Entry And Experimentation In Oligopolistic Markets For Experience Goods

Daniel Krähmer*

Wissenschaftszentrum für Sozialforschung Berlin (WZB)

Social Science Research Center Berlin

January 13, 2003

Abstract

We investigate a two-period Bertrand market in which one seller introduces an experience good. The new product competes with an alternative good of known quality. Ex ante neither sellers nor consumers know the value of the new product. While consumers can learn their valuation for the new good by actual consumption (*experimentation*), sellers cannot observe experimentation outcomes. Thus, asymmetric information arises if the buyer experiments. As a result, the equilibrium is inefficient, and too little entry occurs.

Keywords: Entry, experimentation, asymmetric information, bandit problem, Bertrand competition..

JEL Classification: L13, L15, D82, D83

*Daniel Krähmer, Research Unit: Market Processes and Governance, Wissenschaftszentrum für Sozialforschung Berlin, Reichpietschufer 50, 10785 Berlin, Germany. Phone: ++49-30-25491-439. Fax: ++49-30-25491-400. E-mail: `kraehmer@medea.wz-berlin.de`

1 Introduction

When an experience good is introduced, buyers may learn their valuation for the new good through actual consumption (Nelson, 1972). Sellers, however, can often not observe the consumers' private experience with the new good. If buyers purchase repeatedly, this gives rise to asymmetric information between sellers and buyers in later periods.

The paper explores the consequences of this informational asymmetry for a duopolistic Bertrand market in a simple two period framework. In each period, a risky seller who initially introduces a new experience good competes with a safe seller who offers an alternative good of known quality. The provider of the new product may be an entrant into a previously monopolistic market, or the new good may be a re-launch of an outdated design. A buyer learns about her valuation for the new product by actual consumption. The buyer's valuation for the safe good may be known from previous experience or from an exogenous recommendation, for instance, a consumer guide book.

From the consumer's perspective, a seller is therefore one arm of a two-armed bandit. On the one hand, by trying out the risky good (*experimentation*), the buyer learns about her valuation. On the other hand, the safe good provides a safe benefit but does not give any information about the new good. Thus the buyer faces the well-known intertemporal trade-off between exploration and exploitation.

The aim of the analysis is to determine when experimentation or, equivalently, entry occurs in equilibrium and to study the welfare properties of the equilibrium. We say that experimentation, or entry, occurs if the buyer chooses the risky seller in the first period in equilibrium.

Our results are as follows: first, the market equilibrium is in general inefficient with social welfare in equilibrium lower than in the social planner solution. Second, in equilibrium there is too little entry compared with the planner solution.

What drives these results is that experimentation outcomes are private information. If the buyer experiments in period 1, sellers cannot condition their prices on the buyer's valuation in the second period. Therefore with positive probability, the buyer chooses the seller for whom her valuation is actually lowest. As a result, a welfare loss occurs.

The inefficient use of information reduces the benefits of the risky choice. This impairs the risky seller's competitive position relative to the safe seller as the welfare loss rises. As a consequence, when a planner is indifferent between the risky and the safe arm, the risky seller's competitiveness is too weak to induce the buyer to experiment. The risky seller therefore stays out of the market although entry would be socially desirable. In this sense, there is too little entry.

Moreover, we discuss whether the buyer is willing to pay ex ante for information about the risky good. We shall identify a trade-off between information and competition. An informed buyer enjoys an information rent relative to her uninformed counterpart, but information transforms the pricing game in such a way that sellers' market power is increased, competition is weakened, and higher prices are charged.

Our paper is a first step towards extending the market experimentation literature to the case of privately observed experimentation outcomes. Several other papers (Bergemann and Välimäki, 1996, 1997, 2002; Keller and Rady, 2003) study the relationship between entry, experimentation and oligopolistic competition, but in all of them experimentation outcomes are public information. Except for the finite time horizon, our model is closest to Bergemann and Välimäki (1996). Because experimentation outcomes are public information in their setting, in each stage sellers can condition their prices on the buyer's valuation. In contrast to our results, the market solution and the social planner solution are then identical.

