

Continuous Differentiation: Hotelling Revisits the Lab

Curtis Kephart and Daniel Friedman*

October 21, 2014

Abstract

We investigate experimentally the impact of continuous time on a four-player Hotelling location game. The static pure strategy Nash equilibrium (NE) consists of firms paired-up at the first and third quartiles of the linear city. In repeated simultaneous move games (discrete time grid), we fail to obtain convergence to this distinctive NE, as have previous studies. However, with asynchronous moves in continuous time treatments, the NE clearly emerges.

1 Introduction

In his seminal paper on spatial competition, Hotelling (1929) analyzed the behavior of two sellers of a homogenous product choosing price and location in a bounded, one-dimensional marketplace. The Hotelling model has since been expanded to allow numerous sellers to interact strategically in more general marketplaces. As the preeminent model of spatial competition it has been widely applied, e.g., in industrial organization to analyze geographic competition and product differentiation (Netz and Taylor, 2002), and in political economy as a tool to analyze voting dynamics (Downs, 1957).

We investigate the dynamic foundations of this static model. When firms can quickly adjust product characteristics, and quickly respond to competitors' repositioning, where (if anywhere) does behavior settle down? The question is theoretical but quite relevant

*Kephart: Economics Department, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064 (email: Curtis.Kephart@googlemail.com); Friedman: Economics Department, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064 (email: dan@ucsc.edu)

to 21st century applications. New technology and processes allow – and competitive pressures arguably force – firms to reduce the time between product versions. New marketing techniques allow firms to quickly and continuously modify how consumers perceive a product. Now more than ever, firms compete on their agility in bringing new versions of products to market.

In some applications the adjustment rate, although rapid, remains bounded. To capture this possibility, we impose “speed limits” in some treatments, and allow arbitrarily rapid adjustment in other treatments. We apply these treatments as well as the standard discrete time (finitely repeated) treatment to a four player model of location choice (no price adjustment) with a uniform customer distribution over a bounded line segment. Our study is empirical and includes agent-based simulations as well as a human subject laboratory experiment.

Previous empirical studies of the Hotelling location model, apart from voter analysis, have been scarce (Brenner, 2011). In product differentiation, the lack of agreement as how to best measure product position has hampered field studies. Also, firms have been understandably reluctant to allow public access to their data on product positioning and sales volume. Laboratory experiments can overcome the data access difficulties, but so far the most relevant lab results have been disappointing.

Brown-Kruse et al. (1993) and Brown Kruse and Schenk (2000) investigate a duopoly model with varying customer densities over a bounded one-dimensional action space. Collins and Sherstyuk (2000) look into a three-agent model with inelastic demand and uniform prices. There exists no pure-strategy Nash equilibrium in this set-up (Eaton and Lipsey, 1975). Shaked (1982) finds a unique mixed strategy equilibrium in which players randomize uniformly over the second and third quartiles, but Collins and Sherstyuk find little support for it in their laboratory study.

In laboratory work most closely related to ours, Huck et al. (2002) investigate a four-player implementation of the location-only model. Eaton and Lipsey show that the

unique pure-strategy Nash equilibrium has two players located back-to-back at the first quartile, and the other two players similarly located at the third quartile. Like Collins and Sherstryuk, Huck et al. find little support of the Nash equilibrium prediction. Subjects instead exhibit a “W-shaped” distribution of locations, with significant clustering near the second quartile (the median) as well as near the first and third quartile.

We focus on the four-player location-only game because we believe it to be diagnostic. The central location is an obvious choice in a two player game, and deviations can be written off as endogenous idiosyncratic conventions, *folie à deux*. The three player game has no Nash equilibrium in pure strategies, and it is relatively easy to dismiss support (or lack of support) for the mixed equilibrium. By contrast, the four player Hotelling location game has a distinctive and unique Nash equilibrium. It is in pure strategies and contrary to most people’s initial intuition: the equilibrium locations consist of two widely separated back-to-back pairings. Discovering when (if ever) this equilibrium arises can suggest broad insights into spatial equilibration.

We run preliminary simulations of an agent-based model with agents pursuing the myopic best response. Behavior fails to converge to the NE when the four agents make simultaneous choices, but when they choose asynchronously they eventually converge precisely to the static pure Nash equilibrium just described. Encouraged by that finding, we use novel software to conduct human subject laboratory experiments with the four agent, location-only Hotelling model. As explained in Section 4, we too find that subjects fail to converge to the Nash equilibrium in discrete time (simultaneous choice), but also find good convergence to the NE in continuous time (asynchronous) treatments with and without speed limits.

