

An Experimental Investigation of Price Dispersion and Cycles

Timothy N. Cason

Purdue University

Daniel Friedman

University of Essex and University of California Santa Cruz

Ed Hopkins

University of Edinburgh

We report a continuous-time experiment studying the Burdett-Judd model, whose unique Nash equilibrium (NE) features dispersed prices. Adaptive dynamics predict that the NE is stable for one of our parameter sets and unstable for another. The empirical price distribution is close to the NE distribution for the stable parameter set, but for the unstable parameter set it skews toward higher prices in its NE support interval. We offer an empirical definition of price cycles in terms of changes over time in robust measures of central tendency and dispersion, by which the data exhibit persistent cycles in both treatments but larger cycles for the unstable parameters. Results are roughly similar for professional and student sellers and for limited-information treatments.

We are grateful for financial support from the National Science Foundation (SES-1357867) and for programming assistance from Matt Jee, Emily Hockel, Joshua Pena, and Morgan Grant. For valuable comments we thank two anonymous referees and an editor of this journal; Evan Calford, Jacob Goeree, Dan Kovenock, Heinrich Nax, Mareen Hallier, Lise Vesterlund, Alistair Wilson, and Zhijian Wang; participants at the Learning, Evolution and Games Workshop at Lund (June 2018), the Economic Science Association Conference in

Electronically published January 15, 2021

[*Journal of Political Economy*, 2021, vol. 129, no. 3]

© 2021 by The University of Chicago. All rights reserved. 0022-3808/2021/12903-0004\$10.00

I. Introduction

Few laws are disobeyed more frequently than the Law of One Price (LOP). For a multitude of goods ranging from alphabet blocks to zenon software, at a given point in time one can find a broad range of prices for apparently identical units (e.g., Baye, Morgan, and Scholten 2004). Jensen (2007) illustrates the major consequences such price dispersion can have for market efficiency and welfare. At least since Stigler (1961), economists have pointed to search frictions as the underlying cause. The now-canonical equilibrium search model of Burdett and Judd (1983) obtains a dispersed price distribution as a (mixed-strategy) best response by sellers to the search behavior of buyers, who in turn are best-responding to the dispersed prices. Varian's (1980) search model likewise explains LOP violations in terms of a mixed-strategy Nash equilibrium (NE), albeit with less strategic buyers.

It is not clear how to take such equilibrium predictions to the data. Subtle heterogeneities (e.g., in shipping, reputation, or individualized discounts) may be unobservable and are difficult to account for properly. More fundamentally, field behavior typically consists of asynchronous price adjustments in real time, and the proper connection to an equilibrium mixture of prices is debatable.

Since Smith (1962), economists have used data from simple laboratory markets to help resolve ambiguities in interpreting field data. Such markets can directly implement the primitives of a theoretical model, and the model's prediction successes and failures can help guide subsequent analysis of data from more complex ongoing markets in the field. Here we use new tools to study price dispersion dynamics in experimental markets, varying the search frictions in a controlled manner while eliminating confounds such as subtle heterogeneities and synchronized price changes. Thus, our setting is ideally suited for initial clean tests of dynamics and equilibrium in dispersed prices.

Broadly speaking, in a series of experimental markets inspired by the Burdett-Judd model, we find that the equilibrium distribution is not a bad predictor for the observed overall (time-collapsed) price distribution. However, we uncover a major role for dynamics. Individual sellers continually move their prices around within the dispersion. Beyond that, the dispersion itself moves over time and exhibits strong and persistent cycles that move consistently in the direction predicted by adaptive dynamics.

Our project also makes several methodological contributions of interest in their own right. We introduce nonparametric techniques for identifying

Berlin (June 2018), and the Tsinghua Conference on Behavioral, Experimental and Theoretical Economics (July 2018); and seminar audiences at Amsterdam, Caltech, Exeter, Melbourne, Middlebury (MiddExLab webinar), Monash, New York, Pittsburgh, and Vienna Universities. None of these organizations or individuals is responsible for remaining errors or idiosyncrasies. Data are provided as supplementary material online.

cycles in price distributions; these may be more useful than standard spectral (or time-series) techniques when cycle frequencies and amplitudes are irregular in lab or field data. Second, in contrast to the usual discrete-period lab software that enforces simultaneous pricing decisions, we use continuous-time software that enables asynchronous pricing decisions in real time. The importance of that distinction is exemplified by Wang (2009), who finds that introducing enforced simultaneous price setting for gasoline in Australia significantly changes behavior relative to the natural asynchronous setting.

Third, our treatments are multidimensional. In terms of search technology, we examine stable (S) versus unstable (U) search parameter vectors. As explained below, S theoretically promotes convergence to equilibrium if sellers follow simple adaptive-learning rules, while U theoretically prevents convergence. We also examine immediate versus speed-limited price changes and high versus limited information on counterfactual profits and on other sellers' prices. A final treatment variation is in the subject pool. Some sessions use the usual pool of undergraduate college students, while others use professionals—executive MBAs, many of whom have experience in setting prices for their company's products. Those sessions constitute an artifactual field experiment (Harrison and List 2004) and provide an anchor point for a bridge to future investigations of price dispersion and cycles using field data.

Before positioning our work within existing literature, we offer a few remarks on external validity. Our primary goal is to provide clean initial evidence on foundational issues regarding price dispersion and dynamics, giving theory its best shot; in the terminology of List (2020), ours is a WAVE1 study and establishes that the range of applicability is not empty. No single study, whether field or lab, can in itself guarantee generalization to any significantly different field or lab setting. However, we would argue that price competition gives strong basic incentives that will be the same in any other market that has the same general features, in particular, consumers who vary in their price sensitivity. Further, a study reporting robust results can increase confidence that the range of applicability is broad. Thus, our robust findings for multidimensional treatments—information conditions that bracket what might be seen in the wider world, as well as various adjustment speeds and experience levels of subjects—show that the basic incentive structure continues to determine behavior and bodes well for generalizability.

The literature on price cycles goes back at least to Edgeworth (1925), who describes gradually declining prices in an undercutting phase and upward jumps after capacity constraints begin to bind. He regards the succession of phases as predictable, but not the time spent in each phase. Capacity constraints are not relevant to our work, and Edgeworth's cycles have little dispersion at most points of time, but the intuition may nevertheless be helpful.

Maskin and Tirole (1988) obtain equilibrium cycles in an alternating-moves Bertrand duopoly model on a discrete-price grid. Their dynamic pricing model has many Markov-perfect NEs, one of which is turn taking as sellers undercut each other before returning to a high price. The resulting cycles are very regular, with minimal dispersion. It is not clear how to generalize this equilibrium beyond two sellers. The equilibrium model of Deneckere and Kovenock (1992) includes capacity constraints, and price cycles are also featured in the general equilibrium literature on the Scarf example, for example, Anderson et al. (2004) and Goeree and Lindsay (2016). None of these papers has a role for dispersed prices.

Considerable evidence exists for cyclical pricing behavior in real-world oligopolistic markets; see, for example, Slade (1998), Noel (2007), Wang (2009), or Byrne and de Roos (2019). Some of this evidence is reminiscent of Edgeworth cycles, with episodes of slowly falling prices often followed by a large upward jumps.

Regarding equilibrium price dispersion, a crucial question is whether in practice the equilibrium is approachable. Given the strategic nature of the Burdett-Judd setting and the complexity of the equilibrium in question, this might seem unlikely. Nonetheless, Hopkins and Seymour (2002) show that simple adaptive-learning rules can result in equilibrium play emerging over time, even in this setting. In fact, depending on exogenous parameters describing search efficacy, the Burdett-Judd equilibrium dispersion can be either stable or unstable under adaptive learning. In the stable case, we would therefore predict convergence (at least at the population level) toward the equilibrium price distribution, but it is unclear what to expect in the unstable case.

Our experiment is guided by that prediction and open question, along with previous experimental work by ourselves and others. Cason and Friedman (2003) study the comparative static predictions of Burdett-Judd model with the best tools then available, that is, a discrete-time experiment with enforced simultaneous moves. Further, they use only unstable parameter sets. Cason, Friedman, and Wagener (2005) reexamine the same data with an eye on dynamics and find suggestions of Edgeworth-like price cycles; for example, small price changes are more likely to be downward than upward. Using this older, discrete-time and simultaneous-move experimental technology, however, only one or two complete cycles could be observed in a typical run of 30 trading periods. Morgan, Orzen, and Sefton (2006) take the Varian (1980) model to the lab and find support for its static predictions of price dispersion. Brown Kruse et al. (1994) earlier examined mixed-strategy equilibrium in an oligopoly setting with quantity constraints similar to Edgeworth's.

To prepare for the present study, we examined static equilibrium and dynamics in simple bimatrix games. Cason, Friedman, and Hopkins (2010) use standard discrete-time protocols and find that a point prediction called

TASP (time-average Shapley polygon) outpredicts NE in a set of 4×4 bimatrix games. Cason, Friedman, and Hopkins (2014) study asynchronous real-time choices and introduce new techniques for detecting cycles in symmetric 3×3 bimatrix (rock-paper-scissors-style) games. In the present study, we extend those techniques for cycle detection to much larger action spaces of a far more realistic and complex environment. Our study thus addresses questions about price dynamics and price dispersion that have intrinsic interest, as well as testing predictions of learning theory.

New technology for conducting continuous-time experiments has gained favor recently, in part because it can enable long-run behavior to emerge in much less clock time than with standard discrete time periods and in part because it provides sharper and richer data. Oprea, Henwood, and Friedman (2011), Benndorf, Martínez-Martínez, and Normann (2016), and Stephenson (2019) find substantial support for the long-run stability predictions of evolutionary game theory, albeit for simple games with only two alternative actions. Experiments with larger action sets are reported by Calford and Oprea (2017) and Stephenson and Brown (2017). However, we believe that ours is the first continuous-time experiment of oligopolistic price competition and the first to study cycles in dispersed prices. Importantly, the continuous-time technology allows observation of hundreds of cycles per experimental session, instead of the small handful of cycles seen in discrete time.

The next section summarizes existing theory on equilibrium in the Burdett-Judd model and on its dynamic stability. A new discussion follows about the nature of cycles in high- (or infinite-) dimensional spaces, and the section concludes with a list of testable hypotheses. Section III details the experimental design, which features a novel videogame-like screen for subjects, who serve as sellers in a six-firm oligopoly played in continuous time. Experimental treatments are explained, including stable (S) and unstable (U) search parameter vectors, and price adjustment speeds (almost instantaneous, medium, and slow) that potentially affect the nature of dynamic adjustment and convergence. The section also describes the subpool of professional subjects and the high-versus limited-information treatments.

Section IV presents the results. At the aggregate level, we compare overall empirical distributions by treatment to the equilibrium distributions and find impressively close agreement in some treatments (especially S parameters, student subjects, and full information) and looser agreement in others. Regressions confirm several significant treatment effects on medians (our preferred measure of central tendency) and some marginally significant effects on dispersion (measured as the interquartile range [IQR]). The section then documents the presence and persistence of cycles in (median, IQR) space and reports support for most (but not all) of the testable hypotheses.

Section V summarizes and interprets the main findings and suggests possible future research. Appendix A supplements theoretical analysis, appendix B supplements the data analysis, and appendix C presents the experiment instructions provided to subjects.

II. Theoretical Considerations

This section reviews equilibrium in the sequential search price dispersion model of Burdett and Judd (1983) and the stability analysis of Hopkins and Seymour (2002). A discussion follows on how one might detect cycles in distributions supported on an interval of prices. The section concludes with a list of predictions to be tested.

A. Equilibrium

Burdett and Judd (1983) explain how price dispersion—a range of prices for a homogeneous good—can be supported in equilibrium. A continuum of firms chooses prices, and a continuum of consumers searches sequentially for low prices. An equilibrium is a distribution of prices such that all prices in its support earn the same (maximal) profit and such that consumer search behavior is optimal, given the distribution.

Consumer search is both sequential and noisy in the Burdett-Judd model. Each consumer initially receives a random number of price quotations and then either buys at the cheapest among existing quotations or else sinks search cost $c > 0$ to obtain a fresh sample of price quotations. The optimal policy is to adopt a particular reservation price p^* and to search again if and only if no current price quote is equal to or below p^* .

Sellers thus face buyers who vary ex post in the quality of their information. Some buyers will randomly receive multiple quotations and will buy from a lower-priced seller, while others will see only one price and will take it even if it is only slightly below p^* . This has two consequences. First, there is a nontrivial range of prices $[0, \underline{p})$ that sellers will never choose because these prices are dominated by p^* , a price at which demand is positive because of the positive fraction of buyers who see only one price. Second, demand is decreasing in price, because low prices attract relatively informed consumers as well as the lucky uninformed, while high prices attract only unlucky, relatively uninformed consumers. With decreasing demand (and, for analytical convenience, zero production cost), sellers get the same payoff at a high price in $[\underline{p}, p^*]$ that yields a few highly profitable sales as they do at a lower price that yields a larger volume of less profitable sales.

This leads to the main result of Burdett and Judd (1983): given a positive fraction q_1 of consumers who see only one price, there exists a unique equilibrium. The relevant equal-payoff condition yields a specific distribution of prices with continuous cumulative distribution function (CDF) $F(p)$

with support interval $[\underline{p}, p^*]$. The main exogenous variables, which determine the shape and support of the equilibrium distribution, are the cost of search $c > 0$ and the number of quotations received by consumers. In particular, assume that each consumer with probability q_k receives $k = 1, 2,$ or 3 independent and identically distributed quotations from F . Then the information conditions can be summarized by a point in the two-dimensional simplex, that is, by a nonnegative vector $q = (q_1, q_2, q_3)$, with $q_3 = 1 - q_1 - q_2$.

It is not hard to see that a seller of rank $u \in [0, 1]$ in the price distribution (the highest-priced seller has rank $u = 1$) faces demand

$$D(u) = \mu[q_1 + 2q_2(1 - u) + 3q_3(1 - u)^2], \tag{1}$$

where μ is the number of buyers per seller. Imposing the equal-payoff condition (see app. A), one can then show that the inverse equilibrium distribution is

$$F^{-1}(u) = \frac{p^* q_1}{q_1 + 2q_2(1 - u) + 3q_3(1 - u)^2} \tag{2}$$

on the support interval $[\underline{p}, p^*]$, with

$$\begin{aligned} \underline{p} &= F^{-1}(0) = \frac{p^* q_1}{q_1 + 2q_2 + 3q_3} \quad \text{and} \\ p^* &= F^{-1}(1) = \frac{c}{1 - q_1}. \end{aligned} \tag{3}$$

Note that in equilibrium sellers are indifferent over choosing any price in the equilibrium support $[\underline{p}, p^*]$, since they all generate the same profit, $p^* \mu q_1$. Note also that in this model, all buyers consume a single indivisible unit of the good, so total demand is the same for all prices (below p^*); there is no welfare loss from demand restriction, as in textbook models of monopoly. Higher prices simply imply higher transfers from buyers to sellers.

B. Stability

The Burdett-Judd model implies strategic interaction among sellers, with the best response of each seller depending on the distribution of prices among other sellers. In such complex situations of strategic uncertainty, it is unclear whether players realistically can play equilibrium in a one-shot setting. Evolutionary game theory and adaptive-learning theory investigate whether equilibrium can be reached if players continually play the same game, adjusting their strategies over time using simple rules.

Hopkins and Seymour (2002) identify parameters for which this equilibrium distribution is dynamically stable or unstable under such evolutionary or adaptive dynamics. We can explain intuitively, in terms of the relative size of wins, draws, and losses. In our oligopoly games, a “win” is where a seller has the lowest price and attracts lots of customers ($D(0)$ in eq. [1]), a “loss” is where a seller has the highest price and can sell only to buyers who receive a single price quote ($D(1)$), and a “draw” is where the seller has the median price among sellers ($D(0.5)$).

