Centre for Economic Policy Research

[redacted]

Experimentation in Two-Sided Markets*

Martin Peitz† Sven Rady‡ Piers Trepper§

May 9, 2011

Abstract

In this pap r, w study optimal xp rim ntation by a monopolistic platform in a two-sid d mark t fram work. Th platform provid r fac s unc rtainty about th siz of th xt rnality ach sid is x rcising on th oth r. It maximiz s xp ct d profit in a continuous-tim infinit -horizon fram work by s tting participation f s or quantiti s on both sid s. W comput th boundari s of xp rim ntation giv n by th optimal polici s of a myopic and an infinit ly pati nt platform provid r. Th y r v al that an xp rim nting platform provid r will always incr as both quantiti s but, d p nding on param t rs, may n v rth l ss rais th pric on on sid of th mark t.

Keywords: Two-Sid d Mark ts, Platform Mark t, Monopoly Exp rim ntation, Bay sian L arning, Optimal Control, N twork Eff cts
JEL classification: D42, D83, L12

---

*Our thanks for helpful discussions and comments are due to participants of the Conference on Platform Markets, Mannheim, 31 May - 2 June 2010. We would like to thank the Studienzentrum Gerzensee for its hospitality during the 2008 and 2009 European Summer Symposia in Economic Theory (ESSET). Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 is gratefully acknowledged.

†Department of Economics, University of Mannheim, D-68131 Mannheim, Germany; email: martin.peitz@googlemail.com.

‡Department of Economics, University of Munich, Kaulbachstr. 45, D-80539 Munich, Germany; email: sven.rady@lrz.uni-muenchen.de.

§Department of Economics, University of Munich, Kaulbachstr. 45, D-80539 Munich, Germany; email: piers.trepper@lrz.uni-muenchen.de.

1 Introduction

In many real-world markets, transactions are intermediated through platforms. This paper studies a monopolistic platform in a two-sided market framework. The platform provider is uncertain about the size of the externality each side of the market is exerting on the other and, thus, may want to experiment in order to learn about the externality parameters. Its aim is to maximize expected lifetime profit in a continuous-time infinite-horizon world.

The platform provider's choice variables (prices or quantities) determine its current profit in every instant of time as well as the amount of information received. Thus, there is a trade-off between maximizing current profit and extracting as much information as possible in order to increase future profits. The higher the rate at which future profits are discounted, the more important maximizing current profit becomes, up to the extreme of infinite discounting when information acquisition is completely ignored. Reversely, the benefit of information increases if the discount rate decreases, up to the opposite extreme of no discounting when maximal weight is put on learning. We determine the optimal policy for each of the two extremes. These policies are the boundaries of experimentation for any finitely patient platform provider.

We consider two variants of the model, one in which the platform provider sets prices and learns from quantities, and one in which the platform provider selects quantities and learns from prices. Prices take the form of membership or subscription fees. While the price-setting version of the model seems more realsitic, the quantity-setting version turns out to be more tractable.

Restricting our analysis to positive quantities, we find that an infinitely patient quantity-setting platform provider will always increase quantities relative to the myopic optimum. The effect of quantity experimentation on the expected price levels depends on market characteristics, however. There are two cases. In the first case, expected prices are lower than their myopic counterparts on both sides of the market. In the second case, the expected price is lower than its myopic counterpart on one side of the market, but higher on the other. Such a one-sided price increase may occur on a market side that itself exerts a low externality on the opposite side while benefiting strongly from participation on that side. Thus, it might be optimal to strongly increase participation on the opposite side and to collect part of the high additional benefit through a higher price on this side.

We obtain qualitatively similar results on the direction of experimentation when we allow the platform provider to set prices and observe quantities. In this scenario the optimal prices set by an infinitely patient platform become technically untractable, even if we assume the uncertainty to prevail only on one side of the market. Nevertheless, we can state the first-order conditions these prices have to satisfy and obtain some insights from these.

Two-sided markets have received a lot of attention in industrial economics recently. In general a market is said to be two-sided whenever potential participants care about the number of counterparts on the other side of the market—i.e., when each side exerts an externality on the other side, be it positive or negative. Potential interactions take place on some platform or by means of some vehicle, allowing the provider of such a platform or

vehicle to charge participants for services and to manage usage on both sides.

Real world examples and applications of two-sided markets are manifold. Examples with positive externalities include payment systems (where card holders will want to hold a card if many merchants accept it, while merchants will be willing to accept cards that many customers hold), game consoles (players, software developers), nightclubs and matching agencies (men, women), shopping malls, supermarkets, and department stores (where consumers are interested in a large diversity of products, and producers in a large number of customers).

Examples for negative externalities on one side of the market include newspapers and other media platforms, where advertisers prefer to promote their products on platforms with many readers/watchers, while readers/watchers prefer fewer ads to more. One can even think of situations with negative externalities on both sides. Imagine, for example, rivaling fans in a football stadium, where each group would prefer fewer fans from the other group. Note that in this case the intrinsic value of the platform (watching the football match) has to be large enough to compensate both sides for the disutility generated by the opposite side.

Seminal papers on two-sided markets are Rochet and Tirole (2003, 2006) and Armstrong (2006). For a theoretical investigation of media platforms see, in particular, Anderson and Coate (2005). A general model of monopoly platforms is analyzed by Nocke, Peitz, and Stahl (2007). Empirical work includes Rysman (2004) and Kaiser and Wright (2006). For a selective survey, see Rysman (2009). None of the existing literature treats two-sided markets in a setting of uncertainty where it is unclear how strong the relevant externalities are, and where the platform provider might benefit from experimenting with prices or quantities in order to learn about the true state of the world. Relative to the existing literature on two-sided markets, our contribution is to introduce uncertainty and learning into the set-up proposed by Armstrong (2006).

The economics literature on optimal experimentation by a single Bayesian decision maker starts with the work of Prescott (1972) and Rothschild (1974); a brief overview of this literature can be found in Keller and Rady (1999). Our contribution here is to extend the analysis of optimal experimentation to two-sided markets and, building upon the infinite-horizon continuous-time model of Keller and Rady (1999), to provide a tractable framework for it. To the best of our knowledge, ours is the first experimentation model in which the decision maker has more than one instrument (i.e., two quantities or two prices) with which to trade off exploration versus exploitation. Because of this, even a platform provider primarily concerned about information acquisition can still pursue the secondary goal of current profit maximization: from all pairs of actions generating the same amount of information, the optimal policy selects the pair with the highest current profit.

Our analysis suggests that markets characterized by indirect network effects of uncertain size provide incentives for the experimenting platform provider to initially increase the amount of participation. This improves the learning. Whether it implies lower prices on both sides of the market, depends on the market characteristics, as will be elaborated on in the main text. Since at least one price is lowered in expectations initially, our theory provides a new rationale for price discounts in dynamic two-sided markets.[1]

[1]An alternative explanation could be dynamic consumer behavior which might make a platform provider strive to build up a critical mass. We exclude this channel by assuming that participants can revise their participation decision in each period at no cost.

The remainder of the paper is structured as follows. Section 2 presents the model for the price-setting platform provider with uncertainty on both sides of the market and characterizes the evolution of beliefs. Section 3 analyzes the optimal pricing strategy, and also considers the case of uncertainty about one externality only and the case of an externality on one side only. The optimal policy for a quantity-setting platform provider is analyzed in Section 4. Section 5 concludes. Technical proofs are relegated to the appendix.

2 The Model

We follow Armstrong (2006) and focus on participation decisions. For tractability reasons, we assume linear demand functions on both sides of the market. We refer to the two sides as A and B. Depending on the application, these may be buyers and sellers, advertisers and readers, or men and women. The novelty is to introduce uncertainty with respect to the size of the network effect. Arguably, such uncertainty is an important feature of network industries: a platform provider typically cannot perfectly foresee how strongly one side reacts to the number of users on the other side and has to infer this from market outcomes which noisily reveal the true state of the world.

2.1 The price-setting platform pro ider

In each period, there is a continuum of participants on both sides of the market. Invoking a uniform distribution over the value of the outside option (on a support that is sufficiently large such that aggregate demand is price elastic when positive) gives rise to linear demand functions. The platform provider can set membership fees (M_A, M_B), but no usage fee.[2] Suppose that the total mass of potential participants is such that demand n_i on side $i = A, B$ satisfies $dn_i/dM_i = -1$. The resulting masses of participants n_A and n_B are then characterized by the system of linear equations

$$n_A = u_0 + \tilde{u} n_B - M_A, \qquad (1)$$

$$n_B = \pi_0 + \tilde{\pi} n_A - M_B, \qquad (2)$$

where u_0 and π_0 are the intrinsic platform values, and $\tilde{u}$ and $\tilde{\pi}$ are externality parameters. For the sake of concreteness, we assume positive intrinsic values and positive externalities. While the intrinsic values are common knowledge, the externality parameters are known to market participants, but not to the platform provider.[3] The provider only knows that $(\tilde{u}, \tilde{\pi}) \in \{(\underline{u}, \underline{\pi}), (\overline{u}, \overline{\pi})\}$ with $0 < \underline{u} < \overline{u} < 1$ and $0 < \underline{\pi} < \overline{\pi} < 1$. We denote the ex ante probability of the realization $(\overline{u}, \overline{\pi})$ by p_0 and assume that this is the prior belief held by the platform provider.[4]

[2] Our notation closely follows Belleflamme and Peitz (2010).

