D_1(\Delta^h + x) - c(x, D_1(\Delta^h)) + \delta V_1^{t+1}(\Delta^h + x) \\
& \geq D_1(\Delta^h + x^l) - c(x^l, D_1(\Delta^h)) + \delta V_1^{t+1}(\Delta^h + x^l) \\
& \geq D_1(\Delta^l + x^l) - c(x^l, D_1(\Delta^l)) + \delta V_1^{t+1}(\Delta^l + x^l) \\
& = V_1^t(\Delta^l)
\end{aligned}$$

where the inequality follows from the fact that D_1 is increasing and c is decreasing in D_1 as well as the induction assumption that V_1^{t+1} is increasing.

To show that $V_2^t(\Delta_{t-1})$ is decreasing, recall that $\Delta_t = \Delta_{t-1} + x(\Delta_{t-1})$ is increasing in Δ_{t-1} and therefore

$$\begin{aligned} V_2^t(\Delta^h) &= D_2(\Delta^h + x(\Delta^h)) + \delta V_2^{t+1}(\Delta^h + x(\Delta^h)) \\ &\leq D_2(\Delta^l + x(\Delta^l)) + \delta V_2^{t+1}(\Delta^l + x(\Delta^l)) = V_2^t(\Delta^l) \end{aligned}$$

since D_2 and V_2^{t+1} are decreasing.

If t is even, the proof is analogous. □

A.4. Proof of Lemma 2

Take some Δ_t as given. Assume $\Delta_{t'} \geq 1$ for some $t' > t$. Then $\Delta_{t'}$ is also above 1 for all $\Delta'_t > \Delta_t$. This follows directly from Lemma 1 as a higher Δ_t leads to a higher Δ_{t+1} , which leads in turn to a higher Δ_{t+2} ... which leads to a higher $\Delta_{t'}$.

This implies the following: Whenever for a given Δ_t we have $\Delta_{t'} \geq 1$, for some $t' > t$, then the same is true for all higher Δ_t . Clearly, we can obtain the same result for -1 : Whenever for a given Δ_t we have $\Delta_{t'} \leq -1$ for some $t' > t$, then the same is true for all lower Δ_t . These two statements imply the corollary. □

A.5. Proof of Lemma 3

Lemma 3 is true for period T ; see equations (4) and (5). We will argue via induction that it is true for any t . From the definition of I^t and lemma 1, it is clear that any Δ_t in the interior of I^t leads in equilibrium to a Δ_{t+1} in the interior of I^{t+1} . Then the first-order conditions of the investing firm, (7) or (8) respectively, implies that the optimal investment is linear in Δ_{t-1} as D is linear, c is quadratic and V_i^{t+1} is by the induction hypothesis quadratic.

Consider an odd t where firm 1 invests. Then using the implicit function theorem on (7) yields

$$\frac{dx^t}{d\Delta_{t-1}} = \frac{\alpha/2 + \delta V_1^{t+1}''}{\gamma - \delta V_1^{t+1}''} > 0 \tag{A.2}$$

where the inequality follows because V_i^{t+1} is convex by the induction hypothesis and the denominator is positive by the second-order condition of the maximization problem. Note

that x is linear as V_i^{t+1} is quadratic by the induction hypothesis.³⁷ Similarly, if t is even then

$$\frac{dx^t}{d\Delta_{t-1}} = \frac{-\alpha/2 - \delta V_2^{t+1''}}{\gamma - \delta V_2^{t+1''}} < 0. \quad (\text{A.3})$$

Now we want to derive $V_1^{t''}$ in odd t . By the envelope theorem,

$$V_1^{t'}(\Delta_{t-1}) = \frac{1}{2} + \alpha x^t \frac{1}{2} + \delta V_1^{t+1'}(\Delta_{t-1} + x^t)$$

which after differentiating yields

$$V_1^{t''} = \alpha \frac{dx^t}{d\Delta_{t-1}} \frac{1}{2} + \delta \left(1 + \frac{dx^t}{d\Delta_{t-1}}\right) V_1^{t+1''}. \quad (\text{A.4})$$

We conclude that $V_1^{t''} \geq 0$ as $V_1^{t+1''} \geq 0$ by the induction hypothesis.

Furthermore, V_1^t is quadratic as x^t is linear and V_1^{t+1} is quadratic. Next we consider $V_2^{t''}$ for odd t :

$$\begin{aligned} V_2^{t'}(\Delta_{t-1}) &= -\frac{1}{2} + \delta \left(1 + \frac{dx^t}{d\Delta_{t-1}}\right) V_2^{t+1'}(\Delta_{t-1} + x^t) \\ V_2^{t''} &= \delta \left(1 + \frac{dx^t}{d\Delta_{t-1}}\right)^2 V_2^{t+1''} \end{aligned} \quad (\text{A.5})$$

where the last step utilizes that x^t is linear (hence its second derivative is zero). We obtain that $V_2^{t''} \geq 0$ as $V_2^{t+1''} \geq 0$ by the induction hypothesis. Also V_2^t is quadratic as x^t is linear and V_2^{t+1} is quadratic. The result for even t is derived analogously. \square

