Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

MYOPIA IN DYNAMIC SPATIAL GAMES

Shane Auerbach

Marketplace Labs
Lyft Inc.
185 Berry St., #5000
San Francisco, CA 94107, USA

Rebekah Dix

Department of Economics
University of Wisconsin-Madison
1180 Observatory Drive
Madison, WI 53706, USA

ABSTRACT

We design an experiment to evaluate behavior in a dynamic spatial game representing the incentives faced by drivers on a ridesharing platform while waiting to be matched with a rider. The design is unique in that it allows us to observe not only participants' choices, but also the considerations that went into those choices. The results of the experiment show that a large majority of player choices are consistent with myopic best responding—a myopic best response maximizes a player's flow payoff at the time of the decision but ignores strategic considerations regarding the future choices of opponents. Given this finding, we develop agent-based models of spatial competition built upon myopic agents. Myopic behavior in our model results in quite efficient outcomes, suggesting that ridesharing platforms may benefit from sharing with drivers the locations of other nearby drivers to allow them to compete spatially.

1 INTRODUCTION

When Lyft (or Uber) drivers are waiting for ride requests, they often sign out of the driver app and into the passenger app. They do this because they want to see where nearby drivers are, and Lyft shows nearby drivers only to passengers. If in the passenger app a driver finds that she is surrounded by other drivers in near proximity, she knows she may have to wait a while for a request or move to a different location—when passengers request rides, they are matched, roughly speaking, with the nearest available driver, and therefore a driver's catchment region is small when she is surrounded by others.

An efficient spatial allocation of idle drivers would minimize the expected wait-time for passengers, which we assume proportional to the expected distance to nearest driver. Supposing passengers are uniformly distributed over a unit disk and drivers travel as the crow flies, Figure 1 shows three possible spatial allocations, each with six drivers. The points represent drivers, the black borders represent the catchment regions, and the shading represents the distance from a point to its nearest driver. The allocation on the left yields an expected distance of about 0.364, while the center and right allocations yield expected distances of about 0.285 and 0.282, respectively. In this sense, the allocation on the right is the most efficient—in fact, it is the optimal allocation of six drivers in this space.

Why do ridesharing platforms make it difficult for drivers to see the locations of their competitors? Intuition might suggest that competitive markets would yield reasonably efficient spatial allocations—if there were high expected demand in an area with no driver in close proximity, an idle driver would reduce their expected idle time by moving to this area. This intuition has been both supported and questioned in theoretical literature dating back to Hotelling (1929).

We develop an experiment to examine behavior in a dynamic spatial game. In the experiment, we provide players with calculator software to compute flow payoffs for any possible spatial allocation. As players use the calculator to consider choices before making them, we observe not only their choices in the actual game but also, in the calculator data, indications of the allocations that they consider in making those choices. We find that a large majority of choices are consistent with myopic best responding. Of

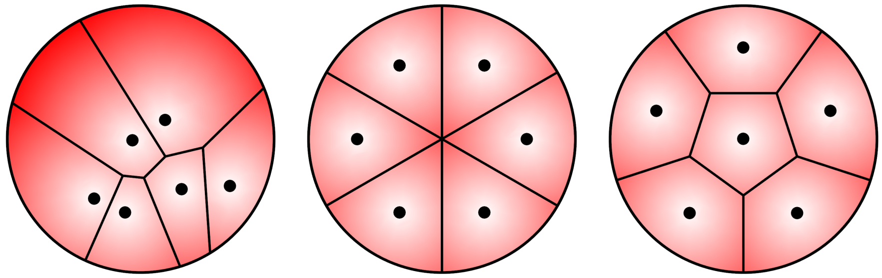

Figure 1: Three different 6-driver spatial allocations of varying efficiency.

the choices that are not, the calculator data suggests that most are attributable to error. Furthermore, the number of times a player chooses a myopically optimal move has a statistically significant positive effect on a player's payment, suggesting that myopic optimization is a good rule of thumb and that learning may reinforce it.

Agent myopia is a behavioral assumption in dynamic environments under which an agent views the positions of her opponents as fixed when deciding whether or not she would profit from changing her own position—that is, she simply maximizes her instantaneous payoff flow. To define agent myopia, consider a dynamic spatial game where each agent periodically makes decisions. Assume that at any given moment at most one agent makes a decision and that each agent can revise her decision in a future time period. In this setting, if an agent takes the locations of opponents as given, she is effectively not strategic—she simply maximizes her instantaneous payoff flow given the current state at the moment of her decision. We say that this agent follows a *myopic best response* (MBR) dynamic.