We impose this for the sake of tractability. If side A, say, does not know the strength of the externality it exerts on the other side either, it has to form a belief about it. This, in turn, has to be taken into account by the platform provider who then must form a belief about the true strength of the externalities as well as about the belief of side A. We leave the analysis of such a model for future work. In the present set-up, only the platform provider holds beliefs and learns.

[4] The assumption that the externality parameters are perfectly correlated is clearly restrictive. Imperfect correlation leads to a much more complicated situation with two (correlated) beliefs.

As $\tilde{u}\,\tilde{\pi} \neq 1$, the system (1)–(2) has a unique solution, given by

$$\begin{aligned} n_A(M_A, M_B, \tilde{u}, \tilde{\pi}) &= \frac{u_0 - M_A + \tilde{u}(\pi_0 - M_B)}{1 - \tilde{u}\tilde{\pi}} \,, \\ n_B(M_A, M_B, \tilde{u}, \tilde{\pi}) &= \frac{\pi_0 - M_B + \tilde{\pi}(u_0 - M_A)}{1 - \tilde{u}\tilde{\pi}} \,. \end{aligned}$$

This constitutes the unique Nash equilibrium of the anonymous game that potential participants play for given membership fees.

In every period $t \in [0, \infty[$, the platform provider sets prices (M_A^t, M_B^t) and then observes noisy signals of the quantities $n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})$ and $n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})$. More precisely, the provider observes the cumulative quantity processes N_A^t and N_B^t with increments given by

$$\begin{aligned} dN_A^t &= n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_A dZ_A^t, \\ dN_B^t &= n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_B dZ_B^t, \end{aligned}$$

where Z_B^t and Z_A^t are independent standard Brownian motions and the constants σ_A and σ_B are positive. Note that, using normally distributed shocks, we cannot restrict the observed quantities dN_A^t and dN_B^t to be positive. We will, however, only allow the platform provider to choose prices such that, in expectation, demand is (weakly) positive. Later, when we use quantities as choice variables, we can explicitly rule out negativity.

The platform provider's revenue increment is

$$\begin{aligned} dR_t &= M_A^t dN_A^t + M_B^t dN_B^t \\ &= M_A^t \left[n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_A dZ_A^t \right] + M_B^t \left[n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi})\, dt + \sigma_B dZ_B^t \right]. \end{aligned}$$

We normalize costs to zero. Hence, the platform provider's total expected profits (expressed in per-period terms) are

$$E_{p_0} \left[\int_0^\infty r e^{-rt} dR_t \right],$$

where $r > 0$ is the discount rate. By the martingale property of the stochastic integral with respect to Brownian motion, this expectation reduces to

$$E_{p_0} \left[\int_0^\infty r e^{-rt} \left\{ M_A^t\, n_A(M_A^t, M_B^t, \tilde{u}, \tilde{\pi}) + M_B^t\, n_B(M_A^t, M_B^t, \tilde{u}, \tilde{\pi}) \right\} dt \right].$$

Let p_t be the subjective probability at time t that the platform provider assigns to the realization $(\overline{u}, \overline{\pi})$. Invoking the law of iterated expectations, we can rewrite total expected profits as

$$E_{p_0} \left[\int_0^\infty r e^{-rt} R(M_A^t, M_B^t, p_t)\ dt \right] \tag{3}$$

where

$$R(M_A, M_B, p) \;=\; E_p \left[M_A\, n_A(M_A, M_B, \tilde{u}, \tilde{\pi}) + M_B\, n_B(M_A, M_B, \tilde{u}, \tilde{\pi}) \right]$$

is the expected *current* profit. Using the shorthand notation $\overline{n}_A(M_A, M_B) = n_A(M_A, M_B, \overline{u}, \overline{\pi})$ (and similar definitions for $\underline{n}_A$, $\overline{n}_B$ and $\underline{n}_B$), we can rewrite expected current profit as

$$\begin{aligned} R(M_A, M_B, p) &= M_A\left[(1-p)\underline{n}_A(M_A, M_B) + p\overline{n}_A(M_A, M_B)\right] \\ &\quad + M_B\left[(1-p)\underline{n}_B(M_A, M_B) + p\overline{n}_B(M_A, M_B)\right], \end{aligned}$$

where the terms in square brackets are the expected numbers of participants on the two sides of the market.

If the platform provider were myopic (corresponding to $r = \infty$), it would maximize expected current profit at each instant. Under our parameter restrictions, expected current revenue is strictly concave, thus optimal myopic prices are unique. As a function of the belief p, they are given by

$$\begin{aligned} M_A^\mu(p) &= \tfrac{2(a+b)[(a+b)u_0+(a\underline{u}+b\overline{u})\pi_0]-[a(\underline{\pi}+\underline{u})+b(\overline{\pi}+\overline{u})][(a+b)\pi_0+(a\underline{\pi}+b\overline{\pi})u_0]}{4(a+b)^2-[a(\underline{\pi}+\underline{u})+b(\overline{\pi}+\overline{u})]^2}, \\ M_B^\mu(p) &= \tfrac{2(a+b)[(a+b)\pi_0+(a\underline{\pi}+b\overline{\pi})u_0]-[a(\underline{\pi}+\underline{u})+b(\overline{\pi}+\overline{u})][(a+b)u_0+(a\underline{u}+b\overline{u})\pi_0]}{4(a+b)^2-[a(\underline{\pi}+\underline{u})+b(\overline{\pi}+\overline{u})]^2} \end{aligned}$$

with $a = \frac{1-p}{1-\underline{\pi u}}$ and $b = \frac{p}{1-\overline{\pi u}}$. Owing to our assumptions, the denominator of $M_A^\mu(p)$ and $M_B^\mu(p)$ is strictly positive for all $p \in [0,1]$, so the myopically optimal fees are well-defined and vary continuously with p.

For future reference we denote the myopically optimal revenue by

$$R^\mu(p) = \max_{M_A, M_B} R(M_A, M_B, p) = R(M_A^\mu(p), M_B^\mu(p), p)$$

and the *ex ante full-information pay-off* by

$$\overline{R}(p) = pR^\mu(1) + (1-p)R^\mu(0) .$$

2.2 The e olution of beliefs

We define

$$S(M_A, M_B) = \left[\frac{\overline{n}_A(M_A, M_B) - \underline{n}_A(M_A, M_B)}{\sigma_A}\right]^2 + \left[\frac{\overline{n}_B(M_A, M_B) - \underline{n}_B(M_A, M_B)}{\sigma_B}\right]^2 .$$

Lemma 1 *The beliefs of the price-setting platform provider evolve according to*

$$dp_t \sim N\big(0, p_t^2(1-p_t)^2 S(M_A^t, M_B^t)\, dt\big) . \tag{4}$$

PROOF: See the appendix. □

In the expression for the infinitesimal variance of the change in beliefs, $S(M_A^t, M_B^t)$ measures the information content of the demand observations obtained after setting prices (it is the sum of the squared signal-to-noise ratios of these observations). The more informative the observations are, the more strongly the beliefs react to them. The information content can be shown to be globally convex in the two prices. So moving away from uninformative prices will always generate additional information.

The beliefs $p = 0$ and $p = 1$ are clearly absorbing—if the platform provider is subjectively sure about the true state of the world, no further learning is possible. Non-degenerate beliefs are invariant if and only if both

$$\begin{aligned}\overline{n}_A - \underline{n}_A &= \frac{u_0 - M_A + \overline{u}(\pi_0 - M_B)}{1 - \overline{u\pi}} - \frac{u_0 - M_A + \underline{u}(\pi_0 - M_B)}{1 - \underline{u\pi}} \\ &= \frac{(\overline{u\pi} - \underline{u\pi})(u_0 - M_A) + [\overline{u} - \underline{u} + \overline{u}\underline{u}(\overline{\pi} - \underline{\pi})](\pi_0 - M_B)}{(1 - \overline{u\pi})(1 - \underline{u\pi})}\end{aligned} \tag{5}$$

and

$$\begin{aligned}\overline{n}_B - \underline{n}_B &= \frac{\pi_0 - M_B + \overline{\pi}(u_0 - M_A)}{1 - \overline{u\pi}} - \frac{\pi_0 - M_B + \underline{\pi}(u_0 - M_A)}{1 - \underline{u\pi}} \\ &= \frac{(\overline{u\pi} - \underline{u\pi})(\pi_0 - M_B) + [\overline{\pi}\underline{\pi}(\overline{u} - \underline{u}) + (\overline{\pi} - \underline{\pi})](u_0 - M_A)}{(1 - \overline{u\pi})(1 - \underline{u\pi})},\end{aligned} \tag{6}$$

equal zero. But as we show in the appendix for any pair of myopically optimal prices both differences are strictly positive. Thus for non-degenerate beliefs even a myopic platform provider will always generate information. Hence:

Lemma 2 *With uncertainty about both sides of the market, the only beliefs for which a price-setting platform provider does not learn and acts myopically optimally are those corresponding to subjective certainty, $p = 0$ and $p = 1$.*

PROOF: See the appendix. □

In other words, there are no potentially confounding actions in the sense of Easley and Kiefer (1988). By well-known results, this means that unless the platform provider starts with a degenerate prior belief $p_0 \in \{0, 1\}$, it will always set fees that generate some new information about the state of the world. As a consequence, there is complete learning: beliefs converge to the truth almost surely as $t \to \infty$.[5]

To determine how the information content of observed quantities changes with the prices charged, we look at the partial derivatives of S with respect to M_A and M_B. Figure 1 visualizes this in (M_A, M_B)-space. The two lines represent the combinations of prices for which the partial derivative of the quantity of information with respect to one of the prices equals zero, i.e., $\frac{\partial S}{\partial M_A} = 0$ and $\frac{\partial S}{\partial M_B} = 0$, respectively. For each line, the respective derivative is positive above the line and negative below. Their point of intersection (u_0, π_0) will never be chosen by the platform provider, because it does not generate information and can not be myopically optimal, as stated in Lemma 2.[6]

We show in the appendix that the line $\frac{\partial S}{\partial M_A} = 0$ is always steeper than the line $\frac{\partial S}{\partial M_B} = 0$.