A.6. Proof of Proposition 1

Note that Proposition 1 holds for $t = T$ by the definition of I^t with $t = T$.³⁸ For a given equilibrium in the game with length T , consider the function that assigns to each Δ_t the resulting Δ_{t+1} , e.g. for even t we have $\Delta_{t+2}(\Delta_t) = \Delta_t + x^{t+1}(\Delta_t) - x^{t+2}(\Delta_t + x^{t+1}(\Delta_t))$. By Lemma 3, this function is linear on the interior of I^t (if both x^{t+1} and x^{t+2} are strictly greater than 0). Using the first-order conditions, it is straightforward to calculate that the slope of

³⁷Here one might consider the possibility of a corner solution $x^t = 0$ if marginal costs, i.e. α , are excessively high. If we assume $\alpha \leq 1/2$ this is impossible (recall that V_1^{t+1} is increasing). For large $T = \infty$, the assumption $\alpha < 1$ would be enough to rule this out: Suppose it $x^t = 0$: Then $\Delta_{t-1} < 0$ as otherwise $\alpha < 1$ implies that firm 1 wants to invest (recall that V_1^{t+1} is increasing). This implies by $\alpha < 1$ that firm 2 will invest a positive amount in $t+1$ and $\Delta_{t+1} < \Delta_{t-1}$. By the convexity of V_1^{t+1} , firm 1 will again find it optimal to invest 0 in $t+2$. Repeating the argument shows that Δ will diverge to -1 which contradicts that $\Delta_t \in I^t$.

³⁸In fact, we could use the analysis of section 3.1 to give the tighter (though somewhat messy) bound $[U_\alpha + (2\gamma - 1 + \alpha)/(2\gamma + \alpha)] / [(1 + \alpha/(2\gamma))^{T-t/2}]$ in the proposition. The proof for this tighter bound is the same.

$\Delta_{t+2}(\Delta_t)$ is (for concreteness, we let t be odd though this has no impact on the final result):

$$\begin{aligned} \frac{d\Delta_{t+2}}{d\Delta_t} &= \frac{d\{\Delta_{t+1} + x^{t+2}(\Delta_{t+1})\}}{d\Delta_t} = \frac{d\{\Delta_t - x^{t+1}(\Delta_t) + x^{t+2}(\Delta_t - x^{t+1}(\Delta_t))\}}{d\Delta_t} \\ &= 1 + \frac{\alpha/2 + \delta V_2^{t''}}{\gamma - \delta V_2^{t''}} + \frac{\alpha/2 + \delta V_1^{t+1''}}{\gamma - \delta V_1^{t+1''}} \left(1 + \frac{\alpha/2 + \delta V_2^{t''}}{\gamma - \delta V_2^{t''}}\right) \\ &= \left(1 + \frac{\alpha/2 + \delta V_1^{t+1''}}{\gamma - \delta V_1^{t+1''}}\right) \left(1 + \frac{\alpha/2 + \delta V_2^{t''}}{\gamma - \delta V_2^{t''}}\right) > (1 + \alpha/(2\gamma))^2 > 1 \end{aligned}$$

If either x^{t+1} or x^{t+2} is zero (say for concreteness $x^{t+2} = 0$), then the slope is only

$$\underline{s} = 1 + \frac{\alpha/2 + \delta V_i^{t+1''}}{\gamma - \delta V_i^{t+1''}} > 1 + \alpha/(2\gamma) > 1.$$

By $\alpha < 1$, it is impossible that both x^{t+1} and x^{t+2} are zero at any quality difference between -1 and 1. From the definition of I^t , it follows that the length of I^t can be at most $\text{length}(I^{t+2})/\underline{s}$. The condition in the proposition iterates this reasoning, e.g. the length of I^t can be at most $\text{length}(I^{t+4})/\underline{s}^2$ etc. Since $1 + \alpha/(2\gamma) > 1$, the maximal length of I^0 shrinks to zero as T becomes large. \square

A.7. Proof of Lemma 4

As the proof of Proposition 1 shows, the upper (lower) bound of I^t will be strictly below 1 (above -1) for $t < T - 1$. This implies that Δ_t is strictly above (below) I^t if $\Delta_t = 1$ (if $\Delta_t = -1$). By the definition of I^t and the monotonicity derived in Lemma 1, it follows that $\Delta_{t'}$ has again to be 1 (respectively -1) in some later period $t' > t$. The monotonicity in Lemma 1 in fact implies that $\Delta_{t'}$ will be above $I^{t'}$ in all $t' > t$ and therefore firm j cannot have full demand in any following period.

