Economic Theory (2007) 33: 211–224
DOI 10.1007/s00199-006-0152-0

SYMPOSIUM

Martin Dufwenberg · Uri Gneezy
Jacob K. Goeree · Rosemarie Nagel

Price floors and competition

Received: 4 August 2005 / Accepted: 18 August 2006 /
Published online: 23 September 2006

Abstract A potential source of instability of many economic models is that agents have little incentive to stick with the equilibrium. We show experimentally that this can matter with price competition. The control variable is a price floor, which increases the cost of deviating from equilibrium. According to traditional theory, a higher floor allows competitors to obtain higher profits. Behaviorally, the opposite result obtains with two (but not with four) competitors. An error model, which

We thank Jan Bouckaert, David Cooper, Simon Gächter, Reinhard Selten, Frank Verboven, a referee, and the participants at several seminars for stimulating discussions and helpful suggestions. The research has been sponsored by the Swedish Competition Authority. Martin Dufwenberg gratefully acknowledges the hospitality of the Laboratory of Experimental Economics at Bonn University, where he was visiting when part of the research was completed. Rosemarie Nagel acknowledges financial support of Spain's Ministry of Education under Grants SEC2002-03403 and SEC2005-08931 and thanks the Barcelona Economics Program of CREA for support.

M. Dufwenberg (✉)
Department of Economics and Economic Science Laboratory,
University of Arizona, Tucson, AZ 85721-0108, USA
E-mail: martind@eller.arizona.edu

U. Gneezy
The Rady School of Management, UC San Diego,
9500 Gilman Dr., MC 0093,
La Jolla, CA 92093-0093, USA
E-mail: ugneezy@ucsd.edu

J. K. Goeree
Division of the Humanities and Social Sciences 228-77, CALTECH,
Pasadena, CA 91125, USA
E-mail: jkg@hss.caltech.edu

R. Nagel
Department of Economics, Universitat Pompeu Fabra, Ramón Trias Fargas 25-27,
08005 Barcelona, Spain
E-mail: rosemarie.nagel@upf.edu

builds on Luce (Individual Choice Behavior, 1959), can adequately describe supra-Nash pricing with a low-floor, but then fails to capture the overall pro-competitive effect of a high-floor seen for duopolies.

Keywords Price competition · Price floors · Bertrand model · Experiment · Luce model

JEL Classification Numbers C72 · C92 · D43 · L13

1 Introduction

Students in microeconomics classes are often baffled by the classical prediction that competitive markets with free entry result in zero long-run profits. They wonder why firms would be willing to produce if they gain nothing by doing so. In addition, their own casual observations suggest that profits are non-negligible even in mature markets with close substitutes. Textbooks mostly try to bridge this gap by pointing out that one or more assumptions underlying the zero-profit result may not be met (e.g., by introducing capacity constraints, market power through product differentiation, costly entry, etc.). However, they rarely address the awkwardness of the zero-profit outcome in the standard case. One interpretation of the students' uneasiness with this prediction is in terms of lack of cost of deviating from the equilibrium. Why should firms exhibit rational equilibrium behavior if at the equilibrium they have no incentive to do so?

The Bertrand model of price competition, one of the most important pillars of modern oligopoly theory, predicts an equally stark outcome even with as few as two firms.[1] When both firms have the same constant marginal costs, their incentives to capture more market will result in cut-throat competition, driving prices down to marginal costs and eliminating all profits. Again the lack of cost of deviating from the equilibrium casts doubt on the predictive power of this result. In equilibrium, a firm's expected profit function is completely flat and *any* price (greater than or equal to the marginal cost) yields the same expected payoff. Moreover, if there is a slight chance that the rival will price above marginal cost, a firm is better off setting a higher non-competitive price as well.

One way to restore a non-negligible cost of deviating from equilibrium is by introducing minimum prices, or price floors, which are regularly employed in a wide spectrum of markets. For example, it is rather common that suppliers impose bounds on the pricing of retailers (see Ippolito (1991)), governments sometimes introduce minimum prices for certain goods (e.g., a minimum wage), and auctioneers may introduce 'bid caps' (like 'minimum bid constraints' in procurement auctions) as part of auction rules. With a price floor, the competing parties still have an incentive to undercut their rivals as in the standard Bertrand game, but if prices spiral downwards as a result, the final price level will be above marginal cost yielding some positive profit. Choosing a sufficiently high price floor thus ensures a non-negligible costs of deviating from equilibrium, and improves the drawing power of the Bertrand-Nash solution.

[1] See Tirole (1994, Chap. 5) for a textbook presentation. The model is named after Bertrand (1883).

Against this background, our goal is to examine the impact of price floors on competition experimentally. We consider four treatments, which differ in terms of

- The size of the price floor (*low-floor* or *high-floor*), and
- Whether there are two or four competitors (*duopoly* or *quadropoly*).