Above either of the two lines, at least one of the implied expected quantities becomes negative, as we again show in the appendix. So only price combinations below both lines are admissible, implying that both partial derivatives of the quantity of information must be negative. We summarize these findings in

[5] If the platform provider were uncertain about the intrinsic platform values (u_0; $_0$) instead of the externalities (u;), computations similar to (5) and (6) would show the quantity of information to be independent of the prices charged. The platform provider would then trivially always set the myopically optimal prices.

[6] Note that revenue equals zero in the intersection point. Thus the platform provider can always increase revenue by marginally lowering one price.

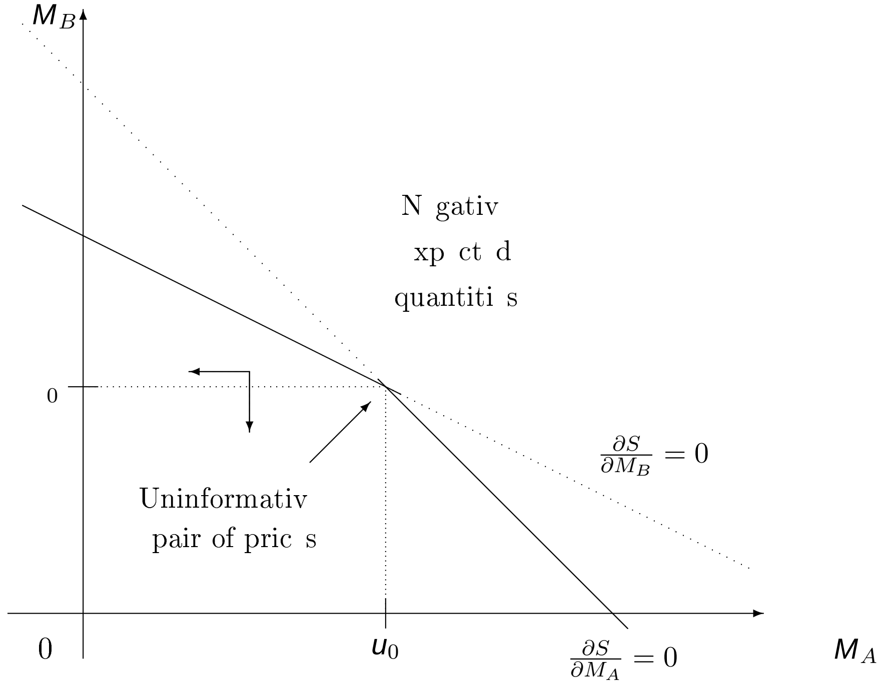

Figure 1: The directions of experimentation in the price plane.

Lemma 3 *Over the admissible range of prices, a price decrease on either side of the market increases the information content of observed quantities, whereas a price increase reduces it.*

For any pair of myopically optimal prices the marginal information content is always negative. This immediately follows from the positivity of $\overline{n}_A - \underline{n}_A$ and $\overline{n}_B - \underline{n}_B$ if evaluated for the myopically optimal prices (see proof of Lemma 2). Thus the myopically optimal prices lie in the lower left area in Figure 1.

3 The Optimal Pricing Strategy

We are now ready to characterize the platform provider's optimal pricing strategy. After addressing the scenario where the provider is uncertain about the size of the externality on both sides of the market, we investigate two simpler scenarios: one in which there is uncertainty about one externality only, and another in which there is only one externality to begin with.

3.1 Uncertainty about both externalities

In view of the objective function (3) and the law of motion (4), standard arguments yield the following Bellman equation for the platform provider's value function:

$$v(p) = \max_{M_A, M_B} \left\{ R(M_A, M_B, p) + \frac{p^2(1-p)^2}{2r} \, S(M_A, M_B) \, v''(p) \right\} \, . \tag{7}$$

We can interpret the second term in the maximand as the value of information, given by the product of the shadow price of information, $p^2(1-p)^2v''/2r$, and the quantity of information, S. For $p \in \{0,1\}$, the value of information is zero, and the platform provider chooses the myopically optimal prices. For all other beliefs, the platform provider experiments, i.e., deviates from the myopic strategy.

Intuitively speaking, the lower the platform provider's discount rate, the farther it might want to deviate from the myopic optimum. To characterize the extent to which this might happen, it is sufficient to establish the optimal behavior in the undiscounted limit where $r = 0$. Unfortunately, this limit turns out to be intractable for a price-setting platform provider. We shall, therefore, consider a quantity-setting platform provider in the next section. However, let us for now state the line of argument we are going to follow several times.

We can restate the Bellman equation (7) as the second-order differential equation

$$\frac{p^2(1-p)^2}{2r}\, v''(p) = \min_{M_A,M_B} \frac{v(p) - R(M_A, M_B, p)}{S(M_A, M_B)} \,. \tag{8}$$

Given p, the set of optimal fees is

$$O(v(p), p) = \arg\min_{M_A,M_B} \frac{v(p) - R(M_A, M_B, p)}{S(M_A, M_B)} \,.$$

Arguing as in Keller and Rady (1999), one shows that the platform provider's value function is the unique solution to (8) subject to the condition that $v(p) = R^\mu(p)$ at $p = 0$ and 1. The value function is decreasing in r at all p in the open unit interval. It converges to the myopic revenue function R^μ as $r \uparrow \infty$, and to the expected full-information payoff $\overline{R}$ as $r \downarrow 0$.

For any given belief p, the set of optimal fees $O(v(p), p)$ thus converges to $O(\overline{R}(p), p)$ as $r \downarrow 0$. As in Keller and Rady (1999), this is the policy correspondence of a platform provider maximizing its undiscounted transient payoff, that is, total expected revenue net of the full-information payoff that it will obtain in the long run; see also Bolton and Harris (2000). Although this policy correspondence does not depend on the value function associated with this objective, the relevant system of first-order conditions is intractable. Thus, we cannot provide explicit solutions for the optimal prices.

3.2 Uncertainty about one externality only ($\overline{\pi} = \underline{\pi}$)

To make the maximal experimentation policy tractable, in this and the following subsection, we impose additional restrictions on our setting.

Let us now assume that only the size of the externality on side A, but not the one on side B, is unknown to the platform provider, i.e., $\underline{\pi} = \overline{\pi}$, for which we simply write π. In this case, the equilibrium demands are

$$\begin{aligned} n_A(M_A, M_B, \tilde{u}) &= \frac{u_0 - M_A + \tilde{u}(\pi_0 - M_B)}{1 - \tilde{u}\pi} \,, \\ n_B(M_A, M_B, \tilde{u}) &= \frac{\pi_0 - M_B + \pi(u_0 - M_A)}{1 - \tilde{u}\pi} \,. \end{aligned}$$

The myopically optimal prices are

$$
\begin{aligned}
M_A^\mu(p) &= \frac{2(c+d)[(c+d)u_0+(c\underline{u}+d\overline{u})\pi_0]-[c(\pi+\underline{u})+d(\pi+\overline{u})](c+d)(\pi_0+\pi u_0)}{4(c+d)^2-[c(\pi+\underline{u})+d(\pi+\overline{u})]^2}\ , \\
M_B^\mu(p) &= \frac{2(c+d)^2(\pi_0+\pi u_0)-[c(\pi+\underline{u})+d(\pi+\overline{u})][(c+d)u_0+(c\underline{u}+d\overline{u})\pi_0]}{4(c+d)^2-[c(\pi+\underline{u})+d(\pi+\overline{u})]^2}
\end{aligned}
$$

with $c = \frac{1-p}{1-\pi\underline{u}}$ and $d = \frac{p}{1-\pi\overline{u}}$.

As $\overline{n}_B - \underline{n}_B = \pi(\overline{n}_A - \underline{n}_A)$, the quantity of information simplifies to

$$
S(M_A, M_B) = \left(\sigma_A^{-2} + \sigma_B^{-2}\pi^2\right)\left[\overline{n}_A(M_A, M_B) - \underline{n}_A(M_A, M_B)\right]^2 .
$$

Non-degenerate beliefs are invariant if and only if $\overline{n}_A - \underline{n}_A = 0$. As already shown in the more general case this difference is always positive for myopically optimal prices. Thus, Lemma 2 extends to the present scenario—there are no potentially confounding actions.

We again employ the (M_A, M_B)-space to visualize the directions of increasing informativeness of observed quantities—see Figure 2.