2 Hotelling Location Model

Each firm $i = 1, 2, \dots, n$ chooses a location $s_i \in [0, 1]$. Firms produce homogenous goods with identical mill prices and linear transport costs. Each of a uniform continuum of consumers inelastically purchase a single unit at the lowest delivered price, i.e., from the closest firm.

Payoffs are determined as follows. Sort the strategy profile (s_1, s_2, \dots, s_n) , so $S_{[1]} = \min\{s_1, s_2, \dots, s_n\}$, $S_{[2]}$ is the second lowest location, ..., and $S_{[n]} = \max\{s_1, s_2, \dots, s_n\}$, so $S_{[1]} \leq S_{[2]} \leq \dots \leq S_{[n]}$. If there are exact ties ($s_i = s_j$), then average the payoffs defined below over all feasible assignments of the tied players.

Normalizing unit profit to 1.0, firm i 's payoff is the length of its territory. As illustrated in Figure 1, that territory (except for the 'edge' players [1] and [n]) extends from the midpoint of the interval $[S_{[i-1]}, S_{[i]}]$ with the firm just below to the midpoint of the interval $[S_{[i]}, S_{[i+1]}]$ with the player just above. Thus the payoff function is

$$\Pi_i = \frac{1}{2}(S_{[i+1]} - S_{[i-1]}), i = 2, \dots, n - 1 \quad (1)$$

with

$$\begin{aligned} \Pi_1 &= S_{[1]} + \frac{1}{2}(S_{[2]} - S_{[1]}) \\ \Pi_n &= (1 - S_{[n]}) + \frac{1}{2}(S_{[n]} - S_{[n-1]}). \end{aligned} \quad (2)$$




Figure 1: Hotelling Linear City action space

Equilibria in this game are sensitive to the number n of competing firms. It is well known that for $n = 4$ there is a unique pure NE.

Proposition. The Hotelling location game for $n = 4$ players has a unique pure Nash equilibrium, up to relabeling of players. The unique sorted equilibrium profile is $S_{[1]} = S_{[2]} = \frac{1}{4}$ and $S_{[3]} = S_{[4]} = \frac{3}{4}$.

That is, in NE, players are paired “back-to-back” at the first and third quartiles. It is easy to check that this is indeed a NE: by equation (2) a deviation to $[0, .25)$ or to $(.75, 1]$ clearly shrinks the deviator’s territory and payoff, while by (1) a deviation to $(.25, .75)$ shifts the deviator’s territory but does not increase payoff. For a complete formal proof, see appendix B of Huck et al. (2002), and for a proof of the uniqueness of pure NE, see Eaton and Lipsey (1975).

2.1 Dynamic Considerations

In previous laboratory examinations of the Hotelling location-only game (Collins and Sherstyuk, Huck et al.), subjects were given random initial positions and allowed to select new actions simultaneously in discrete time, i.e., in a finitely repeated game. Since the game is constant sum (the total payoff is always 1.0) and only finitely repeated, there is no scope for Folk Theorem complications. However, in this standard laboratory implementation of the static Hotelling model, subjects face considerable strategic uncertainty — to chose well, they must accurately predict their opponents’ next location choices. The difficulty of predicting increases considerably as the number of players increases beyond $n = 2$. Thus it is easy to question the relevance of the static pure Nash equilibrium to a game with this type of dynamic structure.

To help understand what we might see in the dynamic game we ran computer simulations. Automated agents played the myopic best response: in period $t + 1$ each player allowed to move chooses a location that would maximize payoff given the period t location profile of the other players. In simulations in which agents moved simultaneously

(everyone allowed to move every period), all agents nearly always converge to the midpoint 0.5 of location space and remain concentrated in the neighborhood. The static model's NE tends to appear only when initial locations are already at the NE, or when locations were quite close to the NE and moves were limited to small increments. By contrast, in asynchronous simulations, only one player gets to move each period. In such simulations agents do tend to converge to the NE from a diverse set of initial locations. These simulations suggest that the static model's predictive power may hinge on the dynamic specifications.