Our paper is also related to the literature on multi-armed bandit problems.[1] In the traditional bandit literature, the pay-off characteristics of an arm are exogenous. By contrast, in our model, as in Bergemann and Välimäki (1996), the bandit is rendered a player such that the bandit characteristics are subject to a player's choice.[2]

The experience goods literature has mainly focussed on how the lemons problem

[1] See, e.g., Rothschild (1974) and Berry and Fristedt (1985).

[2] Active bandits are also considered in Bar-Isaak (2000). In his model, bandits are privately informed sellers who can signal their type.

a la Akerlof, which arises in the context of non-contractible quality, can be overcome. The lemons problem is typically mitigated when firms can build up a reputation[3], or when they can signal high quality[4]. In our model, the lemons problem is not an issue because sellers are uninformed. In a model similar to our's, Kim (1992) studies the pricing behaviour of a monopolist who, as in our case, is uninformed about consumers' tastes and cannot observe experimentation outcomes.

The paper is organized as follows. Section 2 presents the model. In section 3 the model is analyzed and the main results are derived. Section 4 concludes.

2 The Model

There are two time periods, $t = 1, 2$. In each period, two sellers, a *safe seller*, S, and a *risky seller,* R, each produce a unit of a good at zero cost. Goods are also denoted by S, and R respectively. Sellers compete in prices to attract a buyer, B, who buys at most one unit of the goods in each period.

B's willingness to pay for a good equals her expected instantaneous valuation for that good.[5] B's instantaneous valuation for S is known to all players to be $\lambda \in [0, 1]$. B's true instantaneous valuation $x \in [0, 1]$ for R is known neither by sellers nor by B. However, players hold a common belief about x, that is, B's valuation for R is a random variable X with values in $[0, 1]$. Let F be the cumulative density function of X, and let $f = F'$ be the corresponding probability density. We assume that f is continuously differentiable and strictly positive on $[0, 1]$. Suppose that F is common knowledge among all players.[6]

B can learn x with certainty by consuming R. So R is an experience good, and

[3] See, e.g., the influential article of Klein and Leffler (1981), or Shapiro (1982, 1983a, 1983b), Liebeskind and Rumelt (1989).

[4] See, e.g., Allen (1984), Riordan (1986), Hoerger (1993).

[5] Risk-neutrality is assumed for simplicity. Yet, quasi-linearity is a crucial assumption.

[6] F can be interpreted as the distribution of valuations for R among a continuum $[0, 1]$ of consumers where consumer $x \in [0, 1]$ has valuation x for R. This includes that consumers may have heterogeneous tastes not only for R but also for S. For what matters is only a consumer's difference $x - \lambda$ between the valuations for the goods. If λ depends on x, we get the same results by replacing $F(x)$ by $\overline{F}(x) = F(x - \lambda(x))$ and assuming that the valuation for S is 0.

for B each seller is an arm of a two-armed bandit. We say that *experimentation,* or *entry,* occurs when B chooses R in period 1, and we call her observation x the *experimentation outcome.* The experimentation outcome is B's private information. So x is B's subjective taste for R which cannot be observed by the sellers.

The timing is as follows. In $t = 1$, R and S simultaneously set prices r_1 and s_1. Given these prices, B chooses whether to buy a good, and, if so, which one. B learns her true x only if she experiments. Sellers observe consumption decisions but not experimentation outcomes. Call AI (*Asymmetric Information*) the information set reached if B experiments and NI (*No Information*) the information set reached if B does not experiment. In $t = 2$, R and S simultaneously set prices r_2 and s_2, and B, observing these prices, chooses whether to buy a good, and if so, which one.

We assume that players are Bayesian sequentially rational and play a perfect Bayesian Nash equilibrium and that they do not discount stage 2 pay-offs. To introduce a trade-off between learning and instant gratification, we assume $E[X] \leq \lambda$. Otherwise, the risky arm dominates not only in terms of future but also in terms of current benefits.

3 Analysis of the Model

We solve for the perfect Bayesian Nash equilibria of the game by backward induction.

3.1 The Second Stage

3.1.1 Information Set NI: Uninformed Players

At NI no party knows B's true valuation for R. The pricing game is then a classical Bertrand game where B's valuation for R is $E[X]$, and that for S is λ. We assume the following tie-breaking rule when B is indifferent: if there is exactly one seller who could reduce price slightly without making negative profits, B chooses that seller. Otherwise, B chooses S with probability 1.

