

Peitz, Martin; Rady, Sven; Trepper, Piers

Working Paper

Experimentation in Two-Sided Markets

CESifo Working Paper, No. 5346

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Peitz, Martin; Rady, Sven; Trepper, Piers (2015) : Experimentation in Two-Sided Markets, CESifo Working Paper, No. 5346, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/110838>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Papers

www.cesifo.org/wp

Experimentation in Two-Sided Markets

Martin Peitz
Sven Rady
Piers Trepper

CESIFO WORKING PAPER NO. 5346
CATEGORY 11: INDUSTRIAL ORGANISATION
MAY 2015

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

Experimentation in Two-Sided Markets

Abstract

We study optimal experimentation by a monopolistic platform in a two-sided market. The platform provider is uncertain about the strength of the externality each side is exerting on the other. Setting participation fees on both sides, it gradually learns about these externalities by observing actual participation levels. This provides an informational rationale for introductory pricing, with the platform provider charging a fee below the myopically optimal level on at least one side of the market. If the externality that the other side exerts is sufficiently well known and weaker than the externality it experiences, the platform provider extracts surplus from that side by charging it a fee above the myopically optimal level. This interplay between learning and surplus extraction is crucial to the market outcome and its dynamics.

JEL-Code: D420, D830, L120.

Keywords: two-sided market, network effects, monopoly experimentation, Bayesian learning, optimal control.

Martin Peitz
Department of Economics
University of Mannheim / Germany
martin.peitz@googlemail.com

Sven Rady
Department of Economics and Hausdorff
Center for Mathematics / University of
Bonn / Germany
rady@hcm.uni-bonn.de

Piers Trepper
Department of Economics
University of Munich / Germany
piers.trepper@lrz.uni-muenchen.de

May 2, 2015

Our thanks for extremely helpful comments are due to the Editor Dirk Bergemann, four anonymous reviewers, and to participants of the Conference on Platform Markets, Mannheim, 31 May-2 June 2010, and the CEPR Applied IO School and Conference, Tel Aviv, 24-27 May 2011. We would like to thank the Studienzentrums 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.

1 Introduction

Many real-world markets are two-sided in the sense that potential participants care about the number of counterparts on the other side of the market. Transactions in such markets are often mediated through platforms.¹ It is well understood that the optimal pricing strategy of a provider of such a platform depends delicately on the precise nature and strength of the cross-group externality each side of the market is exerting on the other. But how should the provider set prices if it is uncertain about these externalities and needs to estimate them from observed market outcomes?

Arguably, such uncertainty is an important feature of platform industries: a platform provider typically cannot perfectly foresee how strongly one side reacts to the number of users on the other side. When one side of the market are buyers and the other side are sellers, for example, uncertainty about the externalities reflects uncertainty about some aspect of the buyer-seller relationship: the platform provider may be imperfectly informed about the sellers' production function or advertising technology, or about the buyers' demand function. In such a situation, the market outcome not only determines the platform provider's current profit but also yields information about the true externalities. To the extent that the provider can influence the information content of the market outcome through its choices, it may be worth its while sacrificing some short-term profit so as to extract information that will be beneficial in the long term.

The two-sidedness of a platform market renders this trade-off between earning and learning particularly interesting. To start with, changing the price on one side of the market will alter the level of participation and its information content on *both* sides. Any information extracted from market observations thus always comes as a blend of two signals, one from each side of the market. In turn, this raises the possibility that one and the same amount of information can be generated with very different price combinations. It is far from clear, therefore, how a forward-looking platform provider ought to adjust its prices relative to the myopic benchmark of current-profit maximization.

To shed some light on this dynamic pricing problem, we embed a standard model of two-sided monopoly markets into a canonical Bayesian learning framework. Like Armstrong (2006), we focus on participation decisions, with prices taking the form of access, membership or subscription fees.² Like Keller and Rady (1999), we consider a continuous-time infinite-horizon model in which there are two possible states of the world, the demand function on each side of the market is linear in either state, and observed demand is an imperfect signal of the true state because of additive Brownian noise.³ The platform

¹Examples include payment systems (where shoppers 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 (providing a platform for players and software developers), smart phones (users, application developers), nightclubs and matching agencies (men, women), news media (consumers, advertisers), shopping malls, trade fairs, B2C and B2B platforms (where buyers are interested in a large variety of offerings, and sellers in a large number of customers).

²This framework has become the workhorse model in two-sided markets. A particular example are trade fairs where exhibitors pay stand rental fees and visitors entrance fees. The trade fair company as the platform provider has to decide on these prices in advance of any event; in particular, for new events the platform may face strong uncertainty as to the strength of cross-group external effects which determine the value of interaction on each side of the platform.

³While the focus of Keller and Rady (1999) is on the effects of a changing state of the world on the learning dynamics, here we assume that the state is fixed over time. In Peitz, Rady and Trepper (2011), we also analyze a scenario in which the platform provider selects quantities and learns from prices. This variant of the model turns out to be more tractable, but the price setting version is clearly more relevant

provider is assumed to incur no costs, so that revenue and profit are synonymous.

In this framework, current revenues and the information content of observed participation levels are both quadratic functions of the two fees. This allows us to compute the myopic benchmark in closed form and to characterize the optimal pricing policy of a forward-looking platform provider by means of the linear first-order conditions associated with the provider's Bellman equation. A familiar advantage of our stationary continuous-time setting is that, at any given belief about the state of the world, the provider's incentive to deviate from the myopic benchmark is entirely captured by a single number: the shadow price of information at this belief.⁴ Even though in general there are no closed-form solutions for the value function and the shadow price, we can exploit this fact to derive qualitative predictions as to how the optimal fees will be adjusted relative to the myopic benchmark. In geometric terms, we are tracing out the locus of tangency points between iso-revenue and iso-information curves at the given belief; the higher is the shadow price of information, the farther along this locus we move, attaining higher quantities of information.

For low shadow prices of information,⁵ the direction of price experimentation can be determined by a simple comparison between the slope of the iso-information curve through the myopically optimal fees and two numbers: the common slopes of all iso-revenue curves along a horizontal and a vertical axis through these fees, respectively. This comparison already yields the important insight that, even if lowering either fee makes the market outcome more informative, uncertainty about the cross-group externalities may induce the platform provider to raise one fee above the myopic benchmark. In fact, while the two fees can be complements with respect to the quantity of information, they are always substitutes with respect to current revenue. A lower fee on one side of the market then makes reducing the fee on the other side more attractive from an informational perspective, but less attractive with respect to current revenue, and this second effect may well dominate.

The main part of the paper is concerned with analytical results on optimal price experimentation which do not rely on the shadow price of information being small. The most tractable scenario in our model is that of *symmetric externalities*, meaning that in both states of the world the externality one side exerts on the other is exactly as strong as the converse externality. Symmetric externalities neutralize each other completely in the sense that it becomes optimal for the platform provider to behave as if it were a monopolist in two unrelated markets. The revenue-maximizing fees are independent of the provider's beliefs, therefore, and there is no incentive to deviate from these fees.

The second most tractable scenario is that of *one-sided externalities*. This means that side A , say, benefits from an increase in participation on side B , but side B does not benefit from increased participation on side A .⁶ In this scenario, the platform provider always

in applications.

⁴This shadow price is the product of the subjective variance of the state of the world and the second derivative of the value function, divided by twice the interest rate. As usual in single-agent Bayesian learning problems of this kind, a non-negative value of information translates into a value function that is convex in beliefs, implying a non-negative shadow price of information. Another standard result is that the shadow price decreases in the interest rate.

⁵The shadow price is low if the platform provider is subjectively quite certain about the true state of the world, for example, or if the interest rate is high.

⁶As an example, consider readers whose utility of a magazine is independent of the amount of advertising. As another example, consider a nightclub for heterosexual singles (men and women). Externalities are one-sided if people from one group go there for the music, while people from the other group are also

sets a fee lower than the myopic optimum on side B , and a fee higher than the myopic optimum on side A . In fact, the only way to make the market outcome more informative is to raise participation on side B by lowering the fee there, giving the participants on both sides a larger surplus. As raising the fee on side A does not affect participation on side B , the provider can safely extract part of the extra surplus given to side A by charging it a fee above the myopic optimum.

The same intuition applies to the case of *one-sided uncertainty* where the externality exerted by side A is perfectly known but positive. Again, the platform provider can only increase the amount of information by lowering the fee on side B . Whether part of the extra surplus this creates on side A is extracted through a higher fee on that side now depends on a comparison of indirect price effects, however. If the responsiveness of expected participation on side A to price changes on side B (as measured by the absolute value of the relevant partial derivative) is higher than the responsiveness of expected participation on side B to price changes on side A , the platform provider will again raise the fee on side A .

Moving on two scenarios with uncertainty about both cross-group externalities, we distinguish between *ordered* and *mixed price effects*. We say that price effects are ordered if in one state of the world the responsiveness of expected participation on either side to a change in either fee is always at least as high as in the other state; otherwise we say that price effects are mixed. Price effects are ordered if the two cross-group externalities are positively correlated across states of the world. Price effects are mixed when the externalities are negatively correlated and the spread between the two possible externality parameters on each side is of similar size.

The scenarios described so far are all limiting cases of ordered price effects. Applying a standard continuity argument, we first show that for approximately symmetric externalities, the platform provider always sets both fees below their myopically optimal levels. Intuitively, the direction of experimentation must be the same on both sides in this case, and charging less than the myopic fee clearly makes the market outcome more informative. By the same continuity argument, a scenario of approximate one-sided uncertainty in which one externality is much better known than the other can give rise to a fee above the myopic optimum on one side of the market, with exactly the same intuition as above.

For ordered price effects more generally, we formulate sufficient conditions which ensure fee adjustments of a particular sign. These conditions again have a natural interpretation in terms of the geometry of iso-revenue and iso-information curves. They point to asymmetries in the signal-to-noise ratios on the two sides of the market as another natural reason for fee increases on one side. If the spread between the two possible externality parameters is roughly the same on both sides but the level of noise is considerably smaller on side B , for example, there is more to be learned about the true state of the world by lowering the fee on side A . If the platform provider expects this to induce a sufficiently large increase in the surplus given to side B , it will raise the fee on that side above the myopically optimal level so as to capture part of the extra surplus.

It is straightforward to provide similar sufficient conditions for mixed price effects. We refrain from spelling them out in this paper and instead turn to the very tractable special case of *antisymmetric externalities* with the two sides being symmetric in all other regards.⁷ When the provider is highly uncertain about which side exerts the stronger

interested in getting to know people of the opposite sex.

⁷This includes the limiting case in which the platform provider knows that only one side of the market exerts an externality on the other, but does not know which side it is.

externality, it charges both sides less than in the myopic benchmark; the intuition for this finding is the same as the one suggested for approximately symmetric externalities. When the platform provider is fairly confident in knowing the side that exerts the stronger externality, it learns most effectively by lowering the fee on this side; whether it is optimal to recoup some of the surplus this creates by raising the fee on the other side then depends on the actual strength of the externality. By continuity, these findings carry over to scenarios of approximate antisymmetry.

In the absence of closed-form solutions, we resort to numerical simulations when analyzing the impact of optimal experimentation on market participation, the dependence of optimal policies on the model parameters, and dynamic implications for prices and quantities. More precisely, we compute optimal policies in the undiscounted limit of our model, which provides a tight upper bound on the shadow price of information and hence corresponds to maximal experimentation. It is well known that these policies can be computed without solving for the value function first. In fact, we show that the computation reduces to solving a quadratic equation in one unknown, albeit with coefficients that are too cumbersome for analytical results.

Whenever the platform provider reduces both fees below the myopically optimal level, expected participation on both sides obviously increases whatever the true state of the world. In scenarios where the provider sets a fee above the myopically optimal level on one side, however, expected participation on this side may well decrease in some state. What we see in numerical examples can be summarized as follows. Participation rises on a given side either when the externality exerted on this side is strong (so that any fee increase on this side is more than compensated by the concurrent fee reduction on the other), or when this externality is weak and the platform provider believes it is weak (so that the provider will refrain from raising the fee on this side). Participation falls on a given side when the externality exerted on this side is weak but the platform provider believes it is strong (and so raises the fee on this side in an attempt to capture some of the rents created by lowering the other fee).

Holding the externalities in one state of the world fixed and changing those in the other state so that we move from symmetric externalities to one-sided uncertainty and eventually to antisymmetric externalities, we are gradually increasing the incentives to experiment in the sense that the revenue-maximizing fees become more and more sensitive to the true state of the world. In our numerical examples, this goes hand in hand with an increase in the extent of experimentation that takes place, involving ever larger deviations from the myopic benchmark. Changes in the model parameters can have a surprisingly strong effect on the slope and curvature of the optimal pricing policies, moreover.

This great variability in the shape of optimal policy functions gives rise to rich intertemporal patterns. An optimal fee can be decreasing and concave in beliefs over some part of the unit interval, for example, and increasing and convex over another part; by Ito's Lemma, the sign of the expected infinitesimal change in the fee will depend on the current belief, therefore. In addition, the sample-path properties of prices and quantities can depend quite strongly on the true constellation of cross-group externalities. In one of our numerical illustrations with approximately antisymmetric externalities, for example, one of the fees is expected to rise both in the short and long run; this rise is fairly gentle on average in one state of the world, but involves drastic adjustments in the other.

Broadly speaking, we can distinguish two different regimes when it comes to price dynamics. In the *two-sided experimentation regime*, consumers on both sides initially are charged lower fees than in the myopic benchmark, whereas in the *experimentation and*

exploitation regime one side initially faces a higher fee. In either regime, a price path which starts at a lower fee and then rises more steeply on average amounts to introductory pricing with larger initial discounts for informational reasons.⁸ Despite higher prices on one side for some time, our simulations suggest that experimentation tends to raise participation on both sides most of the time, which tends to benefit the users of the platform. Moreover, experimentation can lead to a non-monotonic time trend in participation on one side, with increasing participation as beliefs move towards the middle of the unit interval where uncertainty is most pronounced, and declining participation as the platform provider becomes more confident of the true state of the world.

The remainder of the paper is structured as follows. Section 2 reviews the related literature. Section 3 presents the model, Section 4 considers the benchmark of myopic behaviour, and Section 5 describes the evolution of beliefs. Section 6 characterizes the optimal pricing policy and its limit as the platform provider becomes infinitely patient. Section 7 presents pricing implications. Section 8 discusses expected quantities, comparative statics and dynamic implications. Section 9 concludes. Auxiliary technical results and all proofs are relegated to the appendix.

2 Related Literature

Pricing strategies in two-sided markets have received a lot of attention in industrial economics. 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). Weyl (2010) proposes the alternative solution concept of insulating tariffs. Empirical work includes Rysman (2004) and Kaiser and Wright (2006). For a selective survey of this literature, see Rysman (2009); a textbook treatment can be found in Belleflamme and Peitz (2010). 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 in order to learn about the true state of the world. 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 cross-group externalities of uncertain size provide incentives for the experimenting platform provider to initially lower at least one price.

We are aware of one other contribution that embeds two-sided markets in a dynamic setting. Cabral (2011) considers a monopoly platform whose users can reassess their participation decisions with some probability in each period. He finds that the dynamic model may have a unique equilibrium even when the static pricing model exhibits multiple equilibria. He also shows that this setting can provide a dynamic foundation for the equilibrium concept of insulating tariffs proposed by Weyl (2010).

The economics literature on optimal experimentation by a single Bayesian decision maker starts with the work of Prescott (1972) and Rothschild (1974); a survey of this literature can be found in Bergemann and Välimäki (2008). Our contribution here is

⁸An alternative explanation for introductory pricing in two-sided markets 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.

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. Like in that model, the maximand in the Bellman equation is the sum of a concave quadratic (expected current revenue) and a convex quadratic (the shadow price of information times the quantity of information to be gathered from the market outcome). We make the novel observation that in the absence of any restrictions on the action space, the shadow price of information must be small enough to render the combined quadratic strictly convex. While there exists an action that leads to a completely uninformative market outcome, this action does not maximize expected revenues at any belief and hence is never chosen. In contrast to the Keller-Rady model with a fixed state of the world, learning is thus always complete in the long run.

