GOVERNANCE AND THE EFFICIENCY
OF ECONOMIC SYSTEMS
GESY

Discussion Paper No. 365

Experimentation in Two-Sided Markets

Martin Peitz*
Sven Rady**
Piers Treppers***

*Department of Economics, University of Mannheim, Germany,
E-Mail: martin.peitz@googlemail.com
**Department of Economics, University of Bonn, Germany,
E-Mail: rady@hcm.uni-bonn.de
***Department of Economics, University of Munich, Germany,
E-Mail: piers.treppers@lrz.uni-muenchen.de

November 2011

Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 is gratefully acknowledged.

Sonderforschungsbereich/Transregio 15 · www.sfbtr15.de
Universität Mannheim · Freie Universität Berlin · Humboldt-Universität zu Berlin · Ludwig-Maximilians-Universität München
Rheinische Friedrich-Wilhelms-Universität Bonn · Zentrum für Europäische Wirtschaftsforschung Mannheim

Speaker: Prof. Dr. Urs Schweizer. · Department of Economics · University of Bonn · D-53113 Bonn,
Phone: +49(0228)739220 · Fax: +49(0228)739221

Experimentation in Two-Sided Markets*

Martin Peitz† Sven Rady‡ Piers Trepper§

November 15, 2011

Abstract

We study optimal experimentation by a monopolistic platform in a two-sided market framework. The platform provider faces uncertainty about the strength of the externality each side is exerting on the other. It maximizes the expected present value of its profit stream in a continuous-time infinite-horizon framework by setting participation fees or quantities on both sides. We show that a price-setting platform provider sets a fee lower than the myopically optimal level on at least one side of the market, and on both sides if the two externalities are of approximately equal strength. If the externality that one side exerts is sufficiently weaker than the externality it experiences, the optimal fee on this side exceeds the myopically optimal level. We obtain analogous results for expected prices when the platform provider chooses quantities. While the optimal policy does not admit closed-form representations in general, we identify special cases in which the undiscounted limit of the model can be solved in closed form.

Keywords: Two-Sided Market, Network Effects, Monopoly Experimentation, Bayesian Learning, Optimal Control
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, and of the CEPR Applied IO School and Conference, Tel Aviv, 24-27 May 2011. 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, Germany, martin.peitz@googlemail.com; also affiliated with CEPR, CESifo, Encore, and ZEW.

‡Department of Economics, University of Bonn, Germany, rady@hcm.uni-bonn.de; also affiliated with CEPR and CESifo.

§Department of Economics, University of Munich, Germany; 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 positive externality each side of the market is exerting on the other and, therefore, 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 setting.

In every instant of time, the platform provider's actions determine its current profit as well as the amount of information received. Thus, there is a trade-off between maximizing current profit and extracting information that will increase future profits. The higher the rate at which future profits are discounted, the more important current profit becomes, up to the extreme of myopic behavior which completely ignores information acquisition. 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 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.[1] Prices take the form of membership or subscription fees. In both versions, we first compute the myopic benchmark, then investigate the optimal experimentation policy of a forward-looking platform provider, and finally consider the undiscounted limit in which experimentation is maximal. Our investigation of the optimal experimentation policy relies on an analysis of the first-order conditions associated with the platform provider's Bellman equation; we show that the second-order conditions for a maximum are always satisfied. In general, there are no closed-form solutions for the platform provider's value function and optimal policy. Turning to the undiscounted limit, by contrast, we are able to identify special cases of the model that yield a maximal experimentation policy in closed form.

In the price-setting version of the model, we first establish that the experimenting platform provider will charge a fee lower than the myopic benchmark on at least one side of the market. This immediately implies that if the two sides are approximately symmetric with respect to the participants' intrinsic platform value, the strength of the externality and the informativeness of observed quantities, the provider will charge fees lower than their myopically optimal counterparts on both sides of the market. In sufficiently asymmetric settings, however, the platform provider may find it optimal to charge a fee higher than the myopic benchmark on one side of the market. More precisely, we show that a price increase may occur on a side that exerts a low externality on the other side, yet itself benefits from a strong externality in the other direction. In such a situation, it is optimal to increase participation on the side that exerts the strong externality by lowering the fee there and to extract part of the additional surplus through a higher fee on the side that exerts the weak externality.

In the quantity-setting variant of the model, we obtain analogous results for expected prices. While the platform provider increases the quantity on both sides of the market relative to the myopic benchmark, this may entail an increase in the expected price on one side if the externality that this side exerts is much weaker than the externality it experiences.

Pricing implications in two-sided markets have received a lot of attention in industrial

[1] The price-setting version of the model seems more widely applicable, but the quantity-setting version turns out to be more tractable.

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 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), smart phones (users, application developers), nightclubs and matching agencies (men, women), shopping malls, supermarkets, and department stores (where consumers are interested in a large variety of products, and producers in a large number of customers).

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). This allows us to analyze how the optimal price structure differs from the myopic benchmark and how it evolves over time. Our analysis suggests that markets characterized by indirect network effects of uncertain size provide incentives for the experimenting platform provider to initially lower at least one price. This provides a new rationale for price discounts in dynamic two-sided markets.[2]

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.

The remainder of the paper is structured as follows. Section 2 presents the model for the price-setting platform provider and characterizes the evolution of beliefs. Section 3 analyzes the directions of optimal experimentation, while Section 4 elaborates on the maximal experimentation policy. The optimal policy of a quantity-setting platform provider is analyzed in Section 5. Section 6 concludes. Technical proofs are relegated to the appendix.

[2]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.

2 The Model

We propose a two-sided market model following Armstrong (2006) to focus on participation decisions. For tractability reasons, we analyze a setting with 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 provider

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 decreasing when positive) gives rise to linear demand functions. The platform provider can set membership fees (M_A, M_B), but no usage fee.[3] 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

$$n_A = u_0 + \tilde{u} n_B - M_A, \tag{1}$$
$$n_B = \pi_0 + \tilde{\pi} n_A - M_B, \tag{2}$$

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.[4] 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 probability that the platform provider intially assigns to the realization $(\overline{u}, \overline{\pi})$ by p_0 and assume that this prior belief is non-degenerate, i.e., $0 < p_0 < 1$.[5]

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

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

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

[4] 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.

[5] The assumption that the externality parameters are perfectly positively correlated is clearly restrictive. Imperfect correlation leads to a much more complicated situation with two-dimensional beliefs. We will see that our results for the quantity-setting scenario carry over to perfect *negative* correlation.

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 non-negative. 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

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

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

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

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

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

where

$$R(M_A, M_B, p) = M_A\, \mathrm{E}^p \left[n_A(M_A, M_B, \tilde{u}, \tilde{\pi}) \right] + M_B\, \mathrm{E}^p \left[n_B(M_A, M_B, \tilde{u}, \tilde{\pi}) \right] \tag{4}$$

is the expected *current* revenue from charging the fees (M_A, M_B) given the posterior belief p.

2.2 The myopic benchmark

If the platform provider were myopic (corresponding to $r = \infty$), it would maximize expected current revenue at each instant. Under our parameter restrictions, this revenue is strictly concave in (M_A, M_B), so the myopically optimal fees,

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

are well-defined.

To compute these fees, we write the expected quantities appearing on the right-hand side of (4) as

$$\begin{aligned}
\mathrm{E}^p\left[n_A(M_A, M_B, \tilde{u}, \tilde{\pi})\right] &= \ell_0(p)[u_0 - M_A] + \ell_A(p)[\pi_0 - M_B], \\
\mathrm{E}^p\left[n_B(M_A, M_B, \tilde{u}, \tilde{\pi})\right] &= \ell_0(p)[\pi_0 - M_B] + \ell_B(p)[u_0 - M_A],
\end{aligned}$$

where

$$\ell_0(p) = \frac{1-p}{1-\underline{\pi}\underline{u}} + \frac{p}{1-\overline{\pi}\overline{u}}$$

and

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

measure the expected direct and indirect effects, respectively, of lowering M_A or M_B.