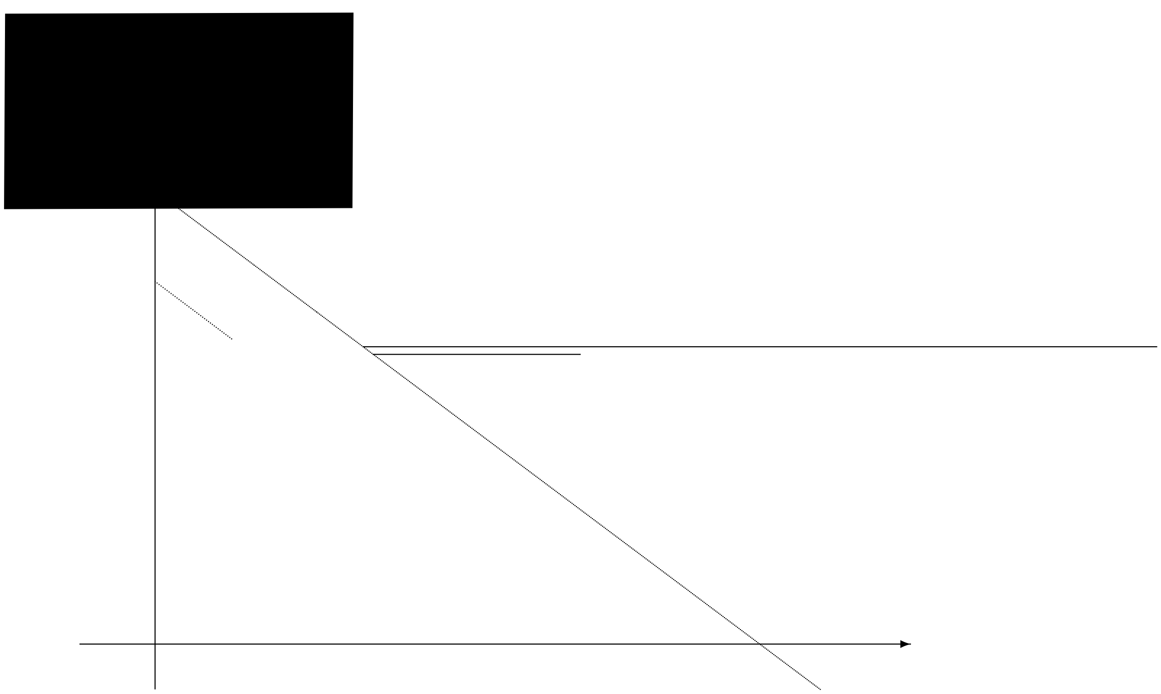

pair of prices fulfilling this reverse inequality cannot be admissible due to the required weak positivity of expected demands. So all admissible price combinations lie in the area where lowering either price increases the information content of observed quantities.

Applying the implicit function theorem to (9), we additionally see that the amount of information is constant for all price pairs fulfilling $M_B = c - \pi M_A$ with some constant c. Hence, all "iso-information curves" are linear and run parallel to the line of no learning (the iso-information curve with an information value of zero).

We show in the appendix that the myopically optimal prices always lie to the left of this line. The experimenting platform provider will deviate from the myopically optimal prices so as to reach an iso-information line that is closer to the origin in Figure 2. This does *not* imply that it will always lower both prices. The reason why this might not be the case is that even the most patient platform provider pursues a secondary goal of revenue maximization in addition to the primary goal of optimal information extraction. To understand the underlying mechanism, recall the optimal policy correspondence for the infinitely patient platform provider,

$$O(\overline{R}(p), p) = \arg\min_{M_A, M_B} \frac{\overline{R}(p) - R(M_A, M_B, p)}{S(M_A, M_B)} \ .$$

The minimand is the smaller, the more information is generated but also the higher current revenue is. Further, for a fixed amount of information (i.e., along an iso-information line) the platform provider will choose prices such as to maximize current revenue (i.e., choose the point on the iso-information line that is tangent to the highest iso-revenue curve).

The iso-revenue curves are convex (ellipses, to be precise). The shape of the locus of tangency points depends on the parameters. Figure 3 illustrates two possible realizations of the locus of optimal price pairs as tangency points of the iso-information lines and the iso-revenue curves. In the left panel the locus points to the lower left – the optimal trade-off between information and current revenue induces a decrease in both prices for increased information. However, if the iso-information lines are rather flat (i.e., π is small) it is optimal to decrease M_B but increase M_A as indicated by the locus of optimal prices in the right panel.

A numerical example comparing myopically optimal prices with optimal prices of an infinitely patient platform provider is given in Table 1. One can see that for low values of π prices on side A are increased by the patient platform provider, while the price for side B is always lowered. Thus the patient platform provider gains additional information on side B relative to the myopic optimum, but might "sacrifice" part of it on side A to achieve the optimal trade-off between revenue and information. For high values of π both prices are lowered.

We have the following result.

Proposition 1 *If the strength of only one externality is unknown, the infinitely patient platform provider optimally decreases the price that generates information relative to the price set by a myopic platform provider. If the known externality is sufficiently weak, the provider optimally increases the other price relative to the myopic benchmark.*

PROOF: The optimal policy of the infinitely patient platform provider is continuous in the known externality parameter π. So the result follows from Proposition 2 below. $\square$

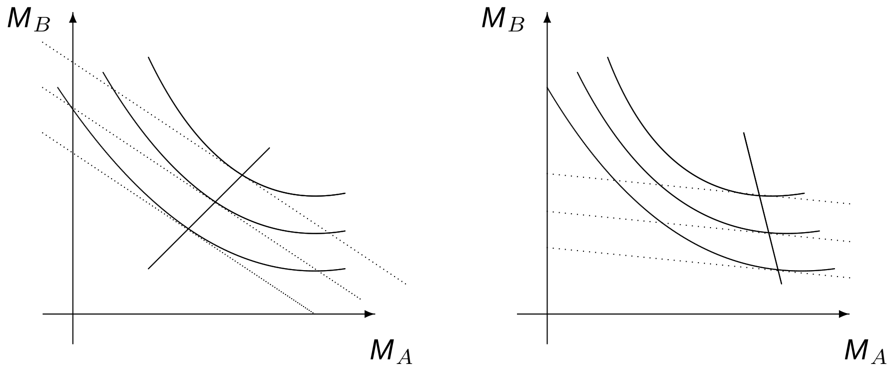

Figure 3: Two examples of iso-information lines (dotted) and iso-revenue curves (solid). The solid line in each case indicates the locus of optimal price combinations.

π	M_A^μ	M_B^μ	M_A^*	M_B^*	$M_A^* - M_A^\mu$	$M_B^* - M_B^\mu$
0.1	0.6459	0.3541	0.6522	0.3241	0.006	-0.030
0.3	0.5943	0.4057	0.5978	0.3782	0.003	-0.028
0.5	0.5217	0.4783	0.5224	0.4557	0.001	-0.023
0.7	0.4094	0.5906	0.4078	0.5766	-0.002	-0.014
0.9	0.2009	0.7991	0.1996	0.7962	-0.001	-0.003

Table 1: Numerical example of optimal price adjustments for varying values of π. All other parameter values are fixed at $u_0 = \pi_0 = \sigma_A = \sigma_B = 1$, $p = 0.5$, $\underline{u} = 0.25$, $\overline{u} = 0.75$.

The intuition for this result is that the platform provider extracts most of the information on one side of the market and recoups part of the rent by increasing the fee on the other side which is only exerting a small externality anyway. We will present a more detailed discussion of this intuition in the next section when we derive explicit solutions for the case $\pi = 0$. For the problem considered in this section there still is no possibility of providing an analytical solution, although the first-order conditions for optimal fees are considerably simpler than in the more general model.

3.3 Only one externality ($\pi = 0$)

We further restrict the setting to the case that only one side is subject to indirect network effects. In particular, we consider the case $\pi = 0$. As an example, consider readers whose utility of a magazine is independent of the amount of advertising.

The additional restriction will substantially facilitate the calculation of optimal policies and generate further important insights. When only one side exerts an externality on the

other, only one of the demands depends on the externality:

$$\begin{aligned} n_A &= u_0 - M_A + \tilde{u}(\pi_0 - M_B) \ , \\ n_B &= \pi_0 - M_B \ . \end{aligned}$$

Uncertainty only enters linearly now, while it entered through fractions before. This significantly reduces complexity.

The myopically optimal prices in this case are

$$\begin{aligned} M_A^\mu(p) &= \frac{2u_0 + u(p)\pi_0}{4 - u^2(p)} \ , \\ M_B^\mu(p) &= \frac{2\pi_0 - u(p)[u_0 + \pi_0 u(p)]}{4 - u^2(p)} \ , \end{aligned}$$

where $u(p) = E_p[\tilde{u}] = p\overline{u} + (1-p)\underline{u}$. The myopic revenue is

$$R^\mu(p) = \frac{\pi_0^2 + u_0^2 + \pi_0 u_0 u(p)}{4 - u^2(p)} \ .$$

The quantity of information now is

$$S(M_A, M_B) = \sigma_A^{-2}(\overline{u} - \underline{u})^2(\pi_0 - M_B)^2.$$

Thus, information is only generated by adjusting the price M_B and observing the quantity n_A. The other price and quantity yield no information on the true size of the externality. The derivative of the quantity of information with respect to M_B is weakly negative for admissible prices $M_B \leq \pi_0$—i.e., reducing M_B increases information. Quantities are uninformative for $M_B = \pi_0$ which once again can be shown never to be myopically optimal, as margiJ/F2211.9552Tf-32e

so for non-degenerate beliefs there is a price increase relative to the myopic optimum. Note that the two price differentials depend linearly and negatively on each other.

The expected quantity on side B clearly increases relative to the myopic optimum as the price M_B goes down. Using the last equation, one can additionally establish that the expected quantity on side A changes by $-\frac{u(p)}{2}[M_B^*(p) - M_B^\mu(p)]$, which is again positive for non-degenerate beliefs. Hence, the platform provider also expects activity on this side to rise relative to the myopic optimum. Overall, therefore, optimal experimentation leads to uniform increases in quantities while price adjustments go in opposite directions.