Pettit et al. (2013) note several reasons why continuous time dynamics may deliver outcomes that differ from the usual simultaneous move finitely repeated implementation. However, they do not emphasize the role of strategic uncertainty, which our simulations suggest may be key to the Hotelling location game. Moves are almost always asynchronous in continuous time, so in that respect it is similar to our turn-based simulations. When a player selects a new position she can take the locations of her opponents as fixed, at least temporarily. And when their opponents do relocate, players can quickly respond with their own relocations.

2.2 Testable Predictions

Our human subject experiment is designed to test two hypotheses suggested by the preceding theoretical discussion.

Hypothesis 1: The static Nash equilibrium prediction will become more accurate over time.

Hypothesis 2: The static Nash equilibrium prediction will be more accurate in continuous time (asynchronous) treatments than in the discrete time (synchronous) treatment.




Figure 2: Hotelling game interface for Continuous Treatments.

3 Experiment

The experiment was programmed in ConG, software designed to implement continuous time economics experiments (Pettit et al., 2013). Subjects choose their target location using their mouse to click or to drag a slider, the black rectangle seen at the bottom of the screen in Figure 2. Subjects may also use the left and right arrow keys to shift locations incrementally. The horizontal position of the large green dot indicates the subject's current location, and other players' current locations are indicated with smaller dots of different colors. The vertical height of each dot indicates the players' current flow payoffs, also indicated by the number displayed by each players' dot. Accumulated flow payoffs are shown in the "Current Points" field, while points earned in all previous paid periods are indicated in the "Cumulative Points" field.

3.1 Treatments

We study three main treatments. The first is discrete time (“Discrete”). Periods are divided into n equal-length subperiods. Within each subperiod, subjects are freely able to move their location target using the mouse or arrow keys, but subperiod payoffs are determined solely by the target location profile chosen at the very end of the subperiod. Only at that point do subjects see the other players’ chosen locations. Subjects see a progress bar (not visible in Figure 2) filling smoothly to indicate when the subperiod will end, and a “Subperiods Left” field counts down until the end of the period. Payoffs for the entire period are the integral across subperiods of the piecewise constant flow payoffs, or equivalently, the average of the lump sum subperiod payoffs. A video of this and the other two treatments may be seen online at <http://youtu.be/NX6L1mV9iII>.

The other two treatments are continuous time. In continuous-time slow (“Slow”) when subjects select a new target location their current position moves toward the target at a constant rate (the “speed limit”). The chosen speed limit is such that it would take 30 seconds to traverse the entire interval.

In continuous-time instant (“Instant”), the subject’s current position moves immediately to the chosen target, with no perceptible delay. Actual latencies in our lab are less than 50 milliseconds; that is, no more than 50 ms elapse from the time a subject clicks a new location until the time when the effect of that click is revealed on all subjects’ displays. Of course, human reaction times are considerably longer than that, and subjects perceive the action as continuous in this treatment.

3.2 Procedures

Treatments are varied across sessions. Subjects are matched and rematched within “silos” of six subjects, and most sessions include two silos. Sessions run for twelve periods, each period lasting three minutes. Discrete time periods are broken into 60 subperiods of 3 seconds each.

Table 1: Matching protocol

Subject	Period											
	1	2	3	4	5	6	7	8	9	10	11	12
1	R1	4p	4p	4p	4p	R2	4p	4p	R3	4p	4p	R4
2	R1	4p	4p	R2	4p	4p	4p	R3	4p	4p	R4	4p
3	4p	R1	4p	4p	R2	4p	4p	4p	R3	4p	R4	4p
4	4p	R1	4p	R2	4p	4p	R3	4p	4p	R4	4p	4p
5	4p	4p	R1	4p	R2	4p	R3	4p	4p	4p	4p	R4
6	4p	4p	R1	4p	4p	R2	4p	R3	4p	R4	4p	4p

Note: Player assignments are shown for each 6-player silo. The four subjects labelled “4p” in each column play the Hotelling location game that period; the other two subjects each play against three automated agents (“robots”) using one of four algorithms (R1 to R4) described in the text.

As shown in Table 1, each period a different subset of six subjects in the silo play the $n = 4$ location game, while the two excluded subjects each play separately against three automated agents (“robots”). Robots reset their target position every 30 seconds following a specified sequence. For example in continuous time treatments subjects matched with robots in periods 1 through 3 face algorithm R1, which begins with the target profile $\{0.25, 0.5, 0.75\}$ for 30 seconds, followed by $\{0.12, 0.33, 0.55\}$, and then four other profiles.