We incorporate the second treatment because previous research has indicated that the predictive power of the Bertrand model may crucially depend on the number of competitors (see e.g., Dufwenberg and Gneezy (2000)), so it seems natural to check whether the anti- or pro-competitiveness of price floors depends on the number of competitors.[2]

We test a decision-error model that incorporates in an intuitive way the aforementioned idea that the drawing power of equilibrium depends on the associated costs of deviations. Following Lopez-Acevedo (1997) and Baye and Morgan (2004), we generalize the classic Luce (1959) model. The Luce model incorporates boundedly rational choice in that players choose better responses with higher probabilities, but not necessarily the best response with probability one. More precisely, the choice probabilities for specific strategies are proportional to the expected payoffs associated with such strategies. We augment this framework by including a free parameter λ that determines how sensitive behavior is with respect to payoffs. Depending on λ, completely random behavior and Nash equilibrium appear as different limiting cases. The model is analytically tractable, and we derive testable predictions concerning the impact of price floors (which are new) and concerning the number of competitors (which appear also in Lopez-Acevedo (1997)).[3]

Section 2 describes the theory in more detail. Section 3 presents the experimental design. Section 4 contains the experimental results. Section 5 is a discussion, including suggested directions for follow-up research.

2 Theory

We consider a simple variant of the classic Bertrand duopoly game, where $n \geq 2$ competing firms simultaneously and independently choose prices for a homogeneous good produced at zero costs. Demand has a 'box' structure; there is demand for one unit of the good for prices up to a reservation value, and beyond that value demand falls to zero. Prices are constrained to be in the range $p \in [p_L, p_H]$, where

[2] Experimental research on price competition goes back to Fouraker and Siegel (1963). See Plott (1989) and Holt (1995)) for surveys and Brown-Kruse et al. (1994), Cason (1995), Cason and Davis (1995), Mason and Phillips (1997), Dufwenberg and Gneezy (2000), Huck et al. (2000), and Selten and Apesteguia (2005) for some more recent work. These studies typically do not explore the role of price floors. Murphy (1966) comes closest, showing that behavior is influenced by whether or not profits are negative in equilibrium. However, the study is not about price floors that guarantee positive profits, and differs in many other ways from our design (like having a repeated game, and having incomplete information about profits). Isaac and Plott (1981) and Smith and Williams (1981) introduce price controls (floors and ceilings) in a different market institution: double auctions. Unlike Bertrand duopoly games, these are well known for their extraordinarily competitive properties in experiments, and some of the price controls considered reduce competition and may be the source of some inefficiency.

[3] The spirit of the Luce model is similar to McKelvey and Palfrey (1995) notion of logit equilibrium. It is harder, however, to derive comparative statics properties of the logit equilibria in the games we consider.

$p_L > 0$ is the price floor and p_H is the reservation value (consumers' maximum willingness to pay). Under the homogeneity assumption, the lowest-price firm sells one unit of the good and higher-price firms sell nothing. If m, with $1 < m \leq n$, firms tie for the lowest price each of these sells one mth of a unit. The Nash equilibrium prediction for this game is easy to derive: Any strategy profile where some firm i chooses a price higher than p_L cannot be an equilibrium; firm i could increase its profit from zero to at least $p_L/n > 0$ by choosing its price equal to p_L. On the other hand, for each firm to choose price equal to p_L is a (strict) equilibrium, in fact the unique equilibrium of the game.

We next derive predictions of the Luce (1959) model, which assumes boundedly rational players that are prone to mistakes with the probability of a mistake being inversely related to its cost. In other words, in the Luce model, choice probabilities are positively correlated with expected payoffs although not perfectly so, i.e., players choose better options more frequently but not necessarily the best one with probability one. There are several ways to parameterize decision-error models of this kind.[4] The original Luce (1959) model, where choice probabilities are proportional to expected payoffs, is arguably one of the simplest possible formulations. Here we consider a one-parameter generalization, which assumes that choice probabilities are proportional to expected payoffs raised to the power of λ.

In contrast to simple decision-making tasks where errors simply add 'noise' around the optimal choice, they can have a compounding effect in interactive contexts such as games. In the Bertrand duopoly pricing game, for example, an upward error by one player makes higher prices by others more profitable, and hence more likely, which reinforces the original error. This way, endogenous errors can cause decisions to be systematically different from Nash predictions. Here we investigate the equilibrium effects of errors in the Bertrand game in the presence of a price floor.