We summarize these findings in

Proposition 2 *If only one side of the market is subject to indirect network effects, the infinitely patient platform provider optimally decreases the price that generates information relative to the price set by a myopic platform provider, while it optimally increases the other price.*

One might wonder why the platform provider changes M_A at all when it does not learn anything from the change. The answer is that it can set prices according to the following sequential procedure: it first adjusts the information-generating price such as to infer the appropriate amount of information (first-order condition with respect to M_B) and, then, consistently sets the price not yielding information in a profit-maximizing way (first-order condition with respect to M_A). As expected quantities increase in response to the decrease in M_B, it is optimal to increase M_A in order to recoup part of the rent: M_A^* is the revenue maximizing price on side A, given $M_B = M_B^*$. This confirms the intuition already given in Section 3.2.

4 The Quantity-Setting Platform Provider

Let the platform provider now choose quantities $(n_A, n_B) \in \mathbb{R}_+^2$ and observe noisy signals of the prices

$$\begin{aligned} M_A(n_A, n_B, \tilde{u}) &= u_0 + \tilde{u} n_B - n_A, \\ M_B(n_A, n_B, \tilde{\pi}) &= \pi_0 + \tilde{\pi} n_A - n_B, \end{aligned}$$

where $\tilde{u} \in \{\underline{u}, \overline{u}\}$ and $\tilde{\pi} \in \{\underline{\pi}, \overline{\pi}\}$. We assume that $|\tilde{u} + \tilde{\pi}| < 2$ to ensure that the optimal quantities are positive and provide a revenue maximum.[7] We impose the natural restriction that the platform provider can only decide to sell non-negative quantities, while fees are not restricted to be positive. Negative prices can be interpreted as subsidies to one side or (temporarily) both sides of the market.

Note that the price on one side of the market does not depend on the externality parameter on the other side. However, as we assume perfect positive correlation between $\tilde{u}$ and $\tilde{\pi}$, any information gained on one side of the market immediately translates into a similar piece of information on the other side. If there is no a priori uncertainty on one side of the market (if $\underline{u} = \overline{u}$, for instance), then any deviation from the expected price on this side must be attributed to noise, and is, thus, uninformative. Hence, in such a case, the platform

[7] For $|\tilde{u} + \tilde{}| > 2$, corner solutions become optimal.

provider can only experiment on the other side of the market. This is the setting already thoroughly discussed in Keller and Rady (1999). In the following we concentrate, therefore, on non-degenerate uncertainty on both sides of the market. We maintain the assumption that costs are zero. As before, we write p for the subjective probability assigned to the realization $(\overline{u}, \overline{\pi})$.

4.1 Re enues and beliefs

In every period $t \in [0, \infty[$, the platform provider chooses quantities (n_A^t, n_B^t) and then observes the increments $M_A(n_A^t, n_B^t, \tilde{u}, \tilde{\pi})\, dt + \theta_A dW_A^t$ and $M_B(n_A^t, n_B^t, \tilde{u}, \tilde{\pi})\, dt + \theta_B dW_B^t$ of two cumulative price processes where W_A and W_B are independent standard Brownian motions and the constants θ_A and θ_B are positive.

The resulting revenue increment at date t is

$$dR_t = n_A^t \left[M_A(n_A^t, n_B^t, \tilde{u})\, dt + \theta_A dW_A^t \right] + n_B^t \left[M_B(n_A^t, n_B^t, \tilde{\pi})\, dt + \theta_B dW_B^t \right]$$

With the notation

$$\begin{aligned} u(p) &= p\overline{u} + (1-p)\underline{u}, \\ \pi(p) &= p\overline{\pi} + (1-p)\underline{\pi} \end{aligned}$$

and

$$R(n_A, n_B, p) = n_A[u_0 + u(p)\, n_B - n_A] + n_B[\pi_0 + \pi(p)\, n_A - n_B], \tag{10}$$

the platform provider's total expected payoff is

$$E_{p_0}\left[\int_0^\infty r e^{-rt} R(n_A^t, n_B^t, p_t)\, dt \right].$$

The current revenue (10) depends on the expected externalities only through the term $[u(p) + \pi(p)] n_A n_B$, so only the *sum* of the two externalities matters here. In particular, the myopically optimal quantities are

$$\begin{aligned} n_A^\mu(p) &= \frac{2u_0 + \pi_0[u(p) + \pi(p)]}{4 - [u(p) + \pi(p)]^2}, \\ n_B^\mu(p) &= \frac{2\pi_0 + u_0[u(p) + \pi(p)]}{4 - [u(p) + \pi(p)]^2}. \end{aligned}$$

The optimal quantities again are unique due to the convexity of current revenue and have a symmetric structure with interchanged intrinsic platform values. If these platform values coincide, optimal myopic quantities will be the same on both sides of the platform. The corresponding expected prices for each group, however, depend on the specific externality the other group is exerting. They are given by

$$\begin{aligned} M_A^\mu(p) &= \frac{\pi_0[u(p) - \pi(p)] + u_0(2 - \pi(p)[u(p) + \pi(p)])}{4 - [u(p) + \pi(p)]^2}, \\ M_B^\mu(p) &= \frac{u_0[\pi(p) - u(p)] + \pi_0(2 - u(p)[u(p) + \pi(p)])}{4 - [u(p) + \pi(p)]^2}. \end{aligned}$$

The current revenue from the myopically optimal quantities is

$$R^\mu(p) = M_A^\mu(p) n_A^\mu(p) + M_B^\mu(p) n_B^\mu(p) = \frac{u_0^2 + \pi_0^2 + u_0\pi_0[u(p)+\pi(p)]}{4 - [u(p)+\pi(p)]^2} .$$

To describe the law of motions of beliefs, we define

$$\Sigma(n_A, n_B) = \left(\frac{\overline{u} - \underline{u}}{\theta_A}\right)^2 n_A^2 + \left(\frac{\overline{\pi} - \underline{\pi}}{\theta_B}\right)^2 n_B^2.$$

This measures the information content of price observations after quantities have been set. It is straightforward to see that Σ is globally convex.

Lemma 4 *The beliefs of the quantity-setting platform provider evolve according to*

$$dp_t \sim N\big(0, p_t^2(1-p_t)^2\, \Sigma(n_A^t, n_B^t)\, dt\big) \tag{11}$$

PROOF: The proof is similar to the price-setting case and therefore omitted. □

Beliefs are stationary if $p = 0$ or 1 or if both quantities are set to 0. According to our next result, the latter cannot be myopically optimal under any belief. So there are no potentially confounding actions, which again implies complete learning in the long run.

Lemma 5 *In the model with quantities as choice variables, there does not exist a non-degenerate belief such that the platform provider does not learn and acts myopically optimally.*

PROOF: See the appendix. □

As the function Σ is increasing in both n_A and n_B, moreover, we obviously have

Lemma 6 *For a quantity-setting platform provider, a quantity increase on either side of the market increases the information content of observed prices, whereas a quantity decrease reduces it.*

4.2 Optimal quantities

Under discounting at rate r, the Bellman equation is

$$v(p) = \max_{n_A, n_B} \left\{ R(n_A, n_B, p) + \frac{p^2(1-p)^2}{2r}\, \Sigma(n_A, n_B)\, v''(p) \right\} .$$

By standard arguments, $v'' > 0$, so the maximand is the sum of a concave quadratic form, R, and a convex one, Σ. This sum is concave if and only if the shadow price of information, $\frac{p^2(1-p)^2}{2r}\, v''(p)$, is sufficiently small.

We can rewrite the Bellman equation as

$$\frac{p^2(1-p)^2}{2r}\; v''(p) = \min_{n_A, n_B} \frac{v(p) - R(n_A, n_B, p)}{\Sigma(n_A, n_B)} , \tag{12}$$

leading to the optimal policy correspondence

$$O(v(p), p) = \arg\min_{n_A, n_B} \frac{v(p) - R(n_A, n_B, p)}{\Sigma(n_A, n_B)} \ .$$

In the undiscounted limit, this correspondence becomes $O(\overline{R}(p), p)$ where $\overline{R}(p) = pR^\mu(1) + (1-p)R^\mu(0)$ is once more the expected full-information payoff. In general, the associated first-order conditions involve third-order polynomials in the two policy variables n_A and n_B.

For $r > 0$, equation (12) and the fact that $v \leq \overline{R}$ allow us to conclude that

$$\frac{p^2(1-p)^2}{2r} v''(p) \leq \frac{\overline{R}(p) - R^\mu(p)}{\Sigma(n_A^\mu, n_B^\mu)} \ .$$

So, if the quadratic form

$$R(n_A, n_B, p) + \frac{\overline{R}(p) - R^\mu(p)}{\Sigma(n_A^\mu, n_B^\mu)} \Sigma(n_A, n_B)$$

is concave, optimal quantities are fully determined by first-order conditions, no matter how patient the platform provider is. This concavity obviously holds for beliefs close to subjective certainty, and for arbitrary beliefs if the signal-to-noise ratios $\frac{\overline{u}-u}{\theta_A}$ and $\frac{\overline{\pi}-\pi}{\theta_B}$ are sufficiently small.