Mixed NEs are unstable in games for which draws are almost as good as wins, because then learning dynamics tend to spend more time near pure states, where prices are less dispersed, and less time near the mixed equilibrium. By contrast, a mixed equilibrium tends to be stable where draw payoffs are only slightly better than loss payoffs and are far below win payoffs; here, adaptive learners take on more dispersion in their prices to try to get wins. As shown in appendix A, this analogy allows us to characterize stability in terms of the convexity of the demand function $D(u)$ given in equation (1). Simply put, convexity leads to very high sales from low prices and thus big win payoffs and so stability, whereas concavity implies a big loss penalty due to a significant reduction in sales from having the highest price and hence instability. Following Hopkins and Seymour (2002), the following result is stated in terms of PDA (positive definite adaptive) learning dynamics, a class that contains various forms of reinforcement learning and imitation and evolutionary dynamics as special cases.

PROPOSITION 1. A dispersed-price equilibrium is stable under all PDA learning dynamics if and only if $D(u)$ is log convex; it is unstable under all PDA learning dynamics if and only if $D(u)$ is log concave.

Figure 1 illustrates this result. Recall that $q = (q_1, q_2, q_3)$, a point in the 3-simplex, summarizes the information available to consumers. The function $D(u)$ is log convex in the region labeled “Stable” near the $q_1 = 1$ corner, where losses and draws are less painful. In this region, all PDA dynamics converge to the unique Nash dispersed-price equilibrium. In contrast, $D(u)$ is log concave in the region labeled “Unstable,” which includes a large neighborhood of the $q_2 = 1$ corner; here, all PDA dynamics move away from the equilibrium. In the unlabeled intermediate region, $D(u)$ is neither log concave nor log convex, and consequently the equilibrium could be stable under some PDA dynamics but unstable under others.

Our experiment employs two particular parameter vectors highlighted in the figure and summarized in table 1. These were chosen for their simplicity and focality: the U vector was the principal case analyzed in Cason and Friedman (2003), and we chose the S vector to be as similar as possible while lying well within the Stable region identified in figure 1.

Unstable parameter vector.—This vector is represented by the dot U in figure 1 at $q_U = (2/3, 1/3, 0)$; consumer search cost is $c = 0.6$. Figure 2 shows the equilibrium price CDF $F_U(p) = 2 - 9/5p$ on support interval




FIG. 1.—Regions of search parameter space for which the unique equilibrium of the Burdett-Judd model is stable or unstable under PDA dynamics. The labeled dots show the parameter values *U* and *S* used in the experiment. Adapted from fig. 1 of Hopkins and Seymour (2002). A color version of this figure is available online.

$[\underline{p}, p^*] = [0.9, 1.8]$. Here the mean price is 1.25, and expected profits are 120 per seller.

Stable parameter vector.—This vector is represented by the dot *S* in figure 1 at $q_s = (0.8, 0.1, 0.1)$, with $c = 0.36$. The equilibrium CDF is $F_s(p) = 1 - (\{-p + p^{1/2}[24(p^* - p) + p]^{1/2}\}/3p)$ on $[\underline{p}, p^*] = [1.11, 1.8]$, with mean price 1.47 and expected profits 144 per seller.

Figure 2 shows that the two equilibrium distributions are somewhat different, despite (by construction) having the same upper bound, $p^* = 1.8$. The unique mode of F_U is at its lower bound, $\underline{p} = 0.9$, while F_S has a mode

TABLE 1
NASH EQUILIBRIUM (NE) PREDICTIONS FOR STABLE (*S*) AND UNSTABLE (*U*) PARAMETERS

Variable	Symbol	<i>U</i>	<i>S</i>
Distribution of price quotes	<i>q</i>	(2/3, 1/3, 0)	(.8, .1, .1)
Search cost	<i>c</i>	.60	.36
Buyers per seller	μ	100	100
Maximum NE price	p^*	1.80	1.80
Minimum NE price	\underline{p}	.90	1.11
Average NE price	\bar{m}	1.25	1.47
Median NE price	med	1.20	1.48
NE price IQR	IQR	.41	.37
NE price standard deviation	σ	.25	.21
Average seller NE profit	π	120	144




FIG. 2.—Nash equilibrium price density functions (*top*) and CDFs (*bottom*) for stable and unstable parameters. A uniform distribution (dashed line) is provided for contrast. A color version of this figure is available online.

at its upper bound as well as at its lower bound, $\underline{p} = 1.1$. Further, F_S stochastically dominates F_U . Finally, the two equilibrium distributions are quite distinct from a simple uniform distribution, which is also plotted in figure 2 for contrast. Thus, equilibrium entails a very specific distribution of prices in each case.

C. Cycles

As explained in appendix A, price choices of a continuum of agents at a given moment t (or, in our experiment, mixed choices of a finite number of agents) are represented by a point F^t in the space \mathcal{F} of all cumulative distributions over the relevant price interval. Over time $t > 0$, agents adjust prices, and the distribution F^t moves around in that infinite-dimensional space. The movement could be quite complicated, even chaotic. In such a high-dimensional space, how could we tell whether the price distribution eventually approached the equilibrium distribution or converged to a nice limit cycle?

Figure 3 illustrates a classic technique, due to Poincaré (1890), for detecting cycles and convergence in finite- ($n \geq 2$) or infinite-dimensional spaces. Suppose that we have a point P that may be on a limit cycle and a hyperplane through P that is transversal to the cycle. (For example, the hyperplane could be the vectors orthogonal to the tangent vector N at P .) The intersection PS of that hyperplane with a ball of radius ε around P is called a Poincaré section, and it is our cycle detector.

The detector works as follows. At a sequence of times $0 \leq t_1 < t_2 < \dots$, the evolving price distribution $\{F^t : t \geq 0\}$ intersects PS . At each such




FIG. 3.—Detecting limit cycles in high dimensions. Given a point P on a stable limit cycle with tangent vector N , take the intersection of the hyperplane orthogonal to N with a ball of radius ε around P . This Poincaré section, $PS = \{z \in B(P, \varepsilon) : z \cdot N = 0\}$, is the cycle detector.

time t_k , assign the sign $+1$ if the price path goes in the predicted direction (e.g., if the tangent to $\{F^t : t \geq 0\}$ at $t = t_k$ has a positive dot product with N) and -1 if it goes in the opposite direction. If, for $k = 1, 2, \dots$, the points $x_k = F^{t_k} \in \text{PS}$ all have the same sign and converge to some point $F^* \in \text{PS}$, then we have a (stable) limit cycle.

Appendix A explains how to deal with some technical issues in applying this method, but practical difficulties remain. How do we find a point P on (or near) a limit cycle if there is noise in the system? Which dynamics should we use to compute the tangent vector at P ? How big should ε be? With noise, we would see too few observations x_k if ε is too small, but if it is too large we might pick up other pieces of a limit cycle, going the wrong way.

In the background is a metaphysical question: What is a cycle? Physicists and chemists traditionally think of cycles as oscillators with regular frequencies that can be analyzed nicely using spectral techniques. Oscillation may persist at a finite amplitude or may be damped and ultimately disappear. By contrast, economists (and some biologists) tend to think of cycles as processes that never settle down but have distinct phases that recur in a particular sequence. They would likely exclude damped oscillations but might include somewhat irregular recurrence, as in some noisy or chaotic systems.¹ We suspect that most economists would not recognize as cycles any process with subtle periodic movements in the price distribution that change neither the central location nor the degree of dispersion.

From that perspective, it makes practical as well as intuitive sense for economists to work not with the infinite-dimensional space \mathcal{F} but rather with a two-dimensional projection that tracks the central tendency of the price distribution and its degree of dispersion. Our data analysis therefore looks for cycles in median price (a robust measure of central tendency) and IQR (75th minus 25th percentile prices, a robust measure of dispersion). The obvious alternatives, mean and variance, are less robust to extreme behavior by an individual subject. Nevertheless, as we document in appendix B, all of our qualitative conclusions regarding cycles still remain when we use mean and variance.

Figure 4 uses simulations to illustrate what we might see if players gradually adapt their prices toward the current best response. The line surrounding point U tracks the median price and the IQR as the distribution evolves under perturbed best-response dynamics for U parameters. (Specifically, the simulations use eq. [A9] in app. A: logit-perturbed best-response dynamics with precision parameter $\lambda = 4$.) Starting at an initial distribution with median price near 1.45 and IQR near 0.55, the dynamics for U soon bring the median price and IQR down to near their NE values of

¹ Zhijian Wang (private communication) reminds us that in the past 25 years, since the fluctuation theorem became widely recognized, physical scientists have also taken a more statistical approach to describing cycles.




FIG. 4.—Simulated (perturbed) best-response dynamics. Sample trajectories under slightly smoothed best-response dynamics for U and S parameters generate clockwise cycles in (median price, IQR) space. Labeled dots show NE values. A color version of this figure is available online.

(1.20, 0.41). But they do not stay there. Instead, the distribution's median and IQR cycle clockwise around a tilted trapezoid surrounding the NE. Starting at the top-right corner of the trapezoid, prices fall as the best response is to undercut slightly, and the IQR decreases as prices cluster more tightly. Then, near the bottom-left corner of the trapezoid, the best response jumps to 1.8, the maximum price. Dispersion increases as more low-price sellers move quickly toward high prices. When most sellers charge high prices, dispersion again decreases. The line heads back toward the top-right corner of the trapezoid, and then undercutting resumes.

The other line in figure 4 shows, starting from a nearby initial distribution, a simulated time path for S parameters, using the same perturbed best-response dynamics. This time path indeed converges to (1.48, 0.37), the NE median and IQR.

Note that Edgeworth's best-response cycle would not look the same. In discrete time, if all sellers were simply to play a best response to prices of the previous period, then prices would steadily descend in small steps from 1.8 to 0.9 in the U case, or to 1.11 in the S case, and then jump to 1.8. Price dispersion would almost always be nearly zero.

D. Testable Hypotheses

Since a continuum of sellers cannot fit in a laboratory, we follow Cason and Friedman (2003) in implementing the Burdett and Judd (1983) model

with a finite number of sellers, in this case, six. One can write $F(p) = u = 1 - (r_i - 1)/5$, or $1 - u = (r_i - 1)/5$, where $r_i \in \{1, 2, \dots, 6\}$ is the rank of firm i , with the convention that the highest-priced seller has rank 1 and the lowest has rank 6. Recall that demand $s(p_i) = 0$ whenever $p_i > p^* = 1.8$; to avoid trivialities, we allow only choices $p_i \leq p^*$. Thus, by equation (1), seller i faces demand

$$d(p_i) = \mu \left[q_1 + 2q_2 \left(\frac{r_i - 1}{5} \right) + 3q_3 \max \left[0, \left(\frac{r_i - 1}{5} \right) \left(\frac{r_i - 2}{4} \right) \right] \right], \quad (4)$$

where $\mu = 100$ is the ratio of (automated) buyers to sellers in our experiment. If two or more sellers post exactly the same price, then demand is shared by those sellers. For example, if two sellers tie for lowest price, then they both obtain the average of the sales of rank 6 and rank 5. Since production costs are zero and $q_U = (2/3, 1/3, 0)$, sellers' payoff function is

$$\pi_U(p_i, p_{-i}) = 100p_i \left[\frac{2}{3} + \frac{2}{3} \left(\frac{r_i - 1}{5} \right) \right] \quad (5)$$

in the Unstable treatment and

$$\pi_S(p_i, p_{-i}) = 100p_i \left[\frac{4}{5} + \frac{1}{5} \left(\frac{r_i - 1}{5} \right) + \frac{3}{10} \max \left[0, \left(\frac{r_i - 1}{5} \right) \left(\frac{r_i - 2}{4} \right) \right] \right] \quad (6)$$

in the Stable treatment, where $q_S = (0.8, 0.1, 0.1)$. This six-player game approximates the Burdett-Judd model; indeed, it can be shown that its unique symmetric NE is for each player to use the mixed strategy described by $F_U(p)$ in the Unstable treatment and by $F_S(p)$ in the Stable treatment.²

It has been suggested to us that the static NE is implausible in this setting, since the calculations are simply too demanding. We have some sympathy with this view, but we point out that the Burdett-Judd model is far from unique in relying on complex equilibria. To give the original theory a fair shot, our high-information treatment gives subjects extensive information about prices and profits. All our treatments give subjects substantial opportunities to learn from experience, as explained in the next section.

If they do learn, then learning theory makes several predictions. First, as set out in section II.B, behavior in the stable and unstable settings should be very different. Second, our simulations of adaptive dynamics (fig. 4) show clockwise motion in (median, IQR) price space. Thus, we predict cycles in our experiment also to be clockwise and not counterclockwise. Finally,

² Given the results of Baye, Kovenock, and De Vries (1992) for a related model, we suspect that the six-player game also has numerous asymmetric NEs whose composite distributions approximate $F_U(p)$ (or $F_S(p)$). Since our focus is the predictive power of the original Burdett-Judd model, we do not pursue that conjecture.

adaptive dynamics, in theory, can work in both high- and low-information settings, so there is no predicted effect from changing the information given to subjects.

Another consideration is that human subjects might try to collude; for example, all sellers might simply charge p^* . This concern led to our decision to have groups of six, where collusion is more difficult than in smaller groups.

Initially, we had another behavioral concern—in continuous time, subjects might make so many large and frequent price changes that the payoff landscape would seem a random blur, discouraging systematic adaptive learning or evolutionary dynamics. Therefore, we include adjustment speed as treatment variable, as detailed in the next section. Along with fast (virtually instantaneous) price change, we consider slow and medium-speed limits, which may affect the focus and quality of subjects' learning processes. We have no sound theoretical basis for predicting how that treatment will affect overall price distributions or the strength of cycles, but the treatment clearly affects the minimum time to complete a cycle and thus the maximum frequency.

These theoretical considerations lead to the following testable hypotheses.

1 (NE). For all subject pools (professionals and students), all information conditions, and all adjustment speeds, the empirical distributions of prices will be stationary and equal to the NE distribution, $F_S(p)$ for Stable parameters and $F_U(p)$ for Unstable parameters. Table 1 gives numerical values.

2 (Price cycles). For all subject pools, all information conditions, and all adjustment speeds,

- a. cycles will be clockwise in (median, IQR) space;
- b. cycle amplitude will converge to 0 for S parameters;
- c. cycle amplitude will converge to a positive number bounded away from 0 for U parameters;
- d. cycle frequency will be highest for the fast adjustment speed and lowest for the slow adjustment speed.

III. Experimental Design

In our within-subjects design, each session of the experiment has some market periods with the S parameters and other periods with the U parameters specified in table 1. For example, $q_U = (2/3, 1/3, 0)$ means that drawing one price has twice the likelihood of drawing two prices with U parameters, while with S parameters $q_S = (0.8, 0.1, 0.1)$, a single price draw is very likely, but drawing three prices is also possible. Given those parameters, in each period a continuum of buyers is simulated to behave

optimally, given the search costs and their price draws. The continuum eliminates sampling variance, so current payoff is a deterministic function of sellers' price profile.

Between 12 and 24 human subjects participate as sellers in each session. Each subject is assigned permanently to a market with six sellers, and each session has at least two such markets. Sellers are physically interspersed and spread widely throughout a 30-station lab with partitions that obscure subjects' identities and assignments.

We used only business professionals as sellers in nine of our 39 markets. These subjects were 25–54 years old (average age 36) and averaged 13.4 years of work experience. Nearly half (42.4%) report price-setting experience in their occupation.³ The remaining 30 markets (180 subjects) used only subjects from the usual pool of undergraduate students.