Hence, the unique Bayesian Nash pricing equilibrium is $r^{NI} = 0, s^{NI} = \lambda -$

$E[X]$,[7] and B chooses S in equilibrium. B's expected pay-off is $u_{NI} = E[X]$.

3.1.2 Information set AI: Informed Buyer, Uninformed Sellers

At AI the buyer has private information about her valuation x for R. To derive the equilibrium, we go through players' decisions in more detail.

Buyer: Given prices r and s, the buyer chooses R if

$$x - r > \lambda - s, \tag{1}$$

and she chooses S if the reverse inequality holds. If B is indifferent, we use the tie-breaking rule above.[8]

Sellers: Given B's decision, sellers' expected period 2 profits are

$$\pi^R(r,s) = r \cdot P[X - r > \lambda - s] = r \cdot (1 - F(\lambda + r - s)), \tag{2}$$

$$\pi^S(r,s) = s \cdot P[X - r < \lambda - s] = s \cdot F(\lambda + r - s). \tag{3}$$

The necessary first order conditions for a Nash pricing equilibrium are thus

$$1 - F(\lambda + r - s) - r \cdot f(\lambda + r - s) = 0, \tag{4}$$

$$F(\lambda + r - s) - s \cdot f(\lambda + r - s) = 0. \tag{5}$$

The sufficient second order conditions for a solution (r,s) of (4), (5) to be an equilibrium are

$$-2f(\lambda + r - s) - r \cdot f'(\lambda + r - s) < 0, \tag{6}$$

$$-2f(\lambda + r - s) + s \cdot f'(\lambda + r - s) < 0. \tag{7}$$

This condition requires the *curvature* (f'/f) of F to be neither extremely large nor extremely small at the "margin" $\lambda + r - s$. The curvature measures how fast a seller's "market share" reacts to small price changes. If one seller's share reacts very fast, the other seller can profitably reduce prices so as to increase sales drastically.

The latter cannot happen if the curvature is bounded. This is implied by the following condition which guarantees the existence of a pure strategy equilibrium:[9]

$$-\frac{2}{z} \leq -\frac{f'(x)}{f(x)} \leq \frac{2}{z} \qquad \text{for all } x \in (0,1), \tag{8}$$

[7] In this subsection we suppress time subscripts and always mean period 2 prices.

[8] Notice that B is indifferent at AI with probability 0.

[9] For a proof see Bester (1992), p. 438.

where $z = \max\{1/F(\lambda), 1/[1-F(\lambda)]\}$.

Condition (8) is for example satisfied if f is log-concave with $f'(0)/f(0) \leq 2/z$. By log-concavity, f'/f is positive and bounded by $f'(0)/f(0)$. This implies (8).

From now on we shall assume that (8) holds. The equilibrium need not be unique. We assume that there is a focal equilibrium that players coordinate on.

3.1.3 The Value of Information

The second stage can be used to ask whether the buyer is willing to pay for information about the risky good. Denote by u_{AI} the buyer's expected pay-off in the continuation game following AI. Let (r, s) be an equilibrium at AI, and let $\Delta = r - s$. With $u^{NI} = \lambda - s^{NI}$, the willingness to pay for information is

$$\begin{aligned} u_{AI} - u_{NI} &= P[X > \lambda + \Delta](E[X \mid X > \lambda + \Delta] - \lambda) \\ &\quad + P[X < \lambda + \Delta](\lambda - \lambda) \\ &\quad - \left(rP[X > \lambda + \Delta] + sP[X < \lambda + \Delta] - s^{NI}\right). \end{aligned} \tag{9}$$

The first two lines are the *gross* value of information. The sign depends on the size of Δ: if Δ is not too small, the gross value of information will be positive. The third line indicates a price effect: at information set AI both sellers have some market power while at NI they play a Bertrand game with flat demand. Inserting $s^{NI} = \lambda - E[X]$ yields

$$\begin{aligned} u_{AI} - u_{NI} &= P[X > \lambda + \Delta](E[X \mid X > \lambda + \Delta] - E[X]) \\ &\quad + P[X < \lambda + \Delta](\lambda - E[X]) \\ &\quad - rP[X > \lambda + \Delta] - sP[X < \lambda + \Delta]. \end{aligned} \tag{10}$$

Whatever the buyer learns, that is, whether $X > \lambda + \Delta$ or $X < \lambda + \Delta$, the *gross* utility when informed is higher than *net* utility when uninformed. This can be seen as a form of favourable selection. Yet, due to sellers' increased market power at AI, the buyer may have to pay a relatively high price when informed.