Equation (12) reveals an interesting feature of experimentation in two-sided markets.[8] Of all quantity combinations yielding the optimal amount of information, the optimal policy necessarily chooses the pair of quantities that minimizes the distance between the current value of the problem and the current revenue, and, hence, maximizes current revenue. In a one-sided market setting with uncertainty about the slope of a linear demand function, the set of quantities generating a certain amount of information would be a singleton.[9]

4.3 Symmetric signal-to-noise ratios

To obtain more tractable first-order conditions, we assume henceforth that the signal-to-noise

Proposition 3 *Suppose that condition (13) holds and $u_0 \neq \pi_0$. Then the quantities set by an infinitely patient platform provider are*

$$\begin{aligned} n_A^*(p) &= \frac{1}{2(u_0^2-\pi_0^2)[u(p)+\pi(p)]}\Bigg\{\pi_0(\pi_0^2+u_0^2)+4\overline{R}(p)u_0[u(p)+\pi(p)] \\ &\qquad - \pi_0\sqrt{(u_0^2-\pi_0^2)^2+(2u_0\pi_0+4\overline{R}(p)[u(p)+\pi(p)])^2}\Bigg\}, \\ n_B^*(p) &= \frac{1}{2(u_0^2-\pi_0^2)[u(p)+\pi(p)]}\Bigg\{(-u_0(\pi_0^2+u_0^2)-4\overline{R}(p)\pi_0[u(p)+\pi(p)] \\ &\qquad + u_0\sqrt{(u_0^2-\pi_0^2)^2+(2u_0\pi_0+4\overline{R}(p)[u(p)+\pi(p)])^2}\Bigg\}. \end{aligned}$$

PROOF: See the appendix. □

The knife-edge case $u_0 = \pi_0$ will be covered later. Note that for $\pi_0 > u_0$, both numerator and denominator of $n_A^*(p)$ and $n_B^*(p)$ are negative, so the quantities remain positive. One can easily verify that these quantities coincide with the myopically optimal ones for $p = 0$ and $p = 1$.

The quantities in Proposition 3 are the ones an infinitely patient platform provider would choose. An infinitely impatient platform provider would choose the myopic optimum. We have also shown already that any deviations from the myopic optimum will be quantity expansions up to the point where the amount of information acquired is optimal. Arguing as in Keller and Rady (1999), one shows that these deviations decrease with the discount rate. We have thus fully characterized the range of experimentation within which a platform provider with finite positive discount rate will set its quantities.

The expected prices $M_A^*(p)$ and $M_B^*(p)$ given the quantities $n_A^*(p)$ and $n_B^*(p)$ are straightforward to calculate. Comparing them with the myopic optimum allows us to show explicitly what we have seen implicitly in the price-setting model: even if both externalities are smaller than 1 in expectation ($|u(p)| < 1, |\pi(p)| < 1$), there are parameter constellations such that, on one side of the market, the expected price for the patient platform provider is higher than the one for the myopic platform provider, as exemplarily shown for the price M_A in Figure 4.

Numerical examples confirm that the intuition from the case of the price-setting platform provider with $\pi = 0$ qualitatively carries over. If, e.g., the externality side A is exercising is small (i.e., π is close to zero) or the relative externality is small (i.e., $\frac{\pi(p)}{u(p)}$ is small), it is beneficial for the infinitely patient platform provider to strongly increase the number of participants on side B, which leads to reductions in M_B but induces an increase in M_A. Hence, as argued above, the platform provider raises the quantities to optimally gather information, but the increase in n_B outweighs the increase in n_A by such a large amount that M_A also rises. Therefore, the platform provider recoups parts of the rent coming from higher participation on both sides by inducing a higher than myopically optimal price on one side of the market.

To illustrate the difference between the policies of an infinitely patient and a myopic platform provider, we provide some simulation results. The graphs were obtained for the

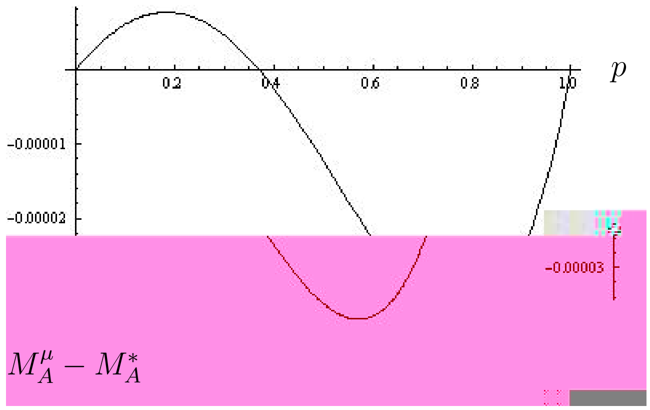

Figure 4: Difference between the prices set by a myopic and an infinitely patient platform provider as a function of belief p for the following parameter constellation: $u_0 = 0.1$, $\pi_0 = 0.7$, $\underline{u} = 0.8$, $\overline{u} = 0.9$, $\underline{\pi} = 0.1$, $\overline{\pi} = 0.2$, $\theta_A = \theta_B = 1$

following parameter constellation: $u_0 = 0.1$, $\pi_0 = 0.8$, $\underline{u} = 0.7$, $\overline{u} = 0.9$, $\underline{\pi} = 0.5$, $\overline{\pi} = 0.7$, $\theta_A = \theta_B = 1$, $\tilde{p}_0 = p_0 = 0.5$ and the "true" values are $\tilde{u} = \overline{u}$, $\tilde{\pi} = \overline{\pi}$.[10]

Figure 5 illustrates that, as predicted, learning is faster for the patient policy—beliefs converge faster to the true state. Figure 6 shows how expected prices evolve on side A. In early stages the patient platform provider induces a lower expected price than in the myopic optimum. So it forfeits revenue but acquires additional information. This pays off in later stages, when expected "patient" prices converge faster to the full-information level. The evolution of expected prices on side B is shown in Figure 7. Here experimentation is quite strong with large deviations from the myopic optimum in early periods. Again the patient platform provider gathers information significantly faster, thus approaching the true value much more rapidly than its myopic counterpart. The per-period revenues depicted in Figure 8 show the benefits of both policies. While the myopic policy creates higher revenues in the early periods, revenues in later periods are higher for the more patient platform provider as its beliefs approach the true state of the world more rapidly.

4.4 Symmetric intrinsic alues

Maintaining the assumption of symmetric signal-to-noise ratios, we now assume that $u_0 = \pi_0$. This may be considered a rather strong assumption since it states that the intrinsic value of the platform is the same for all users. However, this seems to be the appropriate benchmark in a number of examples, such as night clubs and matching agencies.[11] Imposing this additional assumption significantly simplifies the expressions for the optimal quantities.

[10] Simulations were carried out using Wolfram Mathematica 6. Normal shocks were generated by randomly drawing from the normal distribution using the commands "RandomReal" and "NormalDistribution" with mean equal to 0 and variances equalling $_A$ and $_B$ respectively.

[11] We admit that it appears to be an inappropriate benchmark in other examples, such as for merchants and customers in the credit card market.

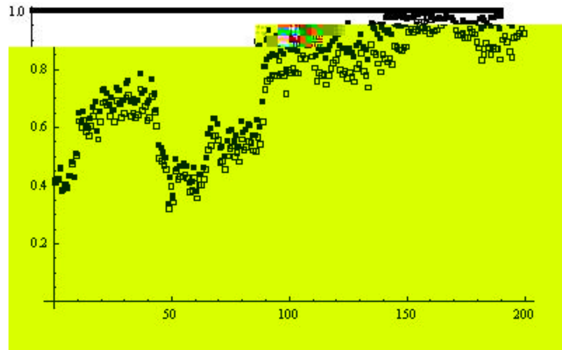

Figure 5: Evolution of beliefs for the myopic policy (white squares) and the infinitely patient policy (black squares), and true state (thick line).

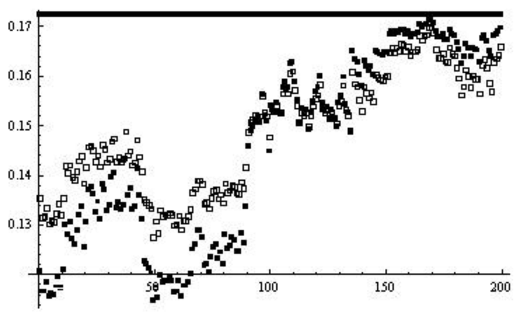

Figure 6: Evolution of expected prices on side A for the myopic policy (white squares) and the infinitely patient policy (black squares), as well as optimal price under full information (thick line).

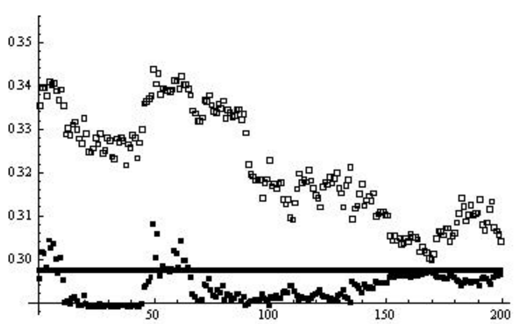

Figure 7: Evolution of expected prices on side B for optimal myopic policy (white squares) and for optimal patient policy (black squares), as well as optimal price under full information (thick line).

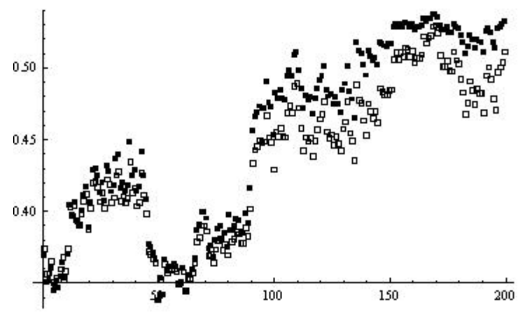

Figure 8: Evolution of per-period revenue for the myopic policy (white squares) and the infinitely patient policy (black squares).