For the infinitely repeated game, V_i will no longer depend on t as the equilibrium is stationary. For concreteness, let firm i be firm 1 and assume that firm 1 has full demand in period t but not in $t-2$ (i.e. t is the first period in which firm 1 has full demand). Let t be odd (this is without loss of generality: if firm 1 has full demand in an even period, it will obviously also have full demand in the directly following odd period). From Lemma 1 it follows that $\Delta_{t+1} \geq \Delta_{t-1}$ is implied by $\Delta_{t-2} < \Delta_t$. This implies, by Lemma 1, that $\Delta_{t+2} \geq \Delta_t$. As firm 1 had full demand in t , firm 1 will have full demand again in $t+2$. This argument can now be iterated to yield the result (i.e. $\Delta_{t+2} \geq \Delta_t$ implies $\Delta_{t+4} \geq \Delta_{t+2}$ etc.). Note that this iteration also shows that Δ_{t+2n+1} is increasing in $n \in \mathbb{N}$. As we consider a stationary equilibrium, nothing depends on time periods per se and the result therefore also holds if firm 1 has full demand in the initial period. An analogous argument works for firm 2. \square

A.8. Proof of Lemma 5

Let the market be weakly tipping and assume for concreteness that firm 1 obtains full demand in some period t'' and that T is even. Lemma 4 (i) (used inductively) implies that firm 1 will also have full demand in period $T - 1$. By $\alpha \geq 1/2$, marginal costs of firm 2 in period T are greater than $1/2$ for every investment $x > 0$ while marginal revenue is – at most – $-D'_2 = 1/2$. Hence, zero investment is optimal for firm 2 in period T and firm 1 will have full demand also in period T . Using $t' = T - 1$ in the definition of absolutely tipping market gives the result. \square

A.9. Proof of Proposition 2

In the interior of I^t , the first-order condition in even periods is

$$1/2 - \gamma'x - \alpha'/2 - \delta V_2^{t+1'}(\Delta_{t-1} - x) = 0.$$

Note that this always yields a positive optimal investment, by $\alpha' < 1$ and $V_2^{t+1'} < 0$. In particular, $(1 - \alpha')/(2\gamma')$ is a lower bound for this investment. Furthermore – as V_2^{t+1} is quadratic in the interior of I^t – the optimal investment is linear in Δ_{t-1} with slope $-\delta V_2^{t+1''}/(\gamma' + \delta V_2^{t+1'}) \leq 0$.

In odd periods, the slope of the optimal investment is as given in (A.2) unless the investment is zero, which is in principle possible if $\alpha > 1/2$ and Δ_{t-1} sufficiently negative. We will now show that the optimal investment of firm 1 has to be strictly positive in the interior of I^t (as $T \rightarrow \infty$) in some periods. Recall that firm 2's investment in even periods is bounded from below by $(1 - \alpha')/(2\gamma')$. If firm 1 invested zero in $4\gamma'/(1 - \alpha')$ consecutive odd periods, then clearly Δ would decrease by more than 2 and, therefore, firm 2 would have captured the whole market, which contradicts the definition of I^t .

Hence, at least in one period out of every time window of $8\gamma'/(1 - \alpha')$ periods, firm 1 will have an interior investment. Let $t + 1$ be such a period with an interior investment by firm 1. Then,

$$\begin{aligned} \frac{d \Delta_{t+2}}{d \Delta_t} &= \frac{d \{ \Delta_{t+1} - x^{t+2}(\Delta_{t+1}) \}}{d \Delta_t} = \frac{d \{ \Delta_t + x^{t+1}(\Delta_t) - x^{t+2}(\Delta_t + x^{t+1}(\Delta_t)) \}}{d \Delta_t} \\ &= 1 + \frac{\alpha/2 + \delta V_1^{t''}}{\gamma - \delta V_1^{t''}} + \frac{\delta V_2^{t+1''}}{\gamma' - \delta V_2^{t+1''}} \left(1 + \frac{\alpha/2 + \delta V_1^{t''}}{\gamma - \delta V_1^{t''}} \right) \\ &= \left(1 + \frac{\alpha/2 + \delta V_1^{t+1''}}{\gamma - \delta V_1^{t+1''}} \right) \left(1 + \frac{\delta V_2^{t''}}{\gamma' - \delta V_2^{t''}} \right) > 1 + \alpha/(2\gamma) > 1. \end{aligned}$$

If firm 1 chooses $x^{t+1} = 0$, then:

$$\frac{d \Delta_{t+2}}{d \Delta_t} = \frac{d \{ \Delta_t + x^{t+1}(\Delta_t) - x^{t+2}(\Delta_t + x^{t+1}(\Delta_t)) \}}{d \Delta_t} = 1 + \frac{\delta V_2^{t+1''}}{\gamma' - \delta V_2^{t+1''}} \geq 1.$$

The effect of increasing Δ_0 by $\varepsilon > 0$ will therefore be greater than 2 in $2\gamma/\alpha * 8\gamma'/(1 - \alpha')$ periods which shows that $\Delta_0 + \varepsilon$ cannot be in I^0 if Δ_0 is. Since $\varepsilon > 0$ is arbitrary, this shows that the length of the interval I^0 is zero if T is sufficiently large. The other results follow immediately from there. \square

A.10. Proof of Corollary 2

Let $\bar{\Delta}_0$ be the infimum of the initial quality differences for which the market tips in favor of firm 1 in case of entry. Part (i) is obviously true as firm 1 will never enter if $F > \bar{F} = 1/(1 - \delta)$. (ii) Recall that V_1^1 is increasing in Δ_0 . As firm 1 enters only if $F \leq V_1^1(\Delta_0)$, this implies that firm 1 enters only if Δ_0 is sufficiently large. For $F \in [V_1^1(\bar{\Delta}_0), \bar{F}]$, firm 1 enters only if it takes over the market eventually, i.e. taking $\hat{F} = V_1^1(\bar{\Delta}_0)$ yields (ii). Part (iii), entry without tipping, occurs if $\Delta_0 < \bar{\Delta}_0$ but nevertheless entry allows a positive profit. For example, assume α' and γ' are very low, say 0 for concreteness, (and α and γ are not low), then firm 2 will react to entry by investing in period 2 sufficiently to force firm 1 from the market. If $F = 0$ and $\Delta_0 > -1$, entry will nevertheless be profitable for firm 1 as positive profits are made in period 1. \square