To see which effect dominates, consider the special case of a symmetric market. In a symmetric market, exactly half of the buyers have a higher valuation for R

than for S, that is, $F(\lambda) = 1/2$. By (4) and (5), an equilibrium is given by $r = s = \min\{1/(2f(\lambda)), \lambda\}$. In particular, $\Delta = 0$. Hence,

$$\begin{aligned} u_{AI} - u_{NI} = {} & P[X > \lambda](E[X \mid X > \lambda] - E[X]) \\ & + P[X < \lambda](\lambda - E[X]) - \min\left\{\frac{1}{2f(\lambda)}, \lambda\right\}. \end{aligned} \tag{11}$$

Thus, the value of information increases in $f(\lambda)$. Intuitively, $f(\lambda)$ measures sellers' market power at equilibrium prices. Sellers trade off price reductions against gains in market share. If $f(\lambda)$ is large, B's valuation for R is relatively concentrated around λ, and price reductions attract relatively many consumers. That is, competition is strong, and equilibrium prices will be low. Thus, the gross favourable selection effect may dominate the market power effect. If $f(\lambda)$ is small, the reverse holds, and the value of information will be low. In this case, ex ante an uninformed buyer may be better off than her informed counterpart and thus unwilling to pay for information.

3.2 The First Stage

We now analyze the entire game. We shall show that—provided sellers do not charge the same prices in period 2—the market equilibrium is inefficient, and that, in a sense made more precise later, this inefficiency gives rise to under-entry.

3.2.1 The Social Planner Solution

We consider the social planner solution in which the planner is as informed as the buyer in the market, that is, if the planner's strategy is to experiment, the planner can observe the experimentation outcome x. This is just the solution to the two-armed bandit problem with one arm paying x with density $f(x)$ and with one arm paying λ with probability 1. We assume that the planner, if indifferent, selects the safe arm with probability 1.[10]

Hence, given experimentation in $t = 1$, the planner prescribes R in $t = 2$ if $x > \lambda$, and S if $x \leq \lambda$. Given no experimentation in $t = 1$, the planner prescribes R in $t = 2$ if and only if $E[X] > \lambda$. The latter is however ruled out by assumption.[11]

[10] This is consistent with the buyer's tie-breaking rule imposed in the market game.

[11] Otherwise the planner solution would always prescribe the risky arm in period 1.

Hence, experimentation yields (first best) utility

$$WR^{FB} = E[X] + P[X > \lambda] \cdot E[X \mid X > \lambda] + P[X < \lambda] \cdot \lambda, \tag{12}$$

whereas selection of the safe arm in $t = 1$ yields

$$WS^{FB} = \lambda + \lambda. \tag{13}$$

Thus the social surplus in the planner solution is $U^{FB}(\lambda) = \max\left\{WR^{FB}, WS^{FB}\right\}$. The planner's strategy is to experiment if, and only if, $WR^{FB} - WS^{FB} > 0$. Define

$$g(\lambda) = WR^{FB} - WS^{FB} \tag{14}$$

as the *first best experimentation incentive*. Notice that g strictly decreases in λ.[12] Hence, there is a unique cut-off point λ^{FB} with $g\left(\lambda^{FB}\right) = 0$, and experimentation is first best for all $\lambda < \lambda^{FB}$.

3.2.2 The Market Solution

We now turn to the market game. Before solving for the equilibrium, it is instructive to think directly about social surplus in the market game. Suppose prices in period 1 are given and B has chosen a particular seller in period 1, be it optimal or not. Suppose further that an equilibrium of the respective continuation game is played. Now consider players' pay-offs along that path in the game tree: due to quasi-linear utility, B's expected (net) life-cycle utility equals B's expected gross life-cycle utility net of the expected life-cycle payments transferred to sellers. Sellers' combined expected life-cycle profits are equal to B's expected life-cycle payments to sellers. Therefore, total social surplus along a particular play path is just equal to B's expected gross life-cycle utility along that path.

