

# INSTABILITY OF ENDOGENOUS PRICE DISPERSION EQUILIBRIA: A SIMULATION

Lucas Herrenbrueck  
Simon Fraser University

November 2015

## Abstract

Models of price posting by firms and search by consumers often feature equilibria with endogenous price dispersion. However, such equilibria are fragile in the sense that an individual firm's decision problem is not concave (or even well-formed) outside of equilibrium. I simulate various procedures firms may use to update their prices sequentially. I find that the benchmark model performs surprisingly well, but that substantial differences in the level, volatility, and dispersion of profits remain even with as many as 100 firms. Certain procedures for price updating lead to "odd" results such as cyclicalities or excessive volatility of prices and profits. When cost dispersion becomes large, prices become less volatile as they are more closely tied to costs.

**JEL codes:** D21, D43, D83

**Keywords:** Simulation; search frictions; Burdett-Judd pricing; Edgeworth cycles; disequilibrium dynamics

**Contact:** herrenbrueck@sfu.ca, +1-778-782-4805  
Simon Fraser University, Department of Economics  
8888 University Drive, Burnaby, B.C. V5A 1S6, Canada

I am grateful to Paul Bergin, Giacomo Bonanno, Athanasios Geromichalos, Martine Quinzii, Katheryn Russ, Burkhard Schipper, Joaquim Silvestre, and Ina Simonovska for their helpful comments and suggestions.

# 1 Introduction

Economic models almost universally focus on equilibrium. This focus is justified to the extent that at least over a sufficiently long time horizon, beliefs converge to facts and best response opportunities have been consummated. But partly, it reflects the limitations of mathematical modeling. Disequilibrium is hard to define and harder to compute. The stability of equilibrium, how disequilibrium converges to equilibrium, and the properties of long-run and large-scale aggregates are therefore important questions of economic theory.

In this paper, I consider a standard model of endogenous price dispersion and simulate it with a finite number of firms. The firms are identical and compete for the sale of a homogeneous good, like in a Bertrand model, which would tend to drive prices down to marginal costs. But consumers face search frictions: at a given time, they will only observe a finite menu of prices, which means that firms can make positive profits by charging a very high price and selling only to consumers who do not have a better option as part of their menu. This tension implies that *price dispersion is present in any equilibrium*, and the distribution of prices will adjust endogenously to make firms indifferent between charging high prices for low sales, and charging low prices for high sales.

However, with a finite number of firms, it is impossible to make firms *exactly* indifferent. Firms would like to charge prices just below but as close as possible to their direct competitors, which, depending on the exact form of the strategy space, could lead to repeated cycles of under-cutting followed by large price rises (Edgeworth cycles). This will make the prices of individual firms volatile, but the same does not need to be true for aggregates like the price distribution, and the level and dispersion of profits.

I therefore simulate various procedures firms may use to update their prices to answer the following questions: (1) Starting from a random price distribution, is there “tâtonnement”, i.e. convergence to an equilibrium? (2) The benchmark model of a continuum of firms predicts a stable equilibrium where firms make equal and constant profits. Does the simulation converge to this benchmark equilibrium? (3) How do the answers to (1) and (2) depend on the number of firms in the simulation?

The results can be summarized as follows. Convergence happens only to an extent, as there seems to be a stationary “attractor set” of price distributions. The average level of profits is higher than in the benchmark model even with very many competing firms, this level is volatile over time, and there is substantial profit dispersion. In the main simulation, the results do conform closer to the benchmark as the number of firms increases, but slowly. In an alternative specification in which firms are extremely myopic, the divergence from the benchmark equilibrium is very stark and does not improve when the number of firms increases.

The concern over the stability and convergence properties of economic equilibrium is not new. The oldest approach is Walrasian tâtonnement.<sup>1</sup> There has also been some work on models with dispersed price equilibria. Hopkins and Seymour (2002) approach a similar problem theoretically and find that the dispersed equilibrium can be stable under learning if competition is not too severe. In several simulations and experiments, however, Cason, Friedman, and Wagener (2005) found that firms’ actions did not converge to a dispersed price equilibrium, but exhibited cycles of under-cutting followed by large price rises (Edgeworth cycles), and the amplitude of these cycles did not tend to decrease over time.

The idea that every firm would like to undercut its competitors by the smallest

---

<sup>1</sup> A good summary of the theory of Walrasian tâtonnement is provided in Chapters 17 and 20 of Mas-Colell, Whinston, and Green (1996). In recent work, Crockett, Oprea, and Plott (2011) provide experimental support of the theory.

possible margin, the “Bertrand paradox”, has been studied in a literature too large to survey here. A relevant example is [Maskin and Tirole \(1988\)](#), who demonstrate that Edgeworth cycles can arise as equilibria of dynamic Bertrand games with a small number of firms, supported by a folk theorem. A key difference in the noisy search approach is that marginal cost pricing is not an equilibrium even with infinitely many firms, because positive profits can be made by charging a very high price and attracting a small, but positive, measure of consumers. If the number of firms is finite but the menu of possible prices is a continuum, no steady-state equilibrium exists at all. Edgeworth cycles seem to be a reasonable guess as to how such a market would evolve in practice.