Let $\pi^e(p)$ denote a firm's expected payoff from choosing a price $p \in [p_L, p_H]$. The expected payoff depends on the distribution of the rivals' prices, denoted by $F(p)$, with associated density $f(p)$:

$$\pi^e(p) = p(1 - F(p))^{(n-1)}. \tag{1}$$

In the generalized Luce model, choice frequencies are proportional to expected payoffs raised to the power λ:

$$f(p) = \frac{(\pi^e(p))^\lambda}{\int_{p_L}^{p_H} (\pi^e(y))^\lambda \mathrm{d}y}. \tag{2}$$

The denominator in (2) is a constant independent of p that ensures that the density integrates to one. The exponent λ is a 'precision' parameter that determines how sensitive choices are with respect to expected payoffs. When λ is small, payoff differences are irrelevant and behavior is completely random. At the other extreme, as λ tends to infinity, the decision rule in (2) limits to the perfect-maximization

[4] For instance, Rosenthal (1989) considers a model where choice probabilities are linear in expected payoffs. McKelvey and Palfrey (1995) quantal response equilibrium allows for a general class of choice functions.

rule; the best option is chosen with probability one. Note that (2) is not an explicit solution since the densities $f(p)$ on the left side also appear on the right side through the expected payoff function. Instead (2) defines a first-order differential equation for the distribution $F(p)$, which can be solved analytically (see Lopez-Acevedo (1997); Baye and Morgan (2004)).

Proposition *For* $0 \leq \lambda < 1/(n-1)$,

$$F(p) = 1 - \left(\frac{p_H^{\lambda+1} - p^{\lambda+1}}{p_H^{\lambda+1} - p_L^{\lambda+1}} \right)^{\frac{1}{1-\lambda(n-1)}} \quad \text{for } p \in [p_L, p_H] \tag{3}$$

constitutes a symmetric Luce equilibrium for the Bertrand game with price floor p_{L}. *The distribution of market prices is given by*

$$F_{market}(p) = 1 - \left(\frac{p_H^{\lambda+1} - p^{\lambda+1}}{p_H^{\lambda+1} - p_L^{\lambda+1}} \right)^{\frac{n}{1-\lambda(n-1)}} \quad \text{for } p \in [p_L, p_H]. \tag{4}$$

For $\lambda \geq 1/(n-1)$ *the Luce equilibrium coincides with the Nash equilibrium with both firms choosing prices equal to the price floor* p_{L}, *which is also the market price.*

To see how the Proposition is derived, notice that the denominator in (2) is a constant, $K(\lambda)$, independent of p. Hence, (2) can be written as: $\mathrm{d}F(p)/\mathrm{d}p = K(\lambda) \cdot p^{\lambda} \cdot (1 - F(p))^{\lambda(n-1)}$. It is readily verified that (3) is the unique symmetric solution to this differential equation that satisfy the boundary conditions $F(p_{\mathrm{L}}) = 0$ and $F(p_{\mathrm{H}}) = 1$. The market price is the lowest of the two firms' prices, so its distribution is simply $F_{\mathrm{market}}(p) = 1 - (1 - F(p))^n$.

Note from (3) that when $\lambda = 0$, prices are uniformly distributed on $[p_{\mathrm{L}}, p_{\mathrm{H}}]$. In contrast, as λ tends to $1/(n-1)$, the exponent in (3) diverges to infinity and since the term between the brackets is less than one, the distribution tends to 1 everywhere. In other words, in this limit the Luce model predicts a spike of mass 1 at the Nash prediction $p = p_{\mathrm{L}}$. For higher values of λ, the Nash equilibrium remains the unique symmetric Luce equilibrium.

Note that (in this game) the Nash prediction is not obtained only as a limiting case with perfectly rational players. While the decision rule in (2) limits to the perfect-maximization rule as λ tends to infinity, the Nash prediction requires only that $\lambda \geq 1/(n-1)$. For example, if $n = 2$, the Nash prediction obtains if $\lambda \geq 1/(2-1) = 1$, and $\lambda = 1$ is the case where choice probabilities are proportional to expected payoffs.

Our main interest is in the comparative statics properties of the Luce equilibrium, in particular how changes in the number of players and the price floor affect average and market prices. The explicit solutions in (3) and (4) make a comparative statics analysis straightforward. Note that as n increases the exponents on the right side of (3) and (4) increase, and since the term in the large brackets is less than 1, the distributions $F(p)$ and $F_{\mathrm{market}}(p)$ increase. Likewise, if the price floor p_{L} increases, the term in the large brackets increases, so the distributions $F_{\mathrm{i}}(p)$ and $F_{\mathrm{market}}(p)$ falls. We thus have:

Corollary *For 0*$\leq \lambda < 1/(n-1)$*, the Luce equilibrium predicts that* (i) *an increase in the number of players results in an decrease in average and market prices in the sense of first-degree stochastic dominance, and* (ii) *an increase in the price floor results in an increase in average and market prices in the sense of first-degree stochastic dominance.*

Our experiment tests these comparative statics predictions.

3 Experiment

For the experiment we consider discretized versions of the games described in Sect. 2 (similar to the games of Dufwenberg and Gneezy (2000)). We have a 2×2 design. In low-floor treatments each player (simultaneously) chooses a number from the set $\{1, 2, \ldots, 100\}$; in high-floor treatments each player chooses a number from the set $\{10, 11, \ldots, 100\}$. In duopoly treatments there are two players; in quadropoly treatments there are four players. In each game, the player who chooses the lowest number gets paid in proportion to the number chosen, and the other player(s) get(s) 0. Ties are split. The Nash equilibrium is for each player to choose 1 in the low-floor games,[5] and for each player to choose 10 in the high-floor games.