With the dependence on the belief p suppressed, the right-hand side of (4) now becomes

$$[\ell_0 u_0 + \ell_A \pi_0] M_A + [\ell_0 \pi_0 + \ell_B u_0] M_B - \ell_0 M_A^2 - [\ell_A + \ell_B] M_A M_B - \ell_0 M_B^2.$$

As $0 < \ell_i < \ell_0$ for $i = A, B$ and hence $0 < \ell_A + \ell_B < 2\ell_0$, this quadratic function is indeed strictly concave, and we obtain

$$M_A^\mu = u_0 - \frac{\left[2\ell_0^2 - (\ell_A + \ell_B)\ell_A\right] u_0 - (\ell_A - \ell_B)\ell_0 \pi_0}{4\ell_0^2 - (\ell_A + \ell_B)^2}\,, \tag{5}$$

$$M_B^\mu = \pi_0 - \frac{\left[2\ell_0^2 - (\ell_A + \ell_B)\ell_B\right] \pi_0 - (\ell_B - \ell_A)\ell_0 u_0}{4\ell_0^2 - (\ell_A + \ell_B)^2}\,. \tag{6}$$

As is well known from the literature on two-sided markets, the myopically optimal fee on one side of the market depends on market characteristics on *both* sides. Independent of the values of the externality parameters $\underline{u}, \overline{u}, \underline{\pi}, \overline{\pi}$, the fee on either side is always increasing in the intrinsic platform value on that same side. Whether or not the fee on one side is increasing in the intrinsic platform value on the *other* side depends on the relative strength of the network effects on both sides. To be precise, the fee M_A^μ is increasing in π_0 if and only if $\ell_A - \ell_B > 0$. Broadly speaking, when the externality side A is experiencing is higher than the one it is exerting, it benefits from the higher attractiveness of the platform for participants on side B as the intrinsic platform value π_0 rises, and can thus be charged a higher price; in this sense, side A "subsidizes" side B.

Further, M_A^μ can only exceed the intrinsic platform value u_0 if ℓ_A exceeds ℓ_B by a sufficient amount, and vice versa for M_B^μ and π_0. Thus, at most one fee at a time can exceed the intrinsic platform value and both fees will be lower than the respective intrinsic platform values if the expected externalities are equal ($\ell_A = \ell_B$) or close together.

For future reference, we denote the myopically optimal revenue by

$$R^\mu(p) = \max_{M_A, M_B} R(M_A, M_B, p) = R(M_A^\mu(p), M_B^\mu(p), p), \tag{7}$$

and, suppressing the dependence on p and other variables, rewrite the expected current revenue as

$$R = R^\mu - \ell_0 \left[M_A - M_A^\mu\right]^2 - \left[\ell_A + \ell_B\right]\left[M_A - M_A^\mu\right]\left[M_B - M_B^\mu\right] - \ell_0 \left[M_B - M_B^\mu\right]^2. \tag{8}$$

2.3 The evolution of beliefs

The platform provider revises its beliefs over time. Writing $\overline{n}_A(M_A, M_B) = n_A(M_A, M_B, \overline{u}, \overline{\pi})$ and using analogous definitions for $\underline{n}_A$, $\overline{n}_B$ and $\underline{n}_B$, 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*

$$dp_t \sim N\big(0, p_t^2(1-p_t)^2 S(M_A^t, M_B^t)\, dt\big). \tag{9}$$

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.[6]

We can gain more precise insights into the structure of the function S by noting that

$$\begin{aligned}
\overline{n}_A(M_A, M_B) - \underline{n}_A(M_A, M_B) &= d_0\left[u_0 - M_A\right] + d_A\left[\pi_0 - M_B\right], \\
\overline{n}_B(M_A, M_B) - \underline{n}_B(M_A, M_B) &= d_0\left[\pi_0 - M_B\right] + d_B\left[u_0 - M_A\right],
\end{aligned}$$

where $d_i = \ell_i(1) - \ell_i(0) > 0$ for $i = 0, A, B$, and computing

$$S(M_A, M_B) = s_A\left[M_A - u_0\right]^2 + 2 s_{AB}\left[M_A - u_0\right]\left[M_B - \pi_0\right] + s_B\left[M_B - \pi_0\right]^2$$

with the constants

$$s_A = \frac{d_0^2}{\sigma_A^2} + \frac{d_B^2}{\sigma_B^2}, \quad s_B = \frac{d_A^2}{\sigma_A^2} + \frac{d_0^2}{\sigma_B^2}, \quad s_{AB} = \frac{d_0 d_A}{\sigma_A^2} + \frac{d_0 d_B}{\sigma_B^2}.$$

[6] If the platform provider were uncertain about the intrinsic platform values (u_0, π_0) instead of the externalities (u, π), the quantity of information would be independent of the fees charged. The platform provider would then trivially always set the myopically optimal fees.

Since $s_A s_B - s_{AB}^2 = \sigma_A^{-2}\sigma_B^{-2}(d_0^2 - d_A d_B)^2$ and, as a simple computation reveals, $d_0^2 < d_A d_B$, we see that S is a strictly convex function which assumes its global minimum of zero at $(M_A, M_B) = (u_0, \pi_0)$.

The beliefs $p = 0$ and $p = 1$ are absorbing—if the platform provider is subjectively sure about the true state of the world, no further learning is possible. For a non-degenerate belief p to be invariant under the optimal learning dynamics, two conditions are necessary: the platform must charge the myopically optimal fees at this belief (since this belief will persist forever), and the information content of the resulting demand observations must be zero (so that this belief will indeed persist). Taken together, this requires $S(M_A^\mu(p), M_B^\mu(p)) = 0$ or, equivalently, $(M_A^\mu(p), M_B^\mu(p)) = (u_0, \pi_0)$ which is impossible since the latter fees generate an expected current revenue of zero and marginally lowering one of the fees would improve upon that. There are thus no potentially confounding actions in the sense of Easley and Kiefer (1988). By well-known results, this implies

Lemma 2 *Any optimal pricing policy induces complete learning in the long run: the platform provider's posterior belief p_t converges to the truth almost surely as $t \to \infty$.*

To determine how the information content of observed quantities changes with the fees 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. We have $\frac{\partial S}{\partial M_A} = 0$ along the line $M_B = \pi_0 - \frac{s_A}{s_{AB}}(M_A - u_0)$, and $\frac{\partial S}{\partial M_B} = 0$ along the line $M_B = \pi_0 - \frac{s_{AB}}{s_B}(M_A - u_0)$; as $s_A s_B - s_{AB}^2 > 0$, the former line is steeper than the latter. For either line, the respective partial derivative is positive above the line and negative below. Along the myopically optimal pricing policy, both $\overline{n}_A - \underline{n}_A$ and $\overline{n}_B - \underline{n}_B$ can be shown to be positive, which directly implies that both partial derivatives of S are negative. Thus, the myopically optimal fees lie below both lines in Figure 1.

In fact, this is true for *all* admissible price combinations.

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.*

PROOF: The proof consists in showing that above either line of vanishing marginal information content in Figure 1, at least one of the implied expected quantities becomes negative. See the appendix for details. □

3 The Optimal Pricing Strategy

We are now ready to characterize the pricing strategy. In view of the objective function (3) and the law of motion (9), standard arguments yield the following Bellman equation for the platform provider's value function, v:

$$v(p) = \max_{M_A, M_B} \left\{ R(M_A, M_B, p) + \frac{p^2(1-p)^2}{2r} S(M_A, M_B)\, v''(p) \right\} . \tag{10}$$

Arguing as in Keller and Rady (1999), one shows that v is strictly convex, twice continuously differentiable, and the unique solution to (10) subject to the condition that $v(p) = R^\mu(p)$ at

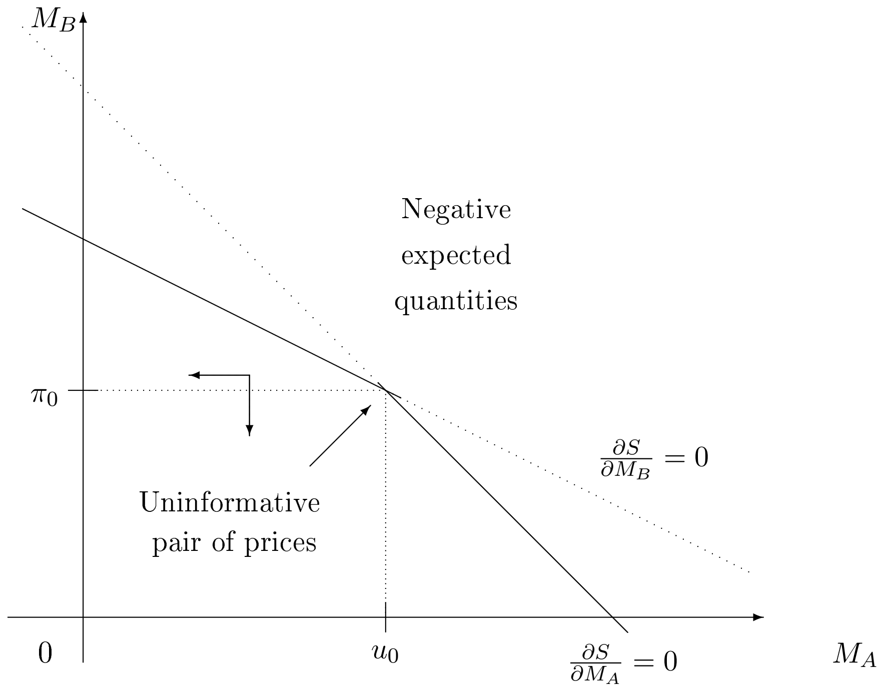

Figure 1: The directions of increasing information in the price plane.