Closely related work on dynamic pricing problems with Brownian information is due to Bergemann and Välimäki (1997, 2002) and Bonatti (2011). These authors study the introduction of a new product whose quality can be either high or low. Buyers and sellers learn about it by observing a statistic that aggregates the experience of all buyers. As the informativeness of this statistic increases with the mass of consumers who try the new product, a positive value of information translates into an incentive to set this product’s price below the myopically optimal level. The causal chain from lower prices to higher participation levels to more informative market outcomes is also present in our model; in sufficiently asymmetric two-sided markets, however, it can be overturned by the incentive to extract surplus on a side that benefits strongly from an experimentation-induced fee reduction on the other side.

Like the platform provider in our model, the monopolistic seller in Bonatti (2011) possesses multiple pricing instruments with which to pursue the conflicting goals of exploitation and exploration—one for each “type” or “group” of customers. In fact, this seller chooses a non-linear tariff so as to screen consumers for their willingness to pay (second-degree price discrimination); different consumer groups are linked through incentive compatibility constraints. In our model, by contrast, the platform can identify to which group a consumer belongs and thus engages in group pricing (third-degree price discrimination); the two groups are linked through the external effects that they exert on each other. These effects are exogenous, and varying their strength allows for a rich analysis of the effects of uncertainty about them on market outcomes.

3 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 .

In each period, there is a continuum (of mass m) of potential participants on both sides of the market. Potential participants are short-lived. The platform provider sets membership fees (M_A, M_B) , but no usage fee. Potential participants observe these fees and then make their participation decisions. A potential participant on side $i \in \{A, B\}$ prefers joining the platform to the outside option if $v_i + e_i^\theta n_j - M_i > u_i^0$ where v_i is the intrinsic platform value, e_i^θ is a cross-group externality parameter in the state of the world $\theta \in \{0, 1\}$,⁹ n_j is the expected mass of participants on the other side of the market,

⁹The assumption that both externalities depend on the same binary state of the world is motivated in the three examples below.

and u_i^0 is the participant-specific valuation of the outside option which is drawn from the uniform distribution on the interval $[0, m]$. For the sake of concreteness, we assume positive intrinsic values, and non-negative externalities of strength less than 1, that is, $0 \leq e_i^\theta < 1$ for $i = A, B$ and $\theta = 0, 1$.¹⁰ Without loss of generality, we further assume $e_A^0 < e_A^1$ and $|e_B^1 - e_B^0| \leq e_A^1 - e_A^0$. Given (M_A, M_B) , the expected masses of active participants n_A and n_B in the current period are characterized by the system of linear equations¹¹

$$n_A = v_A + e_A^\theta n_B - M_A, \quad (1)$$

$$n_B = v_B + e_B^\theta n_A - M_B. \quad (2)$$

While the intrinsic values and the possible externality parameters are common knowledge, the state of the world is known to market participants, but not to the platform provider.¹² Actual participation is expected participation plus some noise term that will be specified below; this noise prevents the platform provider from learning θ instantly.

To understand the assumptions on cross-group externalities, we consider a two-sided platform which hosts buyers (side A) and sellers (side B). We postulate that buyers and sellers cannot bypass the platform, i.e., all trade takes place on the platform. After the participation decision, participating buyers and sellers interact, as we will specify in our micro-foundation below. At the participation stage, the expected surplus a seller obtains from interacting with a buyer is either e_A^0 or e_A^1 ; the expected surplus a buyer obtains from interacting with a seller is either e_B^0 or e_B^1 . In our specification, the indirect utility from cross-group externalities is linear in participation on the other side. The utility of participants is independent of participation on the own side. In terms of the underlying micro model of the buyer-seller interaction this means that the offers by sellers are totally differentiated, i.e., a seller's demand for its product is independent of the pricing decisions of all other sellers.

Two special cases lead to such a demand structure: First, suppose that consumers have a quasi-linear utility function in the products potentially on offer at the platform and an outside good. Consumers view the different products as equally attractive, but independent, and thus have the same demand function D for each product.

Second, suppose that each consumer is interested in exactly one product among all the potential products, while she does not derive any utility from all other products. At the participation stage, consumers do not know which product they like (nor do the sellers know this). Consumers have a quasi-linear utility function in the product they like and an outside good. The platform operates as a matching platform and is able to perfectly match buyers and sellers, i.e., a buyer is matched to the seller whose product she likes whenever this product is available on the platform. If the buyer's net surplus from the successful match is u , her expected surplus is un_B .

Within this framework we provide several examples of buyer-seller interactions and determine the values of e_A^0 , e_A^1 , e_B^0 , and e_B^1 .

¹⁰The upper bound of the strength of the externality guarantees that the equilibrium at the participation stage is unique and stable.

¹¹Linearity obtains because of the uniform distribution of the value of the outside option. We implicitly assume that m is sufficiently large such that $n_A, n_B < m$.

¹²We 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.

Example 1: Suppose monopoly sellers use non-linear prices so as to extract all surplus in each buyer-seller interaction. Suppose that all sellers use a common input with a high or low input price which is uncertain from the viewpoint of the platform, but is privately known to buyers and sellers. This gives rise to constant marginal costs which are either c_H or c_L . In this case,

$$\pi^{pd}(c) = \int_c^\infty D(P)dP \quad \text{and} \quad u^{pd}(c) = 0.$$

The values of the four parameters of our model are thus $e_A^0 = \pi^{pd}(c_H)$, $e_A^1 = \pi^{pd}(c_L)$, $e_B^0 = e_B^1 = 0$.

Example 2: Consider the same setting as in Example 1, but let us postulate that each monopoly seller sets a price P per unit of output.¹³ The buyer's demand function is assumed to be log-concave where positive. Then, at the participation stage, buyers and sellers rationally anticipate that each seller solves

$$\max_P (P - c)D(P),$$

where c is either c_H or c_L . The unique profit-maximizing price is denoted by $P^m(c)$. We obtain

$$\pi^m(c) = (P^m(c) - c)D(P^m(c)) \quad \text{and} \quad u^m(c) = \int_{P^m(c)}^\infty D(P)dP.$$

The values of the four parameters of our model are thus $e_A^0 = \pi^m(c_H)$, $e_A^1 = \pi^m(c_L)$, $e_B^0 = u^m(c_H)$, and $e_B^1 = u^m(c_L)$.

Example 3: Instead of the platform facing uncertainty about the sellers' marginal cost, it may face aggregate demand uncertainty. Suppose that demand is described by a state which is either low or high, $d \in \{d_L, d_H\}$. The demand function is then written as $D(P; d)$. It is assumed to satisfy $D(P; d_H) > D(P; d_L)$ for all P with $D(P; d_L) > 0$. Functions $D(P; d_L)$ and $D(P; d_H)$ are assumed to be log-concave where positive. The realization of the demand state is learnt by buyers and sellers. They rationally anticipate that each seller solves

$$\max_P (P - c)D(P; d),$$

where d is either d_H or d_L . The unique profit-maximizing price is denoted by $P^m(d)$. We then have

$$\pi^m(d) = (P^m(d) - c)D(P^m(d); d) \quad \text{and} \quad u^m(d) = \int_{P^m(d)}^\infty D(P; d)dP.$$

The values of the four parameters of our model are thus $e_A^0 = \pi^m(d_L)$, $e_A^1 = \pi^m(d_H)$, $e_B^0 = u(d_L)$, and $e_B^1 = u(d_H)$. While, by construction, $e_A^1 > e_A^0$, it is not necessarily the case that $e_B^1 > e_B^0$. Thus, this simple example is sufficiently rich to generate positive and negative correlation of the externality parameters across the two sides of the market.

In these examples, sellers set the product price. It is straightforward to generate additional examples where, in addition to setting price, sellers invest in advertising or product

¹³We implicitly assume that the platform either cannot observe the price or, if it does, cannot draw inferences on u and π because it does not understand the buyer-seller interaction.

quality. For instance, sellers may use advertising as a complement as in Becker and Murphy (1993), and the platform is uncertain about the sellers' advertising technology. Alternatively, sellers may invest in product quality, and the platform faces uncertainty with respect to the sellers' investment cost function.

With these simple micro-foundations in place, we return to the stage at which users make their participation decisions. As the product $e_A^\theta e_B^\theta < 1$, the system (1)–(2) has a unique solution, given by

$$n_A(M_A, M_B, \theta) = \ell_0^\theta [v_A - M_A] + \ell_A^\theta [v_B - M_B], \quad (3)$$

$$n_B(M_A, M_B, \theta) = \ell_0^\theta [v_B - M_B] + \ell_B^\theta [v_A - M_A], \quad (4)$$

where

$$\ell_0^\theta = \frac{1}{1 - e_A^\theta e_B^\theta}, \quad \ell_A^\theta = \frac{e_A^\theta}{1 - e_A^\theta e_B^\theta} \quad \text{and} \quad \ell_B^\theta = \frac{e_B^\theta}{1 - e_A^\theta e_B^\theta}$$

measure the direct and indirect effects, respectively, of lowering M_A or M_B in state θ .

For reasons that will become clear very soon, we refrain from imposing a non-negativity constraint on expected participation and interpret the system (3)–(4) as describing participation decisions for *any* prices (M_A, M_B) . In other words, we allow the platform provider to charge arbitrarily high fees. Arbitrarily low fees are unproblematic, by contrast, since negative fees have a natural interpretation as subsidies for participation.

In every period $t \in [0, \infty[$, the platform provider sets prices (M_A^t, M_B^t) and then observes the increments of the cumulative quantity processes N_A^t and N_B^t . These are given by

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

where Z_B^t and Z_A^t are independent standard Brownian motions and the constants σ_A and σ_B are positive. These equations are the continuous-time limit of the equations $\Delta N_A^t = n_A(M_A^t, M_B^t, \theta) \Delta t + \sigma_A \sqrt{\Delta t} \varepsilon_t^A$ and $\Delta N_B^t = n_B(M_A^t, M_B^t, \theta) \Delta t + \sigma_B \sqrt{\Delta t} \varepsilon_t^B$ with $t = 0, \Delta t, 2\Delta t, \dots$ and $\varepsilon_t^i \sim \text{IIN}(0, 1)$. They capture the idea that consumers make occasional mistakes or experience taste shocks that are correlated across consumers and independent over time, either of which are not observed by the other side. Alternatively, these equations can be interpreted to include “noise traders” on each side of the market, who participate or stay away for reasons unrelated to those captured in equations (1)–(2).

Working with normal noise distributions keeps the updating of the platform provider's beliefs tractable; at the same time, their full support implies a positive probability for observed quantity increments and cumulative quantities to be negative—even if the latter do increase in expectation. In line with the literature,¹⁴ we accept this weakness of the Brownian framework for the sake of tractability. This in turn is the reason why we do not insist on expected quantity increments being non-negative at each instant.¹⁵

The platform provider's revenue increment from fees (M_A^t, M_B^t) is

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

¹⁴For instance, Jovanovic (1979), Felli and Harris (1996), Moscarini (2005), Sannikov (2008), Prat and Jovanovic (2014) and Papageorgiou (2014) specify cumulative output with additive Brownian noise.

¹⁵That said, expected participation will turn out to be non-negative at the optimal fees if the platform provider is sufficiently impatient or very certain of the true state, or if the two sides of the market are not too asymmetric.

We normalize platform costs to zero and denote the probability that the platform provider initially assigns to state $\theta = 1$ by p_0 , assuming that this prior belief is non-degenerate, i.e., $0 < p_0 < 1$. Hence, the platform provider's total expected profits (expressed in per-period terms) are

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

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

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

Let p_t be the subjective probability that the platform provider assigns to state $\theta = 1$ at time t . Invoking the law of iterated expectations, we can rewrite total expected profits as

$$\mathbb{E}^{p_0} \left[\int_0^\infty r e^{-rt} R(M_A^t, M_B^t, p_t) dt \right] \quad (5)$$

where

$$R(M_A, M_B, p) = M_A \mathbb{E}^p [n_A(M_A, M_B, \theta)] + M_B \mathbb{E}^p [n_B(M_A, M_B, \theta)] \quad (6)$$

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

4 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 uniquely determined by the (linear) first-order conditions.

To compute these fees, we write the expected quantities appearing on the right-hand side of (6) as

$$\begin{aligned} \mathbb{E}^p [n_A(M_A, M_B, \theta)] &= \ell_0(p)[v_A - M_A] + \ell_A(p)[v_B - M_B], \\ \mathbb{E}^p [n_B(M_A, M_B, \theta)] &= \ell_0(p)[v_B - M_B] + \ell_B(p)[v_A - M_A], \end{aligned}$$

where

$$\ell_i(p) = p\ell_i^1 + (1-p)\ell_i^0 \quad (i = 0, A, B)$$

measures the *expected* direct and indirect effect, respectively, of lowering M_A or M_B given the belief p . We note that $0 \leq \ell_i(p) < \ell_0(p)$ for $i = A, B$ and all p . In the following, it will be convenient to write $\ell_{AB}(p) = \frac{1}{2}[\ell_A(p) + \ell_B(p)]$.

With the dependence on the belief p suppressed, the expected current revenue associated with the fees (M_A, M_B) now becomes

$$[\ell_0 v_A + \ell_A v_B] M_A + [\ell_0 v_B + \ell_B v_A] M_B - \ell_0 M_A^2 - 2\ell_{AB} M_A M_B - \ell_0 M_B^2,$$

and the first-order conditions yield the myopically optimal fees

$$M_A^\mu = v_A - \frac{2(\ell_0^2 - \ell_{AB}\ell_A)v_A - \ell_0(\ell_A - \ell_B)v_B}{4(\ell_0^2 - \ell_{AB}^2)}, \quad (7)$$

$$M_B^\mu = v_B - \frac{2(\ell_0^2 - \ell_{AB}\ell_B)v_B - \ell_0(\ell_B - \ell_A)v_A}{4(\ell_0^2 - \ell_{AB}^2)}. \quad (8)$$

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 $e_A^0, e_A^1, e_B^0, e_B^1$, 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 cross-group externalities on both sides. To be precise, the fee M_A^μ is increasing in v_B 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 v_B 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 v_A if ℓ_A exceeds ℓ_B by a sufficient amount, and vice versa for M_B^μ and v_B . 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 indirect price effects are equal ($\ell_A = \ell_B$) or close together. For $\ell_A = \ell_B$, we actually have $M_i^\mu = v_i/2$ for $i = A, B$; i.e., symmetric expected price effects completely neutralize each other so that the platform provider sets fees as if it were a monopolist in two separate markets.

The level curves of R in (M_A, M_B) -space are concentric ellipses; the farther away such an iso-revenue curve lies from the myopically optimal pair of fees (M_A^μ, M_B^μ) , the lower is the expected revenue. Further details on the function R can be found in Appendix A.1.

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 note that, as the upper envelope of linear functions (one for each fixed pair of fees), R^μ is convex.

5 The Evolution of Beliefs

The platform provider revises its beliefs over time. We define

$$S(M_A, M_B) = \left[\frac{n_A(M_A, M_B, 1) - n_A(M_A, M_B, 0)}{\sigma_A} \right]^2 + \left[\frac{n_B(M_A, M_B, 1) - n_B(M_A, M_B, 0)}{\sigma_B} \right]^2.$$

As the sum of squares of linear functions, S is convex.

Lemma 1 *The platform provider’s beliefs evolve according to*

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

Any pricing policy for which $S(M_A^t, M_B^t)$ is bounded away from 0 induces complete learning in the long run: as $t \rightarrow \infty$ the belief p_t almost surely converges to 1 if the true state of the world is (e_A^1, e_B^1) , and to 0 otherwise.

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).¹⁶ The more informative the observations are, the more strongly the beliefs react to them. If the information content is bounded away from zero, the continuous accrual of information ensures that the truth is learnt eventually. In particular, this is the case for the myopically optimal pricing policy.¹⁷ We shall see shortly that an optimal policy generates no less information than the myopic one, and hence gives rise to complete learning as well.

An explicit expression for the function S and a more detailed discussion of its properties can be found in Appendix A.2. Here, we just note that for $e_B^1 \neq e_B^0$, the function S is strictly convex, and the level curves of S in (M_A, M_B) -space are concentric ellipses. For $e_B^1 = e_B^0$, the iso-information curves are parallel straight lines. The farther away an iso-information curve lies from the uninformative pair of fees (v_A, v_B) , the higher is the amount of information generated by the respective fee combinations.