We now compute B's gross life-cycle utility along all possible paths given equilibrium play in period 2. If B chooses R in $t = 1$, then she gets gross instantaneous utility $E[X]$ in period 1, and play reaches AI where B gets gross period 2 utility

$$P[X > \lambda + \Delta] \cdot E[X \mid X > \lambda + \Delta] + P[X < \lambda + \Delta] \cdot \lambda, \tag{15}$$

[12] By Leibniz' rule it follows immediately that $g' < 0$.

where

$$\Delta = r_2 - s_2 \tag{16}$$

denotes the equilibrium price differential in period 2. Hence, *B's gross life-cycle utility from experimentation* is

$$WR = E[X] + P[X > \lambda + \Delta] \cdot E[X \mid X > \lambda + \Delta] + P[X < \lambda + \Delta] \cdot \lambda. \tag{17}$$

If B chooses S in $t = 1$, she gets λ in $t = 1$, and play reaches NI where B stays with S and gets λ again. Hence, *B's gross life-cycle utility from choosing S in $t = 1$* is

$$WS = \lambda + \lambda. \tag{18}$$

As for efficiency, suppose that the market equilibrium exists, $\Delta \neq 0$, and that experimentation is first best. We show now that in this case the market equilibrium is inefficient. Let U^M be the social surplus in the market equilibrium. That is, $U^M = WR$, if B experiments, and $U^M = WS$ otherwise.

Suppose first that S is chosen in $t = 1$ in the market. Since experimentation is first best, we have $U^{FB} > WS^{FB}$. Thus, $U^M = WS = WS^{FB} < U^{FB}$.

Next, suppose that R is chosen in $t = 1$ in the market. Suppose, for instance, $\Delta > 0$. Then B's choice in period 2 in the market differs from that in the planner solution in period 2 only when experimentation outcomes x are in $(\lambda, \lambda + \Delta)$. Notice that in the market B chooses S in $t = 2$ for all $x \in (\lambda, \lambda + \Delta)$ although her valuation for R exceeds that for S. In contrast, the planner solution prescribes R, for which B has the higher valuation. Therefore with probability $P[\lambda < X < \lambda + \Delta]$ the buyer obtains only gross utility λ in the market, whereas the planner solution yields a gross utility of $E[X \mid \lambda < X < \lambda + \Delta] > \lambda$. We state this finding in Proposition 1.

Proposition 1 *Suppose $\Delta \neq 0$. Then $WR < WR^{FB}$ for all $\lambda \in [0, 1]$.*

Proof: Suppose $\Delta > 0$. (For $\Delta < 0$, the same argument applies.) Then

$$\begin{aligned} WR - WR^{FB} &= P[X > \lambda + \Delta] \cdot E[X \mid X > \lambda + \Delta] + P[X < \lambda + \Delta] \cdot \lambda \\ &- P[X > \lambda] \cdot E[X \mid X > \lambda] + P[X < \lambda] \cdot \lambda \\ &= P[X > \lambda + \Delta] (E[X \mid X > \lambda + \Delta] - E[X \mid X > \lambda + \Delta]) \\ &+ P[\lambda + \Delta > X > \lambda] (\lambda - E[X \mid \lambda + \Delta > X \geq \lambda]) \\ &+ P[X < \lambda] (\lambda - \lambda) \\ &= P[\lambda + \Delta > X > \lambda] (\lambda - E[X \mid \lambda + \Delta > X \geq \lambda]) . \end{aligned}$$

The claim follows since $\lambda < E[X \mid \lambda + \Delta > X > \lambda]$. □

The two previous paragraphs show that the market surplus is *strictly* dominated by the planner surplus whenever experimentation is first best. So the market equilibrium is inefficient for all $\lambda < \lambda^{FB}$. The efficiency loss obtains because the buyer chooses the wrong seller with positive probability in period 2 of the market game. This is due to the presence of a non-zero price differential in period 2 which results from asymmetric information. Proposition 2 summarizes these observations.