$p = 0$ and 1, where the myopically optimal expected current revenue $R^\mu(p)$ has been defined in equation (7).

We can interpret the second term of the maximand in the Bellman equation as the value of information, given by the product of the shadow price of information, $p^2(1-p)^2 v''(p)/2r$, and the quantity of information, $S(M_A, M_B)$. 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 so as to increase the information content of its demand observations.

The maximand in (10) is the sum of two quadratic functions, one of them strictly concave (expected current revenue), the other strictly convex (value of information). As the value function is bounded, so must be the maximum on the right-hand side of (10); and as admissible fees are unbounded below, the shadow price of information must actually be small enough for the combined quadratic function to be strictly concave (the precise argument is in the appendix).

This ensures that optimal fees are fully characterized by the (linear) first-order conditions for the maximization problem in (10). Using the representation of expected current revenues in (8), writing

$$V(p) = \frac{p^2(1-p)^2}{2r}\, v''(p)$$

for the shadow price of information, and suppressing the dependence on p, we compute the optimal pair of fees as

$$M_A^* = M_A^\mu + \frac{2V}{h(V)} \left\{ 2(\ell_0 - s_B V) S_A^\mu - (\ell_A + \ell_B - 2 s_{AB} V) S_B^\mu \right\}, \tag{11}$$

$$M_B^* = M_B^\mu + \frac{2V}{h(V)} \left\{ 2(\ell_0 - s_A V) S_B^\mu - (\ell_A + \ell_B - 2 s_{AB} V) S_A^\mu \right\}, \tag{12}$$

where

$$h(V) = 4(\ell_0 - s_A V)(\ell_0 - s_B V) - (\ell_A + \ell_B - 2 s_{AB} V)^2$$

is the determinant of the Hessian matrix of the maximand in (10) and

$$S_A^\mu = \frac{\partial S}{\partial M_A}(M_A^\mu, M_B^\mu) = s_A(M_A^\mu - u_0) + s_{AB}(M_B^\mu - \pi_0) < 0,$$

$$S_B^\mu = \frac{\partial S}{\partial M_B}(M_A^\mu, M_B^\mu) = s_{AB}(M_A^\mu - u_0) + s_B(M_B^\mu - \pi_0) < 0$$

are the partial derivatives of the quantity of information S at the myopically optimal fees.[7] Strict concavity of the maximand in (10) means $\ell_0 - s_A V > 0$ and $h(V) > 0$, which in turn implies $\ell_0 - s_B V > 0$.

Our first result on the platform provider's optimal pricing strategy is

Proposition 1 *At any non-degenerate belief, the platform provider charges a fee lower than the myopic benchmark on at least one side of the market.*

Proof: Suppose that $M_A^* \geq M_A^\mu$. By (11), this implies $\ell_A + \ell_B - 2 s_{AB} V > 0$ and

$$S_B^\mu \leq \frac{2(\ell_0 - s_B V)}{\ell_A + \ell_B - 2 s_{AB} V} S_A^\mu .$$

As a consequence,

$$2(\ell_0 - s_A V) S_B^\mu - (\ell_A + \ell_B - 2 s_{AB} V) S_A^\mu \leq \frac{h(V)}{\ell_A + \ell_B - 2 s_{AB} V} S_A^\mu < 0,$$

and so $M_B^* < M_B^\mu$ by (12). In exactly the same way, $M_B^* \geq M_B^\mu$ implies $M_A^* < M_A^\mu$. □

The intuition for this result is clear. The purpose of deviating from the myopic optimum is to increase the information content of observed demands. As higher fees mean less information (see Lemma 3), at least one fee must be reduced relative to the myopic benchmark.

This has an obvious consequence for approximately symmetric setups.

Proposition 2 *For $(u_0, \underline{u}, \overline{u}, \sigma_A)$ sufficiently close to $(\pi_0, \underline{\pi}, \overline{\pi}, \sigma_B)$, the platform provider always sets both fees below their myopically optimal levels.*

[7] The argument why both of them are negative was given in Section 2.3.

PROOF: For $(u_0, \underline{u}, \overline{u}, \sigma_A) = (\pi_0, \underline{\pi}, \overline{\pi}, \sigma_B)$, we have $M_A^\mu = M_B^\mu$ by (5)–(6), and $M_A^* - M_A^\mu = M_B^* - M_B^\mu \geq 0$ by (11)–(12) and Proposition 1, with a strict inequality, and the expression in curly brackets bounded away from 0, on the open unit interval. The result thus follows by continuous dependence of the value function and its second derivative on $(u_0, \underline{u}, \overline{u}, \sigma_A)$.□

The analysis of asymmetric settings is more complicated. A lower fee on one side of the market makes reducing the fee on the other side more attractive from an informational perspective (the cross-partial derivative of the quantity of information with respect to prices, s_{AB}, is positive), but less attractive as far as expected current revenue is concerned (its cross-partial derivative, $-(\ell_A + \ell_B)$, is negative). The overall effect is ambiguous.

A different way to see this is to think of the platform provider as following a two stage-procedure. At the first stage, it determines the combination of fees that maximizes current expected revenue subject to the constraint that a certain quantity of information be achieved. This amounts to identifying points of tangency between iso-information and iso-revenue curves in the (M_A, M_B)-plane. At the second stage, the provider then chooses the optimal quantity of information. Depending on the geometry of the iso-information and iso-revenue curves, this may lead it to charge a fee *higher* than in the myopic benchmark on one side of the market, as we shall see below.

To identify the directions of optimal experimentation in some asymmetric settings, we insert the expressions for S_A^μ and S_B^μ into (11)–(12) and collect the terms in $M_A^\mu - u_0$ and $M_B^\mu - \pi_0$, respectively:

$$\begin{aligned} M_A^* &= M_A^\mu + \frac{2V}{h(V)} \Big\{ \left[2\ell_0 s_A - (\ell_A + \ell_B) s_{AB} - 2(s_A s_B - s_{AB}^2)V\right] (M_A^\mu - u_0) \\ &\qquad + \left[2\ell_0 s_{AB} - (\ell_A + \ell_B) s_B\right] (M_B^\mu - \pi_0) \Big\}, \qquad (13) \\ M_B^* &= M_B^\mu + \frac{2V}{h(V)} \Big\{ \left[2\ell_0 s_{AB} - (\ell_A + \ell_B) s_A\right] (M_A^\mu - u_0) \\ &\qquad + \left[2\ell_0 s_B - (\ell_A + \ell_B) s_{AB} - 2(s_A s_B - s_{AB}^2)V\right] (M_B^\mu - \pi_0) \Big\}. \qquad (14) \end{aligned}$$

Proposition 3 *Let $s_A < s_{AB} < s_B$. Whenever both myopically optimal fees are lower than the respective intrinsic values, the platform provider lowers the fee on side B relative to the myopically optimal level.*

PROOF: It is enough to show that in equation (14), the coefficients of $M_A^\mu - u_0$ and $M_B^\mu - \pi_0$ in the expression in curly brackets are positive. As $s_{AB} > s_A$ and $\ell_A + \ell_B < \ell_0$, this is obvious for the coefficient of $M_A^\mu - u_0$. Regarding the coefficient of $M_B^\mu - \pi_0$, we distinguish two cases. If $(\ell_A + \ell_B)/(2s_{AB}) < \ell_0/s_B$, the positivity of $s_A s_B - s_{AB}^2$ and the fact that $V < \ell_0/s_B$ imply that the coefficient of $M_B^\mu - \pi_0$ exceeds $2\ell_0 s_B - (\ell_A + \ell_B)s_{AB} - 2\ell_0(s_A s_B - s_{AB}^2)/s_B$, which is positive. If $(\ell_A + \ell_B)/(2s_{AB}) \geq \ell_0/s_B$, we have $V < (\ell_A + \ell_B)/(2s_{AB})$ and the coefficient of $M_B^\mu - \pi_0$ is no smaller than $2\ell_0 s_B - (\ell_A + \ell_B)s_{AB} - (\ell_A + \ell_B)(s_A s_B - s_{AB}^2)/s_{AB}$, which is again positive. □

The situation assumed in this proposition is one where the marginal informational benefit of lowering the fee is so much larger on side B than on side A that the platform provider will definitely lower the fee on side B. This situation arises naturally when the strength of the externality that side A exerts on side B is relatively well known, i.e., when $\overline{\pi}$ and $\underline{\pi}$ are relatively close to each other. More precisely, as $\overline{\pi}$ and $\underline{\pi}$ tend to a common value π, the ratios d_B/d_0, d_0/d_A, s_{AB}/s_B and s_A/s_{AB} all converge to π, which implies $s_A < s_{AB} < s_B$ for sufficiently small differences $\overline{\pi} - \underline{\pi}$.