As illustrated in figure 5 and in the instructions included in appendix C, sellers select and adjust their prices in continuous time, using a horizontal slider restricted to the interval $[0, 1.80]$. The corresponding profit flow given in equations (5) and (6) is shown as the vertical height of the circle above the current slider position.

In the “high-information” condition, we provide two additional visual cues. Each seller sees differently colored circles on the horizontal axis, indicating the price choices (but not the profit flow) of the other five sellers in the market. This corresponds to the perfect price monitoring available in some markets (e.g., Byrne and de Roos 2019). The same display also includes a black line showing the seller's own counterfactual profit flow at every feasible price, given the prices currently offered by the other sellers. The highest point on this line shows the player the best response to the current price distribution; the sawtooth shape seen in figure 5 is typical. Of course, the line is constantly in motion as sellers continuously adjust their prices. The purpose is to give subjects an impression of the profit opportunities available to them, and we believe that our subjects find it far more useful than detailed mathematical formulas. Twenty-three of the 39 markets, including all nine markets with professionals, provide this counterfactual profit (or “payoff landscape”) line.

To study how such contemporary information may affect pricing behavior, we also conducted sessions where it was suppressed. In eight “limited info—prices only” markets, the display eliminates the payoff landscape line but still shows other sellers' prices on the horizontal axis. Another eight “limited info—no prices” markets eliminate both cues and provide no information about other sellers' current or past prices. An example screen is included in appendix C.⁴

³ These professionals were enrolled in executive or weekend MBA programs at Purdue University. One market with professionals had only five sellers.

⁴ All 16 of these limited-information markets began with four periods of the high-information display so that sellers could learn how their payoffs depend on all sellers' prices. We do not analyze these “training” periods.

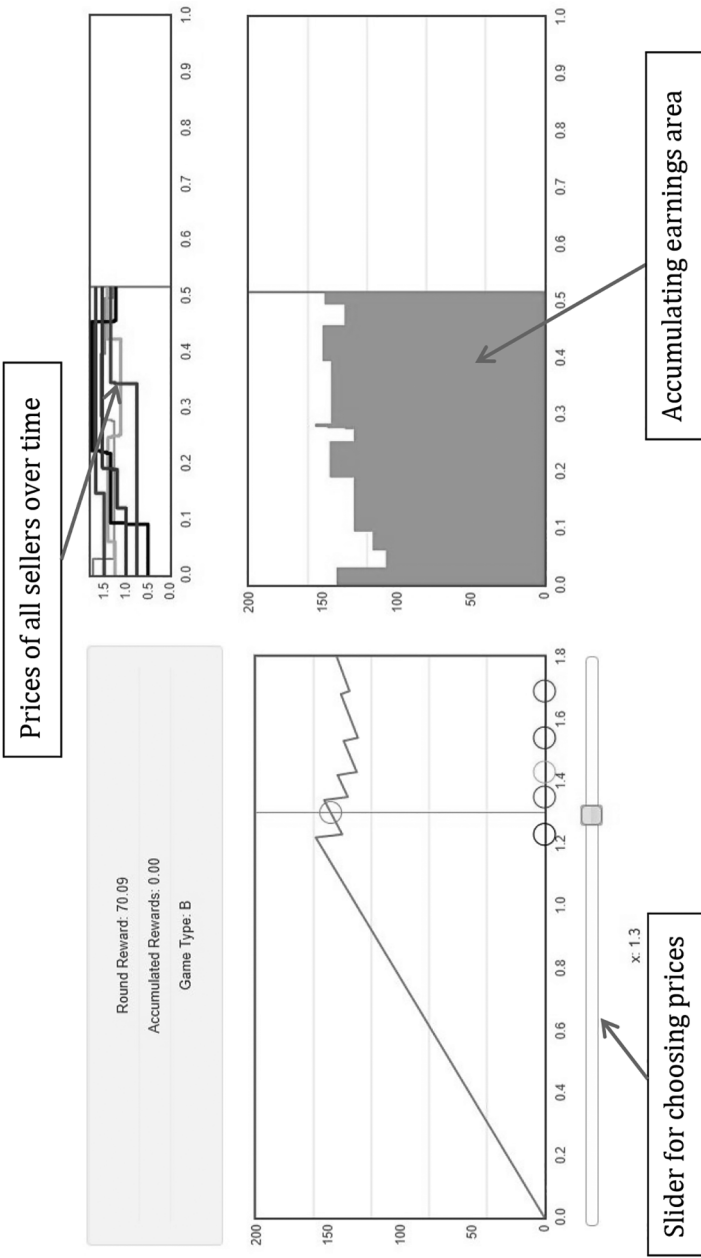


FIG. 5.—Example seller price choice and feedback screen. A color version of this figure is available online.

The two panels on the right of sellers' screens have elapsed time as the horizontal axis. The large lower panel shows the seller's accumulating profit as the area swept out by the vertical height of the circle. The narrow upper panel's line graphs show all six sellers' prices chosen so far; this panel is suppressed in the "no prices" condition.

Besides the primary S - versus U -parameter treatment, we have a secondary within-subject treatment, adjustment speed, with three alternative conditions. When a subject moves the price slider or clicks on a new target price, her actual price moves from its old value to the new target at a constant speed. In the fast condition, the move is virtually instantaneous (0.072 seconds for a typical adjustment spanning 10% of the available price range). In the slow condition, the same adjustment takes 3.6 seconds to complete, and in the medium condition, it takes 0.72 seconds.

Each session with undergraduate student subjects is divided into 32 trading periods. Initial prices are randomly assigned each period. Trading periods each last 90 seconds, except for slow-adjustment periods after the first three sessions. Those periods last 150 seconds, in order to permit more cycles.

The 32 trading periods in each student pool market are divided into eight blocks of four periods each. The within-subjects treatments (i.e., S vs. U parameters and price adjustment speed) are held constant within each four-period block and varied in a balanced fashion across blocks in each session, as detailed in table B12. Because of time constraints for the business professionals, their sessions last only 16 periods, so each market provides data for only four of the six within-subject treatment combinations. Since assignment of subjects to separate six-seller markets is permanent in each session, the 233 subjects provide 36 independent sets of observations for each treatment combination.

Sessions were conducted at Purdue University between October 2017 and February 2020. The student subject sessions took about 80–90 minutes, and earnings per subject (including a \$5.00 participation payment) averaged \$28.95, with a standard deviation of \$1.80. We increased the exchange rate for the business professionals, so their 40–50-minute sessions' earnings averaged \$44.24, with a standard deviation of \$2.31.

IV. Results

Our data consist of prices posted by subjects sampled at 400-ms intervals, that is, 2.5 times per second. Initial prices in the first six markets were drawn randomly without replacement from the set $\{0.50, 0.75, \dots, 1.75\}$; the data analysis drops the first 15 seconds of each period, when sellers are adjusting prices toward the narrower range where they typically operate for the rest of the period. In the remaining 33 markets, we used the narrower set $\{1.10, 1.20, \dots, 1.60\}$ of initial prices, enabling us to reduce the excluded

initial adjustment time to the first 10 seconds of each period. The non-parametric tests reported below all employ the conservative convention that the unit of observation is the independent six-seller market.

A. Overview

The top panels of figure 6 compare overall observed cumulative price distributions in Stable treatments to the NE of the Burdett-Judd model. The prediction seems right on target, except for a clump of prices right at the upper endpoint $p^* = 1.80$ (about 9% of choices with high information but up to 14% with limited information) and a few prices a little below the predicted lower end point, $\underline{p} = 1.11$, of the predicted support interval. The corresponding Unstable-treatment NE prediction shown in the lower half of figure 6 is almost first-order stochastically dominated by each of the empirical cumulative price distributions, mainly because of fewer low prices than predicted. In each panel, adjustment speed seems to have little impact.

Table 2 presents summary statistics for pooled subject pools, with other treatments disaggregated.⁵ The upper-left side shows that the median prices observed for the Stable treatment are remarkably close to the equilibrium prediction of 1.48, and they are not significantly different from this benchmark in the high-information condition.⁶ As predicted, median prices are lower in the Unstable treatment, but they nevertheless are significantly higher than the equilibrium prediction of 1.20. The lower half of table 2 indicates that median prices are higher when sellers have more limited price information and payoff feedback. This is consistent with the price CDFs shown on the right side of figure 6, which have a larger mass of prices on the maximum price.

The right side of table 2 reports predicted and observed dispersion as measured by the IQR of prices. Observed dispersion is often as predicted and sometimes smaller in the Unstable treatment; differences from equilibrium are insignificant, except for some price adjustment speeds in the limited-information condition. In the Stable treatment, observed dispersion is consistently larger than predicted and often significantly larger. This is due mainly to the sprinkling of price choices below $\underline{p} = 1.11$ and the clump at the maximum, $p^* = 1.80$. Table B2 shows that neither the median price nor the IQR dispersion measure change much between the first and second halves of the trading period in any treatment. Thus, the

⁵ Table B1 and figs. B1, B2 summarize data disaggregated by subject pool and information condition.

⁶ These and other unconditional tests in this study employ Holm-Bonferroni p -value correction, as in Andreoni et al. (2019), to adjust for the family-wise error rate (e.g., see List, Shaikh, and Xu 2019).

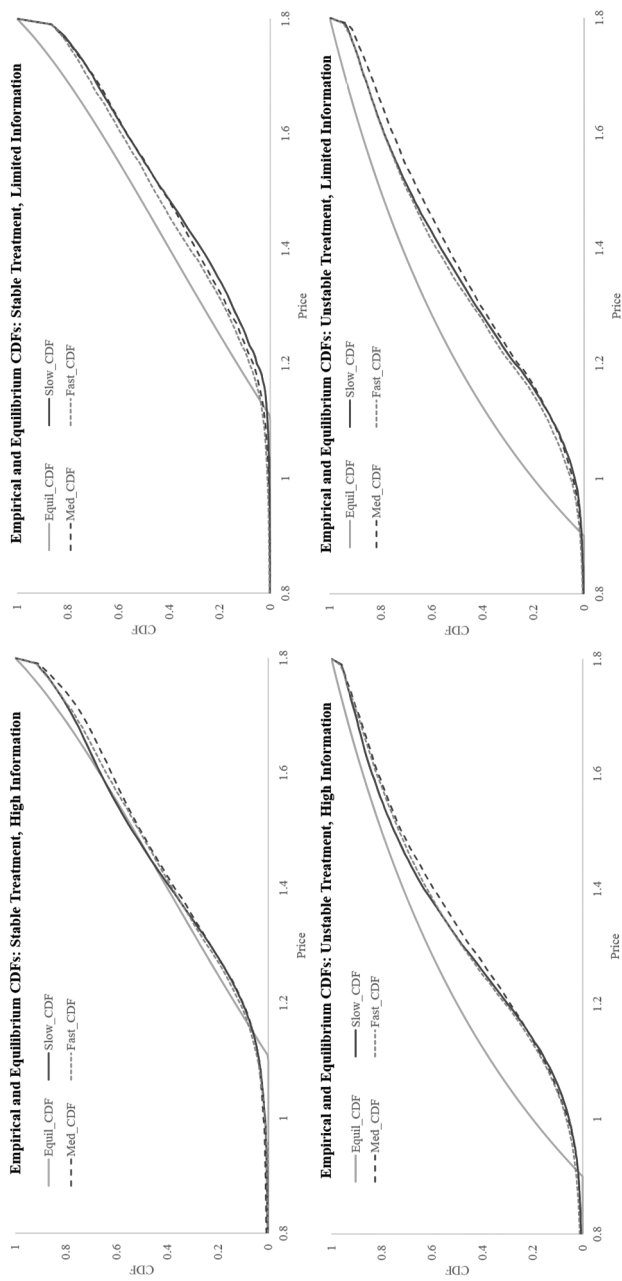


FIG. 6.—Cumulative distributions of observed versus NE prices in stable (*top*) and unstable (*bottom*) treatments. Samples exclude first 10–15 seconds of each period. A color version of this figure is available online.

TABLE 2
 MEAN OF SUMMARY PRICE STATISTICS (Standard Errors [SE] across Periods)

	Median Price	SE	IQR	SE
High Information (Pooled)				
Nash equilibrium (stable)	1.48		.37	
Stable-fast adjustment	1.482	(.007)	.433**	(.006)
Stable-medium adjustment	1.496	(.009)	.447**	(.009)
Stable-slow adjustment	1.469	(.009)	.398*	(.007)
Nash equilibrium (unstable)	1.20		.41	
Unstable-fast adjustment	1.324**	(.008)	.411	(.009)
Unstable-medium adjustment	1.339**	(.011)	.403	(.008)
Unstable-slow adjustment	1.317**	(.010)	.377	(.008)
Limited Information (Pooled)				
Nash equilibrium (stable)	1.48		.37	
Stable-fast adjustment	1.528**	(.006)	.426**	(.007)
Stable-medium adjustment	1.550**	(.007)	.411**	(.006)
Stable-slow adjustment	1.548**	(.006)	.393	(.006)
Nash equilibrium (unstable)	1.20		.41	
Unstable-fast adjustment	1.356**	(.009)	.447*	(.006)
Unstable-medium adjustment	1.387**	(.013)	.439*	(.007)
Unstable-slow adjustment	1.364**	(.011)	.399	(.007)

NOTE.—The first 10 or 15 seconds are omitted from each 90- or 150-second period. Asterisks denote significant differences from Nash equilibrium for two-tailed Wilcoxon signed-rank test, using Holm-Bonferroni p -value correction to adjust for the family-wise error rate.

* Significant at the 5% level.

** Significant at the 1% level.

time-average price distributions seem not to be trending toward a different location or degree of dispersion.

Table 3 reports ordinary least squares (OLS) regressions of median prices and IQR, averaged for all four periods in each block, on a series of dummy variables for exogenous treatments. Consistent with hypothesis 1,

TABLE 3
 OLS REGRESSIONS OF SUMMARY PRICE STATISTICS

Variable	Median Price	SE	IQR	SE
Unstable market	-.163**	(.008)	-.008	(.007)
Professionals	.014	(.015)	.045*	(.020)
No counterfactual projection	.052**	(.012)	.019	(.012)
No information on others' prices	.057**	(.016)	.022*	(.010)
Fast adjustment speed	-.019*	(.007)	.004	(.005)
Slow adjustment speed	-.018*	(.007)	-.034**	(.006)
Intercept	1.494**	(.011)	.413**	(.006)
Observations (clusters)	216 (39)		212 (38)	
R^2	.707		.203	

NOTE.—Standard errors (SE) are adjusted for clustering on markets.

* Significant difference from 0 at the 5% level (two-tailed tests).

** Significant difference from 0 at the 1% level (two-tailed tests).

median prices are significantly lower with the U parameter vector. Relative to high information (the omitted condition), median prices are significantly higher without the counterfactual payoff projection, whether or not display of other sellers' prices is also suppressed. Median prices are slightly lower for the fast and slow adjustment speeds relative to the intermediate (omitted) medium adjustment speed condition. Notably, the use of professional businesspeople as subjects has no significant impact on median prices.

The rightmost column of table 3 shows some small differences in dispersion, with decreases for the slow adjustment speed and increased dispersion for professionals and when price information is suppressed.⁷ Dispersion is not significantly different between the S and U market parameters, but our design deliberately sets similar NE dispersion across the treatments (0.37 for S and 0.41 for U) in order to focus on identification of differences in cycles.

In sum, we have mixed support for hypothesis 1. NE correctly predicts the comparative statics for median prices. Despite slightly underpredicting IQR, it is reasonably accurate in predicting the overall distribution in the Stable treatment, at least when sellers can observe other sellers' contemporaneous prices. The NE prediction in the Unstable treatments is less accurate; although roughly correct, particularly for the level of dispersion, it significantly underpredicts the observed median price.