## 2 Benchmark model

The model is a partial equilibrium version of [Burdett and Judd \(1983\)](#).<sup>2</sup> There is a unit measure of firms, and a measure of consumers described later. Consumers value a good that each firm can produce at constant marginal cost  $c$ . Consumers have one unit of money, and they will spend all of it unless the price exceeds a reservation price  $\bar{p}$ .<sup>3</sup> Meetings between consumers and firms are subject to search frictions: firms post prices before meetings take place, consumers observe a random number of prices (“receive quotes”) and spend all of their money on the firm with the lowest price, unless that price exceeds the reservation price.

The random number of price quotes,  $K$ , can be described by the probability mass function  $q_k = Pr\{K = k\}$  with support  $\{0, \dots, \bar{k}\}$ . For example, [Burdett and Judd](#)

---

<sup>2</sup> The approach taken by Burdett and Judd, and by [Burdett and Mortensen \(1998\)](#) in a related paper, is often termed *random search* because the probability that a buyer meets (or more appropriately in this context, observes) a particular seller is independent of seller characteristics, such as price, capacity, or queue length. [Geromichalos \(2014\)](#) studies the Bertrand paradox, and a general class of resolutions of the paradox, in a model based on the alternative paradigm, *directed search*.

<sup>3</sup> This demand curve arises optimally in [Head and Kumar \(2005\)](#). Here, I use it for simplicity.

(1983), Head et al. (2012), and Wang (2014) assume that consumers receive either 1 or 2 quotes. Others, following Mortensen (2005), assume that  $K$  follows a Poisson distribution. For reasons of generality, I prefer the flexible form where  $K$  follows a negative binomial distribution with parameters  $1/\rho$  and  $\eta\rho/(1 + \eta\rho)$ . This means that the expected number of quotes is  $\eta$ , and higher  $\rho$  increases the dispersion of both quoted and transaction prices.<sup>4</sup> For example, the limit  $\rho \rightarrow 0$  corresponds to Poisson, and  $\rho = -0.5$  with  $\bar{k} = 2$  corresponds to the original Burdett-Judd model.

In order to normalize the measure of matched consumers to 1, assume that the overall measure of consumers is  $(1 - q_0)^{-1}$ .

Let the CDF of posted prices be  $F(p)$ . No firm would find it optimal to charge more than the reservation price in, which implies  $F(\bar{p}) = 1$  in equilibrium. Therefore, every consumer will make a purchase, and will do so from the firm with the lowest price. The CDF of transaction prices will be  $J(F(p))$ , where  $J$  is the function:

$$J(F) = \sum_{k=0}^{\infty} q_k [1 - (1 - F)^k] = \frac{1 - (1 + \rho\eta F)^{-1/\rho}}{1 - (1 + \rho\eta)^{-1/\rho}} \quad (1)$$

If  $F$  has no mass points, a firm charging price  $p \leq \bar{p}$  can expect to sell to  $J'(F(p))$  consumers. Recall that each customer has unit elastic demand so that the *expenditure* is fixed, as long as the price does not exceed the reservation price. Therefore, the profits of a firm with price  $p \leq \bar{p}$  satisfy:

$$\pi(p) = \left(1 - \frac{c}{p}\right) J'(F(p)) \quad (2)$$

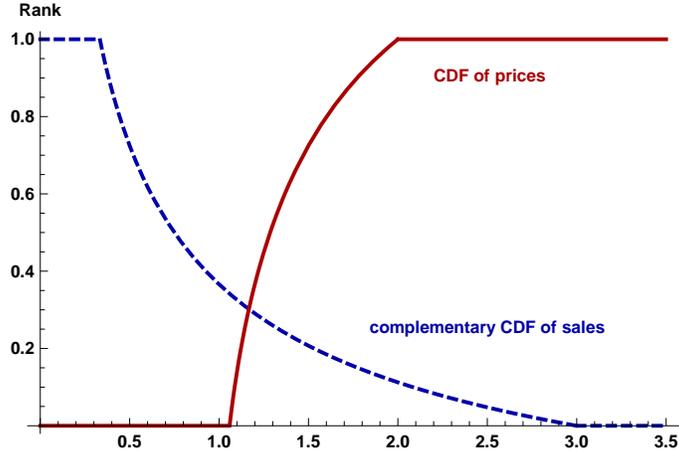
Burdett and Judd (1983) prove that endogenous price dispersion is the only equilibrium as long as both  $q_1 > 0$  and  $q_2 > 0$ , and that some firms do charge the reser-

---

<sup>4</sup> The parameters need to satisfy  $\eta > 0$  and  $\rho > -1/\eta$ . If  $\rho \geq 0$ ,  $\bar{k} = \infty$ , otherwise  $\bar{k} = -1/\rho$ . This class of functional forms for the matching process is derived in detail in Herrenbrueck (2015).

Figure 1: The endogenous price and sales distributions

For firms of any rank, profits (equal to sales multiplied by the profit margin) are constant.



vation price ( $F(p) < 1$  for  $p < \bar{p}$ ). But because all firms are the same ex-ante, they must be indifferent between charging any price on the support of  $F$ ; in particular, they must be indifferent to charging  $\bar{p}$  and making the minimal sales  $J'(1)$ .

$$\left(1 - \frac{c}{p}\right) J'(F(p)) = \left(1 - \frac{c}{\bar{p}}\right) J'(1)$$

Using this equation, we can solve for the (partial) equilibrium  $F(p)$ , taking  $c$  and  $\bar{p}$  as given:

$$F(p) = \frac{1 + \rho\eta}{\rho\eta} \left[ \frac{1 - c/p}{1 - c/\bar{p}} \right]^{\frac{\rho}{1+\rho}} - \frac{1}{\rho\eta} \quad (3)$$

for  $p \in [\underline{p}, \bar{p}]$ , where  $\underline{p} = \arg \sup_p \{F(p) = 0\}$ .