A.11. Proof of Lemma 6

By Proposition 2, the length of I^0 is zero. That is, it contains at most a single point. Consequently, there is a $\bar{\Delta}_0$ such that the market tips in favor of firm 1 (2) if $\Delta_0 > \bar{\Delta}_0$ (if $\Delta_0 < \bar{\Delta}_0$). By the stationarity of the equilibrium, $\bar{\Delta}_0 = \bar{\Delta}_t$ for all t , where $\bar{\Delta}_t$ is the quality difference such that the market will eventually tip in favor of firm 1 if the quality difference is above $\bar{\Delta}_t$ after period t .

Now we will show that $\Delta_{t+2} > \Delta_0$ implies that $\Delta_0 > \bar{\Delta}_0$. To do so we show that Δ is increasing over time if $\Delta_0 > \bar{\Delta}_0$ and decreasing if $\Delta_0 < \bar{\Delta}_0$. To see this, let $\Delta_0 > \bar{\Delta}_0$ or, more generally, $\Delta_t > \bar{\Delta}_t$. If $\Delta_{t+2} < \Delta_t$, then—by the monotonicity shown in Lemma 1— $\tilde{\Delta}_{t+2} < \Delta_{t+2} < \Delta_t$ for all $\tilde{\Delta}_t < \Delta_t$. In a stationary equilibrium, this implies that Δ_{t+2n} can never be above Δ_t . If $\Delta_t < 1$, this would, however, contradict $\Delta_t > \bar{\Delta}_t$. Hence, for all $\Delta_t \in (\bar{\Delta}_t, 1)$, we obtain that $\Delta_{t+2} > \Delta_t$. By monotonicity, we then get $\Delta_{t+2} \geq 1$ for all $\Delta_t \geq 1$. Similarly, we can obtain the result that, for all $\Delta_t \in (-1, \bar{\Delta}_t)$, we have $\Delta_{t+2} < \Delta_t$. By monotonicity, we then get $\Delta_{t+2} \leq -1$ for all $\Delta_t \leq -1$.

Taking this together we can have market entry and $\Delta_2 > \Delta_0$ only if $\Delta_0 > \bar{\Delta}_0$. \square

A.12. Proof of Lemma 7

The Lemma is true for $t = T$. Using (backwards) induction and the first-order conditions (7) and (8), it is straightforward to derive the result for $t < T$.

Let us be a bit more precise here: In period T , $V_1^{T'} = 1/2$ on I^T . In other periods, $V_1^{t'} = 1/2 + \delta V_1^{t+1'}$ on I^t , which implies by induction that $V_1^{t'} = \sum_{i=0}^{T-t} \delta^i / 2$ on I^t . Similarly, $V_2^{t'} = -\sum_{i=0}^{T-t} \delta^i / 2$ on I^t . Using the first-order conditions (7) and (8) yields that

$$x^t = \frac{1/2 + \delta |V_i^{t+1'}|}{\gamma} = \frac{1 - \delta^{T-t+1}}{2\gamma(1 - \delta)}$$

on I^t . □

A.13. Proof of Lemma 1 with decaying quality

First consider Lemma 1.(ii). Take two values of Δ_{t-1} ; a high one, Δ^h , and a low one, Δ^l . Denote firm 1's optimal investment by $x(\Delta_{t-1})$. Now suppose – contrary to the Lemma – that $\Delta_t^h = \mu\Delta^h + x(\Delta^h) < \mu\Delta^l + x(\Delta^l) = \Delta_t^l$. We will show that this leads to a contradiction. Optimality of the investment, $x^h = x(\Delta^h)$, requires that investing x^h leads to a higher value than investing $\mu\Delta^l + x^l - \mu\Delta^h$ when the quality difference is Δ^h :

$$\begin{aligned} & D_1(\mu\Delta^h + x^h) - c(x^h, D_1(\Delta^h)) + \delta V_1^{t+1}(\mu\Delta^h + x^h) \\ & \geq D_1(\mu\Delta^l + x^l) - c(\mu\Delta^l + x^l - \mu\Delta^h, D_1(\Delta^h)) + \delta V_1^{t+1}(\mu\Delta^l + x^l) \\ & \Leftrightarrow D_1(\mu\Delta^h + x^h) - D_1(\mu\Delta^l + x^l) + \delta V_1^{t+1}(\mu\Delta^h + x^h) - \delta V_1^{t+1}(\mu\Delta^l + x^l) \\ & \geq c(x^h, D_1(\Delta^h)) - c(\mu\Delta^l + x^l - \mu\Delta^h, D_1(\Delta^h)). \end{aligned}$$