The limiting case in which $\overline{\pi} = \underline{\pi} = \pi$ lends itself to a simple graphical illustration that will prove valuable when it comes to formulating a sufficient condition for the fee M_A to rise relative to the myopic benchmark. In fact, the identity $s_{AB}/s_B = s_A/s_{AB} = \pi$ implies that in the (M_A, M_B)-plane, the level curves of the function S are parallel straight lines with slope $-\pi$. These iso-information lines and the myopically optimal pricing policy are illustrated in Figure 2.

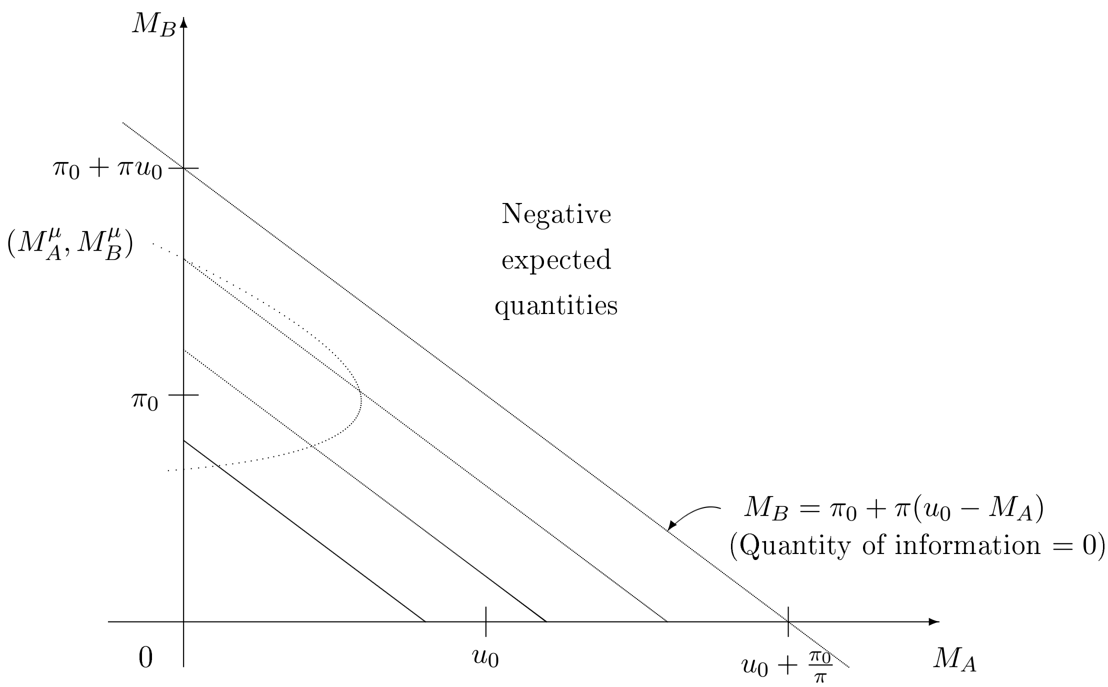

Figure 2: Iso-information lines and myopically optimal policy in the price plane when there is uncertainty about the externality parameter $\tilde{u}$ only ($\overline{\pi} = \underline{\pi} = \pi$).

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. On any iso-information line, it will choose the fees that correspond to a point of tangency with an iso-revenue curve. As Figure 3 illustrates, the slope of the locus of tangency points between iso-information lines and iso-revenue curves (ellipses, to be precise) depends on parameters. In the left panel, this locus slopes upward – the optimal trade-off between information and current revenue induces a decrease in both fees for increased information. However, if the iso-information lines are rather flat (i.e., if π is small), it is optimal to decrease M_B but increase M_A as indicated by the locus of optimal fees in the right panel.

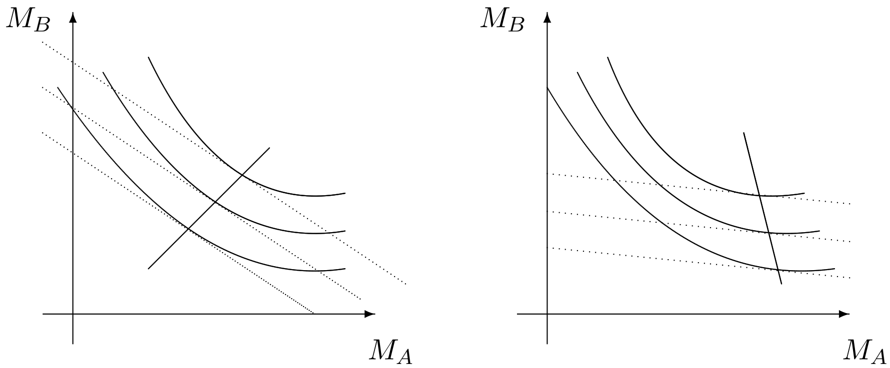

Figure 3: Two examples of iso-information lines (dotted) and iso-revenue curves (solid) for $\overline{\pi} = \underline{\pi} = \pi$. The solid line in each case indicates the locus of optimal fees.

This suggests that for $\overline{\pi}$ different from $\underline{\pi}$ but sufficiently small, we should also be able to see an optimal fee M_A^* that exceeds M_A^μ. Our next result bears this out.

Proposition 4 *For $\overline{\pi}$ sufficiently close to 0, the platform provider increases the fee on side A relative to the myopically optimal level.*

PROOF: For $\overline{\pi} = \underline{\pi} = 0$, we have $\ell_0 = 1$ and $\ell_B = 0$, implying $d_0 = d_B = 0$ and $s_A = s_{AB} = 0$. By (6), moreover, $M_B^\mu - \pi_0$ is negative and bounded away from 0 on the unit interval. Now, the expression in curly brackets in (13) reduces to $-\ell_A s_B(M_B^\mu - \pi_0)$, which is positive and again bounded away from 0. The result thus follows by continuous dependence of the value function and its second derivative on $(\underline{\pi}, \overline{\pi})$. $\square$

We can offer the following intuition for this result. When the externality that side A is exerting on side B is known to be very small, the platform provider learns most by lowering the fee on side B. Side A then benefits from higher participation on side B. Since participation on side A hardly affects participation on side B, the provider can safely extract part of the additional surplus given to side A by charging this side a higher fee.

4 Maximal Experimentation

In the previous section, we were able to analyze the *directions* of optimal experimentation without having to solve for the value function. To establish the precise *extent* of optimal experimentation, one could plug the fees (11)-(12) into the maximand in (10) and numerically solve the resulting second-order ordinary differential equation for the value function.

An alternative route to this differential equation is to write the Bellman equation in the form $0 = \max_{M_A, M_B}\{R - v + \frac{p^2(1-p)^2}{2r}\, S\, v''\}$ and to observe that the maximum remains zero, and the set of maximizers is unchanged, when we divide the maximand by the quantity of information, S.[8] Re-arranging then yields

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

This in turn permits an alternative characterization of the optimal combination of fees as a function of the belief p and the associated value $v(p)$:

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

Arguing as in Keller and Rady (1999), one shows that the value $v(p)$ is decreasing in r at all p in the open unit interval, and that it converges to the *ex ante full-information pay-off*

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

as $r \downarrow 0$. This means that the optimal fees converge to

$$(\overline{M}_A(p), \overline{M}_B(p)) = \arg\min_{M_A,M_B} \frac{\overline{R}(p)-R(M_A,M_B,p)}{S(M_A,M_B)} \,, \tag{15}$$

which is the optimal policy 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 Bolton and Harris (2000).

Intuitively speaking, the lower the platform provider's discount rate, the greater is its incentive to learn, and the farther it might want to deviate from the myopic optimum. Experimentation is maximal when $r = 0$. Once we know the optimal strategy of the infinitely patient provider, therefore, we have fully characterized the range of experimentation in which an impatient provider will set his fees.

Studying the maximal experimentation strategy $(\overline{M}_A, \overline{M}_B)$ has the further advantage that it does not require computation of the value function for the maximization of transient payoffs.[9] While the system of first-order conditions for (15) in general does not permit explicit solutions, it is considerably easier to solve numerically than the differential equation for the value function under discounting. In the next subsection, we will take advantage of this to illustrate the maximal experimentation policy and the associated learning dynamics in a numerical example. Thereafter, we will briefly return to the limiting case $\overline{\pi} = 0$ which does permit a closed-form solution.