The duopoly treatments were conducted at the Technion in Haifa. Students were recruited using posters on campus. We had five sessions for each of the two floor treatments. The number of participants was 12 in all sessions, and an extra student assisted us. That is, in total, 120 students participated in the duopoly games. Each session operated for ten periods. In each period six pairs of participants were grouped together according to a random matching scheme, and then each pair played the relevant game.

In each session, after all 13 students entered the experimental room, they received a standard-type introduction, and were told they would be paid 20 New Israeli Shekels (20 NIS; about \$5 at the time of the experiment) for showing up. Then they took an envelope from a box, which contained 13 envelopes. Twelve of the envelopes contained numbers (A1, ..., A12). These were called "registration numbers". We asked participants not to show their registration number to the others. One envelope was labeled "Monitor", and determined who was the person who assisted us and checked that we did not cheat. That person was paid the average of all other subjects in that session.

Each participant then received the instructions for the experiment (see http://econ.arizona.edu/downloads/working_papers/Econ-WP-04-18-updated.pdf), and ten coupons numbered 1, 2, ..., 10. After reading the instructions and asking questions (privately), each participant was asked to fill out the first coupon with her/his registration number and the chosen number – henceforth referred to as a "price",

[5] Also (2, 2) is a Nash equilibrium in the duopoly low-floor game, but the theoretical underpinning is most secure for (1, 1) which is supported by a variety of differently motivated solution concepts: (1, 1) is the game's only *strict equilibrium*; 1 is the only *evolutionarily stable strategy*, the only survivor of *iterated elimination of weakly dominated strategies*, and the game's only *maxminimizer*; $\{1\}$ is the game's only *fully permissible set* (as defined by Asheim and Dufwenberg (2003)). In terms of economic intuition both profiles (1, 1) and (2, 2) are in line with the Bertrand solution: all players make close to zero profit relative to what is available in principle.

although we did not use that word in the instructions – for period 1. Participants were asked to fold the coupon, and put it in a box carried by the assistant. The assistant randomly took two coupons out of the box and gave them to the experimenter. The experimenter announced the registration number on each of the two coupons and the respective prices. If one price was larger than the other, the experimenter announced that the low price won as many NIS, and the high price won nothing. If the prices were equal the experimenter announced a tie, and said that each person won half as many NIS. The assistant wrote this on a blackboard so that all the participants could see it for the rest of the experiment. Then the assistant took out another two coupons randomly, the experimenter announced their content, and the assistant wrote it on the blackboard. The same procedure was carried out for all the 12 coupons. Then the subsequent periods were conducted the same way. After period 10 payoffs were summed up, and participants were paid privately.

The quadropoly treatments were added later. The results of Dufwenberg and Gneezy (2000) suggest that as the number of competitors increase the predictive power of the Bertrand model improves. We thus had a strong prior [independent of part (ii) of our Corollary in Sect. 2] that if the number of competitors were four then prices would be higher with a (high enough) price floor. However, after seeing the somewhat surprising results for the duopoly treatments (reported in Sect. 4.1 below), we felt slightly less confident about this. So, in order to dispel any fear that our conjecture be invalid, we decided run some sessions involving quadropoly Bertrand games, with low and high floors. These treatments were similar to the duopoly treatments, except that there were four competitors. The sessions were conducted at the University of Chicago, with comparable stakes.[6] We had two sessions for each treatment.

4 Results

We next present our results on the impact of price floors in the duopoly games (Sect. 4.1) as well as a comparison with the quadropoly games (Sect. 4.2). We only present some aggregate statistics, and refer to a working paper version for the complete raw data set; see http://econ.arizona.edu/downloads/working_papers/Econ-WP-04-18-updated.pdf

4.1 The duopoly games

From the viewpoint of traditional theory, an increase in the price floor should raise the level of prices. Recall that Nash equilibrium pricing is at 1 (or possibly 2) in the low-floor treatment and at 10 in the high-floor game. The generalized Luce model allows that predictions are systematically biased away from Nash levels; cf. the Proposition and the case where $\lambda < 1/(n-1)$. However, as part (ii) of our

[6] Our subject pool is divided across Haifa and Chicago, but this problem is mitigated in that our main comparison concerns the impact of different floors within a market structure. For each of the duopoly and quadropoly conditions all subjects were recruited in one and the same location. As pointed out by a referee, it seems unlikely that Chicago students respond to price floors very differently from Haifa students.