6 The Optimal Pricing Strategy

In view of the objective function (5) and the law of motion (9), standard arguments yield the following Bellman equation for the platform provider's value function, v :

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

Arguing exactly as in Keller and Rady (1997, Appendices A-C), one shows that v is convex and twice continuously differentiable with $p^2(1-p)^2 v''(p) \rightarrow 0$ as $p \rightarrow 0$ or 1 ; moreover, v is the only continuous real-valued function on $[0, 1]$ that solves (10) on $]0, 1[$ and coincides with the myopically optimal revenue R^μ on $\{0, 1\}$.

We can interpret the second term of the maximand in the Bellman equation as the value of information, given by the product of the (non-negative) shadow price of information,

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

and the quantity of information, $S(M_A, M_B)$. When $V = 0$, the value of information is zero, and the platform provider chooses the myopically optimal prices. When $V > 0$, the platform provider experiments, i.e., deviates from the myopic strategy so as to increase the information content of its demand observations. As a consequence, any optimal pricing policy has $S(M_A^t, M_B^t)$ bounded away from 0 and thus yields complete learning in the long run by Lemma 1.

The maximand in (10) is the sum of two quadratic functions, one of them strictly concave (expected current revenue), the other convex (value of information). As the value function is bounded, and as fees are unbounded above and below, the shadow price

¹⁶If the platform provider were uncertain about the intrinsic platform values (v_A, v_B) instead of the externalities (e_A, e_B) , the quantity of information would be independent of the fees charged. The platform provider would then trivially always set the myopically optimal fees.

¹⁷ $(M_A^\mu(p), M_B^\mu(p)) \neq (v_A, v_B)$ for all p because the latter fees generate an expected current revenue of zero and marginally lowering one fee would improve upon that.

of information must be small enough for the combined quadratic to be bounded above and hence concave.¹⁸ In fact, we can prove strict concavity.

Lemma 2 *Evaluated at the platform provider's value function, the maximand in the Bellman equation (10) is strictly concave in the fees (M_A, M_B) .*

The optimal fees $(M_A^*(p), M_B^*(p))$ are thus fully characterized by the (linear) first-order conditions for the maximization problem in (10). Explicit representations of these fees in terms of the model parameters, the shadow price of information and the myopically optimal fees are given in Appendix A.3.

These algebraic expressions are closely linked to the geometry of iso-revenue and iso-information curves. To see this, it is useful to think of the platform provider as following a two stage-procedure. At the first stage, it determines the combination of fees that maximizes the quantity of information subject to the constraint that a certain current expected revenue be achieved. This amounts to identifying the points of tangency between iso-information and iso-revenue curves in the (M_A, M_B) -plane; their locus is a curve that can be parameterized by the shadow price of information, starting at the myopically optimal fees for $V = 0$ and moving away from the the myopic optimum as V increases. At the second stage, the provider chooses a point on this locus; it is at this stage that the precise value of V comes into play.

To understand the direction in which this locus leaves the myopic optimum, it is instructive to study the explicit expressions in Appendix A.3 for V close to zero. They show that the differences $M_A^* - M_A^\mu$ and $M_B^* - M_B^\mu$ are of the same sign as $\ell_0 S_A^\mu - \ell_{AB} S_B^\mu$ and $\ell_0 S_B^\mu - \ell_{AB} S_A^\mu$, respectively, where S_A^μ and S_B^μ are the partial derivatives of S at the myopically optimal fees. For concreteness, assume that both are negative, so that lowering either fee increases the information content of observed participation. For small shadow prices of information, we then have $M_A^* < M_A^\mu$ if $-S_A^\mu/S_B^\mu < -\ell_{AB}/\ell_0$, and $M_B^* < M_B^\mu$ if $-S_A^\mu/S_B^\mu > -\ell_0/\ell_{AB}$. While $-S_A^\mu/S_B^\mu$ is the slope of the iso-information curve through the myopically optimal fees, a simple computation reveals that all iso-revenue curves have slope $-\ell_{AB}/\ell_0$ when $M_A = M_A^\mu$, and slope $-\ell_0/\ell_{AB}$ when $M_B = M_B^\mu$ (see Appendix A.1). If $-S_A^\mu/S_B^\mu < -\ell_{AB}/\ell_0$, therefore, the iso-revenue curve through a point (M_A^*, M_B) in a neighbourhood of the myopic optimum is steeper than the iso-revenue curve through this point, so the platform provider can raise both the expected revenue and the information content of observed participation by reducing M_A . Similarly, if $-S_A^\mu/S_B^\mu > -\ell_0/\ell_{AB}$, then the iso-revenue curve through a point (M_A, M_B^*) in a neighbourhood of the myopic optimum is flatter than the iso-revenue curve through this point, and the platform provider can increase both revenue and information by reducing M_B . We know from Section 4 that $-\ell_0/\ell_{AB} < -1 < -\ell_{AB}/\ell_0$. At small shadow prices of information, it is thus optimal to set both fees below their myopically optimal levels when $-\ell_0/\ell_{AB} < -S_A^\mu/S_B^\mu < -\ell_{AB}/\ell_0$. When $-S_A^\mu/S_B^\mu < -\ell_0/\ell_{AB}$ or $-S_A^\mu/S_B^\mu > -\ell_{AB}/\ell_0$, however, it is optimal to set one fee above the myopic optimum. Figure 1 illustrates the three scenarios that can arise this way.

Thus, even if lowering either fee increases the quantity of information, the uncertainty about the externalities between the two sides of the market may induce a learning platform

¹⁸If we had imposed a non-negativity constraint on the expected quantities (3)–(4), the situation would be more complicated. Given the smaller choice set, the shadow price of information would be higher. The combined quadratic could then possess a saddle-point and be strictly increasing towards the boundaries of the set of “admissible” fee combinations.

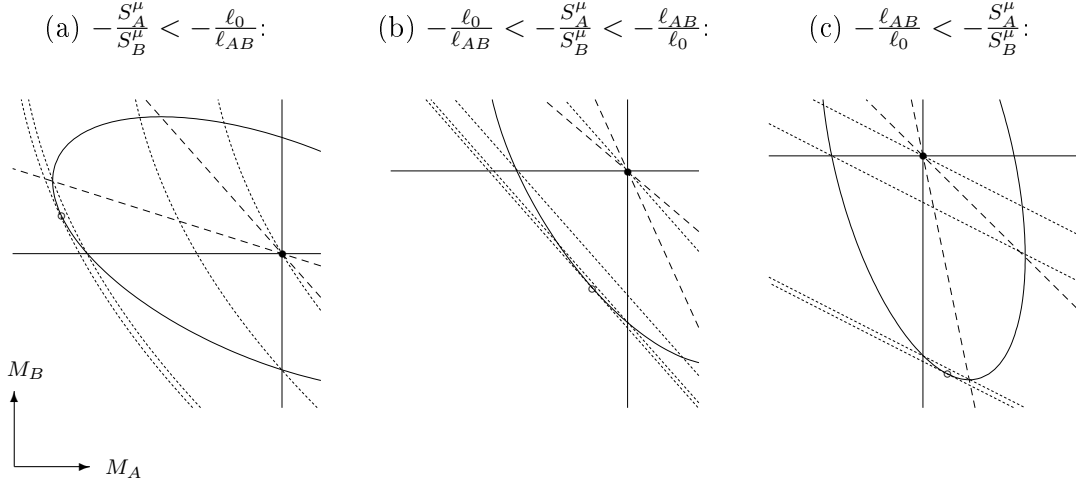


Figure 1: Directions of experimentation for small shadow prices of information and negative partial derivatives of the quantity of information at the myopic optimum. The point marked with a dot indicates the myopically optimal fees (M_A^μ, M_B^μ) , the solid curve is the iso-revenue curve corresponding to 99% of the maximal expected revenue R^μ , and the slopes of the two dashed lines through the myopic optimum are $-\ell_{AB}/\ell_0$ and $-\ell_0/\ell_{AB}$, respectively. The four dotted curves are iso-information curves: one through the myopic optimum, two through the points of intersection of the iso-revenue curve with the lines $M_A = M_A^\mu$ and $M_B = M_B^\mu$, and one tangential to the iso-revenue curve. The point of tangency is marked with a circle. Parameter values: $v_A = v_B = \sigma_A = \sigma_B = 1$, $e_A^0 = 0.1$ and $e_A^1 = 0.9$; (a) $e_B^0 = 0.95$, $e_B^1 = 0.2$ and $p = 0.15$; (b) $e_B^0 = 0.2$, $e_B^1 = 0.7$ and $p = 0.5$; (c) $e_B^0 = 0.1$, $e_B^1 = 0.2$ and $p = 0.75$.

provider to raise one fee. In fact, 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 is positive for the parameters underlying Figure 1), but less attractive as far as expected current revenue is concerned (its cross-partial derivative is always negative). The second effect dominates in Figures 1(a) and 1(c).¹⁹

For shadow prices V further away from zero, the expressions for the optimal fees are harder to analyze, but we shall see in the following section that they nevertheless allow us to investigate (and explain in economic terms) the *directions* of optimal experimentation without precise knowledge of the value function, simply treating V as a parameter.

As to the possible *extent* of optimal experimentation, we can provide a tight upper bound on the shadow price of information. For this purpose, we write the Bellman

¹⁹ Of course, if either S_A^μ or S_B^μ is positive, the directions of experimentation at small shadow prices of information are unambiguous: if $S_A^\mu > 0$, then $M_A > M_A^\mu$ and $M_B < M_B^\mu$; if $S_B^\mu > 0$, these inequalities are reversed. In either case, the tangency between the iso-revenue and iso-information curves occurs at a point where both curves are upward sloping. A positive S_A^μ arises for example when $v_A = 1$, $v_B = 3$, $\sigma_A = 2$, $\sigma_B = 1$, $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = 0.95$, $e_B^1 = 0.2$ and $p = 0.85$; a positive S_B^μ is obtained for the same externality parameters when $v_A = v_B = 1$, $\sigma_A = 3$, $\sigma_B = 1$ and $p = 0.15$.

equation in the form $0 = \max_{M_A, M_B} \{R - v + VS\}$ and note that the maximum remains zero, and the set of maximizers is unchanged, when we divide the maximand by the quantity of information, S .²⁰ Re-arranging then yields

$$V(p) = \min_{M_A, M_B} \frac{v(p) - R(M_A, M_B, p)}{S(M_A, M_B)}. \quad (11)$$

Arguing as in Keller and Rady (1997, Theorem 5.2 and Appendix E.1), 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*

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

as $r \downarrow 0$. This in turn means that the shadow price of information $V(p)$ increases monotonically to

$$\bar{V}(p) = \min_{M_A, M_B} \frac{\bar{R}(p) - R(M_A, M_B, p)}{S(M_A, M_B)} \quad (12)$$

as the platform provider becomes more and more patient. Intuitively speaking, the lower the platform provider's discount rate, the greater is its incentive to deviate from the myopic optimum. The shadow price of information, as a measure of this incentive, is thus maximal for $r = 0$; for any positive discount rate, we have $0 \leq V < \bar{V}$ on the open unit interval.

The transformed Bellman equation (11) further implies that for $r \downarrow 0$, the optimal fees $(M_A^*(p), M_B^*(p))$ converge to

$$(\bar{M}_A(p), \bar{M}_B(p)) = \arg \min_{M_A, M_B} \frac{\bar{R}(p) - R(M_A, M_B, p)}{S(M_A, M_B)}, \quad (13)$$

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 Bergemann and Välimäki (1997) or Bolton and Harris (2000) for details on this performance criterion. We will refer to (\bar{M}_A, \bar{M}_B) as the *maximal experimentation strategy*, reflecting the fact that these fees mark the largest deviation from the myopic optimum.

Besides its important role in delineating the possible range of experimentation, the maximal experimentation strategy has the great advantage of being algebraically computable. In the discounted problem, computing the maximum in the Bellman equation (10) yields a second-order ordinary differential equation for v that generally is cubic in v'' ; see expression (A.7) in Appendix A.3. Alternatively, equation (11) expresses v'' directly as a non-linear function of p and v . Irrespective of the representation one uses, the differential equation has no explicit solution, so one must rely on numerical techniques to compute the value function and trace out the optimal fees as functions of p alone. By contrast, the optimal policy in the undiscounted limit can be computed as the solution to the pointwise optimization problem in (13) which does not involve the value function. Appendix A.4 shows how to reduce this problem to solving a quadratic equation in one variable. The coefficients of this equation are too unwieldy, however, to allow for analytic results.

²⁰As the pair of fees $(M_A, M_B) = (v_A, v_B)$ is clearly suboptimal (yielding zero revenue and zero information), the function S is indeed positive on the relevant domain.

7 Pricing Implications

Our next aim is to provide results on optimal price experimentation which do not require the shadow price of information to be small in the sense of the local analysis performed in the previous section. As these results depend crucially on the precise nature of the externalities which the two sides of the market exert on each other, we are lead to consider a number of different scenarios, starting with the most tractable ones.

7.1 Symmetric Externalities

When the externality that side A exerts on side B is exactly as strong as the converse externality in either state of the world, then the myopic pricing strategy is optimal. In fact, for $(e_A^0, e_A^1) = (e_B^0, e_B^1)$, we have $M_A^\mu \equiv v_A/2$ and $M_B^\mu \equiv v_B/2$ by (7)–(8), so information about the true state of the world is clearly worthless and there is no incentive to deviate from the myopic optimum.

Proposition 1 *For $(e_B^1, e_B^0) = (e_A^1, e_A^0)$, the platform provider always sets the myopically optimal fees $(v_A/2, v_B/2)$.*

The case of symmetric externalities is actually the only one in which the myopically optimal revenue R^μ is linear in beliefs. In all other cases, this function is strictly convex, and so is the value function v , implying a positive shadow price of information.

7.2 One-Sided Externalities

The second most tractable scenario in our model is the one where $e_B^1 = e_B^0 = 0$, so side B does not benefit from an increase in participation on side A . Participation on side B is then independent of the fee M_A , and the platform provider can only increase the information content of observed quantities by varying M_B . The direction of optimal experimentation is obvious in this case: the only way to render observed participation on side A more informative about the externality parameter e_A^θ is to raise participation on side B , i.e., to reduce M_B .

Proposition 2 *For $e_B^1 = e_B^0 = 0$, the platform provider always sets a fee lower than the myopic optimum on side B , and a fee higher than the myopic optimum on side A .*

The intuition behind this result is straightforward. Lowering the fee on side B not only makes the participation observed on side A more informative, but also gives the participants on that side a larger surplus. As raising the fee on side A does not affect participation on side B , the provider can safely extract part of the extra surplus given to side A by charging it a fee above the myopic optimum.²¹

The maximal experimentation policy takes a particularly simple form in this case. The minimum of $[\bar{R}(p) - R(M_A, M_B, p)]/(v_B - M_B)^2$ is attained at

$$\begin{aligned}\bar{M}_A(p) &= \frac{v_A v_B + 2e_A(p)\bar{R}(p)}{e_A(p)v_A + 2v_B}, \\ \bar{M}_B(p) &= v_B + \frac{v_A^2 - 4\bar{R}(p)}{e_A(p)v_A + 2v_B}\end{aligned}$$

²¹To make the connection with our earlier analysis for small shadow prices of information, note that $e_B^1 = e_B^0 = 0$ implies $S_A^\mu = 0$, so iso-information curves are horizontal lines and we are in the case depicted in Figure 1(c).

where $e_A(p) = pe_A^1 + (1-p)e_A^0$. Comparing these fees to the myopically optimal ones, we first see that

$$\overline{M}_B(p) - M_B^\mu(p) = \frac{4[R^\mu(p) - \overline{R}(p)]}{e_A(p)v_A + 2v_B}.$$

As $\overline{R}(p) = pR^\mu(1) + (1-p)R^\mu(0)$ and R^μ is strictly convex, the right-hand side is negative for $0 < p < 1$. Thus, in line with Proposition 2, the maximal experimentation policy will indeed decrease the fee that generates information. On the other side of the market, we find

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

so for non-degenerate beliefs, there is a price increase relative to the myopic benchmark, which is again in line with Proposition 2.

7.3 One-Sided Uncertainty

Our next step is to analyze the scenario where the externality exerted by side A is perfectly known, but positive.