Proposition 4 *Suppose that condition (13) holds and $u_0 = \pi_0 = C_0$. Then the optimal policy of an infinitely patient platform provider is symmetric across market sides and linear in the current belief:*

$$n_A^*(p) = n_B^*(p) = \frac{\overline{R}(p)}{C_0} = C_0 \left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1-p}{2 - (\underline{u} + \underline{\pi})} \right].$$

PROOF: See the appendix. □

The intuition for the symmetry of the optimal quantities is as follows. With identical intrinsic platform values, the myopically optimal quantities are symmetric. With identical signal-to-noise ratios, moreover, the incentive to deviate from the myopic optimum is the same in both quantity dimensions.

Owing to the linearity of the infinitely patient policy, the visualization of the range of experimentation is very simple—see Figure 9. It is the area bounded from below by the myopic policy and bounded from above by the line joining the quantities that are optimal under subjective certainty.

Expected prices need not be symmetric. They are

$$\begin{aligned} M_A^*(p) &= C_0 + (u(p) - 1)\frac{\overline{R}(p)}{C_0} = C_0 \left\{ 1 + (u(p) - 1) \left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1-p}{2 - (\underline{u} + \underline{\pi})} \right] \right\}, \\ M_B^*(p) &= C_0 + (\pi(p) - 1)\frac{\overline{R}(p)}{C_0} = C_0 \left\{ 1 + (\pi(p) - 1) \left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1-p}{2 - (\underline{u} + \underline{\pi})} \right] \right\}. \end{aligned}$$

Depending on the current belief, either both prices are lower than the intrinsic platform value or one price is lower and the other one is higher. Owing to the restriction that $|u(p)+\pi(p)| < 2$, there cannot arise situations where both prices are higher than the intrinsic platform value. The ordering of expected prices depends on the sizes of the externalities and on the current belief, and may even change with beliefs. Let $\underline{u} < \underline{\pi} < \overline{\pi} < \overline{u}$, for example. For high values of p, then, $u(p)$ will exceed $\pi(p)$ and side A will have to pay a higher charge than side B, while for low values of p the reverse is true.

As to the comparison between optimal and myopic expected prices, we have

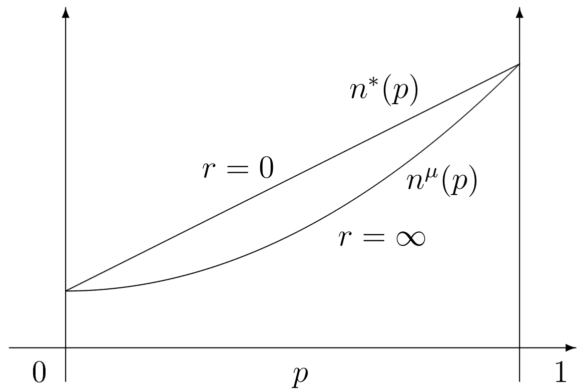

Figure 9: Range of quantity experimentation in the symmetric case.

Corollary 1 *Under the assumptions of Proposition 4, the long-term optimal expected price exceeds its short-term counterpart on a given side of the market if and only if the expected externality that this side experiences is greater than* 1.

PROOF: See the appendix. □

The myopic and the long-term optimal expected prices coincide for beliefs equaling 0 or 1 or if the expected externality equals 1. As $|u(p) + \pi(p)| < 2$, this of course implies that at any time at most one expected price can exceed the myopic benchmark. It also implies that, with equal intrinsic values and for the "standard case" of both externalities smaller than 1, experimentation always means that both expected prices will decrease relative to the myopic optimum.

5 Conclusion

We have studied optimal behavior of a monopolist platform provider in a two-sided market with uncertainty. The demand externalities considered are linear on both sides, fees are charged for participation in the market, but not per transaction. In this respect, our setting is closest to the monopoly setting analyzed in Armstrong (2006).

When the platform provider faces uncertainty about the size of the externality and wants to maximize its expected lifetime profits, it faces the basic trade-off between the conflicting aims of maximizing current payoff and maximizing the information content of the signals it observes. We have characterized the boundaries for optimal policies depending on how much weight the platform provider assigns to future profit. If it does not put any weight on the future ($r = \infty$), it chooses the myopically optimal actions given its current belief. As there is no potentially confounding action, even the myopic platform provider continuously accumulates information about the true state of the world and will, in the limit as time tends to infinity, almost surely learn the true state.

If the platform provider puts some weight on the future, it will deviate from the short-sighted policy and invest in learning. The upper bound on such experimentation is given by the optimal policy of an infinitely patient platform provider ($r = 0$).

If the platform provider chooses quantities, experimentation will always result in quantities higher than in the myopic optimum on both sides of the market. However, the effect on expected prices is ambiguous. Depending on the parameter constellation, either both prices will be lower than in the myopic case or one price will be above and one price below the myopically optimal prices. The price on one market side will go up if the externality this side is exercising is low while the externality it is experiencing is high. The higher price recoups part of the rent created by the higher participation on the other side of the market. These results are in line with our findings in special cases of the version of our model in which the platform provider sets prices.

Our analysis concerns an unrestricted monopoly platform. Future work may want to look at markets with multiple differentiated platforms. As a starting point, it would be interesting to analyze duopoly experimentation in a two-sided market in which there is single-homing on both sides and full observability of actions and outcomes. According to the first assumption, both market sides are symmetric in structure—the model with certainty has been analyzed by Armstrong (2006). According to the second assumption, posterior beliefs of platform providers are identical all the time.

In such a duopoly, a participant acquired by one platform provider is a participant lost for the competitor. Owing to indirect network effects, this makes demand more sensitive to price changes than demand in the monopoly setting with a fixed outside option that has been analyzed in this paper. Therefore, one may conjecture that gaining information about the size of the network effect becomes more important. As has been pointed out in the literature on duopoly experimentation (e.g., Mirman et al. 1994, Harrington 1995, Keller and Rady 2003), however, the public information generated by market signals may have a negative value, in which case the duopolists have an incentive to generate *less* information than in the myopic equilibrium.

Suppose for instance that market participation is perfectly price-inelastic, as is the case in the Hotelling-type model introduced by Armstrong (2006). Then, learning does not increase future equilibrium profits in expectation because profits are linear in beliefs. Since deviations from the myopic best-response are costly, we conjecture that patient platform operators do not behave differently from infinitely impatient ones, and learn only passively. The duopoly setting merits further, more general investigation, and it would be interesting to understand the effect of the degree of differentiation on experimentation in a two-sided market.

Appendix

Proof of Lemma 1

Giv n a pair of pric s $(M_A; M_B)$, th obs rv d quantity incr m nts ar

$$\begin{pmatrix} dN_A \\ dN_B \end{pmatrix} = \begin{pmatrix} \tilde{n}_A \\ \tilde{n}_B \end{pmatrix} dt + \begin{pmatrix} {}_A & 0 \\ 0 & {}_B \end{pmatrix} \begin{pmatrix} dZ_A \\ dZ_B \end{pmatrix}$$

with $\tilde{n}_A = n_A(M_A; M_B; \tilde{u}; \tilde{\ })$ and $\tilde{n}_B = n_B(M_A; M_B; \tilde{u}; \tilde{\ })$.

Giv n th subj ctiv probability p curr ntly assign d to th stat $(\overline{u}; \overline{\ })$, th v ctor of xp ct d d mands is

$$\begin{pmatrix} E_p[\tilde{n}_A] \\ E_p[\tilde{n}_A] \end{pmatrix} = p \begin{pmatrix} \overline{n}_A \\ \overline{n}_B \end{pmatrix} + (1-p) \begin{pmatrix} \underline{n}_A \\ \underline{n}_B \end{pmatrix}$$

with $\overline{n}_A = n_A(M_A; M_B; \overline{u}; \overline{\ })$ tc.

According to Lipts r and Shiryay v (1977), th infinit simal chang in b li fs is giv n by

$$dp = p \begin{pmatrix} \overline{n}_A - E_p[\tilde{n}_A] \\ \overline{n}_B - E_p[\tilde{n}_B] \end{pmatrix} \begin{pmatrix} {}_A^{-1} & 0 \\ 0 & {}_B^{-1} \end{pmatrix} \begin{pmatrix} d\overline{Z}_A \\ d\overline{Z}_B \end{pmatrix}$$

wh r

$$\begin{pmatrix} d\overline{Z}_A \\ d\overline{Z}_B \end{pmatrix} = \begin{pmatrix} {}_A^{-1} & 0 \\ 0 & {}_B^{-1} \end{pmatrix} \begin{pmatrix} dN_A - E_p[\tilde{n}_A] \\ dN_B - E_p[\tilde{n}_B] \end{pmatrix}$$

is th incr m nt of a standard two-dim nsional Brownian motion r lativ to th platform provid r's information filtration.