Similarly, investing x^l must lead to a higher value than investing $x^h + \mu\Delta^h - \mu\Delta^l$ if the quality difference is Δ^l :

$$\begin{aligned} & D_1(\mu\Delta^l + x^l) - c(x^l, D_1(\Delta^l)) + \delta V_1^{t+1}(\mu\Delta^l + x^l) \\ & \geq D_1(\mu\Delta^h + x^h) - c(\mu\Delta^h + x^h - \mu\Delta^l, D_1(\Delta^l)) + \delta V_1^{t+1}(\mu\Delta^h + x^h) \\ & \Leftrightarrow D_1(\mu\Delta^h + x^h) - D_1(\mu\Delta^l + x^l) + \delta V_1^{t+1}(\mu\Delta^h + x^h) - \delta V_1^{t+1}(\mu\Delta^l + x^l) \\ & \leq c(\mu\Delta^h + x^h - \mu\Delta^l, D_1(\Delta^l)) - c(x^l, D_1(\Delta^l)). \end{aligned}$$

Taking these two optimality conditions together, we obtain:

$$c(x^l, D_1(\Delta^l)) - c(\mu\Delta^h + x^h - \mu\Delta^l, D_1(\Delta^l)) \leq c(\mu\Delta^l + x^l - \mu\Delta^h, D_1(\Delta^h)) - c(x^h, D_1(\Delta^h)). \quad (\text{A.6})$$

We will show that this last inequality cannot hold. Note that $\Delta^h > \Delta^l$ implies that $x^h < \mu\Delta^h + x^h - \mu\Delta^l$. Therefore, the strict convexity of c in x implies that:

$$c(x^l, D_1(\Delta^l)) - c(\mu\Delta^h + x^h - \mu\Delta^l, D_1(\Delta^l)) > c(\mu\Delta^l + x^l - \mu\Delta^h, D_1(\Delta^l)) - c(x^h, D_1(\Delta^l))$$

because the difference in x is the same on both sides of the inequality but the cost difference is evaluated at a lower x on the right-hand side. As D_1 is strictly increasing in Δ and $\Delta^h > \Delta^l$,

the assumption $c_{xD_1} < 0$ implies that the right-hand side of the previous inequality is lower when evaluated at $D_1(\Delta^h)$ instead of $D_1(\Delta^l)$ (for this, we use $\mu\Delta^l + x^l > \mu\Delta^h - x^h$, which implies $\mu\Delta^l + x^l - \mu\Delta^h > x^h$). It follows:

$$\begin{aligned} c(x^l, D_1(\Delta^l)) - c(\mu\Delta^h + x^h - \mu\Delta^l, D_1(\Delta^l)) &> c(\mu\Delta^l + x^l - \mu\Delta^h, D_1(\Delta^l)) - c(x^h, D_1(\Delta^l)) \\ &> c(\mu\Delta^l + x^l - \mu\Delta^h, D_1(\Delta^h)) - c(x^h, D_1(\Delta^h)). \end{aligned}$$

But this contradicts (A.6). We can therefore conclude that Δ_t is increasing in Δ_{t-1} .

To prove the robustness of Lemma 1.(i), we have to show that $V_1^t(\Delta_{t-1})$ is increasing in Δ_{t-1} . Consider again $\Delta^h > \Delta^l$ and let x^l be the optimal choice under Δ^l :

$$\begin{aligned} V_1^t(\Delta^h) &= \max_x D_1(\mu\Delta^h + x) - c(x, D_1(\Delta^h)) + \delta V_1^{t+1}(\mu\Delta^h + x) \\ &\geq D_1(\mu\Delta^h + x^l) - c(x^l, D_1(\Delta^h)) + \delta V_1^{t+1}(\mu\Delta^h + x^l) \\ &\geq D_1(\mu\Delta^l + x^l) - c(x^l, D_1(\Delta^l)) + \delta V_1^{t+1}(\mu\Delta^l + x^l) \\ &= V_1^t(\Delta^l), \end{aligned}$$

where the inequality follows from the fact that D_1 is increasing and c is decreasing in D_1 as well as the induction assumption that V_1^{t+1} is increasing.

To show that $V_2^t(\Delta_{t-1})$ is decreasing, recall that $\Delta_t = \mu\Delta_{t-1} + x(\Delta_{t-1})$ is increasing in Δ_{t-1} and therefore

$$\begin{aligned} V_2^t(\Delta^h) &= D_2(\mu\Delta^h + x(\Delta^h)) + \delta V_2^{t+1}(\mu\Delta^h + x(\Delta^h)) \\ &\leq D_2(\mu\Delta^l + x(\Delta^l)) + \delta V_2^{t+1}(\mu\Delta^l + x(\Delta^l)) = V_2^t(\Delta^l) \end{aligned}$$

because D_2 and V_2^{t+1} are decreasing.