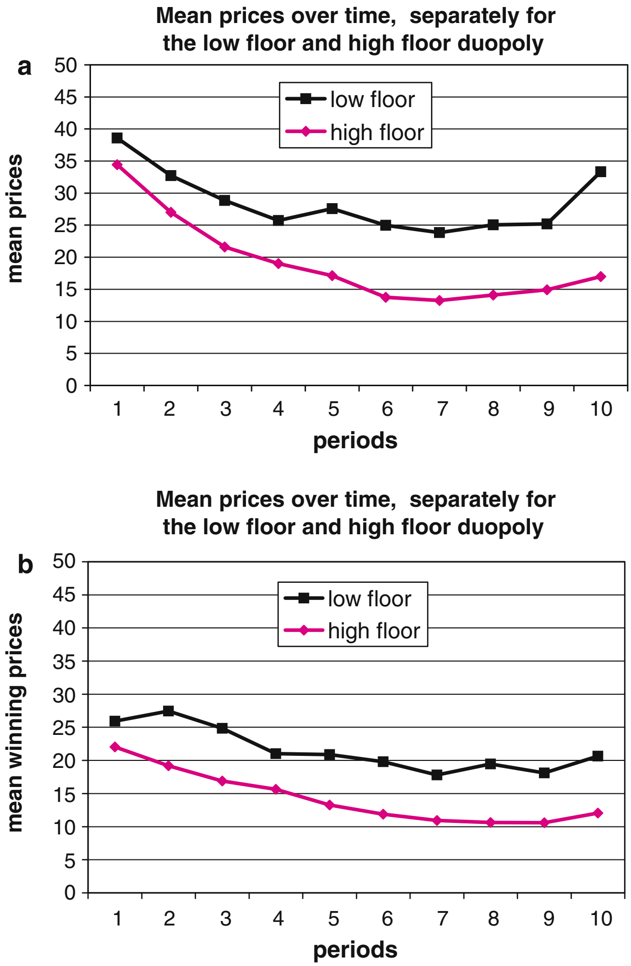

Fig. 1 **a** Mean prices over time in the duopoly treatments. **b** Mean winning prices over time in the duopoly treatments

Corollary shows, the generalized Luce model maintains that a higher price floor leads to less competitive pricing.

Our data do not bear this prediction out. In fact, Fig. 1 suggests the opposite pattern. After the starting periods, all price measures are *higher* in the low-floor treatment than in the high-floor treatment! Moreover, the gaps seem to widen over time.

A closer statistical scrutiny, reflected in Table 1, mainly confirms this picture.[7] The table reports comparisons of mean prices, median prices, mean winning prices, and expected winning price (a recombination measure which takes the average of

[7] p-Values are calculated using the Kruskal–Wallis test instead of the (in experimental economics) more commonly used Wilcoxon test, since the variances in the two treatments are very different (see Siegel and Castellan (1988), pp 137–144).

Table 1 Chi-square results of Kruskal–Wallis test comparing different measures of the duopoly sessions, low- versus high-floor treatments, for each period separately

Measures	Mean price	Median price	Win price	Expected win price
Period 1	0.88	0.011	0.53	0.53
Period 2	0.53	2.73*	2.44	2.45
Period 3	3.15*	3.93**	4.81**	4.81**
Period 4	0	3.93**	1.84	4.81**
Period 5	4.81**	4.81**	4.81**	3.93**
Period 6	6.81***	5.77**	4.81**	5.77**
Period 7	3.93**	5.77**	3.53*	3.15*
Period 8	2.44♣	3.93**	2.45♣	2.45♣
Period 9	3.53*	3.93**	2.45♣	2.45♣
Period 10	4.81**	2.45♣	1.84♣	2.45♣

***Significant at 1%-level; **significant at 5%-level; *significant at 10%-level; ♣$p > 0.1$ but if we add the two low-floor sessions from Dufwenberg and Gneezy (2000) (cf. footnote 8) then there is a significant difference at the 5%-level

winning prices of all possible combinations of matches within a period; cf. Mullin and Reiley (2006)) for the two treatments. For mean and median prices, the differences are not significant in the early periods but mainly become so later on. For mean winning and expected winning prices the tendency is analogous although we get less clear statistical results.[8]

Statistical tests also suggest that the tendency for prices to change over time is different in the treatments. In the high-floor treatment prices decrease in all five sessions ($r = 0.73$ and $p < 0.025$, using Spearman–Rank correlation test). By contrast, in the low-floor treatment, mean prices decrease significantly only in two out of five sessions.

It seems that price floors stimulate competition. This finding becomes even starker if one focuses on the occurrence of prices that are lower than or equal to 10. Figure 2 gives relative frequencies of prices in intervals of 1–10, 11–20, . . ., 91–100, for the final five periods. In the low-floor treatment there are only 9% price choices lower than or equal to 10 (and *none* of those are Nash equilibrium choices of 1 or 2). By contrast, in the high-floor treatment, where price choices of 1–9 are not possible, 40% of the choices are at the equilibrium price of 10. In the final five periods, a whopping 61% of the choices are of 10, while only 13% are at or below 10 in the low-floor treatment.