Proposition 2 *Suppose $\Delta \neq 0$. Then $U^M < U^{FB}$ for all $\lambda < \lambda^{FB}$.*

Remark: With public information, Bergemann and Välimäki (1996) show that the market and the planner solutions are identical. In our simple model it can be easily seen why this is the case. Under public information, sellers can condition stage 2 prices on the experimentation outcome x. Then the unique equilibrium of the continuation game following experimentation in period 1 is

$$(r_2(x), s_2(x)) = \begin{cases} (x - \lambda, 0) & \text{if} \quad x > \lambda \\ (0, \lambda - x) & \text{if} \quad x \leq \lambda \end{cases}, \tag{19}$$

and R will make sales in $t = 2$ if $x > \lambda$, while S will make sales in $t = 2$ if $x \leq \lambda$. Therefore, the efficient seller is chosen in equilibrium, and there is no welfare loss. □

The presence of inefficiencies depends on a non-zero price differential Δ. The following proposition characterizes Δ in terms of the underlying distribution F.

Proposition 3 *$\Delta \neq 0 \Longleftrightarrow F(\lambda) \neq 1/2$, and $\Delta < 0 \Longleftrightarrow F(\lambda) > 1/2$.*

Proof: By adding up (4) and (5) and solving for $F(\lambda + \Delta)$, we obtain

$$F(\lambda + \Delta) = 1/2 - (1/2)\,\Delta f(\lambda + \Delta). \tag{20}$$

Thus, if $\Delta = 0$ solves (20), it follows that $F(\lambda) = 1/2$. Likewise, if $\Delta < 0$, then $F(\lambda + \Delta) > 1/2$. Hence, since F is increasing, $F(\lambda) > F(\lambda + \Delta) > 1/2$.

On the other hand, if $F(\lambda) = 1/2$, then $\Delta = 0$ is a solution to (20). Moreover, if $F(\lambda) > 1/2$ and $\Delta \geq 0$, then the left hand side is strictly larger than 1/2, and the right hand side is weakly smaller than 1/2, a contradiction. $\square$

The proposition says that inefficiencies arise if, and only if, the market is not symmetric. If, say, $F(\lambda) > 1/2$, then after experimentation the buyer is more likely to have a higher valuation for S than for R. This advantage enables S to quote a higher price than R so that $\Delta < 0$.

We now determine the equilibrium. Notice that period 1 prices need not be nonnegative since sellers may make positive profits in period 2. Intuitively, competition squeezes prices until one seller reaches a price below which he is better off when not selling than when selling in period 1. The other seller will then leave the buyer just indifferent between sellers.

More precisely, for $i = R, S$ denote by $\pi_2^i(AI)$, and $\pi_2^i(NI)$ respectively, seller i's period 2 profit in the continuation game following AI, and NI respectively. Then, the smallest price $\widehat{r}_1$ which R is willing to charge in period 1 is

$$\widehat{r}_1 + \pi_2^R(AI) = \pi_2^R(NI). \tag{21}$$

Notice that $\pi_2^R(NI) = 0$, thus, $\widehat{r}_1 = -\pi_2^R(AI)$.

Likewise, the minimum price $\widehat{s}_1$ which S is willing to charge in period 1 is

$$\widehat{s}_1 + \pi_2^S(NI) = \pi_2^S(AI). \tag{22}$$

As above, we use the following tie-breaking rule for B: if exactly one seller sets his minimum price, B chooses the other seller. Otherwise B chooses S.

Due to price competition, in equilibrium one seller must charge his minimum price. Due to our tie-breaking rule, this seller will not make sales in period 1.[13]

Suppose that B experiments in equilibrium. Then it is S who charges $\widehat{s}_1$, and S is indifferent between selling and not selling. Yet, in equilibrium also B is indifferent. Otherwise R could raise period 1 price and increase his profits. Since B and S are indifferent, the surplus B and S achieve together equals the surplus they would achieve if B chose S, that is, their surplus $WS - \pi_2^R(NI)$ from not experimenting. Accordingly, R must extract the rest of the total surplus from experimentation. Hence, because $\pi_2^R(NI) = 0$, R's equilibrium life-cycle profit conditional on selling in period 1 is $WR - WS$. In equilibrium, of course, for experimentation to obtain R's profit from selling in $t = 1$ must be larger than R's profit from not selling in $t = 1$. Hence, it must hold that $WR - WS > 0 + \pi_2^R(NI) = 0$.[14]

By a similar argument, if B does not experiment in equilibrium, then $WR - WS \leq 0$. Thus, B experiments in equilibrium if, and only if, $WR - WS > 0$.