B. Price Cycles

Might the observed skew toward higher-than-NE prices in the Unstable treatments be due to price cycles? Figure 7 illustrates observed price dynamics in (median, IQR) space in two trading periods. Time progresses downward, beginning at second 10 at the top of the vertical axis and ending in the horizontal plane at second 90 (or 150). The "pole" in the center of each panel denotes the time-averaged (median, IQR) for the period. The two panels are representative in that we see rough cycles in all speed and stability treatments, although there is considerable heterogeneity within and across treatments.

To test hypothesis 2 on price cycles, we implement a version of Poincaré's technique in (median, IQR) space. Of course, in any two-dimensional space, a Poincaré section is just a line segment, that is, a "tripwire." As detailed in appendix B, our tripwire connects the center of the (median, IQR) observations to the horizontal axis, and each period we count the number CW of tripwire crossings in the predicted clockwise direction and

⁷ To avoid redefining dispersion, we exclude from the IQR regressions in the one market with five rather than six sellers; thus, these regressions use 212 rather than 216 observations.

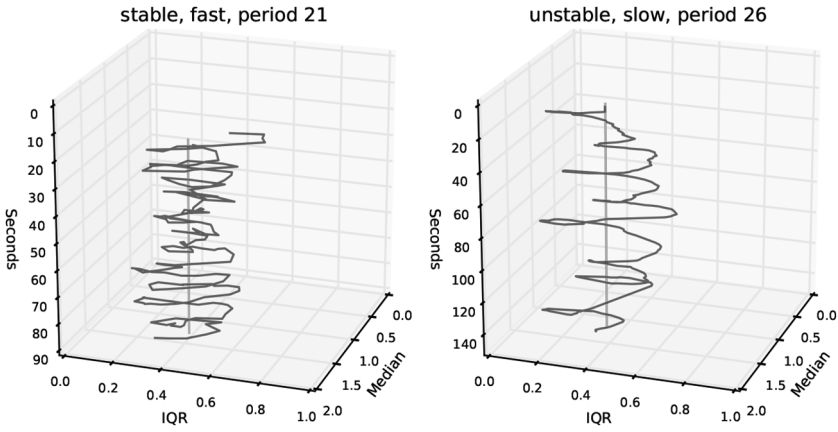


FIG. 7.—Examples of (median, IQR) dynamics in the stable-fast (*left*) and unstable-slow (*right*) treatments. The straight line plots the time-averaged (median, IQR) for the period. Samples exclude first 10 seconds of each period. A color version of this figure is available online.

the number CCW in the opposite (counterclockwise) direction. We define the cycle rotation index for each six-seller market in each period by

$$CRI = \frac{CW - CCW}{CW + CCW}. \tag{7}$$

By construction, $CRI \in [-1, 1]$, with positive values indicating mostly clockwise transits. As noted above, the best-response cycles in (median, IQR) space rotate clockwise around the cycle center, so hypothesis 2a entails positive CRI values.

Table 4 displays the average number of both kinds of transits and the average CRI per period, disaggregated by treatment condition and separately for the high- and limited-information cases. As per hypothesis 2d, transits are less frequent when prices adjust to sellers’ target prices more slowly. This is unsurprising, as the minimum time to complete a given cycle is longer in the slower-adjustment treatments.

More importantly, the CRI is significantly positive in every treatment condition, according to the conservative Wilcoxon signed-rank test, consistent with hypothesis 2a. Cycles are noisy, of course, and not every transit is in the predicted clockwise direction. Cycles are somewhat less frequent when sellers have limited information about their potential payoffs or other sellers’ prices, but the CRI averages across periods are remarkably consistent within information treatments. The CRIs calculated separately for the first half and second half of the trading periods are also very similar; see table B7. Thus, contrary to hypothesis 1, the price distribution remains nonstationary in all treatments. When the adjustment speed and

TABLE 4
MEAN TRANSITS AND CRI

	CW Transits	CCW Transits	CRI
High Information (Pooled)			
Stable-fast adjustment	9.47 (.69)	2.68 (.28)	.59** (.04)
Stable-medium adjustment	7.56 (.48)	2.13 (.17)	.59** (.04)
Stable-slow adjustment	4.18 (.29)	1.25 (.14)	.58** (.04)
Unstable-fast adjustment	8.32 (.49)	1.97 (.32)	.64** (.05)
Unstable-medium adjustment	6.83 (.44)	1.64 (.20)	.65** (.04)
Unstable-slow adjustment	3.80 (.19)	1.05 (.17)	.66** (.05)
Limited Information (Pooled)			
Stable-fast adjustment	6.44 (.49)	2.86 (.27)	.40** (.05)
Stable-medium adjustment	5.67 (.29)	2.81 (.26)	.39** (.05)
Stable-slow adjustment	4.23 (.31)	1.64 (.17)	.48** (.04)
Unstable-fast adjustment	6.17 (.47)	2.58 (.33)	.46** (.06)
Unstable-medium adjustment	4.78 (.37)	2.17 (.25)	.46** (.05)
Unstable-slow adjustment	3.20 (.18)	.84 (.13)	.66** (.05)

NOTE.—Standard error of the means in parentheses.

** Significantly different from zero at the 1% level for a two-tailed Wilcoxon signed-rank test, using Holm-Bonferroni p -value correction to adjust for the family-wise error rate.

information condition are controlled for, however, the CRI is always larger in the Unstable than in the Stable condition. We provide a statistical test for this difference below.

Amplitude is another important aspect of price cycles. As the observed (median, IQR) moves around, its path might typically stay in a smaller or larger neighborhood of the center. We therefore also compare the cycles in terms of their amplitude, defined as the mean distance from the center at which the trajectory crosses the tripwire. Table 5 shows that average amplitude is always greater with the U parameter vector than with

TABLE 5
MEAN (ACROSS PERIODS AND MARKETS) OF MEDIAN (WITHIN-PERIOD) CYCLE AMPLITUDE

	Fast Adjustment	Medium Adjustment	Slow Adjustment
High Information (Pooled)			
Stable	.087 (.006)	.092 (.005)	.103 (.006)
Unstable	.145 (.008)	.122 (.007)	.121 (.007)
p -Value	<.001**	.003**	.050*
Limited Information (Pooled)			
Stable	.085 (.006)	.082 (.006)	.094 (.008)
Unstable	.134 (.010)	.120 (.011)	.110 (.009)
p -Value	<.001**	.011*	.072

NOTE.—Standard error of the means shown in parentheses. Indicated p -values are for one-tailed Mann-Whitney tests of stable vs. unstable, using Holm-Bonferroni p -value correction.

* Significant at the 5% level.

** Significant at the 1% level.

the S vector, holding other treatments constant. The relative differences are substantial, ranging from 17% to 67%, and the differences are highly significant, except for the slow-adjustment treatment (which has fewer cycles). The consistently positive cycle amplitudes in all Stable conditions enable us to reject hypothesis 2b. The systematic, larger-amplitude cycles in the Unstable condition support hypothesis 2c. Appendix B shows that these conclusions are robust to other definitions of the cycle tripwire (tables B5, B6) and for an alternative definition of central tendency and dispersion (mean, standard deviation; see tables B8 and B9).

Table 6 reports OLS regressions of these cycle measures, averaged across all four periods of each treatment within each session, on a dummy variable for the U parameters. These estimates control for the other experimental conditions, using additional dummy variables. Consistent with hypothesis 2, both cycle measures are significantly higher for the U parameters. The use of business professionals lowers the CRI but does not significantly affect cycle amplitude. Other than stability, the only factor significantly affecting amplitude is the availability of information about other sellers' prices. The CRI is also affected by how much information is provided to sellers.

C. Price Changes

Recall that, although sellers could change prices at any time, for data analysis we recorded the actual price vector at 400-ms intervals and at the same time also recorded the vector of prices targeted by each seller. The two price vectors are usually quite close together in the medium- and fast-adjustment data, but they can differ substantially when price adjustment is slow. In this subsection, we focus on the changes in target price between successive observations.

TABLE 6
OLS REGRESSIONS OF CRI AND CYCLE AMPLITUDE

Variable	CRI	SE	Cycle Amplitude	SE
Unstable market	.080**	(.020)	.034**	(.004)
Professionals	-.156**	(.042)	-.012	(.007)
No counterfactual projection	-.110**	(.037)	.003	(.007)
No information on others' prices	-.262**	(.004)	-.024*	(.004)
Fast adjustment speed	-.002	(.003)	.009	(.004)
Slow adjustment speed	.007*	(.003)	.003	(.004)
Intercept	.598**	(.026)	.094**	(.004)
Observations (clusters)	212 (38)		212 (38)	
R^2	.305		.344	

NOTE.—Standard errors (SE) are adjusted for clustering on markets.

* Significant difference from 0 at the 5% level (two-tailed tests).

** Significant difference from 0 at the 1% level (two-tailed tests).

After excluding the initial 10 or 15 adjustment seconds for the 38 markets with six sellers, we have 1,213,674 adjacent 400-ms intervals to consider for possible price changes by individual sellers. Sellers changed their target price in 622,725 of these intervals, or 51%. One reason for this high percentage is that our interface made price changes very easy: when sellers dragged the price slider around, as they often did, our data indicate a constantly moving target price.

Target price changes are most common in the fast-adjustment condition (59.3%) and less frequent in the medium- and slow-adjustment conditions (55.1% and 44.2%, respectively). More importantly, price changes are systematically more frequent in the Unstable markets than in the Stable markets. Overall, price changes occur 55.0% of the time in the Unstable periods, compared to 47.6% for the Stable periods. For 34 of 38 independent markets, target price changes are less frequent in Stable than in Unstable periods. The relevant Wilcoxon signed-rank test establishes that the difference is significant at $p < .001$. However, table B10 shows that the frequency of price changes hardly declines from the early to the late parts of the trading period in any type of market.

Considerable heterogeneity exists in the frequency of price changes across individual sellers. Only 13 of 228 sellers change prices in more than three-fourths of successive observations, and 37 of 228 sellers leave prices unchanged in more than two-thirds of successive observations. The interquartile interval for target price change frequency is [0.399, 0.645] across sellers. Average price change frequency is lower among professionals (0.488) and in markets with limited payoff and price feedback (0.476), relative to markets with student sellers and high information (0.573). This is due, in part, to professionals and sellers in limited-information markets more frequently leaving prices unchanged at the maximum level.

Profits do not vary much across sellers because, as we have seen, the overall price distribution is close to the NE distribution where (within its support) profits are equal for all price choices. The lowest average earning rate among sellers is 126.3, and the highest is 164.4, with an IQR of [135.3, 144.2]. Despite this compressed range, we find that sellers who changed their target price frequently earned greater profits; the Spearman rank correlation of average payoff rate with price change frequency across sellers is $\rho = 0.37$. This correlation is highly significant, based on a cross-sectional regression that clusters standard errors at the market level ($p < .001$).

How do target prices compare to the contemporaneous best-response (BR) price? More often than not, sellers move their targets in the right direction: 61% of price changes are upward when the BR exceeds the current price, and 58% are downward when it is below. There is, however, an interesting asymmetry in how often sellers keep price targets unchanged: 43% of the time when the BR exceeds the current price, but 53% of time

when it is below. This relative reluctance to respond to the incentive to lower prices might be due to the fact that the local profit gradient is typically positive. That is, as can be visualized in figure 5, a seller's counterfactual profit declines as she reduces price until her rank in the price distribution falls below a rival seller's. In other words, while simple (myopic) payoff-based learning would facilitate directional learning (Selten and Stoecker 1986) in the case of upward adjustments, it does not in the case of downward adjustments, as subjects need to have some patience/foresight.⁸

Finally, we note that when sellers move their target toward the BR price, they tend to undershoot; see table B11. Overshooting (setting the target more than 0.05 beyond the BR, as a seller might do in anticipation that the BR will soon change) is rare, less than 5% of cases, and setting the target within 0.05 of the BR is definitely less common than setting it short of that. This behavioral inertia might suggest lack of foresight, but it might also simply be a mechanical artifact of the user interface, since most sellers drag their slider toward their true target rather than clicking on it directly.

V. Discussion

Our experiment investigates price dispersion in the sequential search model of Burdett and Judd (1983). That model predicts parameter-specific distributions of prices, and our aggregate data largely bear out those predictions. For "Stable" (S) parameters, student subjects, and high-information treatments, the prediction is remarkably close to the empirical distribution that we observe when aggregating over sellers and over time. Predicted distributions are in the right ballpark for all treatments, and key comparative-statics predictions are correct; for example, median observed price is indeed lower with the S parameters.

Disaggregated data reveal much more activity than suggested by static equilibrium predictions. In all treatments, most individual sellers continually adjust prices throughout the trading period. Our data thus rule out several plausible hypotheses: that individual sellers' prices stabilize, that price dispersion fades away, and that dispersion is not systematically related to search technology.

In our view, the most important finding is that individual sellers' price adjustments do not tend to cancel out, as in some stochastic equilibrium (in which the equilibrium distribution holds at the population level but

⁸ The asymmetry in the rate of price changes when the BR is above, rather than below, the current price is larger in the slow- and medium-adjustment conditions (11%) than with fast adjustment (8%), which is consistent with this patience/foresight conjecture. With fast price adjustment, patience is not required, since movement to the (recent) BR is nearly instantaneous.

See Nax and Pradelski (2015) and Nax et al. (2016) for recent theory and experimental evidence of myopic directional learning.

not at the individual level); instead, there are strong and persistent cycles in the price distribution. Consistent with learning and evolutionary dynamics, we see that, starting at relatively high values, dispersion and median price both tend to fall as sellers get pulled into intense price undercutting. Dispersion continues to fall as the median price begins to stabilize at a low level. Then a few sellers abandon undercutting and jump to high prices. As more follow, the median increases sharply, and dispersion also increases until the undercutting phase resumes. This process produces irregular but persistent clockwise cycles in (median, IQR) space, where IQR is a robust measure of dispersion, implying that price dispersion is higher when prices are rising.

In each treatment allowing timely price adjustment, we see, on average, 5–10 clockwise cycles per trading period. Of course, cycles are less frequent when we enforce very gradual price adjustment (slow treatment), but the CRI still shows very significant clockwise cycling. Cycle amplitudes are significantly smaller with *S* than with *U* parameters, but (contrary to our initial hypothesis) they remain significantly positive.

One might ask how is it possible for the overall empirical distribution to be close to the equilibrium, even while persistent cycles keep the distribution some distance away from equilibrium. To explain, consider someone constantly driving on a ring road outside of a city. An average of this driver's GPS coordinates would give the center of the ring somewhere in downtown, but the driver never was downtown, just as our subjects are never in equilibrium. Theory formally outlining this possibility is in Benaïm, Hofbauer, and Hopkins (2009). Further, this phenomenon, persistent cycles but averages close to equilibrium, is also reported in Cason and Friedman (2003) and, earlier, in Brown Kruse et al. (1994).

Our results are robust in some important directions. First, to conduct an artificial field experiment, we ran additional sessions with a different subject pool, business professionals enrolled in executive MBA programs, many of whom had field experience setting prices. One might conjecture that such subjects would be more successful in finding equilibrium or, alternatively, in managing to collude or otherwise altering dynamic behavior. Although we detect some interesting differences—cycles are slightly more dispersed and slightly less regular—all major conclusions are confirmed when business professionals set prices.