4.2 Comparing with quadropoly

Figure 3 shows mean prices and mean winning prices over time in the quadropoly treatments. We ran only two sessions of each treatment, which is too little to admit

[8] In the last periods, in one low-floor session the mean winning price is below 10, thus lower than any high-floor session and therefore this session gets the lowest rank. Having only five sessions for the two treatments, the Kruskall–Wallis test does not reject the null-hypothesis of no difference. In Dufwenberg and Gneezy (2000) winning prices are far above 10 in their two sessions with a similar low-floor treatment. If we include these two sessions in our analysis (having 5 high-floor and 7 low-floor sessions) we obtain a significant difference between the two treatments at the 5% level.

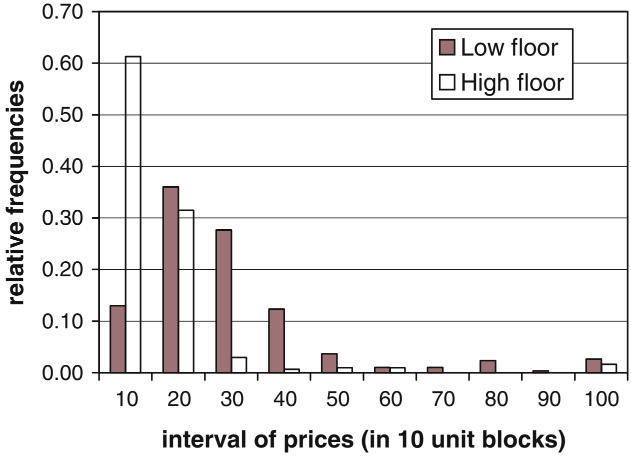

Fig. 2 Relative frequencies of prices for the last five periods in the duopoly treatments

statistical testing.[9] However, after round seven, average winning bids in the low-floor treatment are below 10. We see no reason to revise our strong prior that price floors in quadropoly markets raise prices.

Note, however, that our quadropoly and duopoly high-floor treatments share the feature of making play at the Nash equilibrium more common than in the low-floor treatments. In this sense, a high price floor seems to foster competition in duopoly and quadropoly alike.

Part (i) of our Corollary predicts that an increase in the number of competitors reduces prices for any level of the price floor. Our data goes well with this prediction (as do the prior results of Dufwenberg and Gneezy (2000)). Since we only ran two session of each quadropoly treatment we do not provide formal test results however.

5 Discussion

From the viewpoint of traditional theory, raising price floors in Bertrand models protects competitors from making low profits, and should thus be anti-competitive. With our experiment we have shown that the opposite can be true: a higher price floor may foster competition and may lead to lower prices under conditions of duopoly.[10]

[9] The justification is that, as explained towards the end of Sect. 3, based on the results of Dufwenberg and Gneezy (2000) we already had a strong prior [independent of part (ii) of our Corollary in Sect. 2] that prices would be higher with than without a floor if the number of competitors were four rather than two.

[10] A caveat here is that this result may depend on other design details. An example concerns the specific nature of the information feedback between rounds. In our design after each round each participant was informated about co-players' choices as well as about all other games in that round of the session. Isaac and Walker (1985), Dufwenberg and Gneezy (2002), and Ockenfels and Selten (2005) present evidence indicating that market outcomes may be more competitive if there is less information feedback.

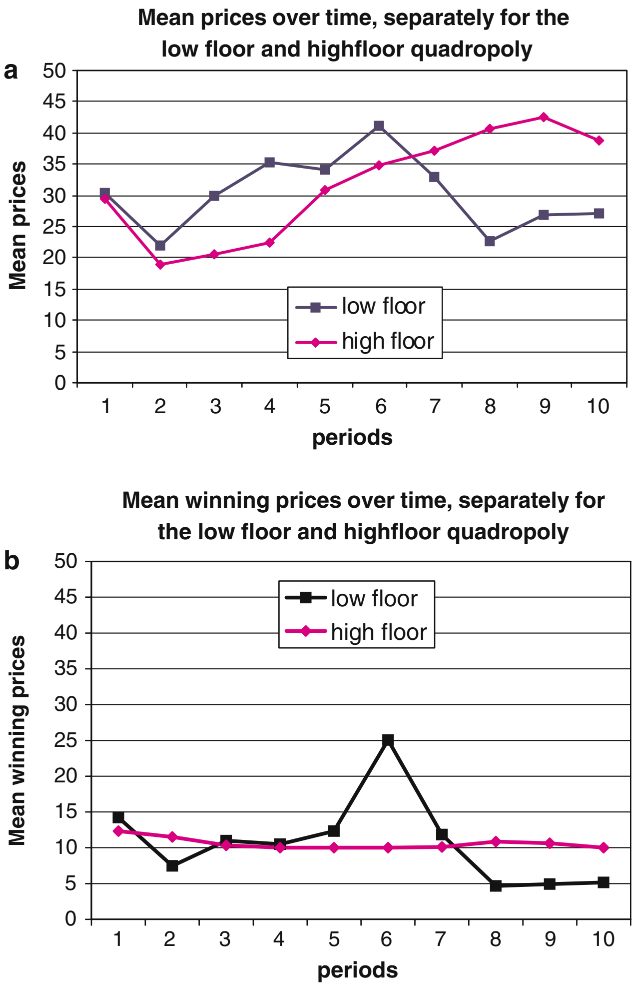

Fig. 3 **a** Mean prices over time in the quadropoly treatments. **b** Mean winning prices over time in the quadropoly treatments

Our results may highlight a possible weakness of economic models in which there is little incentives for decision makers to stick with the equilibrium. Bertrand duopoly competition is just one example. In response, one may be lead to consider models which incorporate the idea that the likelihood of an error is related to the cost of making that error.