Proposition 3 *Suppose $e_B^1 = e_B^0 > 0$. Relative to the myopic optimum, the platform provider then always lowers the fee on side B , and raises the fee on side A if and only if $\ell_A(p) > \ell_B(p)$.*

The intuition for this result is closely related to that for Proposition 2. When the externality exerted by side A is known, the platform provider can only increase the amount of information by lowering the fee on side B . Side A then benefits from higher participation on side B . When $\ell_A > \ell_B$, the price effect that side A is expected to have on side B is weaker than the price effect in the other direction, and the platform provider can again extract part of the additional surplus given to side A by charging this side a higher fee.

Figure 2 illustrates the two situations that can arise for a known externality parameter e_B . Iso-information curves are parallel straight lines with slope $-e_B$. In Figure 2(a), the locus of tangency points between iso-revenue curves and iso-information lines is an upward sloping line – the optimal trade-off between information and current revenue induces a decrease in both fees. In Figure 2(b), the locus of tangency points is a downward sloping line; here, the trade-off between information and current revenue leads to a decrease in M_B but an increase in M_A .²²

Proposition 3 implies in particular that for a known externality parameter $e_B < e_A^0$, the platform provider always sets a fee above the myopic optimum on side A , exactly as in the case where $e_B = 0$. For $e_B > e_A^1$, it lowers both fees relative to the myopic benchmark. For $e_A^0 < e_B < e_A^1$, finally, it sets $M_A^*(p) > M_A^\mu(p)$ for p above some threshold \hat{p} .

²²It is again straightforward to relate these findings with our earlier analysis for small shadow prices of information. For a known externality parameter e_B , we have $\ell_B = e_B \ell_0$. When $\ell_A > \ell_B$, therefore, $-\ell_{AB}/\ell_0 < -\ell_B/\ell_0 = -e_B = -S_A^\mu/S_B^\mu$ and we are in the case depicted in Figure 1(c). When $\ell_A < \ell_B$, then $-\ell_0/\ell_{AB} < -e_B = -S_A^\mu/S_B^\mu < -\ell_{AB}/\ell_0$ and we are in the case depicted in Figure 1(b).

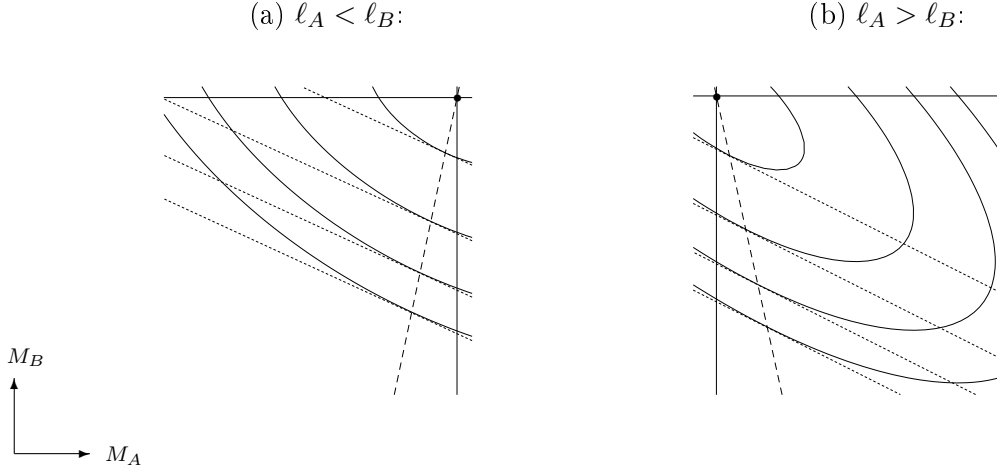


Figure 2: Directions of experimentation when $e_B^1 = e_B^0 > 0$. The point marked with a dot indicates the myopically optimal fees (M_A^μ, M_B^μ) ; the solid curves are the iso-revenue curves corresponding to 99%, 95%, 90% and 85% of the maximal expected revenue R^μ ; the dotted lines are the iso-information lines tangential to these iso-revenue curves; the dashed line is the locus of all tangency points between iso-revenue curves and iso-information lines. Parameter values: $v_A = v_B = \sigma_A = \sigma_B = 1$, $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = e_B^1 = 0.6$; (a) $p = 0.2$; (b) $p = 0.8$.

7.4 Ordered Price Effects

We now turn to scenarios with two-sided uncertainty, meaning that $e_B^1 \neq e_B^0$. In these scenarios, the platform provider can manipulate the information content of observed participation on both sides, and can do so by varying either fee.

To make progress on this problem, we first restrict the model parameters in a way that encompasses the three scenarios analyzed so far (symmetry, one-sided externalities and one-side uncertainty) as limiting cases. We say that price effects are *ordered* if the direct and indirect effects of lowering either fee are at least as strong in state $\theta = 1$ as in state $\theta = 0$; otherwise we say that price effects are *mixed*. As our standing assumptions imply $\ell_A^1 > \ell_A^0$, price effects are ordered if and only if $\ell_0^1 \geq \ell_0^0$ and $\ell_B^1 \geq \ell_B^0$. In particular, this is the case when the externality parameters are positively correlated across states, that is, when $e_B^1 > e_B^0$. For $e_B^1 < e_B^0$, price effects are ordered as long as e_B^1/e_B^0 remains sufficiently large (see Appendix A.2, which also establishes that $\ell_B^1 \geq \ell_B^0$ automatically implies $\ell_0^1 > \ell_0^0$).²³

Our first two results on ordered price effects concern situations which approximate symmetry and one-sided uncertainty, respectively.

Proposition 4 *For (e_B^1, e_B^0) sufficiently close to, but different from, (e_A^1, e_A^0) , the platform provider always sets both fees below their myopically optimal levels.*

²³Price effects are ordered in Figures 1(b) and 1(c); they are mixed in Figure 1(a) and in the two examples given in Footnote 19.

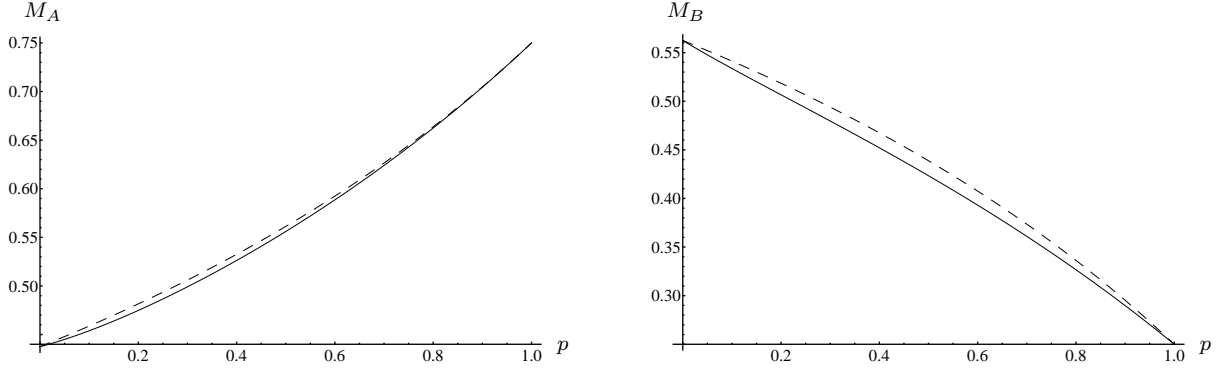


Figure 3: Fees (M_A, M_B) under the maximal experimentation strategy (solid curves) and in the myopic benchmark (dashed curves) for approximately symmetric externalities; $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = 0.3$, $e_B^1 = 0.7$, $v_A = v_B = \sigma_A = \sigma_B = 1$.

This result is illustrated in Figure 3. The intuition for it is clear. Under approximately symmetric externalities, the direction of experimentation must be the same on both sides, and charging less than the myopic optimum makes observed participation unambiguously more informative.

Proposition 5 *For $e_B^1 < e_A^1$ and e_B^0 sufficiently close to e_B^1 , the optimal fee on side A exceeds its myopic benchmark at beliefs close to 1, while the fee on side B is always below its myopically optimal level.*

This result is illustrated in Figure 4. The intuition here is essentially the same as for Proposition 3. If the externality exerted by side A is subject to moderate uncertainty only, the platform provider optimally experiments by lowering the fee on side B. And when the price effect that side A is expected to have on side B is weaker than the price effect in the other direction, the platform provider again extracts surplus from side A by charging it a higher fee. If $e_B^1 < e_A^0$, we have the stronger result that for e_B^0 sufficiently close to e_B^1 , the platform provider always charges side A more than M_A^μ .

Our next aim is to provide sufficient conditions for price experimentation to work in a particular direction, and to cover the ground ‘in between’ symmetry and one-sided uncertainty. Besides the expected price effects ℓ_0 , ℓ_A and ℓ_B , these conditions involve the constants

$$s_A = \frac{1}{2} \frac{\partial^2 S}{\partial M_A^2}, \quad s_B = \frac{1}{2} \frac{\partial^2 S}{\partial M_B^2} \quad \text{and} \quad s_{AB} = \frac{1}{2} \frac{\partial^2 S}{\partial M_A \partial M_B},$$

which are all positive when price effects are ordered (see Appendix A.2, where these constants are computed in terms of the primitive parameters of the model). Strict convexity of S further means that $s_A s_B - s_{AB}^2 > 0$.

Proposition 6 *Suppose that price effects are ordered. Consider a belief p for which both myopically optimal fees are lower than the respective intrinsic values. Then, the platform*

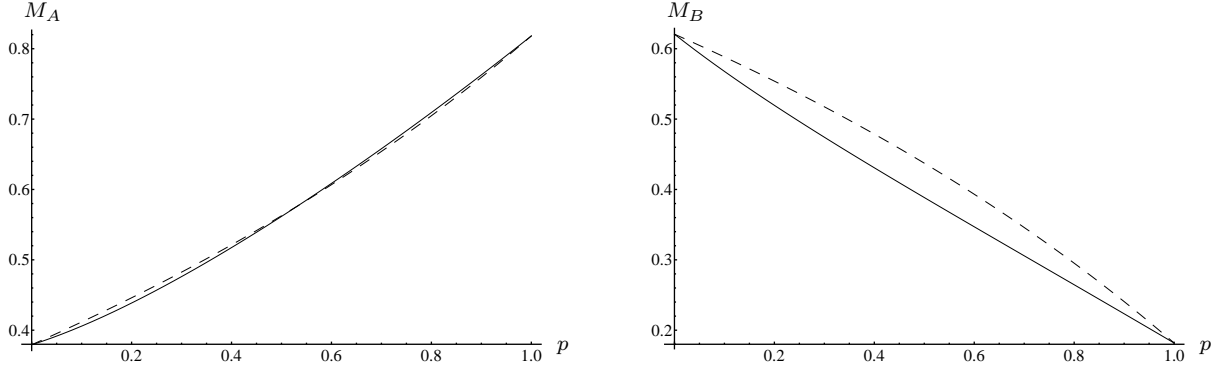


Figure 4: Fees (M_A, M_B) under the maximal experimentation strategy (solid curves) and in the myopic benchmark (dashed curves) for approximately one-sided uncertainty; $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = 0.45$, $e_B^1 = 0.55$, $v_A = v_B = \sigma_A = \sigma_B = 1$.

provider lowers the fee on side A relative to the myopically optimal level if

$$\frac{\ell_0(p)}{\ell_{AB}(p)} > \frac{s_B}{s_{AB}}, \quad (14)$$

and raises it if

$$\frac{\ell_0(p)}{\ell_{AB}(p)} < \frac{s_{AB}}{s_A}. \quad (15)$$

Similarly, the platform provider lowers the fee on side B relative to the myopically optimal level if

$$\frac{\ell_0(p)}{\ell_{AB}(p)} > \frac{s_A}{s_{AB}}, \quad (16)$$

and raises it if

$$\frac{\ell_0(p)}{\ell_{AB}(p)} < \frac{s_{AB}}{s_B}. \quad (17)$$

It is straightforward to give a geometric intuition for conditions (14)–(17) in terms of iso-revenue and iso-information curves in the (M_A, M_B) -plane. We have already seen that the slope of the iso-revenue curve through any point on the vertical line segment $\mathcal{L}_A^\mu = \{(M_A, M_B): M_A = M_A^\mu, M_B \leq v_B\}$ is $-\ell_{AB}/\ell_0$. A simple computation further shows that the slope of the iso-information curve through any point on the line segment $\mathcal{L}_A^v = \{(M_A, M_B): M_A = v_A, M_B \leq v_B\}$ is $-s_{AB}/s_B$; when $M_A^\mu < v_A$ and $M_B^\mu < v_B$, the elliptic shape of the iso-information curves implies that the slope of the iso-information curve through any point on \mathcal{L}_A^μ is strictly smaller than $-s_{AB}/s_B$. Under condition (14), this in turn is strictly smaller than $-\ell_{AB}/\ell_0$, so in each point on \mathcal{L}_A^μ the iso-information curve declines more steeply than the iso-revenue curve and, exactly like in Figures 1(a) and 1(b), the platform provider can raise both its expected revenue and the information content of observed participation by setting M_A below its myopically optimal level.

Similarly, the slope of the iso-information curve through any point on the line segment $\mathcal{L}_B^v = \{(M_A, M_B): M_A \leq v_A, M_B = v_B\}$ is $-s_A/s_{AB}$; when $M_A^\mu < v_A$ and $M_B^\mu < v_B$, the

elliptic shape of the iso-information curves implies that the slope of the iso-information curve through any point on \mathcal{L}_A^μ is strictly larger than $-s_A/s_{AB}$. Under condition (15), this in turn is strictly larger than $-\ell_{AB}/\ell_0$, so in each point on \mathcal{L}_A^μ the iso-information curve declines less steeply than the iso-revenue curve and, as in Figure 1(c), the platform provider can raise both its expected revenue and the information content of observed participation by setting M_A above its myopically optimal level.

In an analogous manner, conditions (16) and (17) translate into inequalities between the slopes of the iso-revenue and iso-information curves along the line segment $\mathcal{L}_B^\mu = \{(M_A, M_B): M_A \leq v_A, M_B = M_B^\mu\}$.

Note that conditions (15) and (17) – which jointly would imply fees above the myopic optimum on *both* sides – cannot hold at the same time. If they did, we would have

$$\left(\frac{\ell_0}{\ell_{AB}}\right)^2 < \frac{s_{AB}^2}{s_A s_B} < 1$$

– a contradiction to the fact, established in Section 4, that $0 < \ell_{AB} < \ell_0$.

The ratio ℓ_0/ℓ_{AB} is either decreasing in p or constant (see Appendix A.1). If condition (14) is met at a belief \hat{p} , therefore, it will be met at all $p < \hat{p}$. The reason for this is straightforward: a lower p implies a flatter iso-revenue curve and thus makes it ‘cheaper’ (in terms of expected revenue) to lower the fee on side A for informational purposes. By the same line of argument, (16) also becomes easier to satisfy as p decreases, while (15) and (17) become harder to satisfy. Thus, a sufficient condition for uniformly lower fees on both sides of the market is that (14) and (16) both hold at $p = 1$, and a sufficient condition for a uniformly higher fee on one side is that either (15) or (17) hold at $p = 0$.²⁴ As the left-hand sides of (14) and (16) exceed 1, the inequality $s_i \leq s_{AB}$ is also sufficient for a uniform fee decrease on side $i = A, B$.

One can easily formulate weaker sufficient conditions by exploiting the fact (established in Appendix A.2) that the ratios s_A/s_{AB} and s_{AB}/s_B are bounded below and above by terms involving only the price effects $(\ell_i^\theta)_{i=0,A,B;\theta=0,1}$. When $e_B^1 > e_B^0$, for example, the inequality $\ell_i^1 - \ell_i^0 \leq \ell_0^1 - \ell_0^0$ for $i = A$ or B already implies a uniform fee decrease on side i , irrespectively of the noise parameters σ_A and σ_B . In line with Proposition 5, this holds on side B , for instance, when e_B^1 is only somewhat higher than e_B^0 , so that the externality which side A exerts on side B is rather well known from the outset and there is much more to be learned from lowering the fee on side B .

Finally, Proposition 6 implies a fee above the myopic benchmark on side B at beliefs close to 1 when (i) $e_B^1 - e_B^0$ is very close to $e_A^1 - e_A^0$, (ii) e_B^1 exceeds e_A^1 by a sufficient amount, and (iii) σ_A/σ_B is sufficiently large. In such a situation, there is more to be learned from lowering the fee on side A . Given enough confidence in the state $\theta = 1$, this is expected to have a comparatively strong effect on the surplus given to side B and hence makes it optimal to charge this side a fee above the myopically optimal level.