4.1 An example

We assume the following parameters: $u_0 = 0.4$, $\pi_0 = 0.1$, $\underline{u} = 0.1$, $\overline{u} = 0.9$, $\underline{\pi} = 0.1$, $\overline{\pi} = 0.2$, $\sigma_A = \sigma_B = 1$, $p_0 = 0.5$, and the "true" values are $(\overline{u}, \overline{\pi})$. These parameters translate into expected direct and indirect price effects of $\ell_0(p_0) = 1.11$, $\ell_A(p_0) = 0.60$, and $\ell_B(p_0) = 0.17$,

[8] As the admissible pair of fees (u_0, π_0) is clearly suboptimal (yielding zero revenue and zero information), the function S is indeed positive on the relevant domain.

[9] This is crucial for the characterization of Markov perfect equilibria in Bolton and Harris (2000), for example.

respectively. In particular, the externality that side B is expected to exert on side A, $\ell_A(p_0)$, is assumed more than three times as large as the expected opposite externality, $\ell_B(p_0)$. Also note that $s_A = 0.06$, $s_{AB} = 0.24$, and $s_B = 1.04$, hence $s_A < s_{AB} < s_B$.

The optimal fees set by a myopic and an infinitely patient platform provider are depicted in Figure 4. It is straightforward to check that both myopically optimal fees are lower than the respective intrinsic values at all beliefs. In line with Proposition 3, the maximal experimentation policy reduces the fee on side B relative to the myopic benchmark at any non-degenerate belief.[10] The fee set on side A under the maximal experimentation policy is lower than the myopic benchmark at all beliefs below a threshold that approximately equals 0.275, and higher than the myopic benchmark at all beliefs above that threshold. Thus, in accordance with Propositions 2 and 4, the fee on side A is reduced when the externalities are of similar expected size (for beliefs close to 0, the large difference between $\overline{u}$ and $\overline{\pi}$ does not matter much), but is increased when $\ell_A(p)$, the expected strength of the externality that side A experiences, is sufficiently larger than $\ell_B(p)$, the expected strength of the externality that side A exerts.

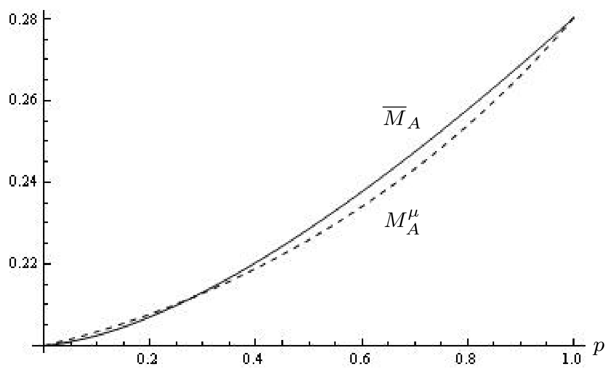

Figure 4: Optimal myopic fees (dashed line) and maximal experimentation fees (solid line) on market side A (left) and B (right) as a function of the belief.

Figure 5 illustrates that the infinitely patient provider learns faster – its beliefs converge more quickly to the true state.[11]

Figure 6 shows a sample path for the optimal fee on side A. At any given belief, the experimenting monopolist increases the fee relative to the myopic monopolist. In later periods, the increase is reinforced by the better information driving the fee towards the high optimum more rapidly. The evolution of optimal fees on side B is shown in Figure 7. Maximal experimentation fees are consistently below their myopic counterparts.

[10]Note that for the given set of parameters, the optimal fee under full information on side B is actually negative, i.e., participants on side B receive a payment from the platform provider. Monetary payments to participants on one side may not always be feasible. However, as pointed out in the two-sided market literature, in-kind payments can often substitute for monetary payments.

[11]Simulations were carried out using Wolfram Mathematica 8. Normal shocks were generated by random draws from the normal distribution using the commands "RandomReal" and "NormalDistribution" with mean equal to 0 and variances equalling σ_A and σ_B respectively.

The expected per-period revenues depicted in Figure 8 show the advantages of each policy. While the myopic policy creates higher revenues in the very early periods, revenues in later periods are higher for the patient platform provider as its belief approaches the true state of the world more rapidly.

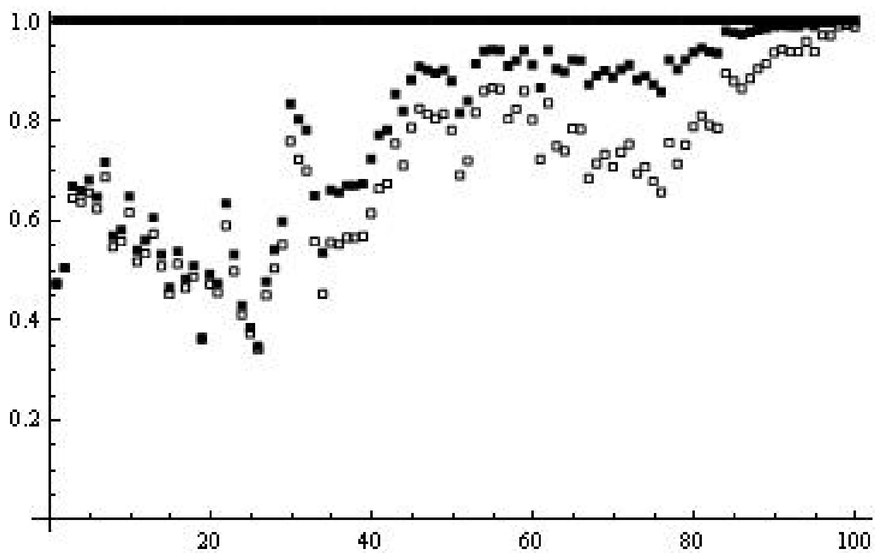

Figure 5: Evolution of beliefs for the myopic policy (white squares) and the infinitely patient policy (black squares), and true state (thick line).

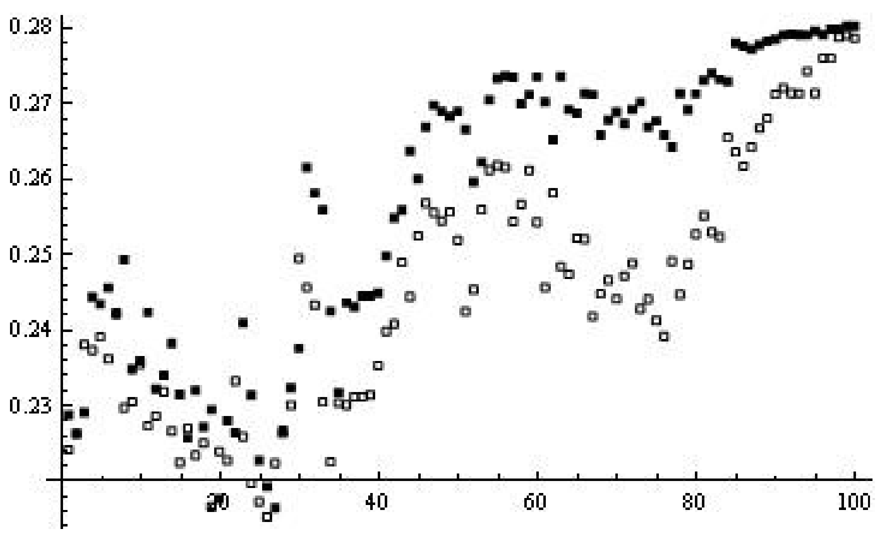

Figure 6: Evolution of fees on side A for the myopic policy (white squares) and the infinitely patient policy (black squares).

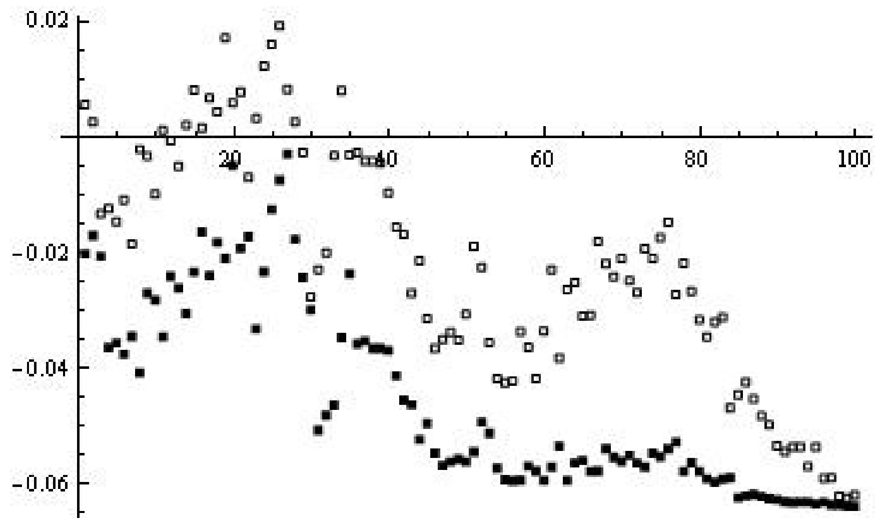

Figure 7: Evolution of fees on side B for the myopic policy (white squares) and the infinitely patient policy (black squares).