Our experiment results support the modeling of spatial competition as a dynamic game with agents that optimize myopically. With this assumption, a dynamic game can be modeled as a sequence of static, individual optimizations, allowing for simulation and agent-based models even in complex environments. In cases where agents following a myopic best response dynamic converge to a fixed point, that fixed point is a Nash equilibrium of the corresponding static game. Where there is no convergence, the dynamic path itself serves as a prediction in that we can evaluate measures of inefficiency at different moments in time and compare averages or dynamics. In the presence of complexity, agents make choices using heuristics and rules of thumb. Agent-based models involve identifying these and building them into a model to generate predictions through simulation and computation. Because of the assumed relative simplicity of MBR, we can work in much richer environments than those used in game-theoretic analyses without losing tractability.

The rest of this paper is organized as follows: Section 2 reviews related literature on spatial competition, depth of reasoning, and agent-based models. In Section 3, we present our model of a dynamic spatial game. In Section 4, we present our experiment results on the prevalence of myopic best-responding. In Section 5, we develop agent-based simulations both in our experimental environment and on an actual transportation network to show that ridesharing platforms may benefit from allowing each driver to see the locations of other drivers. We conclude in Section 6.

2 RELATED LITERATURE

This paper relates to three distinct literatures: experimental analyses of spatial competition, behavioral analyses on depth of reasoning, and simulation with agent-based models.

Brown-Kruse et al. (1993) tests Hotelling's linear-city model in a repeated game with two firms and examines the role of communication between players. Without communication, the two firms locate near the center of the market. With communication, they locate one-fourth and three-fourths of the way along the linear market and maximize joint profits. Kruse and Schenk (2000) extends Brown-Kruse et al. (1993)

to consider non-uniform customer distributions. Players generally chose symmetric strategies with uniform, unimodal, and bimodal distributions, but, even with communication, they struggled to reach the profit-maximizing allocation with non-uniform customer distributions. In order to study the role of complexity in location games, Kruse and Schenk (2000) also considers Hotelling's linear city model with a simplified decision environment where they only allowed players to choose between two locations: the center or one edge of the market. Players in the simplified decision environment reached the profit-maximizing allocation more often than those with the full range of potential locations, but allowing communication between players, even in the more complex environment, led more often to the profit-maximizing allocation. Collins and Sherstyuk (2000) considers a similar model to that in Brown-Kruse et al. (1993) but with three firms and finds that players randomize locations and avoid both the center and edges of the market, highlighting the role of risk aversion as agents choose low-risk locations instead of the risk-neutral equilibrium predictions.

In relation to depth of reasoning, the MBR assumption is similar to assuming that all agents are *level-1* in that they fail to reason to any extent about future opponent play. Depth of reasoning is typically studied within a Keynesian beauty contest, first described in Keynes (1936). Nagel (1995) proposed the *level-k* model of depth of reasoning and experimentally identified heterogeneity in this depth among the experiment's participants. Halpern and Pass (2015) develop a framework for reasoning about strategic agents performing possibly costly computation. Alaoui and Penta (2016) offers a model of endogenous agent depth of reasoning motivated by an axiomatized cost-benefit analysis. *Level-k* models are typically applied to static games, though Rampal (2017) is a recent extension to dynamic games. While assuming *level-1* behavior would be a very strong assumption in a Keynesian beauty contest, our spatial game on the Euclidean plane is far more complex due to the underlying geometry.

We build agent-based models upon the myopic best responding that we find in the behavioral experiment. An agent-based model (ABM) is a computational model for simulating the interactions of autonomous agents to assess their effects on the system. Crooks et al. (2008) and Crooks and Heppenstall (2012) look at the particular challenges of spatial ABM, and economists are starting to apply these models to Hotelling-like environments (Van Leeuwen and Lijesen 2016). The consequences of MBR dynamics on a lattice with local interactions are studied in Blume (1993) and Blume (1995). The MBR dynamic is also common in evolutionary game theory, although that literature focuses on population games with many players.

3 MODEL

Let $\mathscr{I} = \{1,\dots,I\}$ denote the set of players. A graph G is given by the pair $G = (\mathscr{N}, A)$, where $\mathscr{N} = \{1,\dots,N\}$ is a set of nodes and $A \in [0,\infty]^{N\times N}$ is a weighted adjacency matrix. If there is an edge between nodes m and n, $a_{mn} \in \mathbb{R}_{\geq 0}$ denotes the distance or weight. Otherwise, $a_{mn} = \infty$. We embed G in $\mathbb{R}^2$ and assign each $n \in \mathscr{N}$ coordinates $(x_n, y_n) \in \mathbb{R}^2$. G is an undirected graph so that $a_{nm} = a_{mn}$ for all nodes $m,n \in \mathscr{N}$. We also assume that G is connected so that every two nodes in G have a path between them and that each node has a self-loop, $a_{mm} = 0$, so that players may remain in their current positions across periods. Node n's set of neighbors $\mathscr{B}_n = \{m \in \mathscr{N} : a_{nm} < \infty\}$ includes all nodes with which n shares an edge.