A second robustness check was to relax our full-information baseline condition, which was intended to give NE its best shot. Sellers in baseline always see other sellers' current prices and can easily find the best response to those prices. Sellers in many real-world situations do not have optimization tools and/or timely access to rivals' prices, so we tested the impact of going to the other extreme. One limited-information treatment removed access to counterfactual payoffs (including the best response), and a further treatment also removed access to other sellers' current prices.

The impact turns out to be less than we expected. Under either of these limited-information conditions, median prices are somewhat higher and cycles are slightly less consistent in direction and slightly lower in amplitude. However, in all settings we still observe strong cycles that show no tendency to fade away.

What might the results tell us about pricing in the wider world? Online sellers can easily and quickly adjust price offers whenever they desire, although actionable feedback typically takes hours or days to accumulate. Taking into account the different effective timescales, we see no insurmountable barrier to generalizing the theory and results presented here. Certainly, field data present major challenges—for example, in accounting for subtle heterogeneities in products and for endogenous selection of search technology—but robust results from tightly controlled settings such as ours can help isolate and cope with those additional challenges. In particular, empirical economists might consider looking for irregular cycles (perhaps using Poincaré sections) in field data. They could investigate whether the degree of dispersion and the amplitude and frequency of cycles vary with changes in the search technology, such as recommender systems or search engines, in a manner consistent with our laboratory findings.

Many questions also remain for future lab experiments. What happens if periods run for a very long time—would the cycle amplitude eventually fall in the Stable condition? How exploitable are the cycles—could an algorithm that tracks cycles make substantial excess profit? Would sellers' behavior change if buyers were human subjects who could adapt their search behavior or could substitute intertemporally?

We hope that our results also provoke new theoretical developments. Appendix A raises as many questions as it answers about how economists might identify cycles in higher dimensions. Our results are broadly consistent with simple models of adaptive learning, but they also highlight issues in applying existing learning models to realistically rugged payoff landscapes. For example, in our setup (and in some real-world situations), a seller's payoff typically increases with small price increases, but the global best response often is to jump to a much lower price. How should economists model and identify adjustment dynamics in that sort of environment?

Appendix A

Theory

A1. *Equilibrium*

We now review some details of equilibrium calculations in the Burdett and Judd (1983) sequential search model with a continuum of sellers and a continuum of buyers. Each buyer receives $k \geq 1$ random price quotations with given probabilities

q_k , with $q_k = 0$ for $k > 3$, and purchases at the lowest of these prices if at or below the endogenous reservation price \underline{p}^* .

For an arbitrary CDF of seller prices F , each individual seller therefore faces demand

$$D(p) = \mu[q_1 + 2q_2(1 - F(p)) + 3q_3(1 - F(p))^2] \quad (\text{A1})$$

on $[0, \underline{p}^*]$ and zero elsewhere, where μ is the buyer/seller mass ratio. The mixed equilibrium distribution $F(p)$ solves $\pi(p) \equiv pD(p) = \underline{p}^* \mu q_1 \equiv \pi(\underline{p}^*)$ for all $p \in [\underline{p}, \underline{p}^*]$. That is, the profit at each price in the support interval of the equilibrium distribution is equal to the profit earned when charging the reservation price, where the seller's customers see only one price quotation. One can write $D(p)$ as $D(u)$ where $u = F(p)$, one's rank in the price distribution, giving equation (1) in the main text. The equilibrium condition therefore can be rewritten as

$$F^{-1}(u)D(u) = \underline{p}^* D(1). \quad (\text{A2})$$

This implies that the inverse equilibrium price distribution has the explicit form

$$F^{-1}(u) = \frac{\underline{p}^* D(1)}{D(u)}. \quad (\text{A3})$$

This leads directly to equation (2) for F^{-1} and equation (3) for $\underline{p} = F^{-1}(0)$. The reservation price \underline{p}^* equates the incremental cost of another search to the expected price reduction; that is, it solves

$$c = \int_0^1 (\underline{p}^* - F^{-1}(u)) dG(u) = \int_0^1 G(u) \frac{dF^{-1}(u)}{du} du = \underline{p}^*(1 - q_1), \quad (\text{A4})$$

where $G(u) = q_1 u + q_2(1 - u)^2 + q_3(1 - u)^3$, the distribution of lowest prices. This gives the result that $\underline{p}^* = c/(1 - q_1)$.

In the main text, we claimed that with a finite number of sellers, there is a symmetric mixed NE that is equal to the equilibrium distribution above. More precisely, in the finite-sellers case, if all other sellers adopt the same mixed strategy $F(p)$, then the expected sales at a price p are equal to the demand function $D(p)$ as given in equation (A1)—and thus the condition for equilibrium is identical. To see this, note first that each seller sells to μq_1 uninformed consumers (in the finite case $\mu = C/S$, where C is number of consumers and S the number of sellers). Second, a consumer who is part of the proportion q_2 who see two prices will buy from a seller if, first, it is one of the two sellers that she draws (without replacement, so that she always sees two distinct sellers) and, second, it has the lower price of those two. The probability of the first event is $2/S$, and that of the second is $1 - F(p)$, giving the second term in the demand function. The final term can be derived in a similar manner.

A2. Stability

Following Hopkins and Seymour (2002), we now give some results on stability under adaptive dynamics. Let the price range $[\underline{p}, \bar{p}]$ be approximated by an evenly

spaced finite grid of size n .⁹ That is, the strategy set is $[\underline{p}, \underline{p} + \varepsilon, \underline{p} + 2\varepsilon, \dots, \bar{p}] = [p_1, \dots, p_n]$, where $\varepsilon = (\bar{p} - \underline{p})/(n - 1)$. Let x be a point in the n -simplex so that it represents the relative frequencies of strategies, and let $X(p_i) = X_i = \sum_{j \leq i} x_j$ be the associated cumulative distribution. From this we construct a finite approximation to the Burdett-Judd continuous profit function

$$\pi_n(x) = (p_1 D(X_1), \dots, p_n D(1)). \tag{A5}$$

We then consider PDA evolutionary dynamics of the form

$$\dot{x} = Q(x)\pi_n(x), \tag{A6}$$

where $Q(x)$ is a positive definite matrix and $\pi_n(x)$ is the vector of payoffs to the different prices. The best-known dynamics of this form are the evolutionary replicator dynamics, where the relative growth rate of frequency of strategies is proportional to relative payoffs. Hopkins and Seymour (2002) find that the linearization of these dynamics at a mixed equilibrium x^* will be

$$Q(x^*)\pi'_n(x^*), \tag{A7}$$

where $\pi'_n(x) = d\pi_n(x)/dx$ is an $n \times n$ matrix. Hopkins and Seymour show that the mixed equilibrium will be asymptotically stable (unstable) under all dynamics of form (A6) if $\pi'_n(x)$ is negative (positive) definite. Proposition 4 in Hopkins and Seymour (2002) then shows that the definiteness of $\pi'_n(x)$ depends on the function

$$\Theta(u) = D'(u)D(u) - D'(u)^2 = 6q_1q_3 - 4q_2^2 - 12q_2q_3(1 - u) - 18q_3(1 - u)^2. \tag{A8}$$

If $\Theta(u) > 0$ (and thus $D(u)$ is log convex) for $u = 0$, then $\Theta > 0$ for all $u \in [0, 1]$, $\pi'_n(x)$ is negative definite, and the equilibrium is asymptotically stable, but if $\Theta(u)$ is negative (and thus $D(u)$ is log concave) for $u = 1$, then $\pi'_n(x)$ is positive definite. The curve in figure 1 surrounding the Stable area is thus $\Theta(0) = 0$, and $\Theta(1) = 0$ bounds the Unstable area. Proposition 4 of Hopkins and Seymour (2002) then implies both proposition 1 and the following.

PROPOSITION 2. The mixed equilibrium F_S in the game S is asymptotically stable under the PDA dynamics (A6), but the mixed equilibrium F_U in the game U is unstable.

To illustrate the intuition, consider $D(0.5)$, the demand from being at the mid-point of the price distribution, to be the tie payoff. Take $D(1)$ and $D(0)$ to be the loss and win payoffs from being the highest- and lowest-priced seller, respectively. So, the win premium is $D(0) - D(0.5)$, and the loss penalty is $D(0.5) - D(1)$. Then one can see in figure A1 that our Stable parameterization gives rise to convex demand, so that the win premium is bigger than the loss penalty. However, in the Unstable case, demand is linear, so that they are equal.

⁹ It is possible to extend these results to allow for prices below \underline{p} and or above \bar{p} , at the cost of additional notation. However, as these prices are dominated strategies, the population shares for these strategies would quickly approach zero.

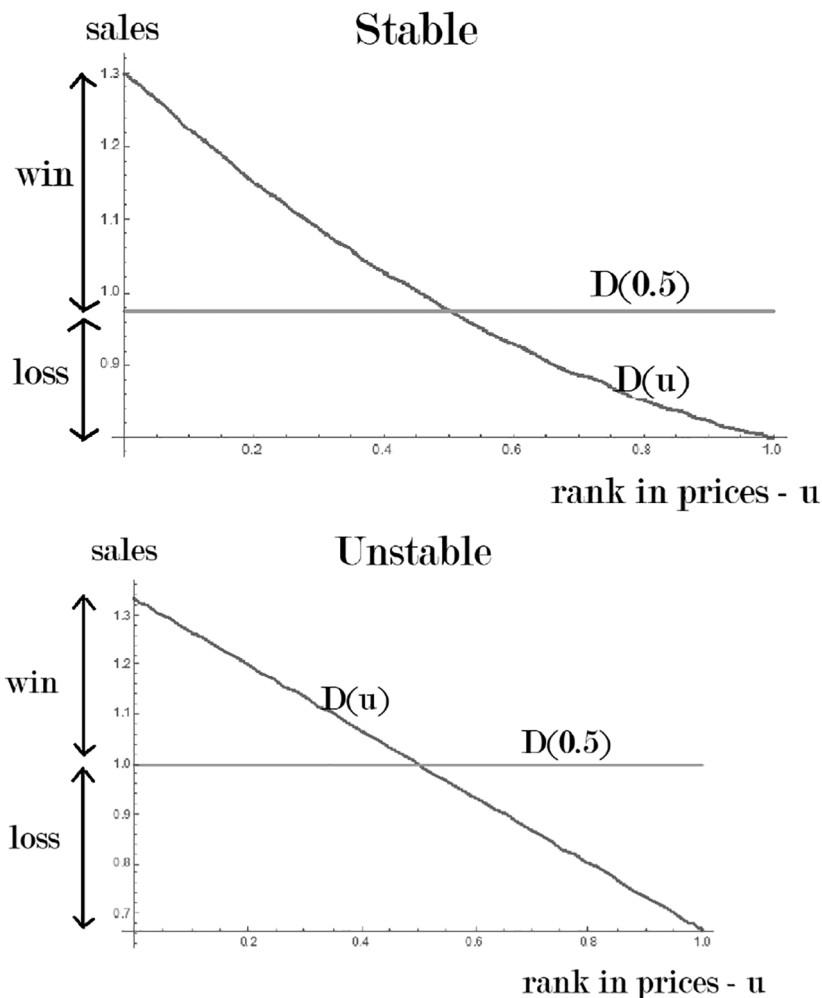


FIG. A1.—Win bonus versus loss penalty. A color version of this figure is available online.

Thus, as Hopkins and Seymour (2002) write, stability requires “sufficient ignorance,” that is, a large number of poorly informed buyers. The Stable game S has some highly informed consumers but many poorly informed consumers. This both raises the payoff to wins (because the highly informed buyers successfully find the lowest-priced seller) and reduces the size of losses (the mass of poorly informed buyers still buy from you even if you are high priced). This is what gives convexity in demand and hence stability. However, the Unstable game U does not have as many poorly informed consumers, and so losses are quite painful.

A3. *(Perturbed-)BR Dynamics*

The text sometimes refers specifically to BR or perturbed-BR (PBR) dynamics, which are described by the ordinary differential equation,

$$\dot{x} = \phi(x) - x, \tag{A9}$$

where the function $\phi(\cdot)$ is a perturbed choice function such as the logit. These PBR dynamics are discussed in detail in Hopkins (1999). The fixed point \hat{x} of the PBR dynamics is not, in general, identical to the NE but instead is a quantal response (QRE) or logit (LE) equilibrium. PBR choice functions such as the logit are usually parameterized with a precision parameter λ , which is the inverse of the noise affecting the individual's choice. Thus, the PBR dynamics approach the (exact) BR dynamics analyzed in, for example, Gaunersdorfer and Hofbauer (1995), as λ becomes large. The results of Hopkins (1999) imply that the linearization of the PBR dynamics at \hat{x} is

$$\lambda Q(\hat{x})\pi'_n(\hat{x}) - I, \tag{A10}$$

where λ is the precision parameter. Proposition 1 of Hopkins (1999) then implies the following (see also Lahkar 2011):

PROPOSITION 3. The perturbed equilibrium (QRE) \hat{x} of the game S is globally asymptotically stable under the PBR dynamics for all $\lambda \geq 0$.

The equivalent result for the unstable equilibrium is somewhat different. The perturbations have a stabilizing effect. The unstable equilibrium is unstable under the perturbed dynamics only if the precision is high enough, equivalent to the level of noise being sufficiently low. The value λ^* is found through numerical calculation.

PROPOSITION 4. In game U , the perturbed equilibrium (LE) \hat{x} is unstable under the logit form of the PBR dynamics for all $\lambda > \lambda^* \approx 0.69$.

A4. *Detecting Cycles*

What sort of behavior should we expect when the fixed point is unstable? In two dimensions, if the unique equilibrium point is dynamically unstable, then the classic Poincaré-Bendixson theorem tells us that any sort of continuous dynamics must have a limit cycle in a generalized sense. The nature of that cycle can depend on the specific dynamics. For example, when the NE is unstable in an asymmetric rock-paper-scissors game, we get convergence to a stable interior Shapley polygon under BR dynamics, but we get convergence to a heteroclinic cycle (i.e., to the edges of the simplex) under replicator dynamics. The Hopf bifurcation is a textbook example of how behavior in two dimensions can transition from convergence to an equilibrium point to regular cycles as parameters change in a two-dimensional system. Continuous-time dynamics in systems of dimension 3 and higher can be much more complicated. Cycles can be quite intricate, and chaos (or convergence to a strange attractor) is possible.

Our dynamical system can be formalized as a system of ordinary differential equations in the infinite-dimensional space \mathcal{F} of probability measures (or CDFs) whose support is contained within a fixed interval $[\underline{p}, \bar{p}]$. One can deal (as in the

previous subsection) with approximations of arbitrarily large finite dimension or (as a literal interpretation of our experiment with six sellers using pure strategies) of fixed finite dimension. In the latter case, one begins with the hypercube $[\underline{p}, \bar{p}]^6$ and then recognizes that the theory treats the sellers symmetrically, so that only the distribution of prices matters, not the identity of the sellers. The proper domain for dynamics (or state space) thus is the $1/6! = 1/320$ slice of the hypercube that satisfies the order statistics condition $\underline{p} \leq x_1 \leq x_2 \leq \dots \leq x_6 \leq \bar{p}$. That state space is still six-dimensional, of course.

Whether the state space is infinite-dimensional or six-dimensional, it seems that very complex behavior is possible. It is also possible that behavior settles down to a nice limit cycle, plus perhaps some residual behavioral noise. To detect nice cycles, we can construct a Poincaré section of codimension 1, as in figure 3, in any state space of dimension 2 or larger, including infinite-dimensional state spaces.

Given data consisting of observations of dispersed prices $(z_1, z_2, \dots, z_t, \dots, z_T)$ sampled on a fine time grid, suppose that we have constructed a Poincaré section PS at point P with tangent vector N .¹⁰ A positive transit (+1) through PS is detected when $z_t \cdot N < 0$ and $z_{t+1} \cdot N \geq 0$, that is, successive dot products go from negative to positive, that is, the trajectory passes through the section normal to the tangent vector between times t and $t + 1$.