### 3 Simulation models

There are  $N$  firms. Each simulation consists of a random initial state, representing these firms' prices, and an updating procedure to map today's state into tomorrow's.

I consider two different updating procedures to allow for flexibility regarding how firms change their prices and what they know. Simulation time is discrete.

Non-atomic firms are an essential ingredient of the Burdett-Judd model, because only a non-atomic price distribution can support equilibrium (corner cases where  $q_1$  equals 0 or 1 aside). The reason is that the *distance* between observed prices does not affect consumer choice: consumers make all their purchases from the cheapest firm that they observe. Consequently, any firm would like to price as close to the next-most expensive firm as possible, but strictly below. This is similar to the Bertrand paradox, but with one wrinkle: because not all consumers observe more than one price, there are always positive profits to be made by charging the reservation price. Unlike in the Bertrand model, there exists no Nash equilibrium with a discrete number of firms.

So in order to make sure every firm updating its price has a unique maximizer, I add some structure to the environment, and I investigate four different simulation models. The first model uses the following approach: every simulation period, a single firm learns what its profits would be if they charged a particular price randomly chosen from a continuum. If these profits exceed their current profits, they switch to the new price, otherwise they keep the current price. The other three models provide variations on this theme. In detail:

### **Model 1**

The set of possible prices is the interval  $[c, \bar{p}]$ . Firms begin with prices randomly selected from the interval. In random sequential order, they draw one potential price from the same interval. They compute the rank of both their current price and of the potential new price (the rank is defined as  $N$  for the cheapest price and 1 for the most expensive), and compare profits according

to the following formula:

$$\pi(p) = \left(1 - \frac{c}{p}\right) \left[ J\left(\frac{N - \text{rank}(p) + 1}{N}\right) - J\left(\frac{N - \text{rank}(p)}{N}\right) \right]$$

Their price rank is well-defined because almost-surely no two firms charge the same price. They switch to the new price if and only if profits would be higher than at the current price.<sup>5</sup>

## Model 2

The environment is the same as in Model 1, but in each time period, all firms make their choices at once, and in a very unstructured manner: they take into account the effects on the price distribution of their own potential switch, but not the fact that other firms may switch, too. (It would be possible to design a procedure where only a few firms switch at a time without taking each other's choices into account, but this wholesale switching procedure provides an extreme point of comparison.)

## Model 3

At the beginning of time, each firm is permanently assigned a cost parameter from the interval  $[c, \bar{c}]$ , where  $\bar{c} < \bar{p}$ . The procedure for updating prices is the same as in Model 1. This setup allows us to see how closely the ranking of prices tracks the ranking of costs in this environment.

## Model 4

All firms have the same cost  $c$ , and they can choose prices freely from a discrete grid of 10 possible prices  $\{p_i\}$ , spaced equally between  $c$  (excluded) and  $\bar{p}$  (included). In random sequential order, firms observe the existing price distribu-

---

<sup>5</sup> This choice is myopic because the updating firm does not anticipate that other firms will be changing their prices, too, before its next opportunity to update.

tion; then for each potential choice  $p_i$ , they compute the *new* price distribution, if they did make that choice:

$$f_j(i) = \text{share of firms charging } p_j \quad \forall j = 1 \dots P$$

$$F_j(i) = \sum_{k=1}^j f_k \quad \forall j = 1 \dots P$$

$$F_0(i) = 0$$

and the implied profits, if they did make that choice:

$$\pi(i) = \left(1 - \frac{c}{p_i}\right) \frac{J(F_i(i)) - J(F_{i-1}(i))}{N f_i(i)}$$