7.5 Mixed Price Effects

The sufficient conditions in Proposition 6 also apply to scenarios with mixed price effects ($e_B^1 < e_B^0$ with $\ell_0^1 < \ell_0^0$ or $\ell_B^1 < \ell_B^0$) as long as s_{AB} remains positive. It is straightforward to formulate analogous conditions when this coefficient is negative (which can happen for $\ell_B^1 < \ell_B^0$); we therefore do not state them in a formal proposition.

²⁴In this discussion, we maintain the assumption of Proposition 6 that both myopically optimal fees are lower than the respective intrinsic values.

Instead, we turn to the very tractable special case of antisymmetric externalities ($e_B^1 = e_A^0$, $e_B^0 = e_A^1$), symmetric intrinsic values ($v_A = v_B$) and symmetric noise intensities ($\sigma_A = \sigma_B$). In this case, $s_A = s_B$ and $s_{AB} = 0$, so the iso-information curves are concentric circles. More importantly, the platform provider's pricing problem is symmetric with respect to the belief $1/2$: the value function and the shadow price of information satisfy $v(p) = v(1-p)$ and $V(p) = V(1-p)$ while the myopic and optimal pricing strategies satisfy $M_A^\mu(p) = M_B^\mu(1-p)$ and $M_A^*(p) = M_B^*(1-p)$.

Proposition 7 *Suppose that $(e_B^1, e_B^0, v_B, \sigma_B) = (e_A^0, e_A^1, v_A, \sigma_A)$. Then there exists a belief $\bar{p} > 1/2$ such that relative to the myopic optimum, the platform provider lowers the fee on side A for beliefs in the interval $]0, \bar{p}[$ and lowers the fee on side B for beliefs in the interval $]1 - \bar{p}, 1[$. The platform provider sets a fee above the myopically optimal level on side A for beliefs close to 1 (and on side B for beliefs close to 0) if and only if the higher of the two externality parameters exceeds $1/2$.*

When the platform provider is highly uncertain about which side exerts the stronger externality, that is, when the belief is close to $1/2$, both sides are charged less than in the myopic benchmark. The intuition for this finding is the same as the one suggested for approximately symmetric externalities (see Proposition 4): given that the fees charged to the two sides must be close to each other, the direction of experimentation must be the same on both sides, and charging less than the myopic optimum makes observed participation unambiguously more informative. For beliefs near a boundary of the unit interval, on the other hand, the platform provider is fairly confident in knowing the side that exerts the stronger externality on the other (if p is close to 1, say, this is side B), and learns most effectively by lowering the fee on this side; whether it is optimal to recoup some of the surplus this creates by raising the fee on the other side (side A for p close to 1) depends on the actual strength of the externality.

By the same continuity argument as in the proofs of Propositions 4 and 5, these findings carry over qualitatively to scenarios of approximate antisymmetry where $(e_B^1, e_B^0, v_B, \sigma_B)$ is close, but not identical, to $(e_A^0, e_A^1, v_A, \sigma_A)$. For the sake of brevity, we omit a formal statement and just present an example in Figure 5. The difference between the maximal experimentation fees and the myopic benchmark is considerably larger here than in Figures 3 and 4. This reflects a stronger incentive to experiment when price effects are mixed: in Figure 5, the differences $M_A^\mu(1) - M_A^\mu(0)$ and $M_B^\mu(1) - M_B^\mu(0)$ are almost twice as large in absolute value as in Figures 3 and 4, so knowing the true state of the world matters much more for optimal pricing. We shall return to this issue when we discuss how the optimal pricing strategy depends on the parameters of the model.

8 Further Findings

8.1 Implications for Quantities

If the platform provider reduces both fees below the myopically optimal level, expected participation on both sides increases irrespectively of the true state of the world. By Proposition 4, this is the case for approximately symmetric externalities.

In scenarios where the platform provider sets a fee above the myopically optimal level on one side, however, expected participation on this side may well decrease in some state of the world. This phenomenon is barely visible in the lower left panel of Figure 6, computed

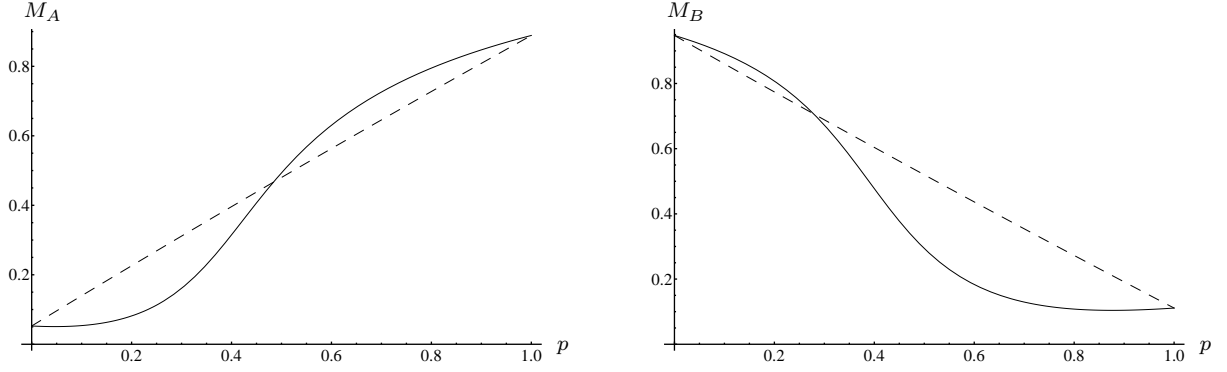


Figure 5: Fees (M_A, M_B) under the maximal experimentation strategy (solid curves) and in the myopic benchmark (dashed curves) for approximately antisymmetric externalities; $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = 0.95$, $e_B^1 = 0.2$, $v_A = v_B = \sigma_A = \sigma_B = 1$.

with the same model parameters as Figure 4 (approximately one-sided uncertainty). In state $\theta = 0$, the externality exerted on side A is considerably weaker than the externality exerted on side B , so when the platform provider raises the fee on side A at high beliefs p , the ensuing negative effect on the participation on side A is not fully compensated by the positive effect of a lower fee on side B . In state $\theta = 1$, by contrast, the externality exerted on side A is considerably stronger than the externality exerted on side B , and the positive effect on n_A of lowering M_B more than compensates the negative effect of raising M_A . In either state, finally, the comparatively large decrease in M_B raises expected participation on side B at all non-degenerate beliefs.

For approximately antisymmetric externalities, decreases in expected participation occur on either side of the market; see Figure 7, which has been computed with the same parameters as Figure 5. On side A , the situation is qualitatively the same as in the example of approximate one-sided uncertainty that we just discussed, and can be explained in exactly the same way; side B is essentially its mirror image with the roles of the two states reversed.

In summary, participation rises on a given side either when the externality exerted on this side is strong, or when this externality is weak and the platform provider believes it is weak; participation falls on a given side when the externality exerted on this side is weak but the platform provider believes it is strong and thus raises the fee on this side in an attempt to capture some of the surplus created by lowering the other fee.

8.2 Comparative Statics

Inspection of the Bellman equation (10) shows that a multiplication of the noise intensities σ_A and σ_B by a common factor γ has the same effect as a multiplication of r by γ^2 . By the results mentioned in Section 6, for $\gamma < 1$ we obtain a higher value v at all non-degenerate beliefs, and a higher shadow price of information V ; by a standard monotone comparative statics argument, this in turn implies a higher quantity of information S at the optimal fees and a lower expected current revenue R . How the optimal fees themselves change in

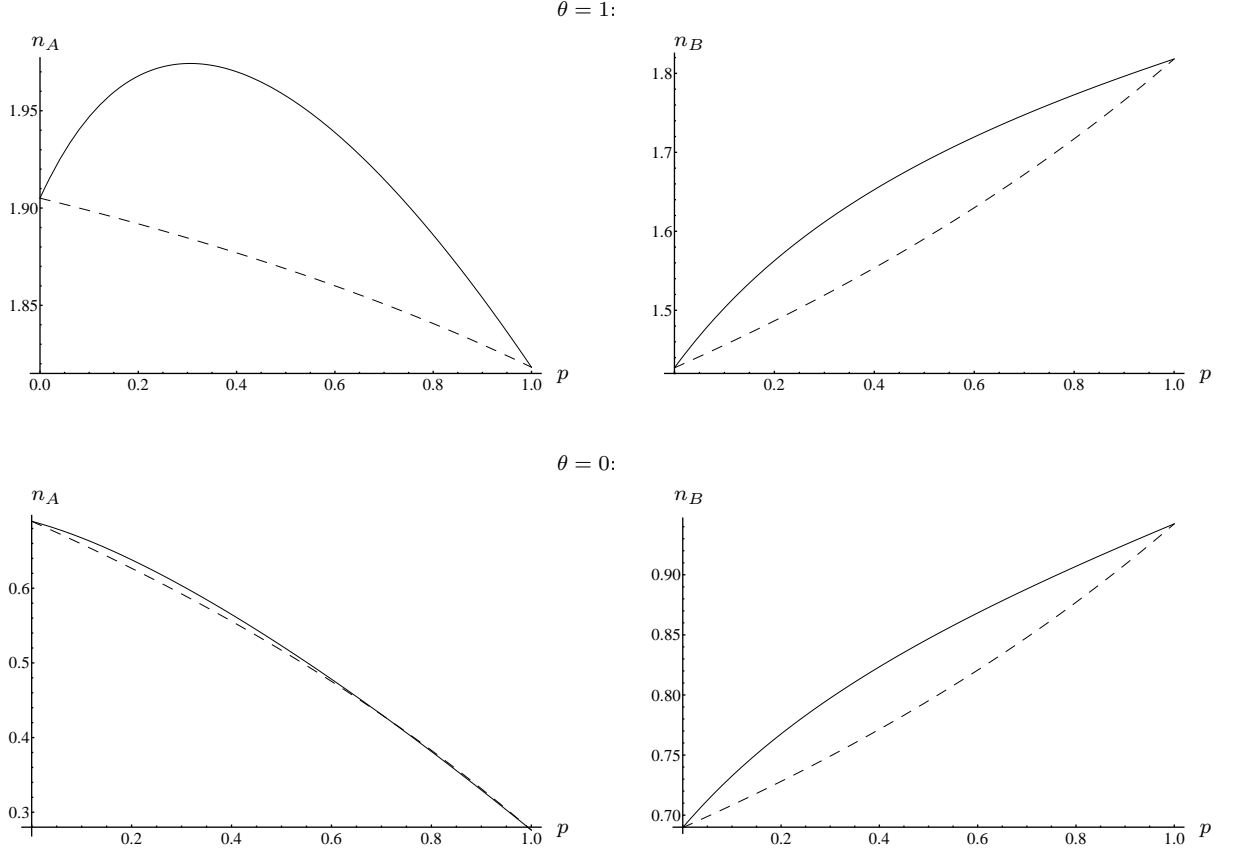


Figure 6: Expected participation $n_A(M_A, M_B, \theta)$ and $n_B(M_A, M_B, \theta)$ under the maximal experimentation strategy (solid curves) and in the myopic benchmark (dashed curves) for approximately one-sided uncertainty; $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = 0.45$, $e_B^1 = 0.55$, $v_A = v_B = \sigma_A = \sigma_B = 1$.

response to this parameter change has been explored in Section 7: whenever the results there imply an unambiguous direction of price experimentation for the fee on a given side, this fee will change further in that direction.

The picture becomes less clear when we change only one noise intensity. By the same comparison argument as in Keller and Rady (1997, Theorem 5.2 and Appendix E.1), a decrease in σ_A , say, again leads to an increase in the value function, but the effect on the shadow price of information can no longer be signed unambiguously because $\min_{M_A, M_B} (v - R)/S$ is increasing in both σ_A and v . In the undiscounted limit, the situation is simpler: the maximal shadow price of information \bar{V} clearly rises as σ_A or σ_B falls. As Figure 8 shows, such a parameter change can have a surprisingly strong effect on optimal fees and participation levels. In each panel, the “leftmost” curve is associated with $\sigma_B = 1$, and hence the same as in Figure 5 and Figure 7, respectively; as σ_B decreases to 0.5 and 0.33, the curves shift to the right and exhibit steeper slopes. A reduction in σ_B improves the signal-to-noise ratio in observed participation on side B and increases the marginal informational benefit of reducing the fee on side A . As a consequence, the platform provider lowers M_A at any non-degenerate belief. In turn, this increases the extra surplus given to side B , so M_B is raised by successively more and over

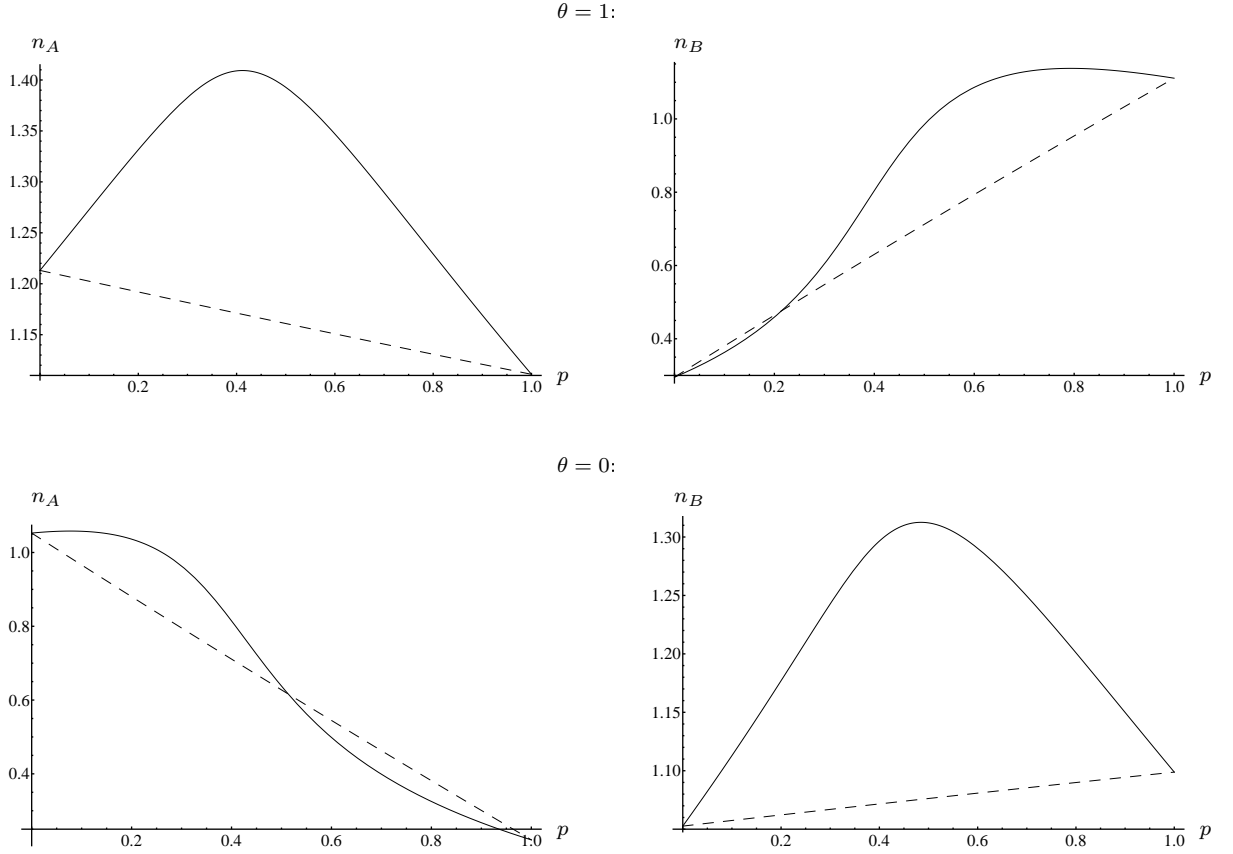


Figure 7: Expected participation $n_A(M_A, M_B, \theta)$ and $n_B(M_A, M_B, \theta)$ under the maximal experimentation strategy (solid curves) and in the myopic benchmark (dashed curves) for approximately antisymmetric externalities; $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = 0.95$, $e_B^1 = 0.2$, $v_A = v_B = \sigma_A = \sigma_B = 1$.

a successively larger range of beliefs. The concomitant changes in expected participation are in line with our discussion of Figure 7.