A major advantage of this approach for us is that it works even if the cycle is a bit irregular or has a small but persistent stochastic component. Better-known spectral techniques deal less effectively with such complications. See Wang, Chen, and Wang (2017) for descriptions of other possible methods, such as velocity fields and angular momentum. For some purposes, our method has the disadvantage that it will not distinguish between limit cycles and some sorts of chaotic behavior (such as Lorenz attractors) that are almost periodic.

As noted in the main text, we chose for practical and metaphysical reasons to work with a two-dimensional projection such as (median, IQR); appendix B notes similar results with (mean, standard deviation). Nevertheless, it should be acknowledged that such a low-dimensional projection can be misleading. Noncyclic behavior can project to cyclic behavior; for example, a helix projected orthogonally to its axis looks like a cycle. Conversely, cyclic behavior can project to something that looks noncyclic; for example, a narrow ellipse projected along its major axis looks like negligible vibration around a fixed point.

Appendix B

Supplementary Data Analysis

This appendix describes the steps taken to document price cycles and also presents supplementary tables and robustness checks for the results presented in the main text.

Figures B1 and B2 reproduce the empirical CDFs for subsets of the data based on subject type (students vs. professionals) and the type of feedback sellers receive.

¹⁰ In the infinite-dimensional case, each z_t is a probability measure with CDF F , so the tangent vector N is defined by some test function ϕ with integral 0 on the price interval $[0, \bar{p}^*]$. The inner product thus is the Stieltjes integral $z_t \cdot N = \int_0^{\bar{p}^*} \phi(x)F(dx, t)$.

Table 2 summarizes the price data (median and IQR) for the entire trading period after excluding some seconds for sellers' initial adjustment. Table B1 reproduces these statistics in the columns labeled "Pooled" and also subdivides the data into different seller types and information feedback conditions. Dispersion tends to be higher with professionals, but other differences do not appear to be systematic. Table B2 pools across all markets and also subdivides the trading periods into first and second halves. Although median prices rise by one or two cents between the early and later phases of the trading period, otherwise the summary statistics are similar across the two phases.

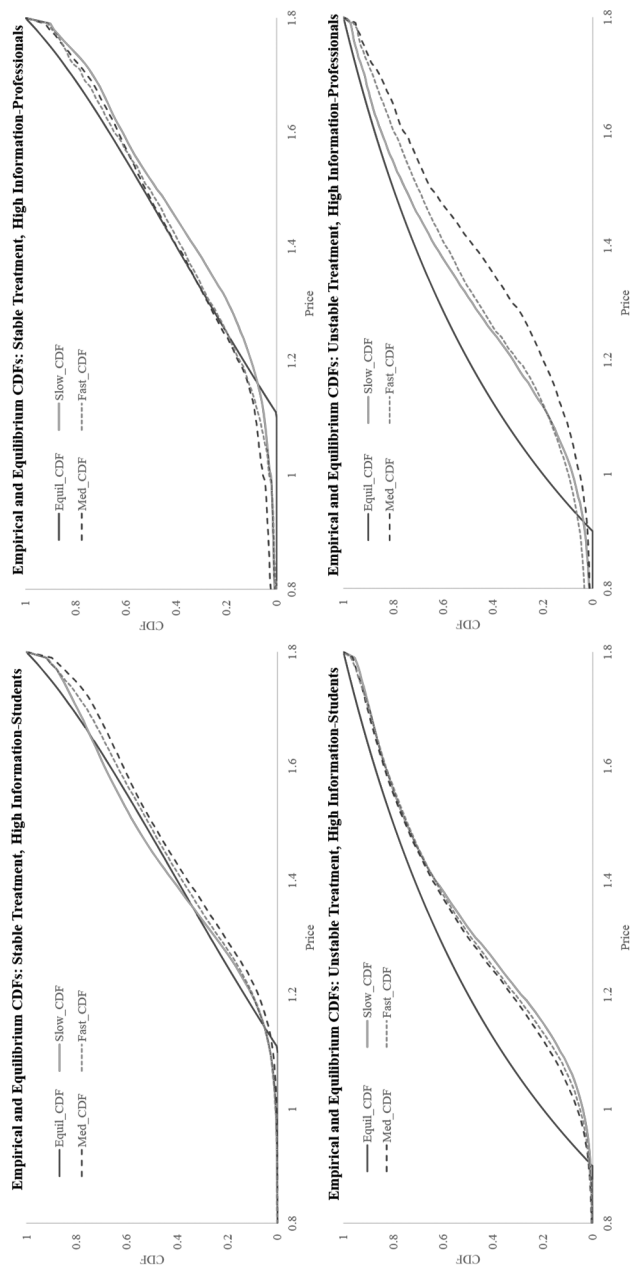


FIG. B1.—Cumulative distributions of observed versus NE prices in treatments with high information, by subject pool. Samples exclude the first 10–15 seconds of each period. A color version of this figure is available online.

TABLE B1
 MEAN OF SUMMARY PRICE STATISTICS FOR DIFFERENT INFORMATION
 FEEDBACK AND SUBJECT TYPES

	MEDIAN PRICE			INTERQUARTILE RANGE		
	(1)	(2)	(3)	(4)	(5)	(6)
High Information						
	Pooled	Students	Professionals	Pooled	Students	Professionals
NE (stable)	1.48	1.48	1.48	.37	.37	.37
Stable-fast adjustment	1.482 (.007)	1.488 (.008)	1.470 (.015)	.433 (.006)	.422 (.006)	.462 (.018)
Stable-medium adjustment	1.496 (.009)	1.505 (.009)	1.475 (.019)	.447 (.009)	.423 (.004)	.505 (.024)
Stable-slow adjustment	1.469 (.009)	1.457 (.010)	1.497 (.017)	.398 (.007)	.394 (.007)	.408 (.017)
NE (unstable)	1.20	1.20	1.20	.41	.41	.41
Unstable-fast adjustment	1.324 (.008)	1.321 (.009)	1.331 (.016)	.411 (.009)	.392 (.007)	.463 (.024)
Unstable-medium adjustment	1.339 (.011)	1.312 (.011)	1.404 (.019)	.403 (.008)	.397 (.007)	.420 (.020)
Unstable-slow adjustment	1.317 (.010)	1.321 (.011)	1.309 (.019)	.377 (.008)	.366 (.006)	.402 (.020)
Limited Information						
	Pooled	No Prices	No Payoff Projection	Pooled	No Prices	No Payoff Projection
NE (stable)	1.48	1.48	1.48	.37	.37	.37
Stable-fast adjustment	1.528 (.006)	1.532 (.009)	1.525 (.007)	.426 (.007)	.444 (.011)	.409 (.009)
Stable-medium adjustment	1.550 (.007)	1.559 (.010)	1.533 (.010)	.411 (.006)	.416 (.008)	.406 (.009)
Stable-slow adjustment	1.548 (.006)	1.549 (.008)	1.546 (.009)	.393 (.006)	.381 (.007)	.405 (.010)
NE (unstable)	1.20	1.20	1.20	.41	.41	.41
Unstable-fast adjustment	1.356 (.009)	1.350 (.013)	1.361 (.013)	.447 (.006)	.450 (.008)	.444 (.010)
Unstable-medium adjustment	1.387 (.013)	1.377 (.019)	1.398 (.018)	.439 (.007)	.448 (.008)	.430 (.011)
Unstable-slow adjustment	1.364 (.011)	1.374 (.016)	1.354 (.014)	.399 (.007)	.387 (.010)	.411 (.011)

NOTE.—Standard errors across periods are shown in parentheses. The first 10 or 15 seconds are omitted from each 90- or 150-second period.

TABLE B2
 MEAN OF SUMMARY PRICE STATISTICS FOR FIRST AND SECOND HALVES
 OF THE TRADING PERIODS

TREATMENT	MEDIAN PRICE			INTERQUARTILE RANGE		
	All Period	1st Half	2nd Half	All Period	1st Half	2nd Half
NE (stable)	1.48	1.48	1.48	.37	.37	.37
Stable-fast adjustment	1.503* (.009)	1.497 (.009)	1.508* (.009)	.430** (.009)	.435** (.010)	.424** (.008)
Stable-medium adjustment	1.518** (.010)	1.509 (.011)	1.527** (.010)	.431** (.008)	.439** (.009)	.423** (.08)
Stable-slow adjustment	1.504 (.010)	1.492 (.010)	1.516* (.011)	.396** (.006)	.397* (.008)	.395** (.006)
NE (unstable)	1.20	1.20	1.20	.41	.41	.41
Unstable-fast adjustment	1.338** (.008)	1.338** (.008)	1.338** (.010)	.427 (.010)	.433 (.011)	.422 (.009)
Unstable-medium adjustment	1.361** (.014)	1.357** (.014)	1.365** (.014)	.419 (.008)	.420 (.009)	.419 (.009)
Unstable-slow adjustment	1.338** (.011)	1.330** (.010)	1.347** (.012)	.387 (.009)	.382 (.010)	.392 (.008)

NOTE.—Standard errors across periods are shown in parentheses. The first 10 or 15 seconds are omitted from each 90- or 150-second period. Asterisks denote significant differences from Nash equilibrium for two-tailed Wilcoxon signed-rank test, using Holm-Bonferroni *p*-value correction to adjust for the family-wise error rate.

* Significant at the 5% level.

** Significant at the 1% level.

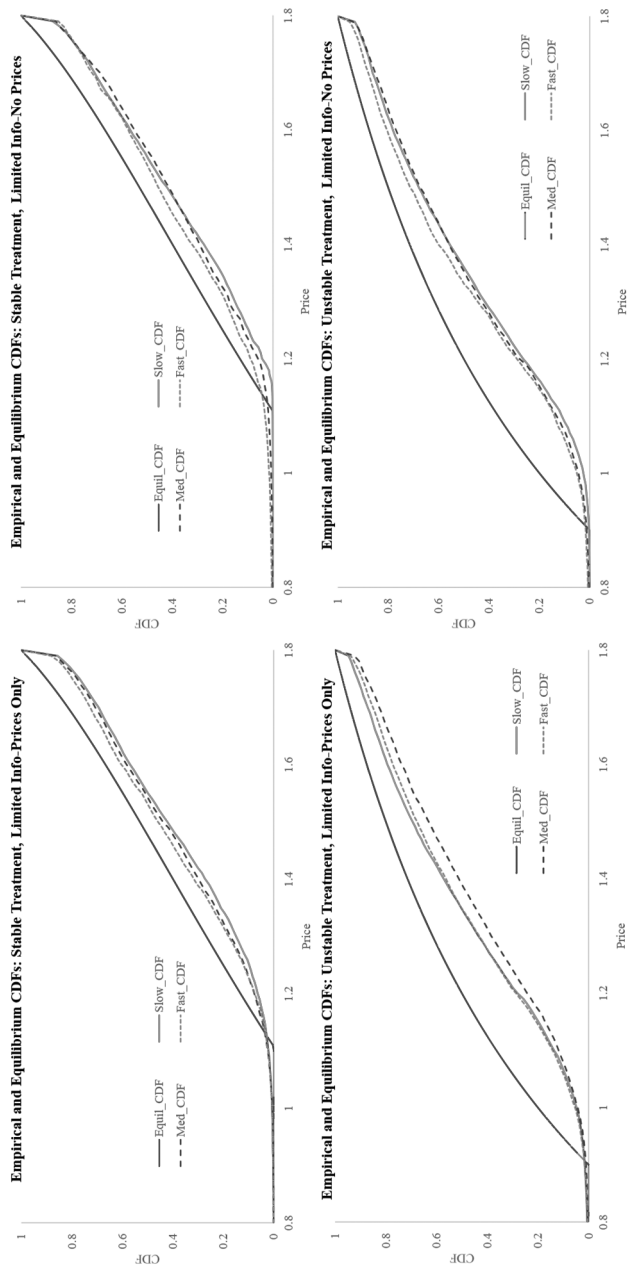


FIG. B2.—Cumulative distributions of observed versus NE prices in treatments with limited information. Samples exclude the first 10 seconds of each period. A color version of this figure is available online.

In section IV.B, we implement the Poincaré procedure for each group of six sellers and each market period as follows.

1 (Sample). Observe the vector of subjects' posted prices at 400-ms intervals, beginning at second 10 (or 15) and ending at second 90 (or 150).

2 (Two-dimensional summary). For each observation t , compute the two-dimensional summary statistic z_t , the median and IQR of the posted price vector.

3 (Center). Compute the mean of z_t . This time average defines the center $C = (C_{\text{med}}, C_{\text{IQR}})$ about which we may observe cycling.

4 (Poincaré sections). Construct the vertical tripwire as the line segment connecting C to the point $V = (C_{\text{med}}, 0)$ on the horizontal axis (median price). To check robustness, also construct the horizontal tripwire as the line segment connecting C to the point $H = (0, C_{\text{IQR}})$ on the vertical axis (IQR). See appendix A for a general discussion of Poincaré sections.

5 (Counting transits). A transit occurs when the line segment connecting successive z_t 's intersects a tripwire. If the crossing is clockwise (e.g., for the vertical tripwire, if $z_t \geq 0 > z_{t+1}$), then increment CW by 1; otherwise, increment CCW by 1.

6 (CRI). This index is defined, following Cason, Friedman, and Hopkins (2014), for each period and market as the difference between CW and CCW, as a fraction of total transits in the period.

7 (Cycle amplitude). At every point where the price median and IQR trajectory crosses a tripwire in the predicted clockwise direction, record the Euclidean distance from the cycle center. For each period and market, define the *cycle amplitude* as the median recorded distance.

Tables 4 and 5 report the average number of cycle transits, the CRI, and the average cycle amplitudes based on the vertical Poincaré section tripwire, defined as the line segment connecting the empirical cycle center to the point $V = (C_{\text{med}}, 0)$ on the horizontal (median price) axis. Tables B3 and B4 report these statistics disaggregated by seller type and information feedback. The lower CRI when sellers receive no information about other sellers' prices indicates that cycles are less regular in this information condition.

TABLE B3
MEAN TRANSITS AND CRIS FOR DIFFERENT INFORMATION FEEDBACK AND SUBJECT TYPES

	(1)	(2)	(3)	(4)	(5)	(6)
High Information						
	Students			Professionals		
	CW Transits	CCW Transits	CRI	CW Transits	CCW Transits	CRI
Stable-fast adjustment	9.66 (.71)	2.48 (.3)	.61 (.04)	8.95 (1.86)	3.25 (.69)	.51 (.08)
Stable-medium adjustment	8.20 (.47)	1.96 (.18)	.65 (.03)	6.08 (.95)	2.50 (.33)	.44 (.08)
Stable-slow adjustment	4.05 (.31)	1.14 (.16)	.60 (.04)	4.55 (.73)	1.55 (.27)	.53 (.10)

TABLE B3 (Continued)

	(1)	(2)	(3)	(4)	(5)	(6)
Unstable-fast adjustment	8.89 (.56)	1.66 (.21)	.70 (.03)	6.70 (.60)	2.85 (1.01)	.47 (.14)
Unstable-medium adjustment	7.11 (.52)	1.55 (.24)	.68 (.04)	6.05 (.81)	1.90 (.43)	.56 (.08)
Unstable-slow adjustment	3.82 (.26)	.84 (.2)	.72 (.05)	3.75 (.26)	1.54 (.23)	.53 (.09)
Limited Information						
	No Prices			No Payoff Projection		
	CW Transits	CCW Transits	CRI	CW Transits	CCW Transits	CRI
Stable-fast adjustment	5.81 (.65)	3.47 (.39)	.27 (.07)	7.06 (.70)	2.25 (.25)	.53 (.05)
Stable-medium adjustment	5.16 (.37)	3.06 (.30)	.28 (.06)	6.19 (.38)	2.56 (.42)	.50 (.07)
Stable-slow adjustment	3.59 (.18)	1.50 (.19)	.47 (.06)	4.88 (.52)	1.78 (.29)	.50 (.06)
Unstable-fast adjustment	6.50 (.68)	3.41 (.43)	.32 (.07)	5.84 (.68)	1.75 (.31)	.60 (.05)
Unstable-medium adjustment	4.50 (.51)	2.59 (.38)	.34 (.05)	5.06 (.55)	1.75 (.28)	.58 (.06)
Unstable-slow adjustment	3.16 (.29)	.72 (.14)	.71 (.05)	3.25 (.25)	.97 (.21)	.60 (.07)

NOTE.—Standard error of the means shown in parentheses.