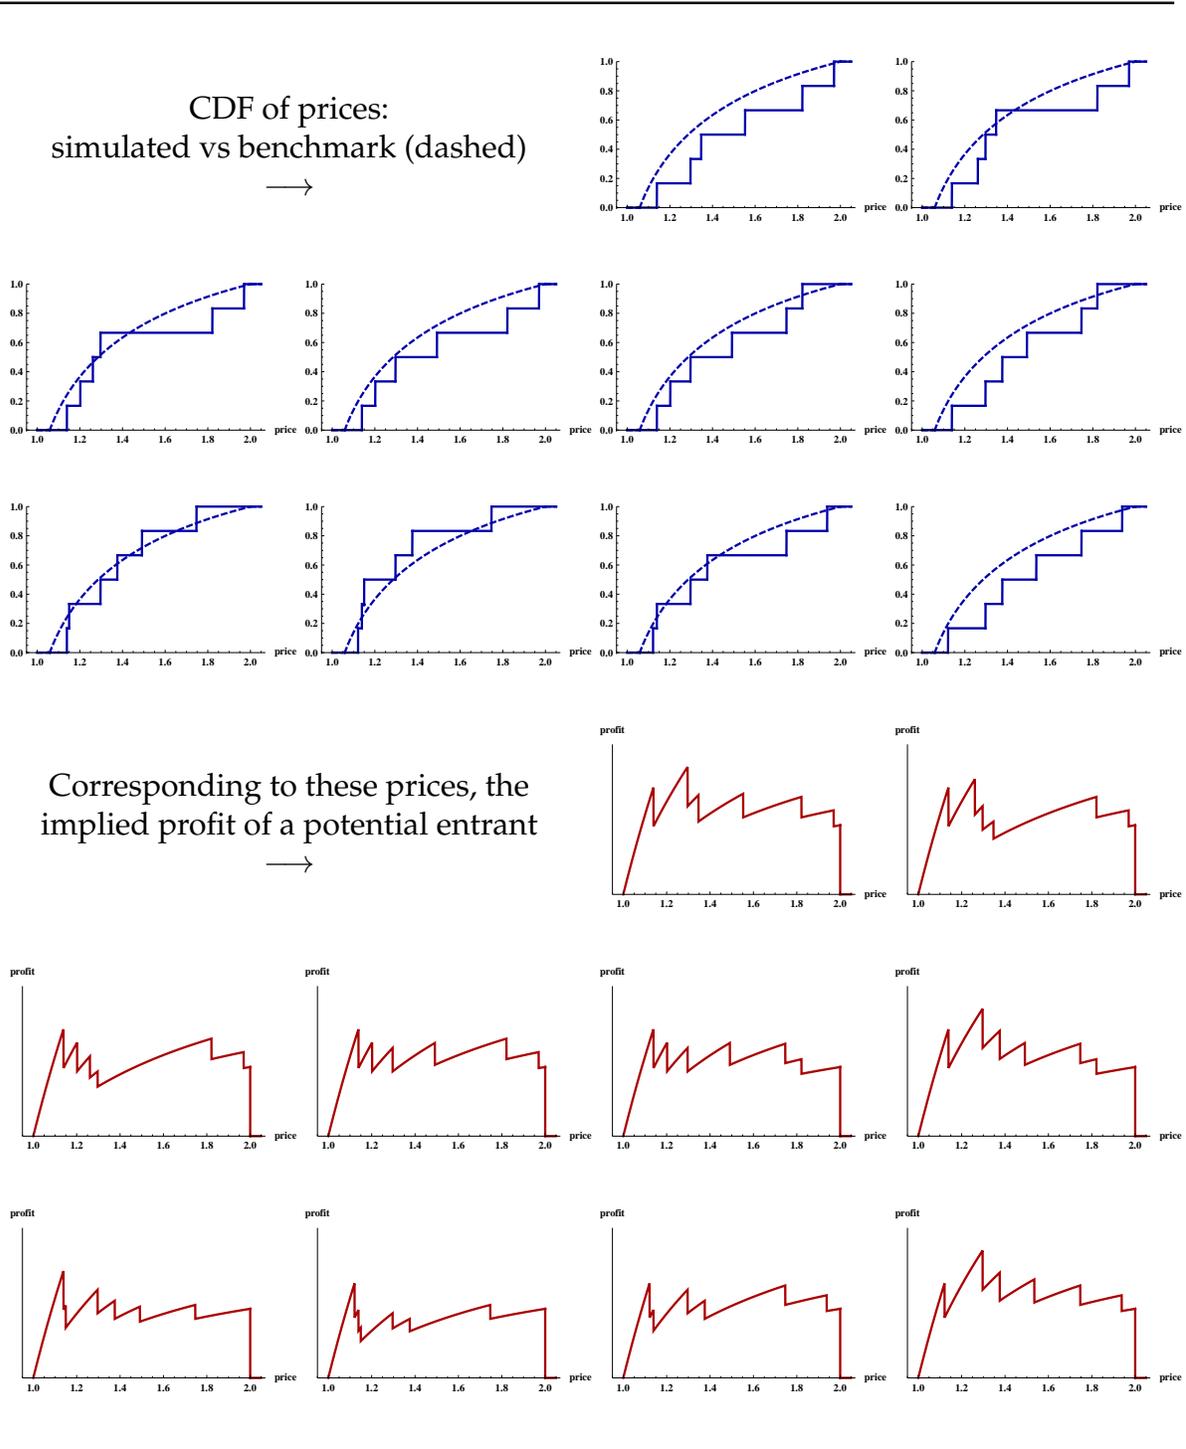
Finally, they simply pick the price with the highest profits  $\max_i \{\pi(i)\}$ .<sup>6</sup>

The reservation price is set to  $\bar{p} = 2$ , and the matching process is parametrized by mean  $\eta = 2$  and shape  $\rho = 1$ . Models 1 and 2 are simulated for  $N = 6$ ,  $N = 16$ ,  $N = 40$ , and  $N = 100$ , with cost  $c = 1$  for all firms. (Model 1 is also simulated for  $N = 625$  to provide an extreme point of comparison.) Model 3 is simulated for  $N = 16$ , with costs equally spaced between  $c = 1$  and one of  $\bar{c} = 1.1$ ,  $\bar{c} = 1.4$ , and  $\bar{c} = 1.9$ . Finally, Model 4 is simulated for  $N = 16$ ,  $N = 40$ , and  $N = 100$ . Each model is simulated for 21,000 periods, and after discarding the first 1000 to make sure the results do not depend on the random initial state, each reported simulation consists of 20,000 periods.

---

<sup>6</sup> This procedure is a version of “fictitious play”, whereas the benchmark procedure is closer to “reinforcement learning” (minus the reinforcement, of course, because firms do not keep a memory of the relative profitability of older actions).

Figure 2: Excerpt from the simulation



Excerpt from the simulation of Model 1 with  $N = 16$ , showing a particular sequence of 10 consecutive price updates. Simulation steps where a firm chose not to update are omitted.