The robot algorithms are held constant within each three-period set (periods 1-3, 4-6, 7-9 and 10-12), allowing each subject to face it once. These robot periods serve as a diagnostic for individual subjects’ ability to best respond, measured by payoffs against robots relative to other subjects. Of course, we can’t compare across different treatments, since the robot performance and the human opportunities are so different. For example, it is far easier to best respond against robots in Instant (where they are sitting ducks for 30 seconds at a time) than in Slow (where they might be moving steadily much of the time and the human subject is also constrained by the speed limit.)

Several considerations motivate our choice of the matching procedures. Pilot experiments using fixed matchings often produced arbitrary conventions on how to split the territory evenly, some of them consistent with the Nash equilibrium, and others not. Either way, a fixed (“partners”) matching protocol doesn’t seem like the best way to test hypotheses based on static non-cooperative games. At the other extreme, matchings randomized over all subjects in the session produce fewer independent observations, and pilot experiments using this protocol seemed to suffer from “contamination” by confused subjects. That is, in our pilots a few individuals never seemed to understand the game well, as indicated by post-experiment surveys and relatively low session payouts, and some other players exposed to one of these confused individuals would exhibit less systematic behavior later in the session, even when no longer playing against that individual. The six-person silo seems like a good compromise that allows remixing between periods while limiting the number of other players exposed to a confused player. Also, performance against robots provides a measure of an individual’s to best respond, albeit a better measure in some treatments than in others.

Sessions were conducted in the LEEPS laboratory at the University of California, Santa Cruz in February and March 2013. A total of 54 human subjects (randomly assigned into 9 silos) were drawn from the LEEPS Lab subject pool – composed of undergraduates across the curriculum – using the subject recruitment software ORSEE (Greiner, 2004). Sessions lasted about 90 minutes, and began with a standard Holt-Laury risk preference elicitation, then instructions¹ for the Hotelling game, two practice periods and a quiz, followed by the 12 periods specified in Table 1, and finally cash payment.

Our location game is constant sum, and points were scaled so that each period the total score of all four players summed to 400 – or 100 points per-player per-period on average. Subjects received a \$5 show-up fee and between \$0.025 and \$0.034 for each

¹See the web appendix associate with this paper at leeps.ucsc.edu

point earned over the entirety over the session in excess of 1000 points, including both all-human periods and robot periods. This dollar amount was added to the subject's payoff from the Holt-Laury game and rounded up to the nearest quarter. Participating subjects were paid on average \$14.22.

4 Results

Figure 3 offers a concise overview of behavior in the human-only 4 player location games. The three left-side panels show a typical three-minute period from each treatment, plotting each player's actual locations sampled ten times a second in panels b1 and c1, and subperiod-by-subperiod in panel a1. The three corresponding panels on the right side show the overall distribution of locations by treatment in the last half of each session (periods 7-12).


Panel a1 shows that some players (like Blue) tend to jump around from one Discrete subperiod to the next, while others (like Red and Yellow) are inclined to settle down for a while. Overall, there is little trend towards the NE configuration in this sample period. Panel a2 shows for Discrete modest modes near the NE quartile points 0.25 and 0.75, but also a mode near the non-NE quartile (median) point 0.50, and a noticeable aversion to edge locations. The data of Huck et al. has similar properties, but all three modes are sharper and location choices between the modes are less common than in our data.

The slopes in Panel b1 reflect the "speed limit" we imposed in the Slow treatment. Orange and Blue take about 80 seconds to sort themselves out, but after that, the location profile remains fairly close to the NE. Panel b2 shows that we have strong modes in the vicinity of both NE quartile points, and no mode at 0.5; indeed, choices distant from .25 and .75 are rare and transient.