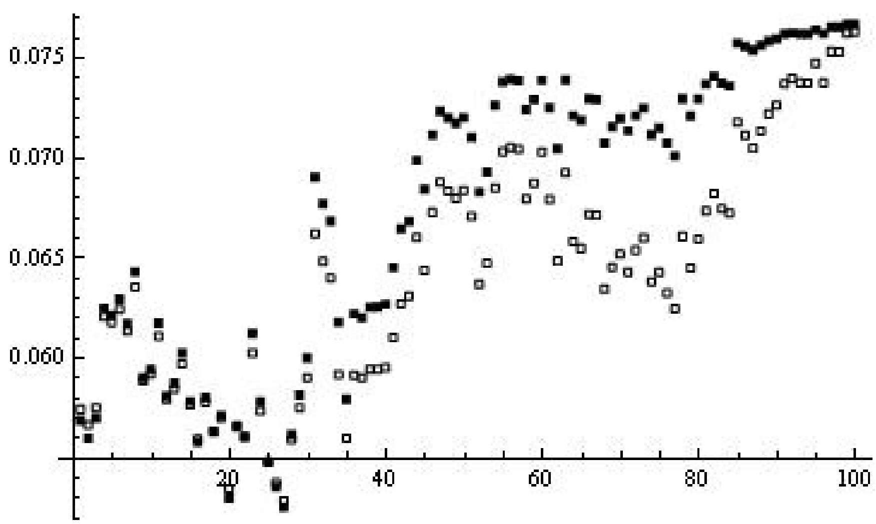

Figure 8: Evolution of expected per-period revenues for the myopic policy (white squares) and the infinitely patient policy (black squares).

4.2 A closed-form solution

We have seen in Proposition 4 above that, for vanishing externality parameter $\overline{\pi}$, the platform provider raises the fee on side A relative to the myopically optimal policy. The limiting case $\overline{\pi} = 0$ turns out to permit a closed-form solution for the maximal experimentation policy.[12]

The myopically optimal fees in this case are

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

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

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

The quantity of information simplifies to $S(M_A, M_B) = \sigma_A^{-2}(\overline{u} - \underline{u})^2(\pi_0 - M_B)^2$, reflecting the fact that only the demand observed on side A is informative.

The minimum of $[\overline{R}(p) - R(M_A, M_B, p)]/(\pi_0 - M_B)^2$ is attained at

$$\begin{aligned} \overline{M}_A(p) &= \frac{\pi_0 u_0 + 2\overline{R}(p)u(p)}{2\pi_0 + u_0 u(p)}, \\ \overline{M}_B(p) &= \pi_0 + \frac{u_0^2 - 4\overline{R}(p)}{2\pi_0 + u_0 u(p)}. \end{aligned}$$

Comparing these fees to the myopically optimal ones, we first see that

$$\overline{M}_B(p) - M_B^\mu(p) = \frac{4[R^\mu(p) - \overline{R}(p)]}{2\pi_0 + u_0 u(p)}.$$

As $\overline{R}(p) = p\, R^\mu(1) + (1-p)R^\mu(0)$ and R^μ is strictly convex, the right-hand side is negative for $0 < p < 1$. Thus, in line with Proposition 3, the infinitely patient platform provider will indeed decrease the fee that generates information.

On the other side of the market, we find

$$\overline{M}_A(p) - M_A^\mu(p) = \frac{2u(p)[\overline{R}(p) - R(p)]}{2\pi_0 + u_0 u(p)} = -\frac{u(p)}{2}[\overline{M}_B(p) - M_B^\mu(p)],$$

so for non-degenerate beliefs, there is a price increase relative to the myopic benchmark, as predicted by Proposition 4.

The expected quantity on side B clearly increases relative to the myopic optimum since the fee M_B goes down. Using the above expression for $\overline{M}_A(p) - M_A^\mu(p)$, one can additionally establish that the expected quantity on side A changes by $-\frac{u(p)}{2}[\overline{M}_B(p) - M_B^\mu(p)]$, 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 expected quantities while price adjustments on the two sides go in opposite directions.

[12] As an example, consider readers whose utility of a magazine is independent of the amount of advertising.

5 The Quantity-Setting Platform Provider

We now assume that the platform provider sets quantities. The quantity-setting assumption seems appropriate in real-world markets where capacity constraints matter. For instance, a shopping mall owner has to decide how much parking space and shop space to provide. If prices are market-clearing, this choice of capacities corresponds to quantity setting.

In standard monopoly, it does not matter (under certainty) whether a price or a quantity is chosen. In two-sided markets, setting quantities means that the platform directly controls the size of the externality, whereas a price setter does so only indirectly. This explains why the quantity-setting case is more tractable: there are no feedback effects to be taken into account when the quantity is changed on one side of the market. As we shall see below, this makes the information content of market observations additively separable across the two sides and implies unambiguous directions of experimentation.

Let the platform provider 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 $\tilde{\pi} \in \{\underline{\pi}, \overline{\pi}\}$ with $0 < \underline{u} < \overline{u}$, $0 < \underline{\pi} < \overline{\pi}$ and $\overline{u} + \overline{\pi} < 2$. As we permit externality parameters exceeding 1, this is somewhat more general than what we assumed in the price-setting case.

We impose the natural restriction that the platform provider can only decide to sell non-negative quantities, while prices are not restricted. Negative prices are interpreted as subsidies to one side or (temporarily) both sides of the market, as discussed earlier. 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.[13]

As before, we write p for the subjective probability assigned to the realization $(\overline{u}, \overline{\pi})$. We maintain the assumption that costs are zero.

5.1 Revenues 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

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

[13] Notably, all insights of this section carry over to the case of perfect *negative* correlation. Results only depend on expected externalities, exchanging the roles of $\overline{u}$ and $\underline{u}$ is unproblematic, therefore. As to Propositions 6–7 below, it suffices that signal-to-noise ratios coincide in absolute value.

for the expected externalities, and

$$R(n_A, n_B, p) = n_A \left[u_0 + u(p)\, n_B - n_A \right] + n_B \left[\pi_0 + \pi(p)\, n_A - n_B \right]$$

for the expected per-period revenue, the platform provider's total expected payoff is

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

The expected revenue R depends on the expected externalities only through the term $[u(p) + \pi(p)] n_A n_B$, so only the *sum* of the externalities matters here. As $|u(p) + \pi(p)| < 2$, moreover, R is strictly concave in (n_A, n_B). The myopically optimal quantities are

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

They exhibit a symmetric structure with interchanged intrinsic platform values. If these platform values coincide, myopically optimal quantities are the same on both sides.

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 expected 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 the strictly convex function

$$\Sigma(n_A, n_B) = \rho_A n_A^2 + \rho_B n_B^2,$$

where the constants

$$\rho_A = \left(\frac{\overline{u} - \underline{u}}{\theta_A} \right)^2, \quad \rho_B = \left(\frac{\overline{\pi} - \underline{\pi}}{\theta_B} \right)^2$$

are the squares of the marginal signal-to-noise ratios.

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

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

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

Complete learning in the long-run follows from the same arguments as in the price-setting scenario (see Lemma 2 above). As Σ is increasing in both n_A and n_B, moreover, we obviously have

Lemma 5 *For the 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.*

Finally, we note that the marginal informational impact of adjusting the quantity on one side of the market does *not* depend on the quantity set on the other.

5.2 Optimal quantities

Under discounting at rate $r > 0$, the Bellman equation is

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

The maximand is again the sum of a strictly concave quadratic function and a strictly convex one. A simpler version of the argument given in the price-setting case shows that the shadow price of information, $V(p) = p^2(1-p)^2 v''(p)/2r$, is again sufficiently small to make the combined quadratic function strictly concave at all beliefs (we omit the details).

Solving the first-order conditions for optimal quantities and suppressing the dependence on the belief p, we obtain

$$\begin{aligned} n_A^* &= n_A^\mu + \frac{2V}{\chi(V)} \left\{ 2(1-\rho_B V)\rho_A n_A^\mu + 2(u+\pi)\rho_B n_B^\mu \right\}, \\ n_B^* &= n_B^\mu + \frac{2V}{\chi(V)} \left\{ 2(1-\rho_A V)\rho_B n_B^\mu + 2(u+\pi)\rho_A n_A^\mu \right\}, \end{aligned}$$

where

$$\chi(V) = 4(1-\rho_A V)(1-\rho_B V) - (u+\pi)^2$$

is the determinant of the Hessian matrix of $R + V\Sigma$. Strict concavity of this function means $1 - \rho_A V > 0$ and $\chi(V) > 0$, which in turn implies $1 - \rho_B V > 0$. As an immediate consequence, we get

Proposition 5 *At any non-degenerate belief, the quantity-setting platform provider chooses quantities above the myopic benchmark on both sides of the market.*

The intuition behind this result is simple. As the information content of observed prices is increasing in quantities, the optimal deviation from the myopic benchmark must entail a higher quantity on at least one side of the market. This raises the marginal revenue on the other side of the market without affecting the marginal informational benefit of adjusting the quantity there. It is optimal, therefore, to set a quantity above the myopic level on that side as well.