When changing an intrinsic value or an externality parameter, another difficulty arises. Now the above comparison argument no longer applies because the boundary conditions change at the same time as the differential equation; put differently, such a parameter change affects both the myopic benchmark and the optimal deviation from it. Still, the propositions of Section 7 and numerical examples such as those in Figures 4 and 5 allow us to gain some intuitive insights into the link between the *incentives* for experimentation on the one hand, and the actual *extent* of experimentation on the other.

Given intrinsic values (v_A, v_B) , noise intensities (σ_A, σ_B) and externality parameters (e_A^1, e_A^0) , for instance, let us say that the incentives for experimentation are higher for one pair of parameters (e_B^1, e_B^0) than for another if both the differences $M_A^\mu(1) - M_A^\mu(0)$ and $M_B^\mu(1) - M_B^\mu(0)$ are larger in absolute value for this pair than for the other. Similarly, let us say that there is more experimentation for one pair of parameters (e_B^1, e_B^0) than for another if both $\max_p |M_A^*(p) - M_A^\mu(p)|$ and $\max_p |M_B^*(p) - M_B^\mu(p)|$ are larger for this pair than for the other.

It is a plausible conjecture that higher incentives for experimentation imply more

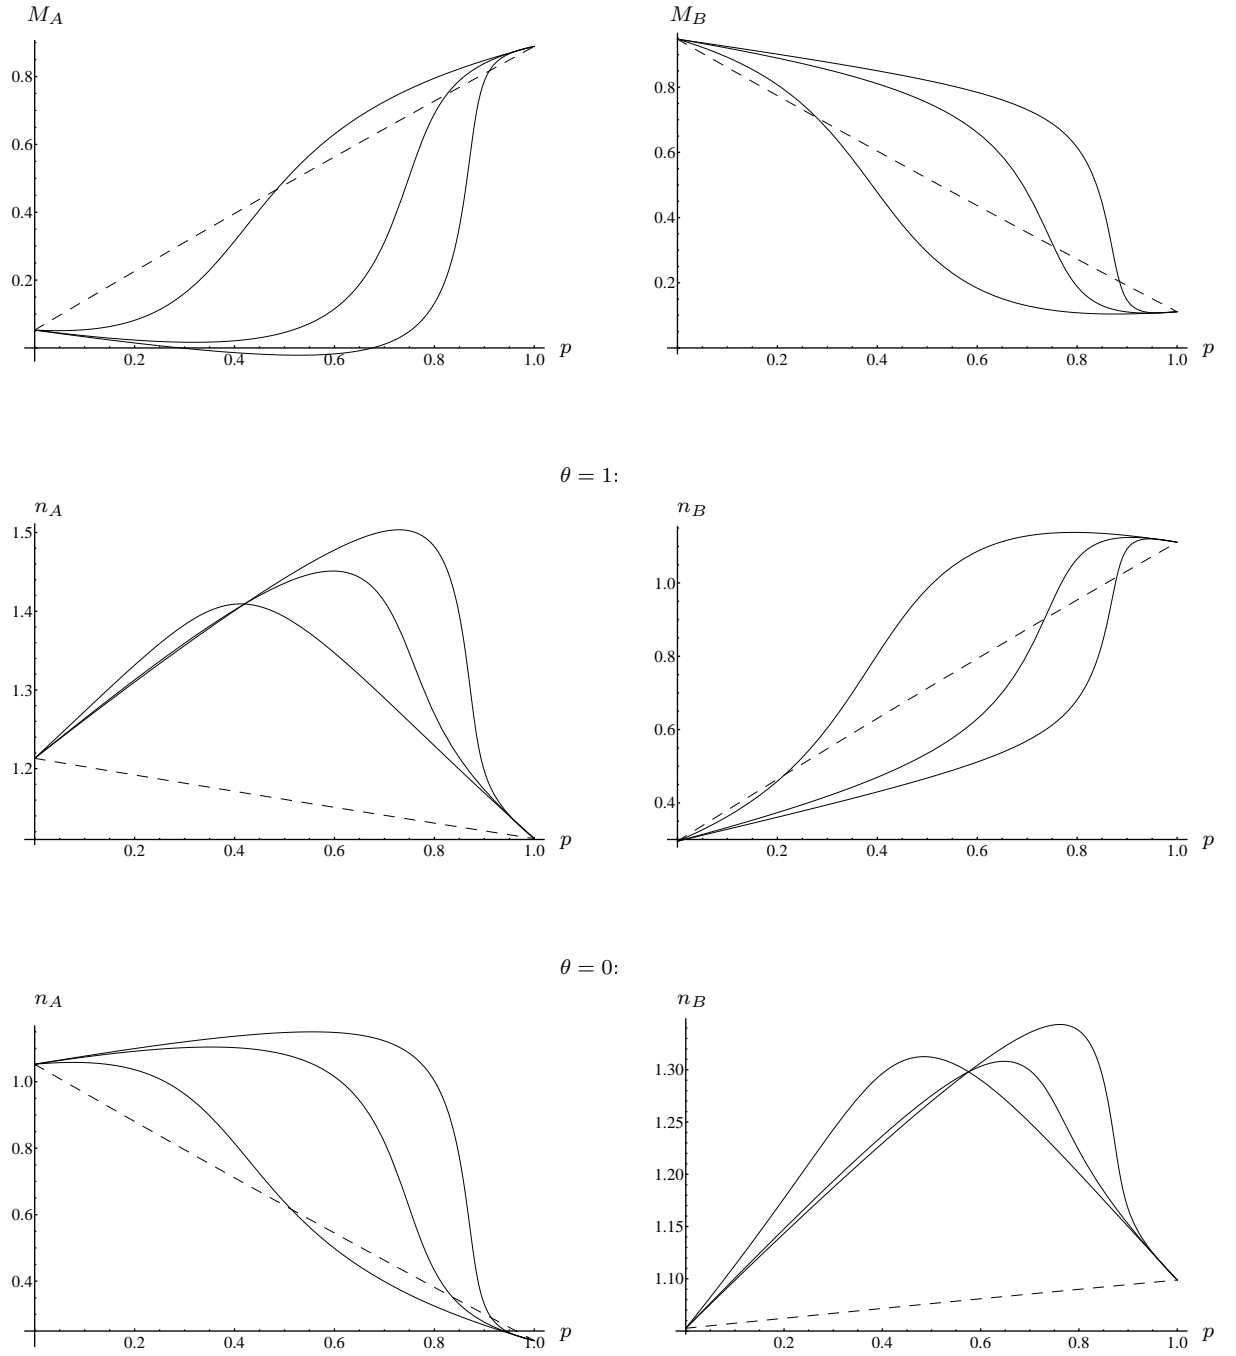


Figure 8: Fees (M_A, M_B) as well as expected participation $n_A(M_A, M_B, \theta)$ and $n_B(M_A, M_B, \theta)$ under the maximal experimentation strategy (solid curves) and in the myopic benchmark (dashed curves) for approximately antisymmetric externalities and different noise intensities σ_B ; $e_A^0 = 0.1$, $e_A^1 = 0.9$, $e_B^0 = 0.95$, $e_B^1 = 0.2$, $v_A = v_B = \sigma_A = 1$, $\sigma_B \in \{1, 0.5, 0.33\}$.

experimentation at the optimum.²⁵ This is trivially true if one of the two situations to be compared involves symmetric externalities $(e_B^1, e_B^0) = (e_A^1, e_A^0)$: these provide no incentives for experimentation at all, and there is indeed no experimentation in this case by Proposition 1; any other pair (e_B^1, e_B^0) is associated with higher incentives and actually features some experimentation. The conjecture is also borne out by Figures 3–5: the example of approximate symmetry in Figure 3 is associated with lower incentives for experimentation than the example of approximately one-sided uncertainty in Figure 4, which in turn is associated with lower incentives than the example of approximately antisymmetric externalities in Figure 5; the extent of experimentation duly increases as we move from one figure to the next.²⁶

This does not mean, however, that the quantity of information generated by the platform provider increases in line with the incentives to experiment. Our examples suggest the contrary, in fact. Evaluated along the maximal experimentation policy, S decreases uniformly as we move from approximate symmetry to approximately one-sided uncertainty and approximate asymmetry: we find $6.46 \leq S(\bar{M}_A, \bar{M}_B) \leq 7.41$ for the parameters underlying Figure 3, $2.02 \leq S(\bar{M}_A, \bar{M}_B) \leq 3.15$ for Figure 4 and $0.59 \leq S(\bar{M}_A, \bar{M}_B) \leq 0.80$ for Figure 5. The reason for this finding is that the quadratic function S becomes uniformly smaller (and less convex) in each step, which reflects a worsening signal-to-noise ratio.²⁷

8.3 Dynamic Implications

Relative to the platform provider's information filtration, the process of beliefs $(p_t)_{t \geq 0}$ is a diffusion without drift (Lemma 1). For any twice differentiable function f on the unit interval, the transformed process $(f(p_t))_{t \geq 0}$ is again a diffusion, and its drift coefficient is of the same sign as $f''(p_t)$ by Ito's Lemma. Relative to the information filtration of an outsider who knows the true state of the world, the process of beliefs has an upward drift if $\theta = 1$, and a downward drift if $\theta = 0$ (see the proof of Lemma 1); by Ito's Lemma, the drift of the transformed process with respect to this filtration depends on both $f'(p_t)$ and $f''(p_t)$. In view of the complicated structure of the optimal fees and the implied participation levels, it seems exceedingly difficult to determine the signs of their first and second derivatives with respect to the belief in general, even in the undiscounted limit. In fact, the slope and curvature of the myopically optimal fees already depend in a rather complicated fashion on the intrinsic values and externality parameter. In Figure 3, for example, M_A^μ and \bar{M}_A are strictly increasing and strictly convex functions of p , while M_B^μ and \bar{M}_B are strictly decreasing and strictly concave; changing (e_B^0, e_B^1) from $(0.3, 0.7)$ to $(0.95, 0.2)$ reverses the signs of all these slopes and curvatures. Figure 8, on the other hand, shows that the slope and curvature of an optimal fee or implied participation level can vary from one belief to another.

Nevertheless, a few robust statements can be made about the dynamics of beliefs, optimal fees and resulting participation levels. Following Bergemann and Välimäki (1997),

²⁵To our knowledge, no conjecture of this type has been proved yet, not even for simpler situations such as a standard monopoly pricing problem with just one decision variable and one outcome process. In the present model, such a proof seems completely out of reach.

²⁶For $v_A = v_B = \sigma_A = \sigma_B = 1$ and $(e_A^1, e_A^0) = (0.9, 0.1)$, actually, one-sided uncertainty with $e_B^1 = e_B^0 = 0.7$ is associated with higher incentives for experimentation than any other scenario with $e_B^1 \geq e_B^0$, and negatively correlated externalities with $e_B^1 = 0.2$ and $e_B^0 \rightarrow 1$ are associated with the highest incentives overall.

²⁷See Appendix A.2 for the formulas necessary to substantiate this claim.

we compute an estimate \hat{p}_t of the posterior belief p_t in a given state of the world by numerically solving the belief dynamics under the maximal experimentation strategy without the stochastic component, focusing on *expected* infinitesimal changes only. We then evaluate the maximal experimentation strategy and the induced participation levels along the path of estimated beliefs. We do so for the state $\theta = 1$ in which, as shown in the proof of Lemma 1, the belief estimate evolves according to $d\hat{p} = \hat{p}(1 - \hat{p})^2 S(\bar{M}_A(\hat{p}), \bar{M}_B(\hat{p})) dt$. We take $p_0 = 0.2$ as the platform provider's prior belief and hence the initial condition for \hat{p} .

The path of \hat{p}_t exhibits the S-shape familiar from Bergemann and Välimäki (1997), first rising slowly, then rapidly, then slowly again as it converges monotonically to 1. Convergence is faster under the maximal experimentation strategy than under the myopic policy. The comparison of quantities of information at the end of Section 8.2 further implies that convergence is fastest in the example of approximately symmetric externalities, slower in the example of approximately one-sided uncertainty, and slowest in the example of approximately antisymmetric externalities. This is reflected in the different time scales used in Figure 9 which shows the induced trajectories of fees and participation levels.²⁸

A couple of dynamic pricing implications emerge from this figure. Note that a myopic platform has an on average increasing price path on one side of the market and a decreasing one on the other. If the platform is forward-looking, fees converge more quickly to their full information limits. We also know from Section 7 that for any given belief about the state of the world, either one or both fees are below their myopically optimal levels. This gives rise to price dynamics that fall into one of two regimes; which regime prevails depends on the configuration of externality parameters and initial beliefs.

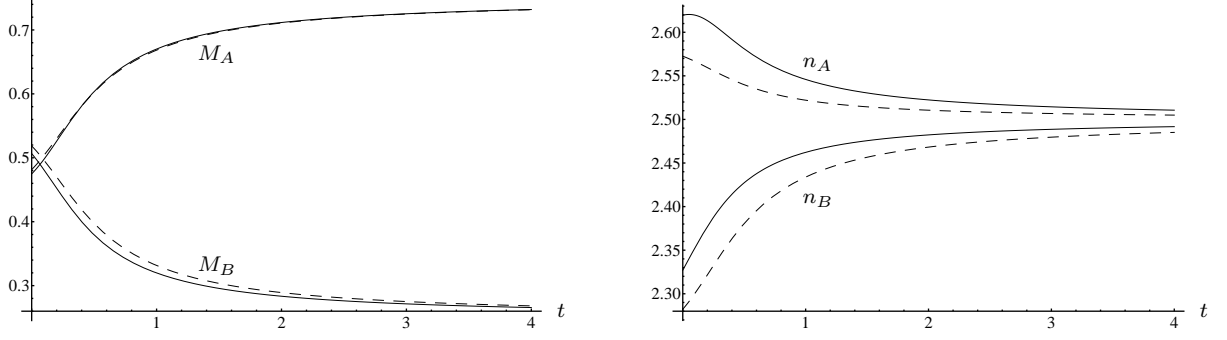
In what we call the *two-sided experimentation regime*, the forward-looking platform initially sets lower fees on both sides of the market than a myopic platform would; see the illustrations with approximate symmetry and approximately one-sided uncertainty in Figure 9. Then any price path that is on average increasing under myopic pricing is initially steeper under forward-looking pricing. In the example of approximately one-sided uncertainty, moreover, a fully informed platform “subsidizes” side B and “monetizes” on side A because side A benefits more strongly from participation on side B . With uncertainty about the state of the world, this subsidization is larger, and the surplus extraction starts earlier, if the platform is forward-looking rather than myopic.

In the other regime, only one fee is below the myopic level at the initial belief. We call this the *experimentation and exploitation regime*; see the illustration with approximate antisymmetry in Figure 9. Since in the illustration we start with a prior belief p_0 close to zero, over time the platform changes its assessment as to which side exerts the stronger externality. Therefore, fees are initially lower on side A and end up being larger; this holds both for the myopic and the forward-looking platform. However, this inversion of the price structure occurs at considerably higher speed if the platform is forward-looking.

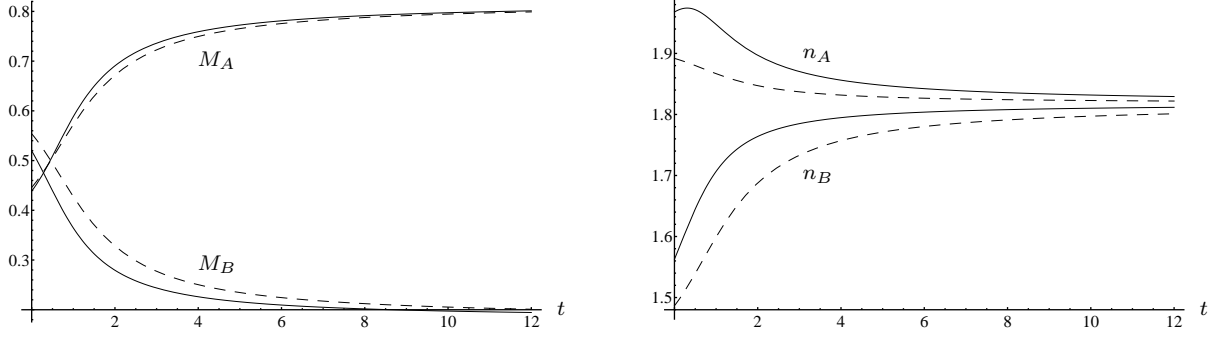
To summarize, while in the two-sided experimentation regime consumers on both sides initially are charged lower fees if the platform is forward-looking rather than myopic, in the experimentation and exploitation regime one side initially faces a higher fee. In either regime, an initially steeper and increasing price path amounts to introductory pricing

²⁸The dashed curves in this figure are generated by solving $d\hat{p} = \hat{p}(1 - \hat{p})^2 S(M_A^\mu(\hat{p}), M_B^\mu(\hat{p})) dt$ and evaluating the myopic policy and corresponding participation levels along this trajectory. That quantities on side A and B converge to a common limit in each case is due to the easily verified fact that $n_A(M_A^\mu(1), M_B^\mu(1), 1) = n_B(M_A^\mu(1), M_B^\mu(1), 1)$.