If t is even, the proof is analogous. □

A.14. Proof of Proposition 3

The proof of lemma 3 goes through. However, the slope of the investment decision in the interior of I^t if firm 1 invests is now

$$\frac{dx^t}{d\Delta_{t-1}} = \frac{\alpha/2 + \mu\delta V_1^{t+1}''}{\gamma - \delta V_1^{t+1}''}$$

and if firm 2 invests we obtain

$$\frac{dx^t}{d\Delta_{t-1}} = \frac{-\alpha/2 - \mu\delta V_2^{t+1}''}{\gamma - \delta V_2^{t+1}''}.$$

Following the proof of Proposition 1, we now have to analyze (concentrating on odd t for concreteness) $\Delta_{t+1}(\Delta_t) = \mu^2\Delta_t - \mu x^{t+1}(\Delta_t) + x^{t+2}(\mu\Delta_t - x^{t+1}(\Delta_t))$. This yields (if both x^{t+1} and x^{t+2} are strictly positive)

$$\frac{d\Delta_{t+2}}{d\Delta_t} = \left(\mu + \frac{\alpha/2 + \mu\delta V_1^{t+1}''}{\gamma - \delta V_1^{t+1}''} \right) \left(\mu + \frac{\alpha/2 + \mu\delta V_2^{t+1}''}{\gamma - \delta V_2^{t+1}''} \right) > (\mu + \alpha/(2\gamma))^2 > \mu(\mu + \alpha/(2\gamma)).$$

If only one of the two investments is positive (say $x^{t+2} = 0$) we obtain the following slope

$$\underline{s} = \mu^2 + \mu \frac{\alpha/2 + \mu\delta V_2^{t''}}{\gamma - \delta V_2^{t''}} > \mu(\mu + \alpha/(2\gamma)).$$

If $\mu(\mu + \alpha/(2\gamma)) > 1$, we can therefore bound the slope Δ_{t+2} as a function of Δ_t from below by a constant strictly higher than 1. The maximal length of I^0 will therefore shrink to zero at exponential speed as $T \rightarrow \infty$ if $\mu(\mu + \alpha/(2\gamma)) > 1$. \square

A.15. Proof of Proposition 4

The proof is done for $\bar{\Delta}$ while the proof for $\underline{\Delta}$ works analogously. First, we show that for sufficiently high Δ investments by both SE are zero. Note that for $\Delta > 1$ the marginal costs of firm 2 are strictly higher than α . Now suppose that $\Delta > 1 + 1/(\alpha - \alpha\delta)$. It is then optimal for firm 2 not to invest: Suppose otherwise, i.e. suppose there is a $\Delta' > 1 + 1/(\alpha - \alpha\delta)$ such that firm 2 invest in equilibrium. This implies that firm 2 must have positive demand eventually in this equilibrium, say firm 2 will have positive demand (for the first time) in t' periods. Firm 2's revenue is then bounded from above by $\delta^{t'}/(1 - \delta)$. Firm 2's investment costs (until period t' , i.e. until $\Delta < 1$) are strictly bounded from below by $\delta^{t'}\alpha/(\alpha - \alpha\delta)$, which equals the upper bound on revenues. Hence, firm 2's value would be negative although it could secure a zero value by not investing ever. This contradicts that there is an equilibrium in which firm 2 invests a positive amount at some $\Delta > 1 + 1/(\alpha - \alpha\delta)$. Given that firm 1 has full demand and firm 2 does not invest for $\Delta > 1 + 1/(\alpha - \alpha\delta)$, firm 1 will also not invest for $\Delta > 1 + 1/(\alpha - \alpha\delta)$. This proves that every $\Delta > 1 + 1/(\alpha - \alpha\delta)$ is a steady state in every Markov equilibrium.

Let \bar{D} be the set of all Δ that are (i) steady states such that firm 2 invests zero in $t + 2$ if $\Delta_{t+1} = \Delta$, (ii) firm 1 has full demand, i.e. $\Delta \geq 1$ and (iii) the steady states are stable in the following sense: There exists an $\varepsilon > 0$ such that $\Delta_{t+2} \geq \Delta_t$ if $\Delta_t \in (\Delta - \varepsilon, \Delta)$ if t is even. By the previous paragraph, this set is non-empty and by (ii) it is bounded from below by 1. Therefore, \bar{D} has an infimum. Let Δ' be this infimum of \bar{D} . We will now show that Δ' is a stable steady state.

Suppose otherwise, i.e. suppose we can find $\Delta'' < \Delta'$ arbitrarily close to Δ' such that $\Delta_{t+2} < \Delta''$ if $\Delta_t = \Delta''$ and t is even. First, note that $V_1(\Delta) = 1/(1 - \delta)$ for $\Delta > \Delta'$ by the definition of \bar{D} . Given that $\Delta_{t+2} < \Delta''$, firm 2 must invest in $t + 2$ more than firm 1 does in $t + 1$. Firm 2 only invests a positive amount if it can enjoy some positive demand in a future period. As firm 2's marginal costs of investment are at least α , the future revenue stream of firm 2 must be at least αx_2^{t+2} and, therefore, firm 1's value at Δ'' has to be less than $1/(1 - \delta) - \delta\alpha x_2^{t+2}$. Now distinguish two cases. First suppose $\Delta' > 1$. Then take $\Delta'' > 1$ and note that firm 2 will in period $t + 2$ expect to make future revenues worth a net

present value of at least $\alpha(\Delta'' - 1)$ (as otherwise $x_2^{t+2} = 0$ would be optimal) and therefore $V_1^{t+1}(\Delta'') < 1/(1 - \delta) - \delta\alpha(\Delta'' - 1)$. By investing $\Delta' - \Delta'' + \varepsilon$ in period t for some $\varepsilon > 0$ arbitrarily small, firm 1 would guarantee itself $1/(1 - \delta)$. For Δ'' sufficiently close to Δ' and ε sufficiently small, this gives 1 a higher value than $1/(1 - \delta) - \delta\alpha(\Delta'' - 1)$, which contradicts the candidate equilibrium.