In the price-setting scenario, by contrast, lowering the fee on one side of the market has an ambiguous effect on the incentives to lower the fee on the other side because the cross-partial derivative of the quantity of information with respect to these fees is not zero and has the opposite sign to the respective derivative of expected current revenue.

5.3 Maximal experimentation

The maximal experimentation strategy is given by

$$(\overline{n}_A(p), \overline{n}_B(p)) = \arg\min_{n_A, n_B} \frac{\overline{R}(p) - R(n_A, n_B, p)}{\Sigma(n_A, n_B)} \ .$$

where $\overline{R}(p) = p\, R^{\mu}(1) + (1-p)R^{\mu}(0)$ is once more the expected full-information payoff. In general, the associated first-order conditions involve third-order polynomials in n_A and n_B. Owing to the simpler structure of the quantity-setting scenario, however, it is easier to obtain closed-form solutions than in the price-setting case, for example by assuming symmetric signal-to-noise ratios.[14]

Proposition 6 *Suppose that $\rho_A = \rho_B$ and $u_0 \neq \pi_0$. Then the quantities set by an infinitely patient platform provider are*

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

PROOF: See the appendix. □

The reason why these quantities do not depend on the common marginal signal-to-noise ratio is simple. For $\rho_A = \rho_B = \rho$, the information content of observed prices simplifies to $\Sigma(n_A, n_B) = \rho\,[n_A^2 + n_B^2]$, so the maximal experimentation strategy minimizes $(\overline{R}(p) - R(p))/(n_A^2 + n_B^2)$. Note that for $\pi_0 > u_0$, both numerator and denominator of $\overline{n}_A(p)$ and $\overline{n}_B(p)$ are negative, so the quantities remain positive. The knife-edge case $u_0 = \pi_0$ will be covered later.

The expected prices $\overline{M}_A(p)$ and $\overline{M}_B(p)$ given the quantities $\overline{n}_A(p)$ and $\overline{n}_B(p)$ are straightforward to calculate. Comparing them with the myopic optimum confirms what we have already seen in the price-setting model: even if the externality parameters $\overline{u}$ and $\overline{\pi}$ are both smaller than 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 myopic benchmark, as exemplarily shown for $\overline{M}_A$ in Figure 9.

[14] With quantities as the choice variables, it is less interesting to consider the limiting case of no uncertainty about the externality on one side of the market. If $\underline{\pi} = \overline{\pi}$, for instance, any deviation from the expected price on side B must be attributed to noise and is, thus, uninformative. The platform can then only experiment on side B, and only by adjusting the quantity n_A. This situation is isomorphic to the one analyzed in Keller and Rady (1999).

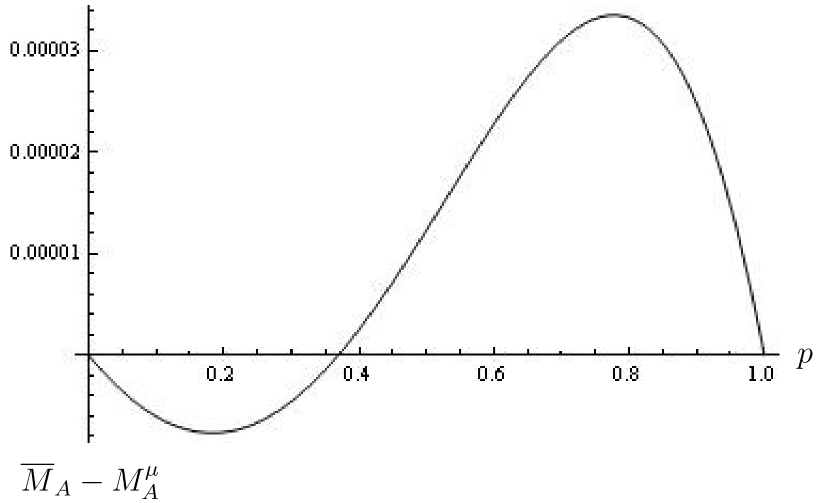

Figure 9: Difference between the expected prices induced by a myopic and an infinitely patient platform provider as a function of beliefs for $u_0 = 0.1$, $\pi_0 = 0.7$, $\underline{u} = 0.8$, $\overline{u} = 0.9$, $\underline{\pi} = 0.1$, $\overline{\pi} = 0.2$, $\theta_A = \theta_B = 1$.

The intuition behind this finding is the one we already gave for Proposition 4 in the price-setting case. If side A is expected to exert a relatively weak externality, i.e., if $\pi(p)$ is small relative to $u(p)$, the platform provider optimally learns by strongly increasing the number of participants on side B, and recoups part of the resulting surplus by inducing a higher than myopically optimal price on side A.

Maintaining symmetric signal-to-noise ratios, we further assume now that the intrinsic value of the platform is the same for all users, i.e., $u_0 = \pi_0$. This admittedly rather strong assumption seems appropriate in a number of examples, such as night clubs and matching agencies.[15] It simplifies the expressions for the optimal quantities considerably.

Proposition 7 *Suppose that $\rho_A = \rho_B$ and $u_0 = \pi_0 = c_0$. Then the optimal policy of an infinitely patient quantity-setting platform provider is symmetric across market sides and linear in the current belief:*

$$\overline{n}_A(p) = \overline{n}_B(p) = \overline{n}(p) = \frac{\overline{R}(p)}{c_0} = c_0 \left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1-p}{2 - (\underline{u} + \underline{\pi})} \right].$$

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.

[15] It is clearly less appropriate in other examples, such as merchants and customers in the credit card market.

The linearity of the maximal experimentation policy makes it easy to visualize the range of quantity experimentation; see Figure 10. It is the area bounded below by the myopic policy and above by the line joining the quantities that are optimal under full information.

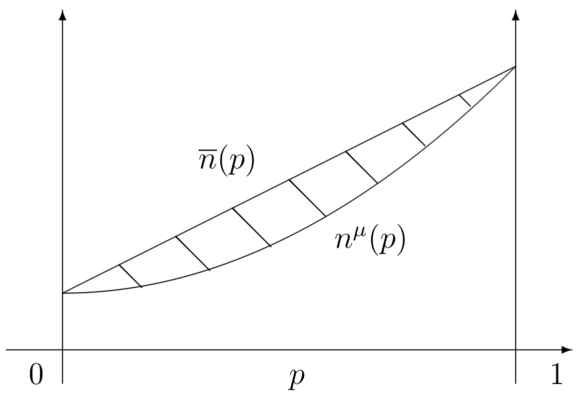

Figure 10: Range of quantity experimentation for symmetric signal-to-noise ratios and symmetric intrinsic values.

Expected prices need not be symmetric. They are

$$\begin{aligned} \overline{M}_A(p) &= c_0 + [u(p) - 1]\,\overline{n}(p), \\ \overline{M}_B(p) &= c_0 + [\pi(p) - 1]\,\overline{n}(p). \end{aligned}$$

As $u(p)+\pi(p) < 2$, either both expected prices are lower than the intrinsic platform value, or one is lower and the other one higher. The ordering of expected prices depends on the size of the externalities and on the current belief, and may 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 price in expectation than side B, while for low values of p the reverse is true.

As to the comparison with the myopic benchmark, we have

Corollary 1 *Under the assumptions of Proposition 7, the expected price induced by an infinitely patient quantity-setting platform provider exceeds its myopically optimal 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 optimal and the myopic expected prices coincide at the beliefs 0 and 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 myopically optimal level. It also implies that for the "standard" case of both externalities smaller than 1, both expected prices will decrease relative to the myopic benchmark.

6 Conclusion

We have studied optimal behavior of a monopolistic platform provider in a two-sided market with uncertainty about the strength of interaction between the two sides. The platform provider either chooses prices or quantities (i.e., participation levels). 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 follows 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 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 ($r = 0$).

The effect of experimentation on (expected) prices is ambiguous. Depending on the parameter constellation, either both prices will be lower than in the myopic benchmark or one price will be above and one price below the myopically optimal prices. The price on one side of the market will go up if the externality this side is exerting is weak while the externality it is experiencing is strong. The higher price recoups part of the surplus created by the higher participation on the other side of the market.

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. 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.