Approximate Symmetry ($e_B^0 = 0.3, e_B^1 = 0.7$):



Approximately One-Sided Uncertainty ($e_B^0 = 0.45, e_B^1 = 0.55$):



Approximate Antisymmetry ($e_B^0 = 0.95, e_B^1 = 0.2$):

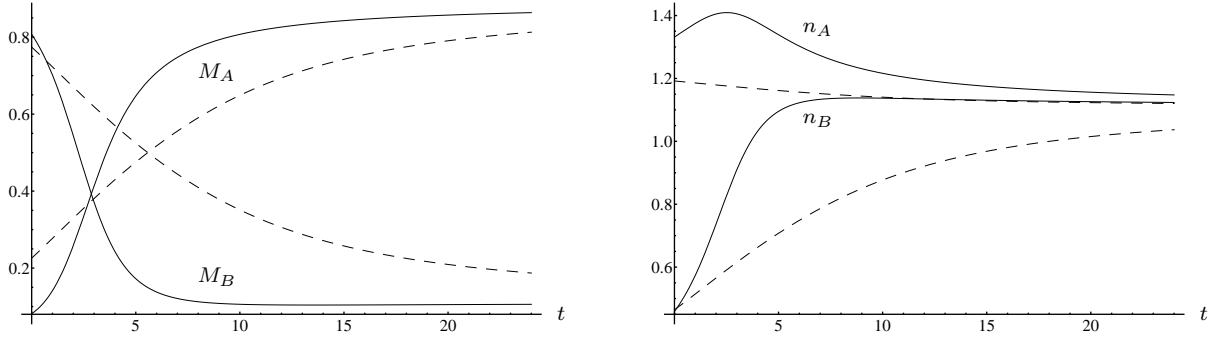


Figure 9: Estimated trajectories of the fees (M_A, M_B) and the expected participation levels $n_A(M_A, M_B, 1)$ and $n_B(M_A, M_B, 1)$ for $\theta = 1$ under the maximal experimentation strategy (solid curves) and in the myopic benchmark (dashed curves); $e_A^0 = 0.1, e_A^1 = 0.9, v_A = v_B = \sigma_A = \sigma_B = 1$.

with larger initial discounts. In all three examples, despite higher prices on one side for some time, participation is larger on both sides if the platform is forward-looking instead of myopic, whenever this difference in participation is pronounced. This tends to benefit participants.

We further note that experimentation can cause the trend in expected participation to be non-monotonic over time. As larger quantities facilitate learning, the forward-looking platform induces particularly high participation levels when the uncertainty about the environment is highest, that is, once beliefs have moved towards the centre of the unit interval; it lets participation levels decline again when it has become fairly confident of the true state of the world.

Performing the same simulations for $\theta = 0$ (and an initial belief $p_0 = 0.8$, say) generates exactly the same insights. If we consider the example in Figure 8 with $\sigma_B = 0.33$, however, we find an interesting difference across the two states of the world. A platform provider starting with the prior belief $p_0 = 0.5$ expects M_A to rise both in the short and in the long term. If the true state is $\theta = 0$, this price increase will be very gentle in expectation. If the true state is $\theta = 1$, by contrast, M_A is expected to rise very quickly when the posterior belief is about 0.85; the associated changes in M_B , n_A and n_B are equally drastic. Experimentation can thus make the variability of prices and quantities along their estimated paths markedly higher in one state than in the other.

9 Conclusion

We have studied a monopolistic platform provider in a two-sided market who is uncertain about the strength of interaction between the two sides. Maximizing expected lifetime profits, the platform provider faces the basic trade-off between the conflicting aims of maximizing current payoff and maximizing the information content of the market outcome. How this trade-off is resolved depends crucially on the precise configuration of cross-group externalities. In particular, the platform provider may raise the fee on one side of the market so as to extract some of the surplus created by the experimentation-induced fee reduction on the other side.

Our model was set in continuous time and allowed the platform provider to adjust the price structure at any moment. In reality, for contractual or other reasons, platforms tend to change prices rather infrequently. We conjecture that the static and dynamic implications of optimal experimentation would be the same in a setting that incorporated such restrictions.

Our analysis concerns a 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 cross-group externalities, 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 strength of the externalities 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. 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 within-group and cross-group externalities. Thus our analysis can possibly be extended to capture experimentation in markets with complementary goods.

Appendix

A.1 Auxiliary Results for the Expected Current Revenue R

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

$$R = R^\mu - \ell_0 [M_A - M_A^\mu]^2 - 2\ell_{AB} [M_A - M_A^\mu] [M_B - M_B^\mu] - \ell_0 [M_B - M_B^\mu]^2. \quad (\text{A.1})$$

A straightforward application of the implicit-function theorem now shows that the slope of the iso-revenue curve at the fee combination (M_A, M_B) is

$$\left. \frac{dM_B}{dM_A} \right|_{R=\text{const.}} = - \frac{\ell_0 [M_A - M_A^\mu] + \ell_{AB} [M_B - M_B^\mu]}{\ell_{AB} [M_A - M_A^\mu] + \ell_0 [M_B - M_B^\mu]}.$$

In particular, we see that

$$\left. \frac{dM_B}{dM_A} \right|_{R=\text{const.}} = - \frac{\ell_0}{\ell_{AB}} \quad \text{for } M_B = M_B^\mu \text{ and } M_A \neq M_A^\mu,$$

and

$$\left. \frac{dM_B}{dM_A} \right|_{R=\text{const.}} = - \frac{\ell_{AB}}{\ell_0} \quad \text{for } M_A = M_A^\mu \text{ and } M_B \neq M_B^\mu.$$

Finally, a simple computation shows that the derivative of the ratio ℓ_0/ℓ_{AB} with respect to p has the same sign as $-(e_A^1 - e_A^0) - (e_B^1 - e_B^0)$. So this ratio is either decreasing in p or constant.

A.2 Auxiliary Results for the Quantity of Information S

We first note that

$$\begin{aligned} n_A(M_A, M_B, 1) - n_A(M_A, M_B, 0) &= d_0 [v_A - M_A] + d_A [v_B - M_B], \\ n_B(M_A, M_B, 1) - n_B(M_A, M_B, 0) &= d_0 [v_B - M_B] + d_B [v_A - M_A], \end{aligned}$$

where

$$d_i = \ell_i^1 - \ell_i^0 \quad (i = 0, A, B)$$

measures the *difference* in direct and indirect price effects between states $\theta = 1$ and 0. Thus, the quotient d_0/σ_A captures the marginal change in the signal-to-noise ratio on side A when the fee M_A is lowered, and d_A/σ_A the marginal change when M_B is lowered. In the same way, d_0/σ_B and d_B/σ_B capture the marginal effects on the signal-to-noise ratio on side B of lowering M_B and M_A , respectively. Our assumptions on the externality parameters imply that d_A is always positive; d_0 is positive if and only if $e_B^1 > (e_A^0/e_A^1)e_B^0$, and d_B is positive if and only if $e_B^1 > e_B^0/[1 + (e_A^1 - e_A^0)e_B^0]$. As $e_A^0/e_A^1 < 1/[1 + (e_A^1 - e_A^0)e_B^0] < 1$, there are three possibly scenarios: either $d_0 > 0$ and $d_B \geq 0$ (in particular, this is the case when $e_B^1 \geq e_B^0$); or $d_0 > 0$ and $d_B < 0$; or $d_0 \leq 0$ and $d_B < 0$. Finally, it is straightforward to check that

$$d_A d_B - d_0^2 = \frac{(e_A^1 - e_A^0)(e_B^1 - e_B^0)}{(1 - e_A^0 e_B^0)(1 - e_A^1 e_B^1)}.$$

Next, we compute

$$S(M_A, M_B) = s_A [M_A - v_A]^2 + 2s_{AB} [M_A - v_A] [M_B - v_B] + s_B [M_B - v_B]^2$$

with the constants

$$s_A = \frac{d_0^2}{\sigma_A^2} + \frac{d_B^2}{\sigma_B^2} = \frac{1}{2} \frac{\partial^2 S}{\partial M_A^2}, \quad s_B = \frac{d_0^2}{\sigma_B^2} + \frac{d_A^2}{\sigma_A^2} = \frac{1}{2} \frac{\partial^2 S}{\partial M_B^2}, \quad s_{AB} = \frac{d_0 d_A}{\sigma_A^2} + \frac{d_0 d_B}{\sigma_B^2} = \frac{1}{2} \frac{\partial^2 S}{\partial M_A \partial M_B}.$$

The coefficient s_A is a measure of how fast the marginal informational gain from lowering the fee M_A increases as M_A falls; it has two components, the first pertaining to demand observations on side A , the second to demand observations on side B . The same structure is evident in the coefficients s_B and s_{AB} . Whereas s_A and s_B are clearly positive, s_{AB} can vanish or become negative when $d_B < 0$. Simple computations reveal that

$$s_A + s_B - 2s_{AB} = \left(\frac{d_0 - d_A}{\sigma_A} \right)^2 + \left(\frac{d_0 - d_B}{\sigma_B} \right)^2$$

and

$$s_A s_B - s_{AB}^2 = \left(\frac{d_A d_B - d_0^2}{\sigma_A \sigma_B} \right)^2.$$

If $e_B^1 \neq e_B^0$, then $s_A s_B - s_{AB}^2 > 0$ and the function S is strictly convex. If $e_B^1 = e_B^0 = e_B$, then the iso-information curves are parallel straight lines of slope $-s_A/s_{AB} = -s_{AB}/s_B = -e_B$.

Finally, a straightforward computation shows that for $s_{AB} \neq 0$, the derivative of s_A/s_{AB} with respect to σ_A^2/σ_B^2 has the same sign as $d_0 d_B (d_A d_B - d_0^2)$, and the corresponding derivative of s_{AB}/s_B the same sign as $d_0 d_A (d_A d_B - d_0^2)$. Letting σ_A^2/σ_B^2 tend to 0 and ∞ , respectively, we see that for $d_0 \neq 0$, the ratios s_A/s_{AB} and s_{AB}/s_B are bounded below by $\min\{d_0/d_A, d_B/d_0\}$ and bounded above by $\max\{d_0/d_A, d_B/d_0\}$.

A.3 The Optimal Fees

Suppressing the dependence on p , we write $S^\mu = S(M_A^\mu, M_B^\mu)$ for the quantity of information at the myopically optimal fees and

$$\begin{aligned} S_A^\mu &= \frac{\partial S}{\partial M_A}(M_A^\mu, M_B^\mu) = 2s_A(M_A^\mu - v_A) + 2s_{AB}(M_B^\mu - v_B), \\ S_B^\mu &= \frac{\partial S}{\partial M_B}(M_A^\mu, M_B^\mu) = 2s_{AB}(M_A^\mu - v_A) + 2s_B(M_B^\mu - v_B) \end{aligned}$$

for the corresponding partial derivatives; as $(M_A^\mu, M_B^\mu) \neq (v_A, v_B)$, at least one of these derivatives is different from 0. The quantity of information can now be rewritten as

$$\begin{aligned} S(M_A, M_B) &= S^\mu + S_A^\mu [M_A - M_A^\mu] + S_B^\mu [M_B - M_B^\mu] \\ &\quad + s_A [M_A - M_A^\mu]^2 + 2s_{AB} [M_A - M_A^\mu][M_B - M_B^\mu] + s_B [M_B - M_B^\mu]^2 \end{aligned} \quad (\text{A.2})$$

Combined with equation (A.1), this yields the following representation for the maximand $R+VS$ in the Bellman equation (10):

$$\begin{aligned} R^\mu + VS^\mu + VS_A^\mu [M_A - M_A^\mu] + VS_B^\mu [M_B - M_B^\mu] - (\ell_0 - s_A V) [M_A - M_A^\mu]^2 \\ - 2(\ell_{AB} - s_{AB} V) [M_A - M_A^\mu][M_B - M_B^\mu] - (\ell_0 - s_B V) [M_B - M_B^\mu]^2. \end{aligned}$$

The first-order conditions are

$$2(\ell_0 - s_A V) [M_A - M_A^\mu] + 2(\ell_{AB} - s_{AB} V) [M_B - M_B^\mu] = VS_A^\mu, \quad (\text{A.3})$$

$$2(\ell_{AB} - s_{AB} V) [M_A - M_A^\mu] + 2(\ell_0 - s_B V) [M_B - M_B^\mu] = VS_B^\mu. \quad (\text{A.4})$$

The determinant of the Hessian matrix of $R+VS$ is $4h(V)$ where

$$h(V) = (\ell_0 - s_A V)(\ell_0 - s_B V) - (\ell_{AB} - s_{AB} V)^2.$$

Strict concavity of $R+VS$ means that $h(V) > 0$ and $\ell_0 - s_A V > 0$, which in turn implies that $\ell_0 - s_B V > 0$.

The optimal pair of fees is the unique solution to the system (A.3)–(A.4):

$$M_A^* = M_A^\mu + \frac{V}{2h(V)} \left\{ (\ell_0 - s_B V) S_A^\mu - (\ell_{AB} - s_{AB} V) S_B^\mu \right\}, \quad (\text{A.5})$$

$$M_B^* = M_B^\mu + \frac{V}{2h(V)} \left\{ (\ell_0 - s_A V) S_B^\mu - (\ell_{AB} - s_{AB} V) S_A^\mu \right\}. \quad (\text{A.6})$$

The maximum $R(M_A^*, M_B^*) + VS(M_A^*, M_B^*)$ can be computed as

$$R^\mu + VS^\mu + \frac{V^2}{4h(V)} \left\{ (\ell_0 - s_B V) (S_A^\mu)^2 - 2(\ell_{AB} - s_{AB} V) S_A^\mu S_B^\mu + (\ell_0 - s_A V) (S_B^\mu)^2 \right\}. \quad (\text{A.7})$$

Inserting the expressions for S_A^μ and S_B^μ into (A.5)–(A.6) and collecting the terms in $M_A^\mu - v_A$ and $M_B^\mu - v_B$, respectively, we see that for $V > 0$, the difference $M_A^* - M_A^\mu$ is of the same sign as

$$[\ell_0 s_A - \ell_{AB} s_{AB} - (s_A s_B - s_{AB}^2) V] (M_A^\mu - v_A) + [\ell_0 s_{AB} - \ell_{AB} s_B] (M_B^\mu - v_B), \quad (\text{A.8})$$

and the difference $M_B^* - M_B^\mu$ of the same sign as

$$[\ell_0 s_{AB} - \ell_{AB} s_A] (M_A^\mu - v_A) + [\ell_0 s_B - \ell_{AB} s_{AB} - (s_A s_B - s_{AB}^2) V] (M_B^\mu - v_B). \quad (\text{A.9})$$

A.4 The Maximal Experimentation Strategy

Fixing p , let $y_i = M_i^\mu(p) - v_i$ for $i = A, B$, and set $\Delta = \bar{R}(p) - R^\mu(p)$. With the new variables $x_i = M_i - M_i^\mu(p)$ for $i = A, B$, the minimization problem in (12) and (13) amounts to minimizing the ratio $F(x_A, x_B)/G(x_A, x_B)$ where $F(x_A, x_B) = \Delta + \ell_0 x_A^2 + 2\ell_{AB} x_A x_B + \ell_0 x_B^2$ and $G(x_A, x_B) = s_A(x_A + y_A)^2 + 2s_{AB}(x_A + y_A)(x_B + y_B) + s_B(x_B + y_B)^2$. We suppress the arguments of these functions from now on and indicate the partial derivative with respect to x_i by means of the subscript " i ".