In a dynamic game of spatial competition on a graph G, players sequentially choose their locations on the graph. We consider games with T periods. A *spatial allocation* s_t is an I-tuple $(s_{t,1}, s_{t,2}, \dots, s_{t,I}) \in \prod_{i\in\mathscr{I}} \mathscr{N}$ that records the locations of players on the nodes of G in period t. Thus, s_t is a list of I nodes, where the i^{th} element, $s_{t,i}$, is the location of player i in period t. Let $\mathscr{S} = \prod_{i\in\mathscr{I}} \mathscr{N}$ be the set of all possible spatial allocations and define $s_{t,-i} \in \prod_{j\in\mathscr{I}\setminus\{i\}} \mathscr{N}$ as the spatial allocation of i's opponents in period t. This notation allows us to consider the movement of player i holding the positions of her opponents fixed. We will refer to the 2-tuple $(n, s_{t,-i})$ as the spatial allocation in which player i is positioned at node n and her opponents are positioned according to $s_{t,-i}$.

We define a multiplicity function, $\psi : \mathscr{N} \times \mathscr{S} \to \mathbb{N}_{\geq 1}$, that returns the cardinality of the set of players who are closest to a particular node: $\psi(n, s_t) = |\arg\min_{i\in\mathscr{I}} \{d(n, s_{t,i})\}|$, where the distance function $d(\cdot)$

finds the length of the shortest path between two given nodes. For example, $\psi(n, s_t) = 2$ implies that the two closest players are equidistant from node n in the spatial allocation s_t. We also allow colocation—where multiple players are located at n, each is a distance of zero from n, and the multiplicity function therefore gives the number of players at n.

In each period t, we use a Voronoi diagram, $\mathrm{Vor}(s_t)$, to calculate each player's market share. A Voronoi diagram on a space and a set of points divides a space into cells, with each cell representing the region that is closer to a particular point than to any other point. $\mathrm{Vor}(s_t)$ partitions $\mathscr{N}$ into Voronoi cells $V_i(s_t) = \{n \in \mathscr{N} \mid d(n, s_{t,i}) \leq d(n, s_{t,j}), \forall j \neq i\}$ for each player $i \in \mathscr{I}$. At period t in the game, the market share of player i is the number of cells to which she is the closest player, including evenly divided shares of cells to which she and other players are equidistant:

$$\pi_i(s_t) = \sum_{n \in V_i(s_t)} \frac{1}{\psi(n, s_t)} .$$

A spatial allocation s_t is a *global* Nash equilibrium of a static location game if for all $i \in \mathscr{I}$ and for all $n \in \mathscr{N}$, $\pi_i(s_t) \geq \pi_i(n, s_{t,-i})$. It is a *local* Nash equilibrium of a static location game if for all $i \in \mathscr{I}$ and for all $b \in \mathscr{B}_{s_{t,i}}$, $\pi_i(s_t) \geq \pi_i(b, s_{t,-i})$. A global Nash equilibrium requires that there is no profitable deviation for any player to any other node, whereas a local equilibrium precludes only profitable deviations for any player to any adjacent node. Obviously the set of local equilibria contains the set of global equilibria.

To measure the *spatial inefficiency* of an allocation for the purposes of the ridesharing motivation, we first calculate the average distance from each node to its nearest driver in the spatial allocation, s_t:

$$\bar{d}(s_t) = \frac{1}{N} \left(\sum_{n \in \mathscr{N}} \min_{i \in \mathscr{I}} \{ d(n, s_{t,i}) \} \right) . \tag{1}$$

Define an optimal spatial allocation s_t^* as one which minimizes (1):

$$s^* \in \arg\min_{s_t \in \mathscr{S}} \bar{d}(s_t) .$$

Then, define the spatial inefficiency of a spatial allocation s_t as the percentage difference between the average distance of an allocation and that of an optimized allocation:

$$\xi(s_t) = \frac{\bar{d}(s_t) - \bar{d}(s^*)}{\bar{d}(s^*)} .$$

4 BEHAVIORAL EXPERIMENT

In this section, we present results from an experiment designed to test the validity of the behavioral assumption that agents myopically best respond (MBR) in the context of a dynamic spatial game. If an agent maximizes her instantaneous flow playoffs at the moment of her decision, ignoring potential future opponent movements, we say that her choices satisfy the MBR assumption.