Another interesting extension is to consider a market for two (or more) goods that are complements. Specifically, suppose that demands are linked through positive network effects. Here we have in mind a situation in which a monopolist sells a product (or technologically related products) to two distinct and distinguishable consumer groups (i.e., the monopolist can practice third-degree price discrimination). If consumers in each group care directly or indirectly about the sum of the total number of buyers in both groups (e.g., because a larger production volume increases average product quality through learning-by-doing), we can rewrite this as a demand system with indirect network effects. Thus our analysis can possibly be extended to capture experimentation in markets with complementary goods.

Appendix

Proof of Lemma 1

Given a pair of prices (M_A, M_B), the observed quantity increments are

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

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

Given the subjective probability p currently assigned to the state $(\overline{u}, \overline{\pi})$, the vector of expected demands is

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

with $\overline{n}_A = n_A(M_A, M_B, \overline{u}, \overline{\pi})$ etc.

According to Liptser and Shiryayev (1977), the infinitesimal change in beliefs is given by

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

where

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

is the increment of a standard two-dimensional Brownian motion relative to the platform provider's information filtration.

Simplifying the expression for dp, we obtain

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

As $d\overline{Z}_A$ and $d\overline{Z}_B$ are normally distributed with mean zero and variance dt, and the infinitesimal covariance $< d\overline{Z}_A, d\overline{Z}_B >$ is zero, the change in beliefs dp is normally distributed with mean zero and variance $p^2(1-p)^2(\overline{n}_A - \underline{n}_A)^2\sigma_A^{-2}\, dt + p^2(1-p)^2(\overline{n}_B - \underline{n}_B)^2\sigma_B^{-2}\, dt$. $\square$

Proof of Lemma 3

We wish to show that in the region where the information content of quantities is increasing in a fee, the expected quantity on at least one side of the market must be negative.

For a partial derivative of S to be positive, at least one of the differences $\overline{n}_A - \underline{n}_A$ or $\overline{n}_B - \underline{n}_B$ has to be negative. This in turn is equivalent to at least one of the following inequalities holding:

$$M_B > \pi_0 + \frac{d_0}{d_A}(u_0 - M_A), \tag{16}$$

$$M_B > \pi_0 + \frac{d_B}{d_0}(u_0 - M_A). \tag{17}$$

For the two expected demands to be non-negative, it is necessary that both $\overline{n}_A$ and $\overline{n}_B$ be non-negative. This requires the following inequalities to hold:

$$M_B \leq \pi_0 + \frac{1}{\overline{u}}(u_0 - M_A), \tag{18}$$

$$M_B \leq \pi_0 + \overline{\pi}(u_0 - M_A). \tag{19}$$

Comparing the coefficients of $u_0 - M_A$ on the right-hand sides of these four inequalities, we see that for $M_A > u_0$, (18) contradicts both (16) and (17), while (19) does so for $M_A < u_0$. For $M_A = u_0$ the contradiction is obvious. □

Strict concavity of the maximand in the Bellman equation

Fixing a belief p and a shadow price of information $V = p^2(1-p)^2 v''(p)/2r$, we write the maximand in the Bellman equation (10) as $R(M_A, M_B, p) + V\,S(M_A, M_B)$ and compute its Hessian, suppressing the variable p from now on:

$$\mathcal{H}(V) = \begin{pmatrix} -2\ell_0 & -(\ell_A + \ell_B) \\ -(\ell_A + \ell_B) & -2\ell_0 \end{pmatrix} + V \begin{pmatrix} 2s_A & 2s_{AB} \\ 2s_{AB} & 2s_B \end{pmatrix}.$$

Its determinant is

$$h(V) = 4(\ell_0 - s_A V)(\ell_0 - s_B V) - (\ell_A + \ell_B - 2s_{AB}V)^2.$$

For global strict concavity of $R + VS$, we wish to show that $\ell_0 - s_A V > 0$ and $h(V) > 0$.

Since the value function, and hence the maximum of $R + VS$, is bounded, the latter is bounded from above along any ray $\{(M_A, M_B) : M_A = u_0 - x,\ M_B = \pi_0 - \beta x,\ x \geq 0\}$ with $\beta \geq 0$ (note that these fees are all admissible). As

$$R(u_0 - x, \pi_0 - \beta x) + VS(u_0 - x, \pi_0 - \beta x) = \left\{u_0\left[\ell_0 + \ell_A\beta\right] + \pi_0\left[\ell_0\beta + \ell_B\right]\right\} x - q(\beta)x^2$$

with the quadratic function

$$q(\beta) = \ell_0 - s_A V + (\ell_A + \ell_B - 2s_{AB}V)\beta + (\ell_0 - s_B V)\beta^2,$$

this implies that q is positive on $[0, \infty[$. Setting $\beta = 0$ yields $\ell_0 - s_A V > 0$.

Next, let $V > (\ell_A + \ell_B)/2s_{AB}$, so that $q'(0) < 0$. As a consequence, $\ell_0 - s_B V > 0$ since q would become negative at high β otherwise. Moreover, q assumes its minimum at

$$\beta^* = \frac{2s_{AB}V - \ell_A - \ell_B}{2(\ell_0 - s_B V)} > 0.$$

This minimum equals

$$q(\beta^*) = \ell_0 - s_A V - \frac{(2s_{AB}V - \ell_A - \ell_B)^2}{4(\ell_0 - s_B V)} = \frac{h(V)}{4(\ell_0 - s_B V)},$$

implying $h(V) > 0$ and concavity of $R + VS$.

As V multiplies the strictly convex function S, concavity of $R + VS$ now also follows for shadow prices $V \leq (\ell_A + \ell_B)/2s_{AB}$. □

Proof of Propositions 6 and 7

For arbitrary ρ_A and ρ_B, the first-order conditions for the fees $\overline{n}_A(p)$ and $\overline{n}_B(p)$ can be written as

$$\begin{aligned}
&(u_0 + [u(p) + \pi(p)]n_B - 2n_A)(\rho_A n_A^2 + \rho_B n_B^2) \\
&\quad + \; 2\rho_A n_A \left[\overline{R}(p) - (u_0 + u(p)n_B - n_A)\, n_A - (\pi_0 + \pi(p)n_A - n_B)\, n_B\right] \; = \; 0, \\
&(\pi_0 + [u(p) + \pi(p)]n_A - 2n_B)(\rho_A n_A^2 + \rho_B n_B^2) \\
&\quad + \; 2\rho_B n_B \left[\overline{R}(p) - (u_0 + u(p)n_B - n_A)\, n_A - (\pi_0 + \pi(p)n_A - n_B)\, n_B\right] \; = \; 0.
\end{aligned}$$

For $\rho_A = \rho_B$, this system simplifies to

$$\begin{aligned}
(u_0 + [u(p) + \pi(p)]n_B)(n_B^2 - n_A^2) + 2(\overline{R}(p) - n_B \pi_0)\, n_A &= 0 \\
(\pi_0 + [u(p) + \pi(p)]n_A)(n_A^2 - n_B^2) + 2(\overline{R}(p) - n_A u_0)\, n_B &= 0.
\end{aligned}$$

For $u_0 \neq \pi_0$, the pair of quantities stated in Proposition 6 constitutes the unique solution to these equations. For $u_0 = \pi_0 = c_0$, setting both quantities equal to $\overline{R}(p)/c_0$ solves the system. □

Proof of Corollary 1

For $u_0 = \pi_0 = c_0$, the myopically optimal expected price on side A simplifies to

$$M_A^\mu(p) = \frac{c_0[1 - \pi(p)]}{2 - [u(p) + \pi(p)]},$$

so the price difference $\overline{M}_A(p) - M_A^\mu(p)$ has the same sign as

$$1 + (u(p) - 1)\left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1-p}{2 - (\underline{u} + \underline{\pi})}\right] - \frac{1 - \pi(p)}{2 - [u(p) + \pi(p)]}.$$

Multiplying with $2 - [u(p) + \pi(p)]$ and simplifying, we see that this in turn has the same sign as

$$(u(p) - 1)\left\{(2 - [u(p) + \pi(p)])\left[\frac{p}{2 - (\overline{u} + \overline{\pi})} + \frac{1-p}{2 - (\underline{u} + \underline{\pi})}\right] - 1\right\}.$$

The expression in curly brackets is strictly concave in p; as it vanishes at $p = 0$ and $p = 1$, it is positive for $0 < p < 1$. The proof for side 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. WRIGHT (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. SHIRYAYEV (1977): *Statistics of Random Processes I.* New York: Springer-Verlag.

MIRMAN, L.J., SAMUELSON, L., and SCHLEE, E. (1994): "Strategic Information Manipulation in Duopolies", *Journal of Economic Theory*, 62, 363–384.

NOCKE, V., PEITZ, M. and K. STAHL (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.

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

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

ROTHSCHILD, 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.