In other words, price competition in stage 1 implies that the buyer chooses the seller for whom her gross life-cycle utility is highest. The buyer's *net* willingness to pay for information, as discussed in section 3.1.1, is irrelevant in equilibrium insofar as sellers' stage 2 profits will be reflected in period 1 prices. If, say, the value of information is negative, the risky seller has to reduce prices in period 1 so as to compensate the buyer up front for her anticipated loss in period 2.

To determine equilibrium prices, consider first the case in which B experiments. Notice that R's life-cycle profit conditional on selling in period 1 is equal to period 1 price plus period 2 profits in the continuation game following AI. Thus, $WR - WS = r_1 + \pi_2^R(AI)$. With $\widehat{r}_1 = -\pi_2^R(AI)$ it follows that R's period 1 price is

$$r_1^* = WR - WS + \widehat{r}_1. \tag{23}$$

Likewise, if B does not experiment in equilibrium, S's period 1 price is

$$s_1^* = WS - WR + \widehat{s}_1. \tag{24}$$

We summarize our findings in Proposition 4.

[13] Except if both sellers charge their minimum price in equilibrium. Then S makes sales in $t = 1$.

[14] The strict inequality results from our tie-breaking rule.

Proposition 4 *Suppose $WR - WS > 0$. Then B experiments and entry occurs. Sellers' equilibrium prices in period 1 are $r_1 = r_1^*$ and $s_1 = \widehat{s}_1$.*

Suppose $WR - WS \leq 0$. Then B does not experiment and entry does not occur. Sellers' equilibrium prices in period 1 are $r_1 = \widehat{r}_1$ and $s_1 = s_1^*$.[15]

We shall now show that there is too little entry in equilibrium. Define

$$h(\lambda) = WR - WS \tag{25}$$

as the *equilibrium experimentation incentive* or the risky seller's *entry incentive.*

We now compare h with the first best experimentation incentive $g = WR^{FB} - WS^{FB}$. Since $WS^{FB} = WS$, the next result follows directly from Proposition 1.

Proposition 5 *Suppose $\Delta \neq 0$, then $h(\lambda) < g(\lambda)$ for all $\lambda \in [0,1]$.*

This implies that there is a cut-off $\lambda^{ME} < \lambda^{FB}$ after which entry will not occur in equilibrium. More precisely, let $\lambda^{ME} = \sup\{\lambda \mid h(\lambda) = 0\}$ be the largest root of h. Thus, $h(\lambda) < 0$ for all $\lambda > \lambda^{ME}$, and, by Proposition 5, $\lambda^{ME} < \lambda^{FB}$.

Therefore, if $\lambda \in [\lambda^{ME}, \lambda^{FB})$, the risky seller stays out of the market although entry is socially optimal. Again, this is because information is used inefficiently in period 2. The informational externality reduces the buyer's gross valuation for the risky seller, WR, compared to the value of the risky arm in the planner solution, WR^{FB}, but does not impair B's valuation for the safe seller, WS. As a consequence, the entry price necessary to induce the buyer to experiment does not allow the risky seller to make positive profits, and he stays out of the market. In this sense, there is too little entry in the market equilibrium. Lemma 1 summarizes our results.

Lemma 1 *Suppose $\Delta \neq 0$. Then the market equilibrium is inefficient for all $\lambda < \lambda^{FB}$ and leads to under-entry for all $\lambda \in [\lambda^{ME}, \lambda^{FB})$.*

4 Conclusion

The model presented in this paper is a first step towards extending the market experimentation literature to the case of privately observed experimentation outcomes.

[15] Though not stated explicitly, this is the unique perfect Bayesian Nash equilibrium.

This leads to asymmetric information and, in contrast to the public information setting, results in a welfare loss and in too little entry.