TABLE B4
MEAN (across Periods and Markets) OF MEDIAN (Within-Period) CYCLE AMPLITUDE,
FOR DIFFERENT INFORMATION FEEDBACK AND SUBJECT TYPES

	ADJUSTMENT SPEED					
	Fast	Medium	Slow	Fast	Medium	Slow
	Students			Professionals		
High information:						
Stable	.087 (.007)	.092 (.005)	.107 (.007)	.086 (.008)	.098 (.006)	.088 (.006)
Unstable	.147 (.007)	.130 (.006)	.125 (.005)	.142 (.021)	.094 (.011)	.109 (.018)
	No Prices			No Payoff Projection		
Limited information:						
Stable	.074 (.003)	.077 (.006)	.074 (.005)	.097 (.009)	.086 (.011)	.114 (.012)
Unstable	.118 (.011)	.107 (.014)	.095 (.014)	.150 (.015)	.133 (.017)	.125 (.008)

NOTE.—Standard error of the means shown in parentheses.

To check the robustness of the conclusions to an alternative tripwire definition, tables B5 and B6 report cycle statistics pooled over all markets, with those based on the vertical tripwire shown on the left and their counterparts based on the horizontal tripwire—the line segment connecting C to the point $H = (0, C_{IQR})$ on the vertical axis (IQR)—on the right. Cycle amplitude for the horizontal tripwire is again greater in the Unstable than in the Stable case, although the differences are smaller, and they are only marginally significant for two of the three adjustment speeds.

TABLE B5
MEAN TRANSITS AND CRIS, POOLED MARKETS

TREATMENT	VERTICAL TRIPWIRE			HORIZONTAL TRIPWIRE		
	CW Transits	CCW Transits	CRI	CW Transits	CCW Transits	CRI
Stable—fast adjustment	8.06 (.50)	2.76 (.20)	.50** (.03)	9.49 (.49)	4.16 (.31)	.40** (.04)
Stable—medium adjustment	6.72 (.33)	2.43 (.16)	.50** (.04)	7.37 (.38)	3.10 (.24)	.44** (.04)
Stable—slow adjustment	4.21 (.21)	1.43 (.11)	.54** (.03)	4.53 (.24)	1.57 (.13)	.54** (.03)
Unstable—fast adjustment	7.34 (.38)	2.25 (.23)	.56** (.04)	10.26 (.41)	5.04 (.32)	.36** (.03)
Unstable—medium adjustment	5.89 (.34)	1.89 (.16)	.56** (.03)	7.14 (.32)	3.04 (.21)	.45** (.03)
Unstable—slow adjustment	3.53 (.14)	.96 (.11)	.66** (.03)	4.28 (.19)	1.42 (.12)	.57** (.03)

NOTE.—Standard error of the means in parentheses.

** Significantly different from zero, at the 1% level, in a two-tailed Wilcoxon signed-rank test using Holm-Bonferroni p -value correction to adjust for family-wise error rate.

TABLE B6
MEAN (across Periods and Markets) OF MEDIAN (Within-Period)
CYCLE AMPLITUDE, POOLED MARKETS

TREATMENT	CROSSING VERTICAL TRIPWIRE			CROSSING HORIZONTAL TRIPWIRE		
	Fast Adjustment	Medium Adjustment	Slow Adjustment	Fast Adjustment	Medium Adjustment	Slow Adjustment
Stable	.086 (.004)	.088 (.004)	.098 (.005)	.110 (.004)	.106 (.004)	.097 (.003)
Unstable	.140 (.006)	.120 (.006)	.116 (.005)	.126 (.006)	.117 (.005)	.109 (.005)
p -value for MW test	<.001**	<.001**	.018*	.008**	.098	.067

NOTE.—Standard error of the means shown in parentheses. Indicated p -values are for one-tailed Mann-Whitney (MW) tests, using Holm-Bonferroni p -value correction.

* Significant at the 5% level.

** Significant at the 1% level.

The main difference between the tripwires is that the horizontal sees consistently more CW transits as well as more CCW transits. This reduces the CRI somewhat, but it always remains very highly significant. This difference may be due to slower movements and/or more noise in the cycle phase near the horizontal tripwire. That is, there seems to be relatively less noise in the cycle phase when median price decreases at low dispersion (near the vertical tripwire) than in the phase when dispersion increases at middling median price (near the horizontal tripwire).

As noted in the main text, the cycle frequency and CRI are similar in the early and late halves of each trading period, providing no evidence of convergence even for the Stable environment. This is documented in table B7 for the vertical tripwire. Comparing each treatment condition across rows does not suggest any systematic change between the early and the late halves, for either the Stable or the Unstable market.

TABLE B7
MEAN TRANSITS AND CRIS, BY PERIOD HALVES

TREATMENT	FIRST HALF			SECOND HALF		
	CW Transits	CCW Transits	CRI	CW Transits	CCW Transits	CRI
Stable-fast adjustment	3.95 (.24)	1.26 (.15)	.58** (.04)	4.13 (.26)	1.50 (.13)	.49** (.05)
Stable-medium adjustment	3.18 (.18)	1.13 (.08)	.54** (.03)	3.54 (.20)	1.30 (.11)	.50** (.04)
Stable-slow adjustment	2.14 (.12)	.81 (.08)	.54** (.04)	2.06 (.13)	.62 (.07)	.63** (.04)
Unstable-fast adjustment	3.38 (.21)	.96 (.11)	.62** (.05)	3.96 (.22)	1.29 (.15)	.57** (.05)
Unstable-medium adjustment	2.88 (.20)	.96 (.11)	.60** (.04)	3.01 (.19)	.93 (.10)	.59** (.04)
Unstable-slow adjustment	1.82 (.18)	.53 (.08)	.66** (.05)	1.71 (.19)	.42 (.14)	.72** (.06)

NOTES.—Standard error of the means in parentheses.

** Significantly different from zero, at the 1% level, in a two-tailed Wilcoxon signed-rank test using Holm-Bonferroni p -value correction to adjust for family-wise error rate.

The cycles are constructed using robust measures of central tendency (median) and dispersion (IQR). As a further robustness check, tables B8 and B9 reproduce the CRI and cycle-amplitude calculations based on an alternative two dimensions of central tendency and dispersion: price mean and standard deviation. Comparison with tables B5 and B6 indicates that all of our conclusions regarding cycles are robust to this alternative. Although the number of transits is slightly lower (table B8), the CRI is very similar for both versions. Cycle amplitudes are lower in the (mean, standard deviation) space (table B9), but the key conclusion continues to hold: cycles are systematically larger in the Unstable than in the Stable case.

TABLE B8
MEAN TRANSITS AND CRIS BASED ON PRICE MEAN AND STANDARD DEVIATION

TREATMENT	VERTICAL TRIPWIRE			HORIZONTAL TRIPWIRE		
	CW Transits	CCW Transits	CRI	CW Transits	CCW Transits	CRI
Stable-fast adjustment	7.12 (.49)	2.69 (.22)	.47** (.04)	8.13 (.45)	3.69 (.25)	.40** (.03)
Stable-medium adjustment	5.35 (.29)	1.93 (.15)	.50** (.03)	6.34 (.28)	2.95 (.21)	.41** (.03)
Stable-slow adjustment	3.18 (.17)	1.09 (.09)	.54** (.03)	3.94 (.19)	1.63 (.11)	.47** (.03)
Unstable-fast adjustment	6.21 (.38)	1.96 (.20)	.55** (.04)	8.93 (.33)	4.32 (.30)	.37** (.04)
Unstable-medium adjustment	4.60 (.27)	1.25 (.13)	.63** (.04)	5.93 (.23)	2.51 (.20)	.46** (.03)
Unstable-slow adjustment	2.58 (.12)	.81 (.09)	.61** (.04)	3.44 (.16)	1.40 (.09)	.49** (.03)

NOTE.—Standard error of the means in parentheses.

** Significantly different from zero, at the 1% level, in a two-tailed Wilcoxon signed-rank test using Holm-Bonferroni p -value correction to adjust for family-wise error rate.

TABLE B9
MEAN (across Periods and Markets) OF MEDIAN (Within-Period) CYCLE AMPLITUDE,
BASED ON PRICE MEAN AND STANDARD DEVIATION

TREATMENT	CROSSING VERTICAL TRIPWIRE			CROSSING HORIZONTAL TRIPWIRE		
	Fast Adjustment	Medium Adjustment	Slow Adjustment	Fast Adjustment	Medium Adjustment	Slow Adjustment
Stable	.034 (.002)	.034 (.003)	.032 (.002)	.072 (.002)	.064 (.003)	.057 (.003)
Unstable	.068 (.004)	.057 (.004)	.053 (.004)	.089 (.004)	.088 (.005)	.083 (.005)
p -value for MW test	<.001	<.001	<.001	<.001	.001	<.001

NOTE.—Standard error of the means shown in parentheses. Indicated p -values are for one-tailed Mann-Whitney (MW) tests, using Holm-Bonferroni p -value correction, and all surpass the 1% significance threshold.

TABLE B10
PRICE CHANGE FREQUENCY (Percent)

	Throughout Trading Period	First Half of Trading Period	Second Half of Trading Period
All treatments	51.3	51.6	51.0
Fast adjustment speed	59.3	59.1	59.4
Medium adjustment speed	55.1	55.4	54.7
Slow adjustment speed	44.2	44.8	43.6
Unstable markets	55.0	55.0	54.9
Stable markets	47.6	48.2	47.0

NOTE.—All percentages exclude the first 10 or 15 seconds of the trading period.

TABLE B11
PRICE TARGET ADJUSTMENTS (% of Cases)

	AWAY FROM BR	TOWARD BR		
		Small toward BR	Near BR	Beyond BR
Fast adjustment speed:				
BR above current price	37.5	44.8	13.4	4.4
BR below current price	39.8	36.3	19.5	4.4
Medium adjustment speed:				
BR above current price	37.4	45.5	13.2	3.9
BR below current price	40.6	36.8	18.4	4.3

NOTE.—“Small toward BR” denotes movements in the direction of the best response but at least 5 cents short of it; “Near BR” denotes target within 5 cents of the best response; “Beyond BR” denotes movements at least 5 cents beyond the best response. This table considers only the fast and medium adjustment speeds, since the actual price often lags considerably behind the target price with slow adjustment.

TABLE B12
TREATMENT CONDITION ORDERING FOR EACH SESSION (39 Total Markets)

SESSION LABEL	BLOCK NO. (Each Block Consists of 4 Trading Periods)							
	1	2	3	4	5	6	7	8
101817	U-CS	S-CI	U-CM	S-D	U-CI	S-CS	U-D	S-CM
110117A	S-CS	U-CM	S-D	S-CM	U-CS	S-CI	U-D	U-CI
110117B	U-CI	S-CS	S-CI	U-D	U-CM	S-D	U-CS	S-CM
021918A	U-D	U-CM	S-CS	S-CI	U-CI	U-CS	S-D	S-CM
021918B	S-CI	S-CS	U-D	U-CM	S-CM	U-CI	U-CS	S-D
022018A	S-CS	U-CM	U-CS	U-D	S-D	S-CM	U-CI	S-CI
022018B	U-D	S-CS	S-D	S-CI	S-CM	U-CI	U-CM	U-CS
060119A	U-CI	S-CM	U-CS	S-CI				
060119B ^a	S-CI	U-CM	S-CS	U-CI				
073019A	U-CI	S-CS	U-CM	S-CI				
073019B	S-CM	U-CS	U-CM	S-CS				
082419A ^a	U-CM	S-CI	S-CM	U-CS				
082419B ^a	S-CM	U-CI	U-CS	S-CS				
011720	U-CS	S-D	U-CM	S-CI	U-CI	S-CS	U-D	S-CM
012320A	U-D	U-CM	S-CS	S-CI	U-CI	U-CS	S-D	S-CM
012320B	S-D	S-CM	U-CS	U-CI	S-CI	S-CS	U-D	U-CM
012320C	U-D	U-CM	S-CS	S-CI	U-CI	U-CS	S-D	S-CM
012320D	S-D	S-CM	U-CS	U-CI	S-CI	S-CS	U-D	U-CM
020620	S-CI	S-CS	U-D	U-CM	S-CM	U-CI	U-CS	S-D
020720A	S-D	U-CM	S-CI	U-CS	U-CI	S-CS	U-D	S-CM
020720B	U-D	S-CM	U-CI	S-CS	S-CI	U-CS	S-D	U-CM

NOTE.—Markets with business professionals had only four blocks (16 periods). Each session configuration included two markets of six sellers each, except for those marked. U denotes unstable, S denotes stable, and CS, CM, CI, and D denote the continuous slow price adjustment, continuous medium price adjustment, continuous fast price adjustment, and discrete pricing rounds, respectively.

^a Session had only one market.

Appendix C

Instructions to Subjects

Welcome! This is an economics experiment. If you pay close attention to these instructions, you can earn a significant amount of money, which will be paid to you in cash at the end of the last round.

Please remain silent, turn off your cell phone, and do not look at other participants' screens. If you have any questions, or need assistance of any kind, please raise your hand and we will come to you. If you disrupt the experiment by talking, laughing, etc., you may be asked to leave and may not be paid. We expect and appreciate your cooperation today.

The Basic Idea

The experiment will be divided into a number of rounds and throughout all rounds you will be anonymously matched with the same five other "counterparts" via the computer. These counterparts are other participants in today's experiment.

You and your counterparts are sellers of a fictitious good, and throughout each round you will select prices to offer to the market. The buyers in today's experiment are simulated by computerized "robots." Your price choices and those of the other sellers determine your rewards for the round. Your production costs are 0, so every sale you make directly adds to your reward, i.e., to the amount of money you will receive at the end of the experiment.

Some trading rounds will be further divided into a number of subperiods. In each subperiod you and the other sellers will privately select prices and at the end of the subperiod the combination of these prices will determine your rewards for the subperiod.

How Buyers Purchase

Buyers want to buy at the lowest price they can, and they will never pay more than \$1.80 per unit. Buyers do not see all six sellers' prices, however. Sometimes they see one price, and as long as that price is no higher than \$1.80 they will purchase at that price. Sometimes buyers see two or three prices, and in those cases they will purchase the lowest-price item they see.

We will not tell you exactly how often buyers see one, two, or three prices, or their cost to search for additional prices, but here are a few facts:

1. Your reward at any time depends entirely on your current price and the other sellers' current prices, and nothing else.
2. The function that determines how your sales quantity (and therefore your rewards) depends on your prices and others' prices will not change over the course of a round. That is, if you and the other sellers choose the same prices at time A as at time B, you will all have the same profits at time A as at time B.
3. The function that determines your sales quantity from the prices chosen can change in different rounds (due to changes in buyers' search costs or how many prices they see).