In this vein, we considered a generalized version of the Luce (1959) model which assumes that the probability of making a particular choice is proportional to the expected payoff associated with that choice raised to the power of a parameter $\lambda \in [0, \infty)$ ('rationality increases' with λ). For appropriate choices of λ the model can predict supra-Nash pricing, but it fails to capture the overall pro-competitive effect of a high floor seen for duopolies.

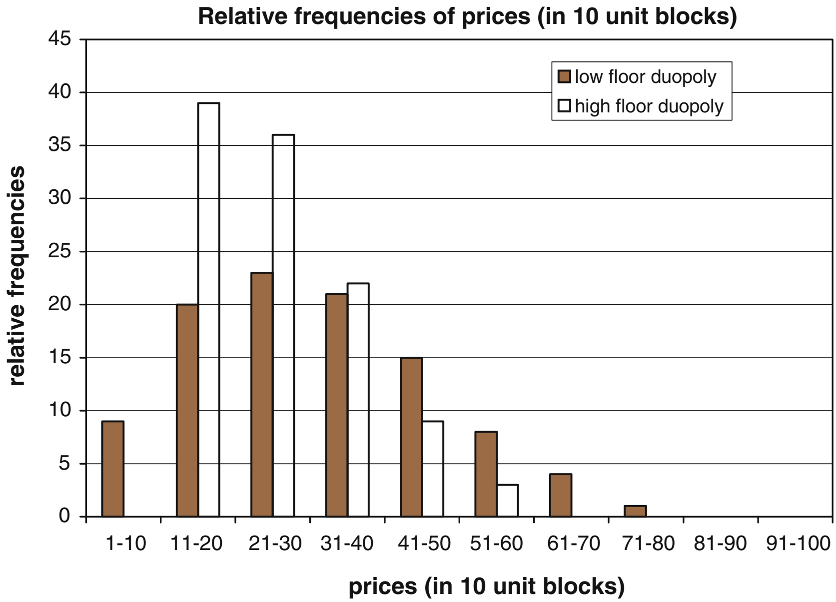

Fig. 4 Relative frequency of prices (in ten period blocks) according to the simulation of the Luce model with λ independently estimated for each duopoly treatment separately

How should one react to this finding? In closing this paper we suggest two rather different lines of future research. The first one attempts to 'fix things' within the generalized Luce framework, whereas the second one addresses a radically different idea.

Reaction 1: *Endogenize* λ

To motivate this approach, let us first illustrate how well the generalized Luce model can describe our data if we allow *different* values of λ across treatments. Applying standard maximum-likelihood techniques to the pricing data of the low-floor duopoly treatment, yields an estimate of $\lambda = 0.85$ (0.01), where the number in parentheses denotes the standard error.[11] The light bars of the histogram in Fig. 4 show the predictions corresponding to this λ estimate. Analogously, for the high-floor duopoly treatment the maximum-likelihood estimate is $\lambda = 0.95$ (0.01).[12] The dark bars in Fig. 4 represent the predictions for the high-floor treatment. Note that the prediction histograms in Fig. 4 are remarkably similar to the actual data histograms in Fig. 2.[13]

The higher λ estimate for the high-floor treatment is not surprising, but the different λ-estimates do raise the question "what determines the precision

[11] The loglikelihood −2309.2 on 600 observations.

[12] The loglikelihood −2340.9 based on 600 observations.

[13] Haile et al. (2004) have suggested that *any* observed data set can be explained by estimating a different λ for each treatment. Goeree et al. (2005) show that this is not true for *regular* decision-error models where choice probabilities are monotone in expected payoffs, as in the Luce model studied here. They show that regular decision-error models impose strong empirical restrictions on the data even without any restrictions on the precision parameters.

parameter?" In most applications of the Luce model (or related decision-error models), λ is assumed to be a constant exogenous parameter. One might argue, however, that λ is better seen as an endogenous variable.[14] For example, experimental data often show that precision rises with experience (learning), i.e., using only data from later periods of the experiment typically yields higher λ estimates. This is no different in our data, where both in the low-floor and the high-floor treatment, estimated precision increases over time. Learning is more difficult, however, in the low-floor treatment where behavior is more volatile, which partly could explain the lower value of λ. Second, subjects may weight the benefits of being more precise against the higher (mental) costs. Since the benefits differ across the treatments, this results in different levels of the optimal precision parameter.