If there are costs associated with complex strategic reasoning, myopia could be rational. Abstracting from these costs, we know that behaving in accordance with the MBR assumption is likely to be suboptimal. Our initial hypothesis was that agents would choose to myopically best respond—it is a reasonably good strategy and does not require complex strategic reasoning. Indeed, we find evidence of this.

Our experiment involves participants playing a two-dimensional, discrete-time, dynamic spatial game. Previewing our results, we observe 2178 decisions and find that only 307 of them are suggestive of higher-order reasoning (SHO) that would violate the MBR behavioral assumption. Further, regression analysis shows that players that make more SHO choices earn no more money than those who make fewer, while players who behave most in accordance with the behavioral assumption do earn more money, suggesting that repetition and learning might work in favor of the assumption's validity.

4.1 Experiment Design

We use a grid in our experimental environment. An $N \times N$ grid is a graph $G = (\mathscr{N}, A)$ whose set of nodes are numbered 1 through N^2. We embed G in $\mathbb{R}^2$ and assign each node coordinates (i, j) with $1 \leq i \leq N$ and $1 \leq j \leq N$. Then, there is an edge between two nodes m and n with coordinates (i, j) and (k,l), respectively, if and only if $|i-k|+|j-l| \leq 1$. For A, this means that $a_{m,n} = 1$ if and only if the previous condition holds, otherwise $a_{m,n} = \infty$. For all $n \in \mathscr{N}$, $a_{n,n} = 0$.

In the experiment, five participants (players, henceforth) played a location game on a 21×21 grid. Players were given the opportunity, one by one, to move one square in any cardinal direction. Turn order was random. Colocation was not allowed. We used the same initial allocation of players, shown in Figure 2, for each session. Players are labeled by number, 1 through 5. There are also eight computer players, each labeled with C, who do not move and are positioned along the perimeter of the grid. We include these static computer players because we want to abstract from issues resulting from the presence of boundaries. Having the static computer players makes the game played by the actual players theoretically similar to that played in a particular region of the unbounded Euclidean plane.

Each player has a color assigned to them. The player's current location is represented by a single cell in the grid with a dark shade of that color and the player's number. The player's Voronoi region, calculated with the ℓ_1 norm (Manhattan distance), is represented by an area of the grid in a lighter shade of the player's color. Black cells are equidistant from two or more players, at least one of which is not a computer. Grey cells are closer to a computer player than any non-computer player.

We conducted the experiment with two pieces of software that we developed. Our main console software shows the players' locations in the current iteration, the player whose turn it is to move, and a grid with the players' Voronoi regions. The grid also highlights up to five move options, including the four cardinal directions and an option to remain at the current location. In the experiment, the main console was projected for all players to view throughout the experimental session.

Our second piece of software, shown in Figure 2, is calculator software that each player used on her own lab computer throughout the experiment. The calculator allows players to enter in an allocation of players, calculate the area of the players' Voronoi regions for that allocation, and see the grid of the players' Voronoi regions. We provided the calculator for two reasons. First, while it is simple arithmetic to work out which Voronoi region a given square belongs to for a given allocation, it is very time-consuming to calculate the area of a Voronoi region by hand. We wanted to alleviate that burden. Second, because players were using the calculator to consider their choices during each turn, and between their turns in many cases, we have data on not only the choices they make in the actual game but also all of the allocations they considered in making their choices.

The game is played as follows. In each iteration, the player number whose turn it is in that iteration is announced. This player has up to two minutes to decide where to move. The player then communicates her decision to the experiment leader and the experiment leader updates the main console accordingly. This process repeats for the duration of the experiment session.

In each turn of the game, a player chooses to move to a square within one unit of her current location. If there are no opponents within one unit of the current player's position and she is not on a boundary, she has five squares to choose from. Positioning the current player in one of these squares within one unit of her current location and keeping the opponents in place creates up to five potential allocations of players for the next iteration. We define each of these potential allocations as a *move option*. We define a move option's flow payment as the area of the current player's Voronoi region after the move is made. In order to classify all move options, we rank the move options by their flow payments: the FP1 move is the move