Hence, we can move to the second case $\Delta' = 1$. Note that firm 1 can, for $\Delta_t = \Delta''$, guarantee itself $1/(1 - \delta) - c(1 - \Delta'' + \varepsilon, D_1(\Delta'')) = 1/(1 - \delta) - \gamma(1 - \Delta'' + \varepsilon)^2/2 - \alpha(1 - \Delta'' + \varepsilon)(1 - \Delta'')/2$ for some $\varepsilon > 0$ arbitrarily small by investing $1 - \Delta'' + \varepsilon$ in $t + 1$. By sticking to its equilibrium investment, firm 1 will get at most $1/(1 - \delta) - \delta(1 - \Delta'')/2$ as firm 2 will have demand of at least $(1 - \Delta'')/2$ in period $t + 2$. But for Δ'' sufficiently close to $\Delta' = 1$ and ε sufficiently small, $\gamma(1 - \Delta'' + \varepsilon)^2/2 + \alpha(1 - \Delta'' + \varepsilon)(1 - \Delta'')/2 < \delta(1 - \Delta'')/2$, which contradicts the optimality of firm 1's equilibrium investment. Hence, Δ' is stable.

Last we show that $\Delta' = 1$. Suppose otherwise, i.e. suppose $\Delta' > 1$. As Δ' is stable and above 1, firm 2's investment when facing quality difference Δ' will pay off only in the period $t + 2$ in which it is made. By $\alpha \geq 1/2$, marginal costs of investing are then higher than marginal revenue even at zero investment. Hence, investing is unprofitable for firm 2 when facing quality difference Δ' . By stability of Δ' , the same is true for all $\Delta \in (\Delta' - \varepsilon, \Delta')$ for $\varepsilon > 0$ sufficiently small. Given that firm 2 invests zero when facing $\Delta \in (\Delta' - \varepsilon, \Delta')$ and given that $\Delta' > 1$, it is optimal for firm 1 to invest zero at $\Delta \in (\Delta' - \varepsilon, \Delta')$ for ε sufficiently small. Hence, these quality differences are stable steady states. But this means that all $\Delta \in (\Delta' - \varepsilon, \Delta')$ are in \bar{D} , which contradicts the definition of Δ' as the infimum of \bar{D} . Hence, $\Delta' = 1$.

To see that $\Delta' = 1$ is a *strictly* stable steady state, note that the arguments two paragraphs above show a profitable deviation in case in case that $\Delta'' < 1$ arbitrarily close to 1 exist such that $\Delta_t = \Delta''$ imply $\Delta_{t+1} = \Delta''$. \square

A.16. Proof of Lemma 8

Consider the completely myopic case: $\delta = 0$ and the firms invest as in the one-shot game. Investments will then be given by (3) and will be denoted by $x_i^*(\Delta)$ for the remainder of this proof. Consider one investment by firm 1 and one investment by firm 2. If neither grabs the whole market, that is, if investments are interior, then the change in Δ is:

$$d(\Delta) \equiv x_1(\Delta) - x_2(\Delta + x_1(\Delta)) = \frac{\alpha(1 + 1/(4\gamma))}{\gamma} \Delta + \frac{1}{4\gamma^2}(1 - \alpha).$$

Clearly, this is a strictly monotone (and linear) function with only one zero; call this zero $\tilde{\Delta}$. Consequently, if firms play myopic and repeatedly, the market will tip for all initial quality differences but $\tilde{\Delta}$. Now consider the profit difference between investing $x^*(\Delta)$ and investing

some x in the current period (assuming that one does not grab the whole market), that is, for firm 1 consider:

$$f(x, \Delta) \equiv \frac{1 + \Delta + x^*(\Delta)}{2} - c(x^*(\Delta), D_1(\Delta)) - \left[\frac{1 + \Delta + x}{2} - c(x, D_1(\Delta)) \right].$$

Clearly, $f \geq 0$ is a quadratic function in x with its minimum at $x = x^*(\Delta)$. Hence, for any $\varepsilon' > 0$, we can find a δ' such that $f(x, \Delta) \geq \delta'/(1 - \delta')$ if $x \notin [x^*(\Delta) - \varepsilon', x^*(\Delta) + \varepsilon']$ (assuming that $\Delta \in [-1, 1]$).