## 4 Simulation results

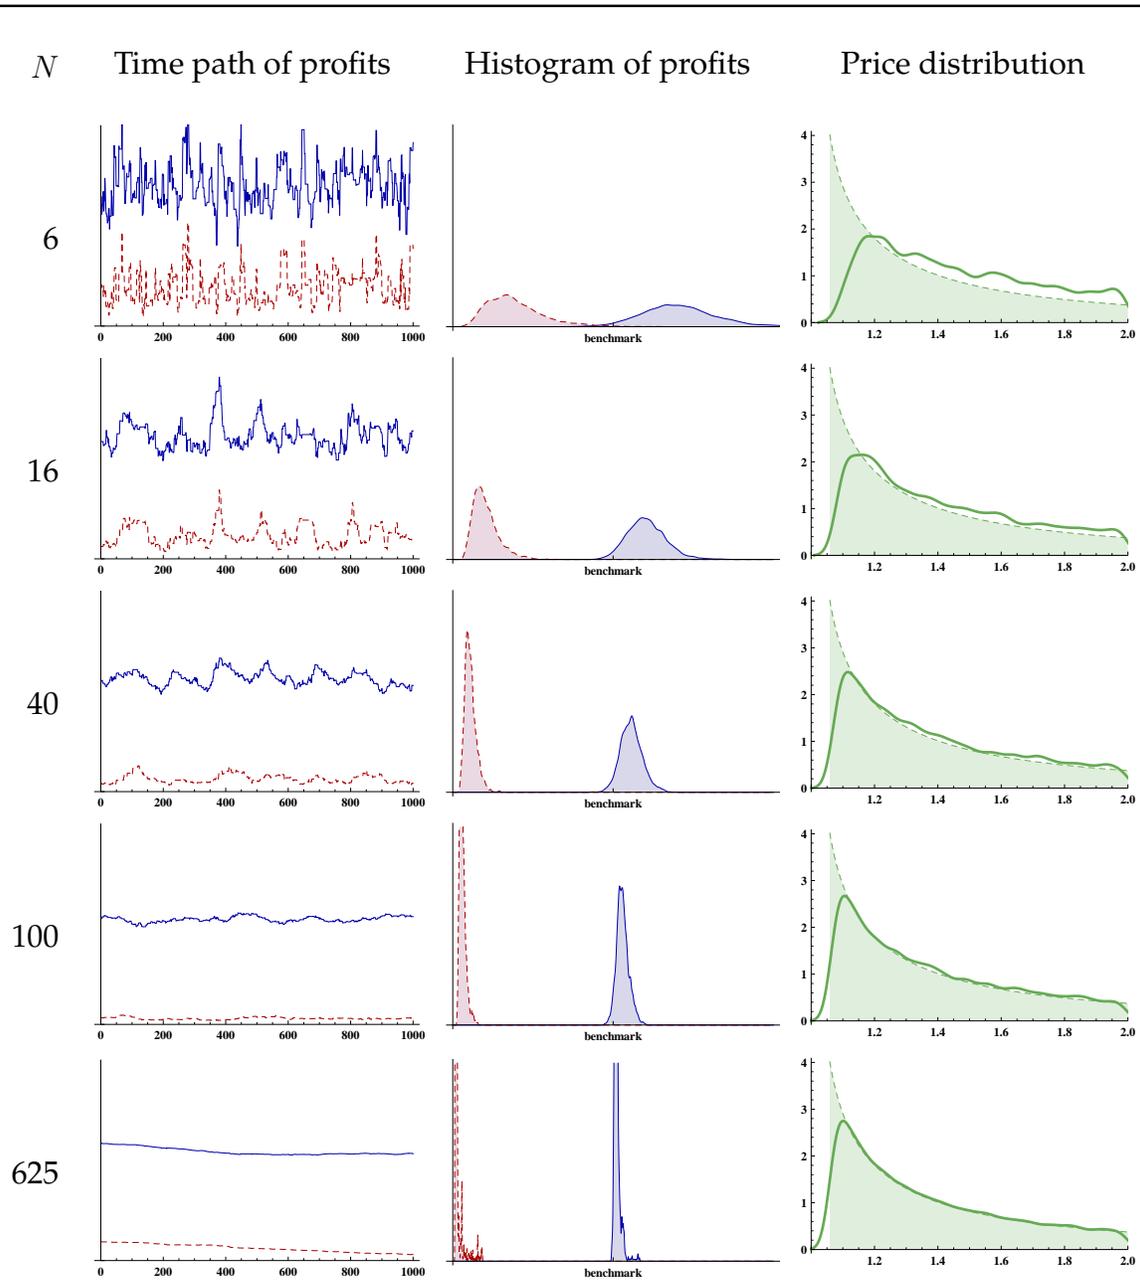
Figure 2 shows a particular excerpt from the simulation of Model 1 with  $N = 16$  to illustrate the principle. We can see that the prices of individual firms never converge because there are always switching opportunities, no matter how infinitesimal. The cumulative distribution functions (CDFs) of prices do not converge to a limit function, either, as there is always churn in the price ranking.

However, it could still be the case that the benchmark model provided an accurate representation of *average* prices and profits over time, even if they differed in any single period. If that was the case, then the long-run price distribution should be well approximated by equation (3). Total profits (since the measure of firms is 1 in the benchmark model) should be approximately equal to  $(1 - c/\bar{p}) J'(1)$  and should not vary much over time. Cross-sectional dispersion of profits should be close to zero in every period.