Denote the minimum of F/G by γ . The first-order conditions are $F_i G - F G_i = 0$ for $i = A, B$ by the quotient rule, so we must also have

$$\frac{F_A}{G_A} = \frac{F_B}{G_B} = \gamma \quad (\text{A.10})$$

at the minimum of F/G . As $F = \Delta + \frac{1}{2}F_A x_A + \frac{1}{2}F_B x_B$ and $G = \frac{1}{2}G_A(x_A + y_A) + \frac{1}{2}G_B(x_B + y_B)$, this in turn implies

$$1 = \frac{F}{\gamma G} = \frac{2\Delta + F_A x_A + F_B x_B}{\gamma G_A(x_A + y_A) + \gamma G_B(x_B + y_B)} = \frac{2\Delta + F_A x_A + F_B x_B}{F_A(x_A + y_A) + F_B(x_B + y_B)}$$

and hence $F_A y_A + F_B y_B = 2\Delta$. Writing out the partial derivatives and re-arranging, we obtain an affine relationship $\alpha x_A + \beta x_B = \Delta$ with $\alpha = \ell_0 y_A + \ell_{AB} y_B$ and $\beta = \ell_{AB} y_A + \ell_0 y_B$. Substituting $x_B = (\Delta - \alpha x_A)/\beta$ into the identity $F_A G_B = F_B G_A$ finally yields a quadratic equation for x_A .

A.5 Proofs

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

Given the subjective probability p currently assigned to the state (e_A^1, e_B^1) , the vector of expected demands is

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

with $\bar{n}_A = n_A(M_A, M_B, 1)$, $\bar{n}_B = n_B(M_A, M_B, 1)$, $\underline{n}_A = n_A(M_A, M_B, 0)$ and $\underline{n}_B = n_B(M_A, M_B, 0)$.

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

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

where

$$\begin{aligned} \begin{pmatrix} d\hat{Z}_A \\ d\hat{Z}_B \end{pmatrix} &= \begin{pmatrix} \sigma_A^{-1} & 0 \\ 0 & \sigma_B^{-1} \end{pmatrix} \begin{pmatrix} dN_A - \mathbb{E}^p[\tilde{n}_A] dt \\ dN_B - \mathbb{E}^p[\tilde{n}_B] dt \end{pmatrix} \\ &= \begin{pmatrix} \sigma_A^{-1} & 0 \\ 0 & \sigma_B^{-1} \end{pmatrix} \begin{pmatrix} \tilde{n}_A - \mathbb{E}^p[\tilde{n}_A] \\ \tilde{n}_B - \mathbb{E}^p[\tilde{n}_B] \end{pmatrix} dt + \begin{pmatrix} dZ_A \\ dZ_B \end{pmatrix} \end{aligned}$$

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

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

Relative to the platform provider's information filtration, $d\hat{Z}_A$ and $d\hat{Z}_B$ are normally distributed with mean zero and variance dt , and the infinitesimal covariance $\langle d\hat{Z}_A, d\hat{Z}_B \rangle$ is zero, so the change in beliefs dp is normally distributed with mean zero and variance

$$p^2(1-p)^2(\bar{n}_A - \underline{n}_A)^2\sigma_A^{-2}dt + p^2(1-p)^2(\bar{n}_B - \underline{n}_B)^2\sigma_B^{-2}dt = p^2(1-p)^2S(M_A, M_B)dt.$$

Now consider a pricing policy with $S(M_A, M_B)$ bounded away from 0, and suppose that the true state is $\theta = 1$. As

$$\begin{pmatrix} d\hat{Z}_A \\ d\hat{Z}_B \end{pmatrix} = \begin{pmatrix} \sigma_A^{-1} & 0 \\ 0 & \sigma_B^{-1} \end{pmatrix} \begin{pmatrix} \bar{n}_A - \mathbb{E}^p[\tilde{n}_A] \\ \bar{n}_B - \mathbb{E}^p[\tilde{n}_B] \end{pmatrix} dt + \begin{pmatrix} dZ_A \\ dZ_B \end{pmatrix},$$

we see that relative to the information filtration of an outside observer who knows the true state of the world, dp is normally distributed with mean

$$p \{ (\bar{n}_A - \mathbb{E}^p[\tilde{n}_A])^2 \sigma_A^{-2} + (\bar{n}_B - \mathbb{E}^p[\tilde{n}_B])^2 \sigma_B^{-2} \} dt = p(1-p)^2 S(M_A, M_B) dt.$$

As this is strictly positive on $]0, 1[$, the process of beliefs is a submartingale with respect to the observer's filtration and, if started at a non-degenerate prior, almost surely converges to its upper bound 1 as $t \rightarrow \infty$. An analogous argument establishes convergence to 0 when the true state is $\theta = 0$. ■

PROOF OF LEMMA 2: We fix a non-degenerate belief p . In Section 6, we saw that $R + VS$ is at least weakly concave in (M_A, M_B) . In Appendix A.3, we computed the determinant of the Hessian matrix of $R + VS$ as $4h(V)$ with $h(V) = (\ell_0 - s_A V)(\ell_0 - s_B V) - (\ell_{AB} - s_{AB} V)^2$ and noted that strict concavity means $h(V) > 0$ and $\ell_0 - s_A V > 0$.

If $R + VS$ is not strictly concave, its Hessian matrix is negative semidefinite but not negative definite (meaning $h(V) = 0$ with $\ell_0 - s_A V > 0$ or $\ell_0 - s_B V > 0$) and the first-order conditions (A.3)–(A.4) collapse into a single equation which defines the locus of all fee combinations along which $R + VS$ achieves its maximum. Thus, we must have

$$\frac{\ell_0 - s_A V}{\ell_{AB} - s_{AB} V} = \frac{\ell_{AB} - s_{AB} V}{\ell_0 - s_B V} = \frac{S_A^\mu}{S_B^\mu}$$

with all numerators and denominators being different from zero. In particular, the inequality $\ell_0 - s_A V > 0$ holds in this case as well.

As $h(0) = \ell_0^2 - \ell_{AB}^2 > 0$ and $h(\ell_0/s_A) \leq 0$, the function h has exactly one root in the interval $]0, \ell_0/s_A]$. We denote it by \hat{V} . Now, recall from Section 6 that V is strictly decreasing in the interest rate r . If there were an $r > 0$ such that $V = \hat{V}$, then we would have $V > \hat{V}$ and $h(V) < 0$ for all interest rates smaller than r , which contradicts what was said in the first two paragraphs of this proof. We can thus conclude that $h(V) > 0$ for all $r > 0$. ■

PROOF OF PROPOSITION 1: The myopically optimal revenue $R^\mu(p) = R(v_A/2, v_B/2, p)$ is linear in p . So the function $v = R^\mu$ trivially solves the Bellman equation (10) with $V \equiv 0$. ■

PROOF OF PROPOSITION 2: For $e_B^1 = e_B^0 = 0$, we have $\ell_0 \equiv 1$, $\ell_B \equiv 0$, $\ell_A(p) = p e_A^1 + (1-p) e_A^0$, $s_A = s_{AB} = 0$ and $s_B = (e_A^1 - e_A^0)^2 \sigma_A^{-2}$. The term (A.8) thus simplifies to $-\ell_A s_B [M_B^\mu - v_B]/2$, and its counterpart (A.9) to $s_B [M_B^\mu - v_B]$. Finally, $M_B^\mu < v_B$ by (8). ■

PROOF OF PROPOSITION 3: For $e_B^1 = e_B^0 = e_B$, the ratios ℓ_B/ℓ_0 , s_{AB}/s_B and s_A/s_{AB} all reduce to e_B . The term (A.9) thus has the same sign as $(\ell_0 - e_B \ell_{AB})[e_B(M_A^\mu - v_A) + M_B^\mu - v_B]$, and a simple computation using (7)–(8) shows that the expression in square brackets is indeed negative. The term (A.8) has the same sign as $(\ell_B - \ell_A)[e_B(M_A^\mu - v_A) + M_B^\mu - v_B]$, and hence is positive if and only if $\ell_A > \ell_B$. ■

PROOF OF PROPOSITION 4: For $(e_B^0, e_B^1) = (e_A^0, e_A^1) = (e^0, e^1)$, we have $\ell_A = \ell_B = \ell_{AB}$, $d_A = d_B$ and $s_A = s_B$, so we shall simply write ℓ , d and s for these objects. We know from Proposition 1 that the Bellman equation (10) is solved by the affine function R^μ , so that $V \equiv 0$. As $M_A^\mu \equiv v_A/2$, $M_B^\mu \equiv v_B/2$, the terms (A.8)–(A.9) are thus of the same sign as $-\ell_0 s - \ell s_{AB} v_A - [\ell_0 s_{AB} - \ell s] v_B$ and $-\ell_0 s_{AB} - \ell s v_A - [\ell_0 s - \ell s_{AB}] v_B$, respectively. If the expressions in square brackets are positive at all beliefs, then both terms are negative and bounded away from 0 on

the unit interval, and the result follows by continuous dependence of the value function and its second derivative on (e_B^0, e_B^1) .

As ℓ_0/ℓ is strictly decreasing in p (see Appendix A.1), and $s_{AB}/s < s/s_{AB} < d/d_0$, it is enough to show that

$$\frac{\ell_0^1}{\ell^1} > \frac{\ell^1 - \ell^0}{\ell_0^1 - \ell_0^0}.$$

This inequality is easily seen to hold for all (e^0, e^1) with $0 \leq e^0 < e^1 < 1$. ■

PROOF OF PROPOSITION 5: We know from the proof of Proposition 3 that for $e_B^0 = e_B^1$, the expression in curly brackets in (A.6) is negative and bounded away from 0 on the unit interval, while the expression in curly brackets in (A.5) is of the same sign as $\ell_A - \ell_B$. The result thus follows by continuous dependence of the value function and its second derivative on (e_B^0, e_B^1) . ■

PROOF OF PROPOSITION 6: Under condition (14), the coefficient of $M_B^\mu - v_B$ in (A.8) is clearly positive, and using the fact that $V < \ell_0/s_B$, we have

$$\begin{aligned} \ell_0 s_A - \ell_{AB} s_{AB} - (s_A s_B - s_{AB}^2) V &> \ell_0 \left(s_A - \frac{s_A s_B - s_{AB}^2}{s_B} \right) - \ell_{AB} s_{AB} \\ &> \ell_{AB} \left[\frac{s_B}{s_{AB}} \left(s_A - \frac{s_A s_B - s_{AB}^2}{s_B} \right) - s_{AB} \right] \\ &= 0, \end{aligned}$$

so that the coefficient of $M_A^\mu - v_A$ in (A.8) is also positive.

Under condition (15), the coefficient of $M_A^\mu - v_A$ in (A.8) is clearly negative, and we have

$$\frac{\ell_0(p)}{\ell_{AB}} < \frac{s_B}{s_{AB}},$$

so that the coefficient of $M_B^\mu - v_B$ in (A.8) is also negative.

The statements about the fee on side B follow in the same way. ■

PROOF OF PROPOSITION 7: We write $(e_B^1, e_B^0) = (e_A^0, e_A^1) = (\underline{e}, \bar{e})$ with $0 \leq \underline{e} < \bar{e} < 1$, and $v_A = v_B = v$. The expected direct price effect ℓ_0 is constant and equal to $1/(1 - \bar{e}\underline{e})$. The expected indirect effects are given by $\ell_A(p) = e(p)\ell_0$ and $\ell_B(p) = e(1-p)\ell_0$ with $e(p) = p\bar{e} + (1-p)\underline{e}$; as a consequence, ℓ_{AB} is constant as well, being equal to $(\bar{e} + \underline{e})\ell_0/2$. We compute $M_A^\mu(p) = [\frac{1}{2} + f(p)]v$ and $M_B^\mu(p) = [\frac{1}{2} - f(p)]v$ with

$$f(p) = \frac{\bar{e} - \underline{e}}{2 + \bar{e} + \underline{e}} \left(p - \frac{1}{2} \right).$$

Appendix A.2 implies $s_{AB} = 0$; given that $\sigma_A = \sigma_B$, moreover, $s_A = s_B$. Writing s for this parameter, we see that the term (A.8) has the same sign as $(\ell_0 - sV + \ell_{AB})f(p) - \frac{1}{2}(\ell_0 - sV - \ell_{AB})$. Appendix A.3 implies $\ell_0 - sV > \ell_{AB}$, so there exists $\bar{p} > 1/2$ such that this expression is negative for all $p < \bar{p}$. The limit of the expression as $p \rightarrow 1$ and $V \rightarrow 0$ is of the same sign as

$$\left(1 + \frac{\bar{e} + \underline{e}}{2} \right) f(1) - \frac{1}{2} \left(1 - \frac{\bar{e} + \underline{e}}{2} \right) = \bar{e} - \frac{1}{2}.$$

Finally, the term (A.9) is of the same sign as $(\ell_0 - sV + \ell_{AB})f(1-p) - \frac{1}{2}(\ell_0 - sV - \ell_{AB})$. ■

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.
- BECKER, G. S. and MURPHY, K. M. (1993): “A Simple Theory of Advertising as a Good or Bad”, *Quarterly Journal of Economics*, 104, 942–964.
- BELLEFLAMME, P. and M. PEITZ (2010): *Industrial Organization: Markets and Strategies*. Cambridge: Cambridge University Press.
- BERGEMANN, D. and J. VÄLIMÄKI (1997): “Market Diffusion with Two-sided Learning”, *RAND Journal of Economics*, 28, 773–795.
- BERGEMANN, D. and J. VÄLIMÄKI (2002): “Entry and Vertical Differentiation”, *Journal of Economic Theory*, 106, 91–125.
- BERGEMANN, D. and J. VÄLIMÄKI (2008): “Bandit Problems,” in: *The New Palgrave Dictionary of Economics* (2nd edition, ed. by S. Durlauf and L. Blume), Basingstoke and New York, Palgrave Macmillan Ltd.
- 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), 53–68, Oxford University Press.
- BONATTI, A. (2011): “Menu Pricing and Learning”, *American Economic Journal: Microeconomics*, 3, 124–163.
- CABRAL, L. (2011): “A Dynamic Theory of Two-Sided Markets”, working paper, New York University, available at <https://www.researchgate.net/publication/228944076>.
- FELLI, L. and HARRIS, C. (1996): “Learning, Wage Dynamics and Firm-Specific Human Capital”, *Journal of Political Economy*, 104, 838–868.
- HARRINGTON, J.E. JR. (1995): “Experimentation and Learning in a Differentiated-Products Duopoly”, *Journal of Economic Theory*, 66, 175–288.
- JOVANOVIĆ, B. (1979): “Job Matching and the Theory of Turnover”, *Journal of Political Economy*, 87, 972–990.
- 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 (1997): “Optimal Experimentation in a Changing Environment”, Research Paper No. 1443, Stanford Graduate School of Business, available at <http://dx.doi.org/10.2139/ssrn.51597>.
- 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., L. SAMUELSON and E. SCHLEE (1994): “Strategic Information Manipulation in Duopolies”, *Journal of Economic Theory*, 62, 363–384.
- MOSCARINI, G. (2005): “Job Matching and the Wage Distribution”, *Econometrica*, 73, 481–516.

- NOCKE, V., M. PEITZ and K. STAHL (2007): “Platform Ownership”, *Journal of the European Economic Association*, 5, 1130–1160.
- PAPAGEORGIOU, T. (2014): “Learning Your Comparative Advantages”, *Review of Economic Studies*, 81, 1263–1295.
- PEITZ, M., S. RADY and P. TREPPER (2011): “Experimentation in Two-Sided Markets”, SFB/TR 15 Discussion Paper No. 365, available at <http://www.sfbtr15.de/uploads/media/365.pdf>.
- PRAT, J. and B. JOVANOVIĆ (2014): “Dynamic Contracts when the Agent’s Quality is Unknown”, *Theoretical Economics*, 9, 865–914.
- 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.
- SANNIKOV, Y. (2008): “A Continuous-Time Version of the Principal-Agent Problem”, *Review of Economic Studies*, 75, 957–984.
- WEYL, G. (2010): “A Price Theory of Multi-Sided Platforms”, *American Economic Review*, 100, 1642–1672.