Panel c1 shows that not all players in the Instant treatment continue to exercise their

Examples of Session Periods


(a1) A Discrete Time Treatment's Period
(period 8 from March 3, 2013 session, all humans)



(b1) A Continuous Time, Slow Speed Limit Treatment Period
(period 8 from February 26, 2013 session, all humans)

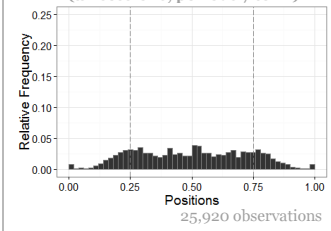


(c1) A Continuous Time, Instant Speed Limit Treatment Period
(period 8 from February 25, 2013 session, all humans)




Histograms of Player Positions


(a2) Discrete time treatments,
(all sessions, periods 7 to 12)



(b2) Continuous time, slow,
(all sessions, periods 7 to 12)



(c2) Continuous time, instant,
(all sessions, periods 7 to 12)



Note: Panels a1, b1, and c1 plot actual locations over time for each of the three treatments from the same randomly selected period (#8) of a randomly selected silo. Panels a2, b2, and c2 plot location distributions in all all-human games from the later half of sessions (periods 7 through 12).

Figure 3: Example periods, and distribution of player locations, by treatment.

ability to jump around. About 20 seconds in, Red settles near 0.75 and scarcely budes after that, and Black, after trying other locations, settles near 0.25 about 70 seconds in. Orange and Blue sometimes play near Nash strategies, but sometimes bounce around to no apparent purpose. Overall, Panel c2 shows sharp modes at the NE quartile points, and rather little activity elsewhere. (The small mode at 0.5 mainly reflects the behavior of a single player who stubbornly occupied that position for a number of periods in one session, despite unimpressive payoffs.)

The histograms in Figure 3 suggest better convergence to NE in Instant than in Slow, and little or no convergence in Discrete. To test our hypotheses more formally, we define a metric for the average absolute distance from NE, as follows. Sort the players' time- t locations as usual: $S_{[1]t}$ is the location of the "left most" player at time t , $S_{[2]t}$ is second from the left, etc., so that $S_{[1]t} \leq S_{[2]t} \leq S_{[3]t} \leq S_{[4]t}$. Then *AvgAbsDist* at time or subperiod t is:

$$AvgAbsDist_t = \frac{1}{4} (|S_{[1]t} - 0.25| + |S_{[2]t} - 0.25| + |S_{[3]t} - 0.75| + |S_{[4]t} - 0.75|) \quad (3)$$

Clearly *AvgAbsDist* = 0 at a NE profile, and is bounded above by 0.5 (achieved when all $S_{[i]t} = 0$ or all = 1). Its expected value for a sample drawn from a uniform random distribution can be shown to be 0.1715, with a median of 0.1646.

The overall median value of *AvgAbsDist* in Discrete is 0.1417, insignificantly different from the random benchmark of 0.1715, while the median values in Slow and Instant, 0.0639 and 0.0461 respectively, are significantly lower. However, looking at the overall average *AvgAbsDist* of each session period we can reject the null hypothesis that the median is 0 in each treatment.

Is there a trend towards NE across periods? Figure 4 plots *AvgAbsDist* in human-only location games averaged period-by-period over all silos in each treatment. In the two Continuous treatments we do see a trend towards 0 (i.e., towards NE) in the first




Figure 4: AvgAbsDist by period and treatment

few periods, but it seems to level off thereafter. In Discrete periods, $AvgAbsDist$ is not much below the random benchmark (the gray dashed line) and has no noticeable downward trend.

Are there trends within a typical period? Figure 5 plots $AvgAbsDist$ in human-only location games averaged second-by-second (or, for Discrete, subperiod-by-subperiod) over all periods in each treatment. The software initializes subjects at uniform random locations, thus at time zero $AvgAbsDist$ is about 0.17, as expected. Again, in the Discrete treatment, there is little trend towards NE within the period, while in the continuous treatments there is a clear trend at first but it stagnates after about 40 seconds.