Detailed results are reported in Table 1 and illustrated in Figure 3. The benchmark model performs surprisingly well in predicting the long-run averages of Model 1. Simulated profits are higher on average than in the benchmark model, and exhibit substantial volatility and dispersion, but all of these statistics fall as the number of firms becomes large, and the average price distribution converges to the one predicted by Equation (3) – especially when the number of firms is extremely large (625 in the simulation).

However, even with 100 firms, average profits remain at least 5% above the benchmark, long-run volatility measured by the coefficient of variation of total profits over time is positive, and so is profit dispersion measured by the average over time of the cross-sectional coefficient of variation. These numbers are even higher for smaller numbers of firms, suggesting that the benchmark model does a poor job

Figure 3: Order (profits and prices in Model 1)



Cross-section **mean** and **standard deviation** (dashed) of profits, truncated to the first 1000 simulation periods.

Smoothed histograms of the cross-section **mean** (right) and **standard deviation** (dashed, left) of profits.

Density functions of prices pooled over all simulation periods, compared with the benchmark equilibrium (shaded area).

of modeling small-market competition, even averaged over time. The correlation over (simulated) time between the cross-sectional mean and coefficient of variation of profits is robustly positive, which suggests that variation in profits is driven by outliers to the top and not to the bottom.

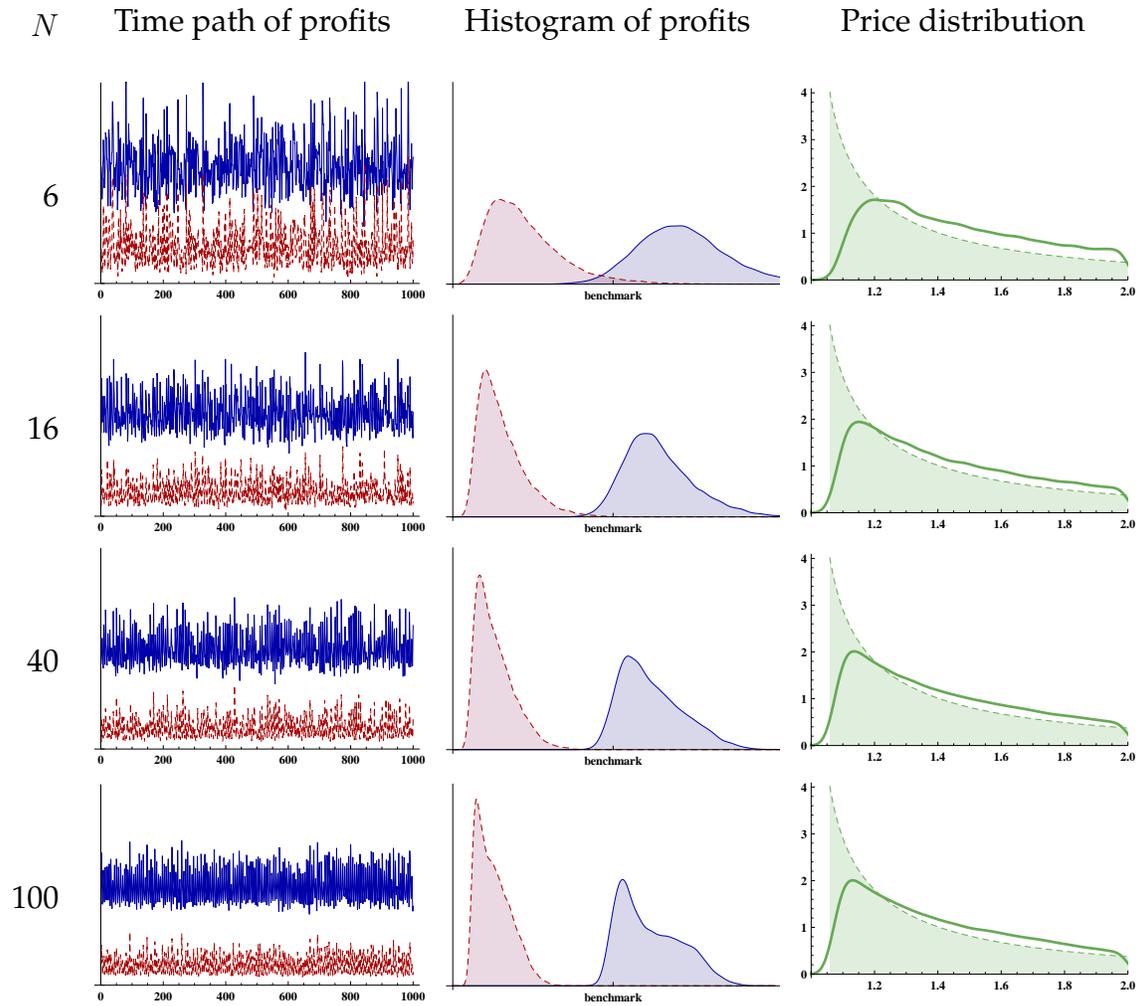
Table 1: **Profits in Models 1, 2, and 4**