Reaction 2: *Examine the available 'punishments'*

An interesting alternative may be to explore what factors other than costs of error that may be influencing choice. Reinhard Selten suggested to us a behavioral idea that could be relevant (cf. also Murphy 1966, who touches on related themes). Perhaps players conceive of the general level of payoffs at some Nash equilibrium as a 'threat', the size of which affects the 'mode' by which they play? In our low-floor duopoly treatment the equilibrium payoffs are low and the threat severe, so the subjects may attempt to avoid entering a competitive mode. With a higher price floor, the equilibrium payoffs are higher, the threat less severe, and the players may see less long-run reason to avoid initiating price wars.

References

Asheim, G.B., Dufwenberg, M.: Admissibility and common belief. Games Econ Behav **42**, 208–234 (2003)

Baye, M., Morgan, J.: Price dispersion in the lab and on the internet: theory and evidence. RAND J Econ **35**, 449–466 (2004)

Bertrand, J.: Review of Theorie Mathematique de la Richesse Sociale and Recherches sur les Principes Mathematique de la Theorie des Richesse. J Savants **48**, 499–508 (1883)

Brown-Kruse, J., Rassenti, S., Reynolds, S., Smith, V.: Bertrand-Edgeworth competition in experimental markets. Econometrica **62**, 343–371 (1994)

Cason, T.: Cheap talk price signaling in laboratory markets. Inf Econ Policy **7**, 183–204 (1995)

Cason, T., Davis, D.: Price communications in a multi-market context: an experimental investigation. Rev Ind Organ **10**, 769–787 (1995)

Dufwenberg, M., Gneezy, U.: Price competition and market concentration: an experimental study. Int J Ind Organ **18**, 7–22 (2000)

Dufwenberg, M., Gneezy, U.: Information disclosure in auctions: an experiment. J Econ Behav Organ **48**, 431–444 (2002)

Fouraker, L., Siegel, S.: Bargaining Behavior. New York: McGraw-Hill (1963)

Goeree, J., Holt, C., Palfrey, T.: Regular quantal response equilibrium. Exp Econ **8**, 347–367 (2005)

Haile, P., Hortacsu, A., Kosenok, G.: On the Empirical Content of Quantal Response Models. Mimeo (2004)

Holt, C.: Industrial organization: a survey of laboratory research. In: Kagel, J., Roth, A. (eds.) Handbook of Experimental Economics, pp. 349–444. Princeton: Princeton University Press (1995)

Huck, S., Normann, H.-Th., Oechssler, J.: Does information about competitors actions' increase or decrease competition in experimental oligopoly markets? Int J Ind Organ **18**, 39–57 (2000)

[14] See, e.g., McKelvey et al. (1997) and Rogers (2000).

Ippolito, P.M.: Resale price maintenance: empirical evidence from litigation. J Law Econ **34**, 263–294 (1991)
Isaac, M., Plott, C.: Price controls and the behavior of auction markets: an experimental study. Am Econ Rev **71**, 448–459 (1981)
Isaac, M., Walker, J.: Information and conspiracy in sealed-bid auctions. J Econ Behav Organ **6**, 139–159 (1985)
Lopez-Acevedo, G.: Quantal response equilibria for posted offer markets. Estud Econ **12**, 95–131 (1997)
Luce, D.: Individual Choice Behavior. New York: Wiley (1959)
Mason, C., Phillips, O.: Information and cost asymmetry in experimental duopoly markets. Rev Econ Stat **79**, 290–299 (1997)
McKelvey, R., Palfrey, T.: Quantal response equilibria for normal form games. Games Econ Behav **10**, 6–38 (1995)
McKelvey, R., Palfrey, T., Weber, R.: Endogenous Rationality Equilibrium. Mimeo (1997)
Mullin, C., Reiley, D.: Recombinant estimation for normal-form games, with applications to auctions and bargaining. Games Econ Behav **54**, 159–182 (2006)
Murphy, J.: Effects of the threats of losses on duopoly bargaining. Q J Econ **80**, 296–313 (1966)
Ockenfels, A., Selten, R.: Impulse balance equilibrium and feedback in first price auctions. Games Econ Behav **51**, 155–170 (2005)
Plott, C.: An updated review of industrial organization: applications of experimental economics. In: Schmalensee, R., Willig, R. (eds.) Handbook of Industrial Organization, vol. II, pp. 1109–1176. Amsterdam: North Holland (1989)
Rogers, B.: Games That Make You Think. Undergraduate Thesis, University of Virginia (2000)
Rosenthal, R.: A bounded rationality approach to the study of non-cooperative games. Int J Game Theory **2**, 65–67 (1989)
Selten, R., Apesteguia, J.: Experimentally observed imitation and cooperation in price competition on the circle. Games Econ Behav **51**, 171–192 (2005)
Siegel, S., Castellan, J.: Nonparametric Statistics for the Behavioral Sciences. New York: McGraw-Hill (1988)
Smith, V., Williams, A.: On non-binding price controls in a competitive market. Am Econ Rev **71**, 467–474 (1981)
Tirole, J.: The Theory of Industrial Organization. Cambridge: MIT Press (1994)