Our main hypothesis tests are based on the following variant of the regression model of Noussair et al. (1995), applied to human-only data from latter half of all sessions (periods 7 through 12):

$$AvgAbsDist_{jt} = \beta_0 \left(\frac{1}{t} \right) + \beta_1 \cdot D_1 \left(\frac{t-1}{t} \right) + \beta_2 \cdot D_2 \left(\frac{t-1}{t} \right) + \beta_3 \cdot D_3 \left(\frac{t-1}{t} \right) + \varepsilon_{jt} \quad (4)$$

where $AvgAbsDist_{jt}$ is the observed per-player absolute distance from the pure Nash

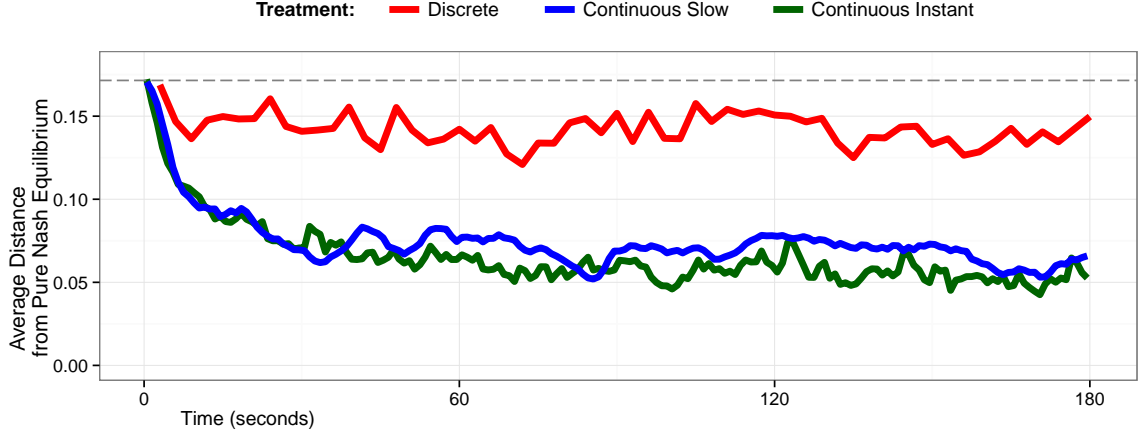


Figure 5: AvgAbsDist by second, by treatment over the period.

equilibrium configuration, in treatment j at time t – each 100 milliseconds in continuous time or subperiod t in the case of discrete time. D_1 is a dummy variable that is equal to 1 if the observation is from treatment $j = 1$, Discrete time. Likewise the dummies D_2 and D_3 are for observations from the Slow and Instant time treatments. The β_0 term has the interpretation as the y-axis intercept, i.e. where the time series starts at time $t = 1$. The coefficient β_j can be interpreted as the value $AvgAvgDist_{jt}$ converges toward as $t \rightarrow \infty$ in treatment j , since $\frac{t-1}{t} \rightarrow 1$ as $t \rightarrow \infty$. Equation (4) imposes a single origin coefficient, and imposes equal convergence target coefficients β_j across silos and periods, but allows those targets to vary across treatments. The ordinary least square estimates are shown in Table 2.

Result 1: Locations in the Discrete time treatment remain far away from the Nash equilibrium profile.

Support for result 1: Coefficient estimate $\hat{\beta}_1 = 0.140$ in Table 2 is statistically and “economically” very different from zero. There is a weak tendency for the distance from NE to decrease over time, inasmuch as $\hat{\beta}_1$ is less than the estimate $\hat{\beta}_0 = 0.264$ of the origin parameter, but the main point is that the estimated asymptotic distance $\hat{\beta}_1$ is not much below the random benchmark of 0.17.

Table 2: Estimates of Equation 4

$\widehat{\beta}_0$: Intercept	0.264*** (0.0030)
<i>Convergence Targets by Treatment</i>	
$\widehat{\beta}_1$: Discrete	0.140*** (0.0020)
$\widehat{\beta}_2$: Continuous Slow	0.063*** (0.0003)
$\widehat{\beta}_3$: Continuous Instant	0.048*** (0.0003)
Observations	65,492
Adjusted R ²	0.616
Residual Std. Error	0.052 (df = 65488)
F Statistic	26,297.220*** (df = 4; 65488)

Note: Dependent variable is AvgAbsDist. Standard errors in parentheses, and *** \Rightarrow $p < 0.01$. Only data from periods 7 through 12 used.

Result 2: Continuous time treatments exhibit convergence toward the pure Nash equilibrium.

Support for result 2: The estimated asymptotic distances from NE in the Slow and Instant treatments are $\widehat{\beta}_2 = 0.0632$ and $\widehat{\beta}_3 = 0.0478$ respectively. These are statistically different from zero, but “economically” speaking they are much closer to zero than to the random benchmark 0.17. Both estimates are also well below the estimate of the origin, $\widehat{\beta}_0 = 0.264$. We conclude that behavior moves decisively towards (but not all the way to) the Nash equilibrium in these treatments.