Model	$N$	Level	Volatility	Dispersion	L-D correlation
1	6	0.236	0.039	0.267	0.56
	16	0.201	0.020	0.167	0.46
	40	0.186	0.011	0.099	0.39
	100	0.177	0.006	0.057	0.43
	625	0.171	0.003	0.026	0.79
2	16	0.214	0.035	0.229	0.63
	40	0.206	0.039	0.205	0.72
	100	0.205	0.032	0.201	0.81
4	16	0.198	0.024	0.126	0.43
	40	0.184	0.013	0.068	0.35
	100	0.177	0.007	0.033	0.31
Benchmark	$\infty$	0.167	0.000	0.000	–

“Level” is the cross-sectional total of firms’ profits in a simulated period, averaged across simulated time. “Volatility” is the time-series standard deviation of profit levels. “Dispersion” is the cross-sectional coefficient of variation of firms’ profits in a simulated period, averaged across simulated time. The “L-D correlation” is the time-series correlation between the level of profits and their dispersion (measured by the coefficient of variation to avoid a trivial correlation between means and standard deviations) in each simulated period.

The simulation of Model 2 exhibits stranger behavior, illustrated in Figure 4. All firms draw a new potential price simultaneously. Without knowing each others’ choices, they decide whether to stay or to switch, and the simulation assumes that they ignore that other firms are also changing their prices. As one would expect, volatility and dispersion are very high in this environment, and they do not decrease with a higher number of firms, because the entire set of firms updates

Figure 4: Chaos (profits and prices in Model 2)



Cross-section **mean** and **standard deviation** (dashed) of profits, truncated to the first 1000 simulation periods.

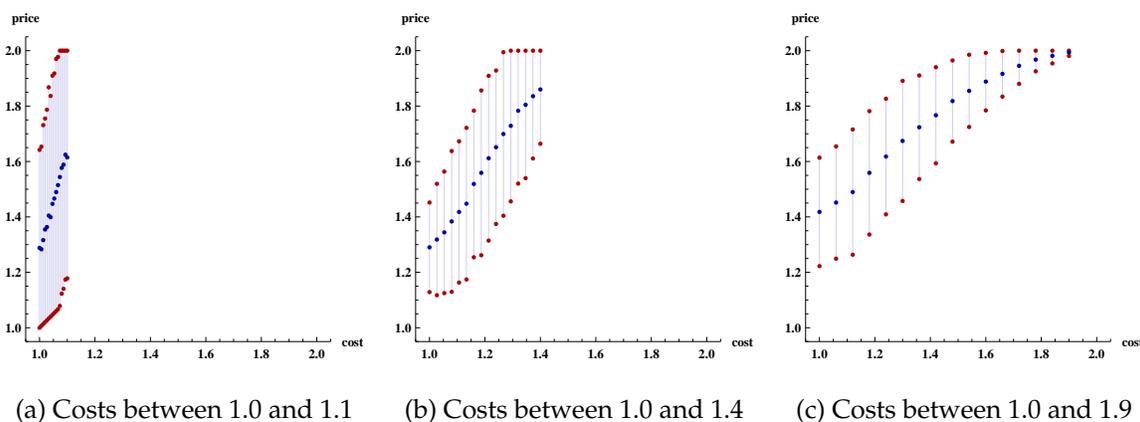
Smoothed histograms of the cross-section **mean** (right) and **standard deviation** (dashed, left) of profits.

Density functions of prices pooled over all simulation periods, compared with the benchmark equilibrium (shaded area).

simultaneously and myopically. In addition, average prices and profits are both negatively autocorrelated when the number of firms is high. This is very much in contrast to Model 1, where autocorrelation is positive and declines to zero smoothly, as we would expect from sequential updating.

To be sure, the head-over-heels updating of Model 2 is even more artificial than the sequential updating of the other Models.<sup>7</sup> However, every firm updating at the same time is just an extreme representation of the idea that firms may not be able to observe their competitors' most recent actions when considering their own choices. And it is worth noting that profits are higher on average (albeit more volatile) in Model 2 than in Model 1, which suggests that even rational firm owners might not mind a myopic management strategy.

Figure 5: The distribution of costs and prices in Model 3



Average prices and two-standard-deviation confidence intervals, truncated at the individually rational bounds: a firm's cost and the reservation price. The intervals are not independent of each other: when some firms charge relatively high prices, others are more likely to choose relatively low prices, and vice versa, which contributes to making price-cost rank reversals very common.

Model 3, which allows for heterogeneous production costs, is different. As [Burdett and Judd \(1983\)](#) proved for a continuum of firms with cost dispersion, firms

<sup>7</sup> All of the models impose a violation of rational expectations, as firms do not attempt to forecast future prices when switching, and they never learn that profit opportunities are ephemeral in these settings. But the violation in Model 2 is very stark.

with higher costs always charge higher prices, so the price ranking is identical to the cost ranking. In fact, if there is such a continuum of firms with a non-atomic cost distribution, a firm’s individual choice problem becomes concave, and while the price distribution is still endogenous in such an equilibrium, the equilibrium itself is much less fragile than when each firm must be indifferent between being at the top or the bottom of the price distribution.

We can clearly see this effect at work in the simulation. When cost dispersion is large, high-cost firms are forced to choose prices close to the reservation price, and consequently, low-cost firms charge low prices and take advantage of the higher sales volumes. As a result, the price distribution tracks the cost distribution closely despite the noise introduced by sequential updating. As the cost dispersion becomes small, on the other hand, the results resemble those of Model 1: all firms charge high prices at some times and low prices at others, and price-cost rank reversals are common. These results are summarized in Table 2 and illustrated in Figure 5.