Simplifying th xpr ssion for dp, w obtain

$$dp = p(1-p)(\overline{n}_A - \underline{n}_A) \; {}_A^{-1} d\overline{Z}_A + p(1-p)(\overline{n}_B - \underline{n}_B) \; {}_B^{-1} d\overline{Z}_B:$$

As $d\overline{Z}_A$ and $d\overline{Z}_B$ 6.551.8f11.4381.689Td(6.551.8fi8lth)82(distribut)T1522. m 8( )1z -375(8(v)-33282(v9-1ar28(c

Slopes of the lines $\frac{\partial S}{\partial M_A} = 0$ and $\frac{\partial S}{\partial M_B} = 0$ in Figure 1

Writing

$$\begin{aligned} e &= \overline{\ }u - \underline{\ u}; \\ f &= \overline{\ } - \underline{\ } + (\overline{u} - \underline{u})\overline{\ }\underline{\ }; \\ g &= \overline{u} - \underline{u} + (\overline{\ } - \underline{\ })\overline{u}\underline{u}; \end{aligned}$$

w s that $\frac{\partial S}{\partial M_A} = 0$ if and only if

$$M_B = {}_0 + \frac{{}_B^2 e^2 + {}_A^2 f^2}{{}_B^2 eg + {}_A^2 ef}(u_0 - M_A);$$

and $\frac{\partial S}{\partial M_B} = 0$ if and only if

$$M_B = {}_0 + \frac{{}_B^2 eg + {}_A^2 ef}{{}_B^2 g^2 + {}_A^2 e^2}(u_0 - M_A):$$

Th diff r nc of th slop co ffici nts is

$$\frac{{}_B^2 e^2 + {}_A^2 f^2}{{}_B^2 eg + {}_A^2 ef} - \frac{{}_B^2 eg + {}_A^2 ef}{{}_B^2 g^2 + {}_A^2 e^2} = \frac{{}_B^2\ {}_A^2 (e^2 - fg)^2}{(\ {}_B^2 eg + {}_A^2 ef)(\ {}_B^2 eg + {}_A^2 ef)} > 0$$

as $e^2 - fg = -(\overline{\ } - \underline{\ })(\overline{u} - \underline{u})(1 - \overline{\ }\overline{u})(1 - \underline{\ u}) < 0$. □

Negativity of quantities above the lines $\frac{\partial S}{\partial M_A} = 0$ and $\frac{\partial S}{\partial M_B} = 0$ in Figure 1

W show that in th r gion wh r th information cont nt of quantiti s is incr asing in pric s, th xp ct d quantity on at l ast on sid of th mark t must b n gativ . For th partial d rivativ s to b positiv ith r $\overline{n}_A - \underline{n}_A$ or $\overline{n}_B - \underline{n}_B$ has to b n gativ . This translat s into

$$M_B > {}_0 + \frac{e}{g}(u_0 - M_A) \tag{14}$$

or

$$M_B > {}_0 + \frac{f}{e}(u_0 - M_A) \tag{15}$$

On th oth r hand positiv d mands $\overline{n}_A$, $\overline{n}_B$ imply

$$M_B < {}_0 + \frac{1}{\underline{u}}(u_0 - M_A) \tag{16}$$

and

$$M_B < {}_0 + \overline{\ }(u_0 - M_A) \tag{17}$$

Comparing th slop s of th s four in qualiti s, for $M_A > u_0$ (16) contradicts (14) and (15), whil for $M_A < u_0$ (17) do s so. For $M_A = u_0$ th contradiction is obvious. From th proof it additionally follows that xp ct d quantiti s might b n gativ v n b low th lin s $\frac{\partial S}{\partial M_A} = 0$ and $\frac{\partial S}{\partial M_B} = 0$. □

Proof of Lemma 5

Quantiti s ar uninformativ if and only if th y both qual z ro. W show that this can n v r b myopically optimal. Assum $n_A^\mu = n_B^\mu = 0$. This translat s into

$$
\begin{aligned}
2u_0 + \quad {}_0[u(p) + \quad (p)] &= 0; \\
2 \quad {}_0 + u_0[u(p) + \quad (p)] &= 0;
\end{aligned}
$$

and adding up yi lds $(2 - [u(p) + \quad (p)])(u_0 - \quad {}_0) = 0$. As $|u(p) + \quad (p)| < 2$, this can only b tru for $\quad {}_0 = u_0$. But th n th first quation abov b com s $u_0(2 + [u(p) + \quad (p)]) = 0$, which is rul d out by th positivity r strictions on u_0, $u(p)$ and $\quad (p)$. □

Proof of Propositions 3 and 4

With $C_A = \quad {}_A^{-1}(\overline{u} - \underline{u})$ and $C_B = \quad {}_B^{-1}(^{-} - {}_{\_})$, th first-ord r conditions for (12) can b writt n as

$$
\begin{aligned}
&(u_0 + [u(p) + \quad (p)]n_B - 2n_A)(C_A^2 n_A^2 + C_B^2 n_B^2) \\
&\quad + \ 2C_A^2 n_A \left[\overline{R}(p) - (u_0 + u(p)n_B - n_A)\, n_A - (\quad {}_0 + \quad (p)n_A - n_B)\, n_B\right] = 0; \\
&(\quad {}_0 + [u(p) + \quad (p)]n_A - 2n_B)(C_A^2 n_A^2 + C_B^2 n_B^2) \\
&\quad + \ 2C_B^2 n_B \left[\overline{R}(p) - (u_0 + u(p)n_B - n_A)\, n_A - (\quad {}_0 + \quad (p)n_A - n_B)\, n_B\right] = 0:
\end{aligned}
$$

Und r condition (13), this syst m simplifi s to

$$
\begin{aligned}
(u_0 + [u(p) + \quad (p)]n_B)(n_B^2 - n_A^2) + 2(\overline{R}(p) - n_B \quad {}_0)\, n_A &= 0 \\
(\quad {}_0 + [u(p) + \quad (p)]n_A)(n_A^2 - n_B^2) + 2(\overline{R}(p) - n_A u_0)\, n_B &= 0:
\end{aligned}
$$

For $u_0 \neq \quad {}_0$, th pair of quantiti s stat d in Proposition 3 constitut s th uniqu solution to th s   quations. For $u_0 = \quad {}_0 = C_0$, on can asily s   that s tting both quantiti s qual to $\frac{\overline{R}(p)}{C_0}$ solv s th syst m. □

Proof of Corollary 1

For $u_0 = \quad {}_0 = C_0$, th myopically optimal xp ct d pric on sid A simplifi s to

$$
M_A^\mu(p) = \frac{C_0[1 - \quad (p)]}{2 - [u(p) + \quad (p)]} ;
$$

so th pric diff r nc $M_A^*(p) - M_A^\mu(p)$ has th sam sign as

$$
1 + (u(p) - 1)\left[\frac{p}{2 - (\overline{u} + ^{-})} + \frac{1 - p}{2 - (\underline{u} + {}_{\_})}\right] - \frac{1 - \quad (p)}{2 - [u(p) + \quad (p)]} :
$$

Multiplying with $2 - [u(p) + \quad (p)]$ and simplifying, w s   that this in turn has th sam sign as

$$
(u(p) - 1)\left\{(2 - [u(p) + \quad (p)])\left[\frac{p}{2 - (\overline{u} + ^{-})} + \frac{1 - p}{2 - (\underline{u} + {}_{\_})}\right] - 1\right\}:
$$

Th   xpr ssion in curly brack ts is strictly concav in p; as it vanish s at $p = 0$ and $p = 1$, it is strictly positiv for $0 < p < 1$. Th proof for sid B is analogous. □

References

Anderson, S. P. and S. Coate (2005): “Market Provision of Broadcasting: A Welfare Analysis", *Review of Economic Studies*, 72, 947–972.

Armstrong, M. (2006): “Competition in Two-sided Markets", *RAND Journal of Economics*, 37, 668–691.

Belleflamme, P. and M. Peitz (2010): *Industrial Organization: Markets and Strategies.* Cambridge: Cambridge University Press.

Bolton, P. and C. Harris (2000): “Strategic Experimentation: the Undiscounted Case," in: *Incentives, Organizations and Public Economics – Papers in Honour of Sir James Mirrlees*, ed. by P.J. Hammond and G.D. Myles. Oxford: Oxford University Press, 53–68.

Easley, D. and N.M. Kiefer (1988): “Controlling a Stochastic Process with Unknown Parameters", *Econometrica*, 56, 1045–1064.

Harrington, J.E. Jr. (1995): “Experimentation and Learning in a Differentiated-Products Duopoly", *Journal of Economic Theory*, 66, 175–288.

Kaiser, U. and J. Wrig t (2006): “Price Structure in Two-Sided Markets: Evidence from the Magazine Industry", *International Journal of Industrial Organization*, 24, 1–28.

Keller, G. and S. Rady (1999): “Optimal Experimentation in a Changing Environment", *Review of Economic Studies*, 66, 475–507.

Keller, G. and S. Rady (2003): “Price Dispersion and Learning in a Dynamic Differentiated-Goods Duopoly", *RAND Journal of Economics*, 34, 138–165.

Liptser, R.S., and A.N. S iryayev (1977): *Statistics of Random Processes I.* New York: Springer-Verlag.

Mirman, L.J., Samuelson, L., and Sc lee, E. (1994): “Strategic Information Manipulation in Duopolies", *Journal of Economic Theory*, 62, 363–384.

Nocke, V., Peitz, M. and K. Sta l (2007): “Platform Ownership", *Journal of the European Economic Association*, 5, 1130–1160.

Prescott, E.C. (1972): “The Multiperiod Control Problem under Uncertainty", *Econometrica*, 40, 1043–1058.

Roc et, J.C. and J. Tirole (2003): “Platform Competition in Two-Sided Markets", *Journal of the European Economic Association*, 1, 990–1029.

Roc et, J.C. and J. Tirole (2006): “Two-Sided Markets: An Overview", *RAND Journal of Economics*, 37, 645–667.

Rot sc ild, M. (1974): “A Two-Armed Bandit Theory of Market Pricing", *Journal of Economic Theory*, 9, 185–202.

Rysman, M. (2004): “Competition between Networks: A Study of the Market for Yellow Pages", *Review of Economic Studies*, 71, 483–512.

Rysman, M. (2009): “The Economics of Two-Sided Markets", *Journal of Economic Perspectives*, 71, 125–143.