Result 3: Continuous time Instant treatments exhibit better convergence toward equilibrium than continuous Slow.

Support for result 3: The coefficient estimate $\widehat{\beta}_3 = 0.0478$ for Instant is significantly less than the estimate $\widehat{\beta}_2 = 0.0632$ for Slow, according to the usual t -test ($Pr(> |t|) < 0.001$).

A more nuanced picture can be obtained by comparing the cumulative distribution functions across treatments of average absolute distance from NE, as in Figure 6. One can see that the Discrete treatment has a vastly different distribution than the two

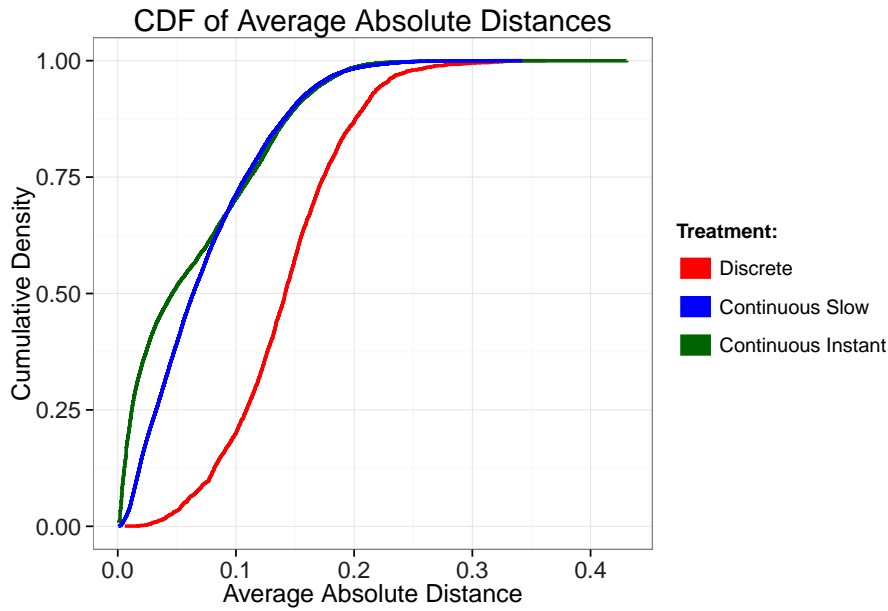


Figure 6

continuous treatments, and that Instant has a much larger fraction of tiny distances (near NE behavior) than Slow. The Figure also shows that the cdfs are close for distances exceeding about 0.07, i.e., the upper tails are similar. A Kolmogorov-Smirnov test confirms that the overall differences in distributions are significant ($p < 0.001$ for all pairwise comparisons).

Table 3 lists the percent of AvgAbsDist observations for which it is implied that subjects are 1 percent and 5 percent away from the pure Nash equilibrium, i.e. the percent of AvgAbsDist observations less than 0.0025 and 0.0125 respectively. We see that in 25.86 percent of observations players in continuous Instant treatments choose to locate within 5 percent from the equilibrium configuration, compared with 6.38 percent in continuous Slow treatments and 0.05 percent in Discrete time treatments. Continuous Instant subjects spend fully 3-5 times as much of their session near equilibrium than continuous slow subjects.

Individual Holt-Laury risk aversion test results are not significantly correlated with

Table 3: Attainment of Near-Equilibrium Location Formations by Treatment

	Discrete	Continuous Slow	Continuous Instant
Periods 1-3	0.0% (0.0%)	0.01% (2.53%)	0.3% (8.76%)
Periods 4-6	0.0% (0.0%)	0.02% (4.33%)	0.41% (22.08%)
Periods 5-9	0.0% (0.19%)	0.01% (6.59%)	8.5% (33.51%)
Periods 10-12	0.0% (0.0%)	0.74% (12.05%)	6.28% (39.06%)
Overall	0.0% (0.05%)	0.20% (6.38%)	3.87% (25.86%)

Note: The table list the percent of observations within 1 percent of Pure NE ($\text{AvgAbsDist} \leq 0.0025$), and within within 5 percent of Pure NE in parentheses ($\text{AvgAbsDist} \leq 0.0125$). Observations are based on subject location configurations sampled ten times a second in continuous time, and each subperiod in Discrete.

payoffs in all-human location games nor with payoffs against robots, nor are they correlated with membership in groups that achieve near-NE outcomes. The only regularity we detected is that subjects who completed the H-L survey nonsensically or who scored as “risk loving” (yielding estimates $-0.49 < r < -0.15$ for $U(x) = x^{1-r}/(1-r)$) did tend to earn lower than average payoffs when playing against follow humans. Nonsensical answerers and risk lovers earned 92 points and 95 points per period on average against humans, respectively, compared to an average of 104 for risk neutral subjects.