Now for a given $\varepsilon > 0$, choose $\varepsilon' > 0$ such that $d(\Delta) > 2\varepsilon'$ if $\Delta \notin [\tilde{\Delta} - \varepsilon/2, \tilde{\Delta} + \varepsilon/2]$. For this ε' , determine δ' as in the previous paragraph. Note that neither firm will deviate from $x^*(\Delta)$ by more than ε' if $\delta \leq \delta'$ as the possible gain in the future is bounded by $1/(1 - \delta)$ (and as it is in the future this potential gain is discounted by δ). Hence, when $\delta \leq \delta'$ and $\Delta \notin [\tilde{\Delta} - \varepsilon/2, \tilde{\Delta} + \varepsilon/2]$, then Δ will change qualitatively as in d and the market will tip eventually. \square

References

- Acquisti, A., C. Taylor, and L. Wagman (2016, June). The economics of privacy. *Journal of Economic Literature* 54(2), 442–492.
- Argenton, C. and J. Prüfer (2012). Search engine competition with network externalities. *Journal of Competition Law and Economics* 8(1), 73–105.
- Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics* 37(3), 668–691.
- Athey, S. and G. Ellison (2011). Position auctions with consumer search. *Quarterly Journal of Economics* 126(3), 1213–1270.
- Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review* 90(4), 1034–1054.
- Besen, S. M. and J. Farrell (1994). Choosing how to compete: Strategies and tactics in standardization. *Journal of Economic Perspectives* 8(2), 117–131.
- Billingsley, P. (2008). *Probability and measure*. John Wiley & Sons, New York.
- Brynjolfsson, E. and A. McAfee (2012). *The Second Machine Age*. Norton and Company, New York.
- Burguet, R., R. Caminal, and M. Ellman (2015). In Google we trust? *International Journal of Industrial Organization* 39, 44 – 55.
- Chen, Y. and C. He (2011). Paid placement: Advertising and search on the internet. *Economic Journal* 121(556), 309–328.
- Chiou, L. and C. E. Tucker (2014). Search engines and data retention: Implications for privacy and antitrust. MIT Sloan Research Paper 5094-14.
- Cyert, R. and M. de Groot (1970). Multiperiod decision models with alternating choice as the solution of the duopoly problem. *Quarterly Journal of Economics* 84, 419–429.
- de Cornière, A. (2016). Search advertising. *American Economic Journal: Microeconomics* 8(3), 156–88.
- de Cornière, A. and R. de Nijs (2016). Online advertising and privacy. *RAND Journal of Economics* 47(1), 48–72.
- Dengler, S. and J. Prüfer (2016). Consumers’ privacy choices in the era of big data. mimeo, Tilburg University.

- Economides, N. (1996). The economics of networks. *International Journal of Industrial Organization* 14(6), 673–699.
- Edelman, B. (2015). Does Google leverage market power through tying and bundling? *Journal of Competition Law and Economics* 11(2), 365–400.
- Edelman, B., M. Ostrovsky, and M. Schwarz (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97(1), 242–259.
- Einav, L. and J. Levin (2014). The data revolution and economic analysis. *Innovation Policy and the Economy* 14(1), 1–24.
- Fudenberg, D. and J. Tirole (1991). *Game theory*. MIT Press Cambridge, MA.
- Goldfarb, A. and C. Tucker (2011). Online display advertising: Targeting and obtrusiveness. *Marketing Science* 30(3), 389–404.
- Hagiu, A. and B. Jullien (2011). Why do intermediaries divert search? *RAND Journal of Economics* 42(2), 337–362.
- Halaburda, H., B. Jullien, and Y. Yehezkel (2016). Dynamic competition with network externalities: Why history matters. NET Institute Working Paper No. 13-10.
- Iskhakov, F., J. Rust, and B. Schjerning (2015). Recursive lexicographical search: Finding all Markov perfect equilibria of finite state directional dynamic games. *Review of Economic Studies* 83(2), 658–703.
- Lianos, I. and E. Motchenkova (2013). Market dominance and search quality in the search engine market. *Journal of Competition Law and Economics* 9(2), 419–455.
- Mailath, G. J. and L. Samuelson (2006). *Repeated games and reputations: long-run relationships*. Oxford University Press.
- Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, i: Overview and quantity competition with large fixed costs. *Econometrica* 56(3), 549–569.
- Mayer-Schönberger, V. and K. Cukier (2013). The rise of big data. *Foreign Affairs*, May/June Issue.
- McAfee, P., J. Rao, A. Kannan, D. He, T. Qin, and T.-Y. Liu (2015). Measuring scale economies in search. <http://www.learconference2015.com/wp-content/uploads/2014/11/McAfee-slides.pdf>, 2 June 2015.

- Rochet, J.-C. and J. Tirole (2006). Two-sided markets: A progress report. *RAND Journal of Economics* 37(3), 645–667.
- Shapiro, C. and H. Varian (1999). *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Review Press.
- Taylor, G. (2013). Search quality and revenue cannibalization by competing search engines. *Journal of Economics & Management Strategy* 22(3), 445–467.
- The Economist (2015). Does Deutschland do digital? November 21, 2015.
- The Economist (2016). Of profits and prophecies. February 6, 2016.
- The New York Times (2012). Big data: Rise of the machines. December 31, 2012.
- The New York Times (2015a). Apple and Google create a buzz at Frankfurt motor show. September 17, 2015.
- The New York Times (2015b). iPhone 6s’s hands-free Siri is an omen of the future. September 22, 2015.
- Wall Street Journal (2015). Uber valued at more than \$50 billion. July 31, 2015.
- Zuboff, S. (2016). The secrets of surveillance capitalism. *Frankfurter Allgemeine Zeitung*, March 5, 2016.