**Table 2: The correlation of costs and prices in Model 3**

$\bar{c}$	Average correlation	Standard deviation
1.1	0.47	0.21
1.4	0.83	0.08
1.9	0.94	0.03
Any continuum of firms with cost dispersion	1.00	0.00

For each simulated time period, I compute Spearman’s rank correlation coefficient for the cross-section of prices and costs, then report the time-series mean and standard deviation. There are  $N = 16$  firms whose costs are evenly spaced between  $c = 1$  and  $\bar{c}$ , and each firm faces the same unit elastic demand curve with reservation price  $\bar{p} = 2$ .

All of the simulations described so far had firms learn what their profits would be with one single, randomly picked price and then decide whether to stay or to

switch. One can imagine an altogether different approach: firms choose prices from a *discrete grid*, *maximizing profits* over the whole set of available prices when it is their turn to update (closely related to the “fictitious play” concept of Fudenberg and Kreps, 1993). I also simulated this alternative approach. The results depended crucially on the order in which firms could update prices. When the order was random, as in Model 4 above, the results were almost identical to those obtained using Model 1 (see Table 1). In a version of Model 4 where the update order was fixed, however, every simulation run eventually reached a stable cycle of period two times the number of firms, implying that each firm ended up alternating between two prices. It is worth noting that not a single simulation run reached an absorbing state – in other words, a pure strategy Nash equilibrium of the simultaneous game – even though such equilibria exist when prices are restricted to a discrete grid.

## 5 Conclusion

Models of endogenous price dispersion feature equilibria where firms are indifferent over a wide range of prices. This makes them different from other models of price setting, such as monopolistic competition, where there is generally a unique price which maximizes profits for each firm. However, if the number of firms is finite, it is impossible to satisfy the indifference condition exactly. This raises the question of how robust an endogenous price distribution can be, and if markets would “find” it when starting from an arbitrary initial state.

Even if the initial state almost perfectly satisfies indifference, two mechanisms will cause instability. First, a firm changing its price to take advantage of arbitrage opportunities will cause another arbitrage opportunity near the price it “vacated” in the price distribution, and the new one might be bigger than the original one.

Second, if firms do not observe the entire sequence of recent price changes, they may attempt to take advantage of profit opportunities which no longer exist.

In the simulation, I show that both mechanisms may contribute to price instability. Edgeworth cycles spontaneously arise and die out, as simulation excerpts show (Figure 2). Qualitatively, the benchmark comparison performs fairly well when prices and profits are averaged over the long run. Quantitatively, however, differences persist even when the number of competitors is very large, and the differences are substantial when the number of competitors is small. Certainly, 6 or 16 competitors seems like a more realistic description of most markets than 100, even of markets for homogeneous goods like gas stations or retailers selling the same branded products. We should expect the mechanisms simulated in this paper to be at work in such markets.

One important qualification might be that in the simulated models, there is no menu cost of changing prices. Even a small menu cost might be enough to lock in an equilibrium in which no firm wants to switch given the opportunity. Starting from random sets of prices, Figures 3 and 4 demonstrate that the dispersion of profits will occasionally become very small. Such states of low dispersion could become absorbing states. Since the level and the dispersion of profits are positively correlated over simulated time, an absorbing state would likely also feature smaller profits than the average of the simulation before it reached the absorbing state. The surprising implication is that positive menu costs could lead to *smaller* average profits than zero menu costs, even net of menu costs themselves, because these would never actually be incurred in an absorbing state.

## References

- Burdett, K. and K. L. Judd (1983, July). Equilibrium Price Dispersion. *Econometrica* 51(4), 955–69.
- Burdett, K. and D. T. Mortensen (1998, May). Wage Differentials, Employer Size, and Unemployment. *International Economic Review* 39(2), 257–73.
- Cason, T. N., D. Friedman, and F. Wagener (2005). The dynamics of price dispersion, or edgeworth variations. *Journal of Economic Dynamics and Control* 29(4), 801–822.
- Crockett, S., R. Oprea, and C. Plott (2011). Extreme walrasian dynamics: The gale example in the lab. *American Economic Review* 101(7), 3196–3220.
- Fudenberg, D. and D. M. Kreps (1993). Learning mixed equilibria. *Games and Economic Behavior* 5(3), 320–367.
- Geromichalos, A. (2014). Directed search and the bertrand paradox. *International Economic Review* 55(4), 1043–1065.
- Head, A. and A. Kumar (2005, May). Price Dispersion, Inflation, And Welfare. *International Economic Review* 46(2), 533–572.
- Head, A., L. Q. Liu, G. Menzies, and R. Wright (2012). Sticky prices: A new monetarist approach. *Journal of the European Economic Association* 10(5), 939–973.
- Herrenbrueck, L. (2015). An open-economy model with money, endogenous search, and heterogeneous firms. Working paper, Simon Fraser University.
- Hopkins, E. and R. M. Seymour (2002). The stability of price dispersion under seller and consumer learning\*. *International Economic Review* 43(4), 1157–1190.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1996). *Microeconomic theory*. Oxford University Press New York.

Maskin, E. and J. Tirole (1988, May). A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles. *Econometrica* 56(3), 571–99.

Mortensen, D. T. (2005, May). A Comment On "Price Dispersion, Inflation, And Welfare" By A. Head And A. Kumar. *International Economic Review* 46(2), 573–578.

Wang, L. (2014). Endogenous search, price dispersion, and welfare. Working paper.