5 Conclusion

To summarize briefly, we are able to replicate the negative result of Huck et al. (2002) that behavior fails to converge to the distinctive Nash equilibrium in 4 player Hotelling location games played in Discrete time. More importantly, we have identified (as far as we know, for the first time) conditions that lead to good NE convergence. In both continuous time treatments, Slow and Instant, we get convergence to a neighborhood of (but not precisely to) NE both within a typical period and across periods. The neighborhoods are tighter under Instant, as confirmed by both parametric tests (in the tradition of Noussair et al. (1995)) and nonparametric tests (Kolmogorov-Smirnov).

We draw two broader conclusions from our results. First, Hotelling models are

more relevant to understanding behavior in the wider world than might previously have been supposed. In particular, in terms of product positioning in a space of features or perceived characteristics, our Instant and Slow treatments capture important aspects of competition in the 21st century. In some situations firms are freely able to reposition their product to any point in the spectrum of differentiation, while in other situations product development and marketing may only gradually adjust the way consumers perceive their product. Our results suggest that Nash equilibrium may have predictive power in both situations.

Second, abstracting from the spatial aspects of our study, we provide an example where the dynamic implementation alters the predictive power of static Nash equilibrium. The predictive power is poor in a conventional synchronous Discrete time implementation, but much better in two different asynchronous continuous time implementations.

Several new avenues of research now come into focus. First, we note that unexplained discrepancies remain between our Discrete time results and those of Huck et al. (2002). They obtain sharper modes than we do, including the contra-NE mode at the center location. A broader avenue is to apply the continuous time treatments to more general Hotelling models, including the no-edge case (the circle), joint decisions of price and location, and different numbers of players. Another broad avenue is to separate the impact of continuous time per se from the impact of asynchronous choice. Our agent-based simulations suggest that taking turns in discrete time suffices to achieve Nash (or near-Nash) equilibrium behavior, but we do not yet know whether humans will agree.

Bibliography

Steffen Brenner. Location (hotelling) games and applications. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.

Jamie Brown Kruse and David J Schenk. Location, cooperation and communication:

- An experimental examination. *International Journal of Industrial Organization*, 18(1):59–80, 2000.
- Jamie Brown-Kruse, Mark B Cronshaw, and David J Schenk. Theory and experiments on spatial competition. *Economic Inquiry*, 31(1):139–165, 1993.
- Richard Collins and Katerina Sherstyuk. Spatial competition with three firms: an experimental study. *Economic Inquiry*, 38(1):73–94, 2000.
- Anthony Downs. An economic theory of political action in a democracy. *The Journal of Political Economy*, pages 135–150, 1957.
- B Curtis Eaton and Richard G Lipsey. The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *The Review of Economic Studies*, pages 27–49, 1975.
- Ben Greiner. An online recruitment system for economic experiments. 2004.
- Harold Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929.
- S Huck, W Müller, and NJ Vriend. The east end, the west end and king’s cross: On clustering in the four-player hotelling game. *Economic Inquiry*, 40(2):231–240, 2002.
- Janet S Netz and Beck A Taylor. Maximum or minimum differentiation? location patterns of retail outlets. *Review of Economics and Statistics*, 84(1):162–175, 2002.
- Charles N Noussair, Charles R Plott, and Raymond G Riezman. An experimental investigation of the patterns of international trade. *The American Economic Review*, pages 462–491, 1995.
- James Pettit, Daniel Friedman, Curtis Kephart, and Ryan Oprea. Software for continuous game experiments. *Experimental Economics*, pages 1–18, 2013.
- Avner Shaked. Existence and computation of mixed strategy nash equilibrium for 3-firms location problem. *The Journal of Industrial Economics*, pages 93–96, 1982.