4. Your rewards are symmetric with the other sellers. That is, the function that determines how prices determine sales quantity is the same for all participants.

The Screen Display

Figure C1 shows the computer display you will use to choose prices and interact with the other sellers. You will receive rewards throughout the round that depend on your price choice and other sellers' price choices, and your rewards will be displayed on this screen.

You select your price by moving a slider (from left to right) or clicking at the bottom of the screen below the left graph. This determines a "price target" that you want to offer. Your actual price being offered to buyers may quickly or slowly adjust to this target, and this adjustment speed may differ across trading rounds. The horizontal axis covers prices from \$0.00 (the far left) to \$1.80 (the far right).

If the round is divided into subperiods, you can freely adjust your tentative price offer by clicking on the screen or dragging the slider. Your actual price for the subperiod is the location of your slider **at the end** of the subperiod. Subperiod endings are designated with vertical lines on the rewards time graph at the right of the screen (see fig. C2).

During the round (or when the subperiod is over) you will see a **circle** displaying your rewards payoff rate at that time. The higher the circle, the faster your reward accumulates. You will also be shown **circles of different colors** at the bottom of the screen showing the other sellers' prices.

In some rounds your display may also show the rewards you could receive at every possible price, based on the current (or previous subperiod) prices of all other sellers. The height of this light gray line indicates the possible rewards amount at other prices you could have chosen.

It is important to keep in mind for rounds divided into subperiods that the other sellers' prices and your reward circle height always refers to outcomes from **the previous subperiod**. You will not learn your rewards or others' prices from the current subperiod until after the subperiod is over.

The small graph on the top right of the screen shows the history of your price choices (green line [colors available online]) and the other sellers' price choices (other colored lines) throughout the current trading round. Higher lines correspond to higher prices.

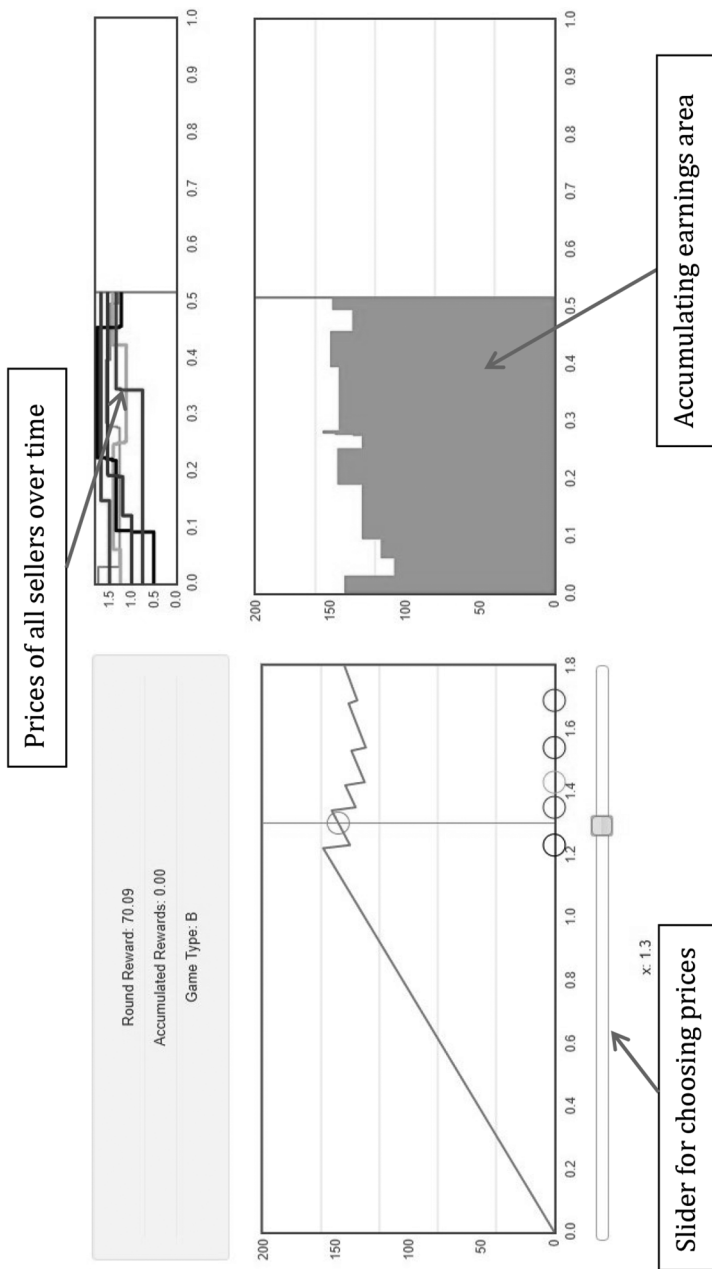


FIG. C.1.—Example seller price choice and feedback screen. A color version of this figure is available online.

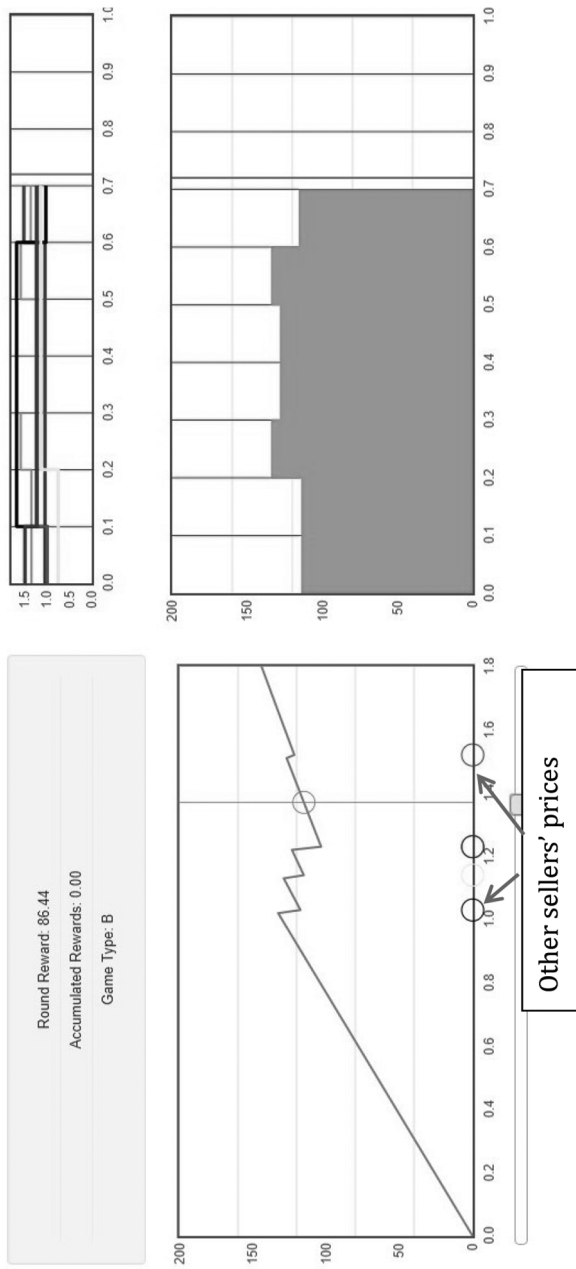


FIG. C2.—Example seller price choice and feedback screen showing subperiods. A color version of this figure is available online.





FIG. C3.—Seller display for the limited information—no prices condition. A color version of this figure is available online.

Rewards

To the right of your price choice graph you will have a display showing your accumulating rewards for the current round. Your rewards are represented by the filled area—the larger the area, the greater your accumulated rewards. The height of the filled area is based on the prices chosen by you and others at that moment (or in that subperiod). So the higher the green line, the faster your rewards are accumulating.

Your rewards will be given in points. Your points will accumulate over the course of the experiment. The upper left of your screen will always display your “Round Rewards” during the round so far and “Accumulated Rewards” earned over all previous rounds.

You will be paid cash for points earned at a conversion rate written on the white board at the front of the room.

Frequently Asked Questions

Q1. Is this some kind of psychological experiment with an agenda you haven’t told us?

Answer. No. It is an economics experiment. If we do anything deceptive or don’t pay you cash as described then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how the market works and show you how you earn money; our interest is simply in seeing how people make decisions.

Q2. What changes from one round to the next?

Answer. The likelihood that the robot buyers see one, two, or three prices might change, and their cost to search for more prices might change, both of which will change the function that determines your sales quantity from the prices chosen. The adjustment speed might also be faster or slower than in the previous round. Also, some rounds may be broken down into subperiods. But often nothing changes between rounds, and you are matched with exactly the same people in exactly the same way as in the last round.

References

- Anderson, Christopher M., Charles R. Plott, Ken-Ichi Shimomura, and Sander Granat. 2004. “Global Instability in Experimental General Equilibrium: The Scarf Example.” *J. Econ. Theory* 115 (2): 209–49.
- Andreoni, James, Michael A. Kuhn, John A. List, Anya Samek, Kevin Sokal, and Charles Sprenger. 2019. “Toward an Understanding of the Development of Time Preferences: Evidence from Field Experiments.” *J. Public Econ.* 177: 104039.
- Baye, Michael R., Dan Kovenock, and Casper G. De Vries. 1992. “It Takes Two to Tango: Equilibria in a Model of Sales.” *Games and Econ. Behavior* 4 (4): 493–510.
- Baye, Michael R., John Morgan, and Patrick Scholten. 2004. “Price Dispersion in the Small and in the Large: Evidence from an Internet Price Comparison Site.” *J. Indus. Econ.* 52 (4): 463–96.
- Benaïm, Michel, Josef Hofbauer, and Ed Hopkins. 2009. “Learning in Games with Unstable Equilibria.” *J. Econ. Theory* 144 (4): 1694–709.

- Benndorf, Volker, Ismael Martínez-Martínez, and Hans-Theo Normann. 2016. "Equilibrium Selection with Coupled Populations in Hawk-Dove Games: Theory and Experiment in Continuous Time." *J. Econ. Theory* 165:472–86.
- Brown Kruse, Jamie, Stephen Rassenti, Stanley S. Reynolds, and Vernon L. Smith. 1994. "Bertrand-Edgeworth Competition in Experimental Markets." *Econometrica* 62 (2): 343–71.
- Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–69.
- Byrne, David P., and Nicolas de Roos. 2019. "Learning to Coordinate: A Study in Retail Gasoline." *A.E.R.* 109 (2): 591–619.
- Calford, Evan, and Ryan Oprea. 2017. "Continuity, Inertia, and Strategic Uncertainty: A Test of the Theory of Continuous Time Games." *Econometrica* 85 (3): 915–35.
- Cason, Timothy N., and Daniel Friedman. 2003. "Buyer Search and Price Dispersion: A Laboratory Study." *J. Econ. Theory* 112 (2): 232–60.
- Cason, Timothy N., Daniel Friedman, and Ed Hopkins. 2010. "Testing the TASP: an Experimental Investigation of Learning in Games with Unstable Equilibria." *J. Econ. Theory* 145 (6): 2309–31.
- . 2014. "Cycles and Instability in a Rock-Paper-Scissors Population Game: A Continuous Time Experiment." *Rev. Econ. Studies* 81 (1): 112–36.
- Cason, Timothy N., Daniel Friedman, and Florian Wagener. 2005. "The Dynamics of Price Dispersion, or Edgeworth Variations." *J. Econ. Dynamics and Control* 29 (4): 801–22.
- Deneckere, Raymond J., and Dan Kovenock. 1992. "Price Leadership." *Rev. Econ. Studies* 59 (1): 143–62.
- Edgeworth, Francis Y. 1925. "The Pure Theory of Monopoly." In *Papers Relating to Political Economy*, vol. 1, 111–42. New York: Burt Franklin.
- Gaunersdorfer, Andrea, and Josef Hofbauer. 1995. "Fictitious Play, Shapley Polygons, and the Replicator Equation." *Games and Econ. Behavior* 11 (2): 279–303.
- Goeree, Jacob K., and Luke Lindsay. 2016. "Market Design and the Stability of General Equilibrium." *J. Econ. Theory* 165:37–68.
- Harrison, Glenn W., and John A. List. 2004. "Field Experiments." *J. Econ. Literature* 42 (4): 1009–55.
- Hopkins, Ed. 1999. "A Note on Best Response Dynamics." *Games and Econ. Behavior* 29:138–50.
- Hopkins, Ed, and Robert M. Seymour. 2002. "The Stability of Price Dispersion under Seller and Consumer Learning." *Internat. Econ. Rev.* 43 (4): 1157–90.
- Jensen, Robert. 2007. "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector." *Q.J.E.* 122 (3): 879–924.
- Lahkar, Ratul. 2011. "The Dynamic Instability of Dispersed Price Equilibria." *J. Econ. Theory* 146 (5): 1796–827.
- List, John A. 2020. "Non Est Disputandum de Generalizability? A Glimpse into the External Validity Trial." Working Paper no. 27535 (July), NBER, Cambridge, MA.
- List, John A., Azeem M. Shaikh, and Yang Xu. 2019. "Multiple Hypothesis Testing in Experimental Economics." *Experimental Econ.* 22 (4): 773–93.
- Maskin, Eric, and Jean Tirole. 1988. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles." *Econometrica* 56 (3): 571–99.
- Morgan, John, Henrik Orzen, and Martin Sefton. 2006. "An Experimental Study of Price Dispersion." *Games and Econ. Behavior* 54 (1): 134–58.

- Nax, Heinrich H., Maxwell N. Burton-Chellew, Stuart A. West, and H. Peyton Young. 2016. "Learning in a Black Box." *J. Econ. Behavior and Org.* 127:1–15.
- Nax, Heinrich H., and Bary S. R. Pradelski. 2015. "Evolutionary Dynamics and Equitable Core Selection in Assignment Games." *Internat. J. Game Theory* 44 (4): 903–32.
- Noel, Michael D. 2007. "Edgeworth Price Cycles, Cost-Based Pricing, and Sticky Pricing in Retail Gasoline Markets." *Rev. Econ. and Statis.* 89 (2): 324–34.
- Oprea, Ryan, Keith Henwood, and Daniel Friedman. 2011. "Separating the Hawks from the Doves: Evidence from Continuous Time Laboratory Games." *J. Econ. Theory* 146 (6): 2206–25.
- Poincaré, Henri. 1890. "Sur le problème des trois corps et les équations de la dynamique." *Acta Mathematica* 13 (1–2): 1–270.
- Selten, Reinhard, and Rolf Stoecker. 1986. "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach." *J. Econ. Behavior and Org.* 7 (1): 47–70.
- Slade, Margaret E. 1998. "Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices." *Rev. Econ. Studies* 65 (1): 87–107.
- Smith, Vernon L. 1962. "An Experimental Study of Competitive Market Behavior." *J.P.E.* 70 (2): 111–37.
- Stephenson, Daniel. 2019. "Coordination and Evolutionary Dynamics: When Are Evolutionary Models Reliable?" *Games and Econ. Behavior* 113:381–95.
- Stephenson, Daniel G., and Alexander L. Brown. 2017. "Predicting Behavior in Disequilibrium in Continuous Space and Time: What Do Equilibrium, Stability Criteria, and Adaptive Models Get Right?" Working paper, Dept. Econ., Texas A&M Univ.
- Stigler, George. 1961. "The Economics of Information." *J.P.E.* 69 (3): 213–25.
- Varian, Hal R. 1980. "A Model of Sales." *A.E.R.* 70 (4): 651–59.
- Wang, Yijia, Xiaojie Chen, and Zhijian Wang. 2017. "Testability of Evolutionary Game Dynamics Based on Experimental Economics Data." *Physica A* 486: 455–64.
- Wang, Zhongmin. 2009. "(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation." *J.P.E.* 117 (6): 987–1030.