The qualitative conclusions of the analysis remain valid if the buyer is assumed to obtain only an imperfect signal about the risky alternative, if the buyer is assumed to be risk averse, if the safe alternative is rendered risky as well, and if the number of sellers is increased. The model can easily be generalized along those lines. Essential for our argument is the assumption of quasi-linear utility. However, we conjecture that our qualitative conclusions still hold for more general utility functions because this essentially amounts to a transformation of the random variable X.

What is, however, crucial for our results is the restriction to a two-period time horizon. If the time horizon is extended, the logic of the model alters considerably. Consider the decision of an informed buyer at the second stage. With more periods, this decision contains an additional strategic element since the buyer's choice reveals information. In particular, there will be a ratchet effect. Suppose the buyer learns in period 1 to have a high taste for the risky good. This should increase the incentive to choose the risky seller in period 2. However, if she chooses the risky seller in period 2, the risky seller learns about the buyer's high valuation and can therefore set a high price in later periods. This reduces the buyer's incentive to choose the risky seller in period 2. How this effect feeds back to the experimentation incentive in the initial period, is a question for future research.

In a further direction of future investigation points the literature on multi-player bandit games in which several players play the same bandit and can observe the pay-offs obtained by other players (see Bolton and Harris, 1999). Bergemann and Välimäki (2000) study a multi-player bandit game in a market context in which bandits set prices. In Kamp (1998) first period buyers transmit their experience to second period buyers. Extensions of these approaches could shed more light on how the spread of information may affect the performance of experience goods markets.

Acknowledgements

This paper is based on Chapter 2 of my PhD dissertation written at the Free University Berlin. I wish to thank the editor Simon Anderson and two anonymous referees for very helpful suggestions. I also wish to thank Helmut Bester, Jürgen Bierbaum, Kai Konrad, and Roland Strausz for useful comments.

References

Allen, F., 1984, Reputation and product quality, RAND Journal of Economics 15, 311 - 327.

Bar-Isaak, H., 2000, Good advisers survive despite bad luck, mimeo, LSE.

Bergemann, D. and J. Välimäki, 1996, Learning and strategic pricing, Econometrica 64, 1125 - 1149.

Bergemann, D. and J. Välimäki, 1997, Market diffusion with two-sided learning, RAND Journal of Economics 28, 773 - 795.

Bergemann, D. and J. Välimäki, 2000, Experimentation in markets, Review of Econonomic Studies 67, 213 - 234.

Bergemann, D. and J. Välimäki, 2002, Entry and vertical differentiation, Journal of Economic Theory 106, 91 - 125.

Berry, D. A. and B. Fristedt, 1985, Bandit problems (Chapman and Hall, New York).

Bester, H., 1992, Bertrand equilibrium in a differentiated duopoly, International Economic Review 33, 433 - 448.

Bolton, P. and C. Harris, 1999, Strategic experimentation, Econometrica 67, 349 - 374.

Hoerger, T., 1993, Two-part pricing for experience goods in the presence of adverse selection, The International Journal of Industrial Organization 11, 451 - 474.

Kamp, B., 1998, Experience dependent information diffusion and product quality, Economic Inquiry 36, 161 - 170.

Keller, G. and S. Rady, 2003, Price dispersion and learning in a dynamic differentiated-good duopoly," RAND Journal of Economics, forthcoming.

Kim, J., 1992, Experience goods, expectations and pricing, Economic Record 68, 7 - 15.

Klein, B. and K. Leffler, 1981, The role of market forces in assuring contractual performance, Journal of Political Economy 89, 615 - 641.

Liebeskind, J. and R. Rumelt, 1989, Markets for experience goods with performance uncertainty, RAND Journal of Economics 20, 601 - 621.

Nelson, P., 1972, Information and consumer behavior, Journal of Political Economy 78, 311 - 329.

Riordan, M., 1986, Monopolistic competition with experience goods, Quarterly Journal of Economics 101, 265 - 279.

Rothschild, M., 1974, A two-armed bandit theory of market pricing, Journal of Economic Theory 9, 185 - 202.

Shapiro, C., 1982, Consumer information, product quality and seller reputation, Bell Journal of Economics 13, 20 - 35.

Shapiro, C, 1983a, Optimal pricing of experience goods, Bell Journal of Economics 14 , 497 - 507.

Shapiro, C., 1983b, Premiums for high quality products as returns to reputations, Quarterly Journal of Economics 98, 659 - 679.