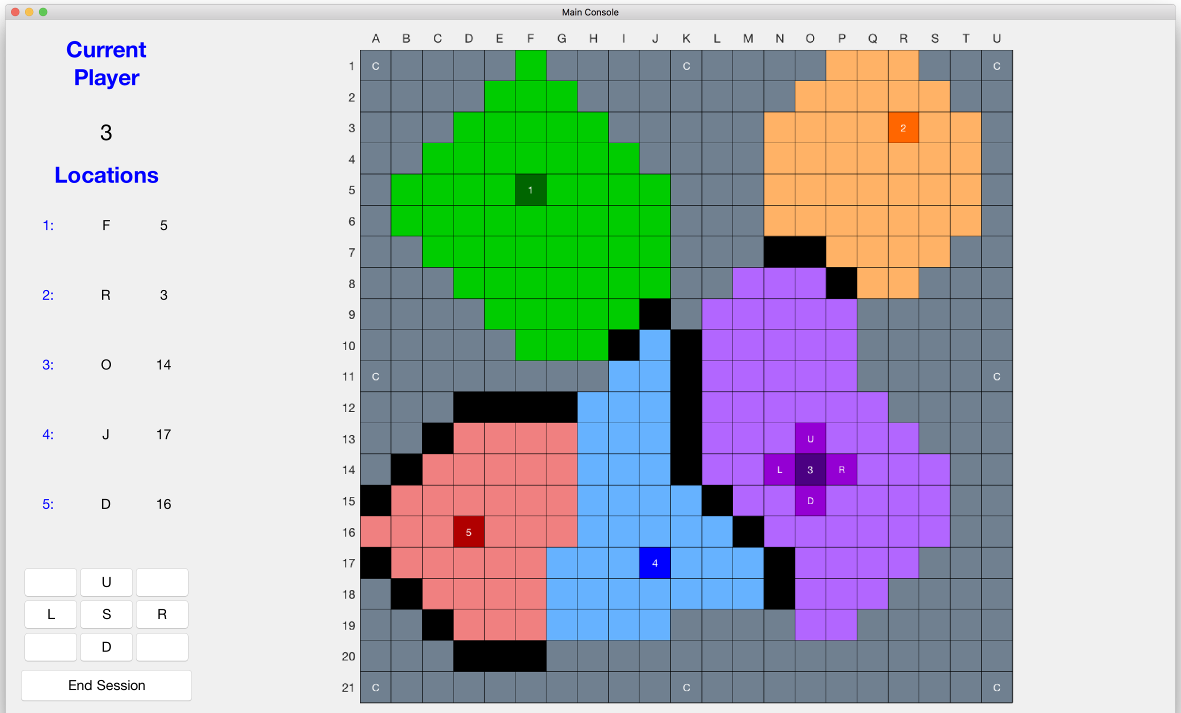

Figure 2: Main console software.

option with the highest flow payment, the FP2 move is that with the second-highest flow payment, etc. An MBR agent would always choose the FP1 move.

We ran the experiment at the Behavioral Research Insights Through Experiments (BRITE) Lab at the University of Wisconsin-Madison in June, 2017. Players were recruited from a pool of students maintained by the BRITE Lab. We conducted 18 experimental sessions, each with 5 players, for a total of 90 players. Players were shown an instructional video at the beginning of each experimental session to explain how the game is played and how their payments would be calculated.

Players' payments were proportional to the average area of their Voronoi regions over all iterations in the experimental session, where an iteration is defined by the allocation at a given time and a new iteration is entered upon every selected move. The number of iterations per session varied based on speed of play. Turn order was determined randomly. Experimental sessions were scheduled to last 90 minutes, including time for the instructional video. The average time spent playing the game was 68 minutes, and the average number of iterations per experimental session was 121. The mean player payment was $20.

4.2 Results

There were a total of 2178 main console moves across the 18 sessions. Table 1 shows the distribution of flow payment move rankings for all moves as well as for those made after the first ten minutes of each session. The percentage of FP1 moves increases with the exclusion of the first ten minutes, suggesting that there was some noise in the beginning of each session as players learned how to play the game and how to use the calculator.

Result 1 Distribution of players chose the FP1 move 60% of the time.

This does not imply that a majority of players behaved in accordance with the MBR behavioral assumption. Nor does it suggest that a majority of moves were made by players behaving in accordance

Table 1: Distribution of move FP rankings.

	All moves		**After 10**	
FP ranking	**Freq.**	**%**	**Freq.**	**%**
1	1315	60	1103	63
2	439	20	339	20
3	229	11	174	10
4	126	6	74	4
5	69	3	50	3

with the MBR assumption. Players selecting a FP1 move may be engaging in sophisticated consideration of anticipated future movements that happens to motivate them to pick the same choice as an MBR player. Alternatively, an agent might just select a move randomly and happen to select the FP1 move. Similarly for the moves selected that were not FP1, the analysis above does not suggest the reasoning that motivated these choices. However, we use data from players' calculators to learn about reasoning processes.

If players were engaging in higher-order reasoning on their opponents' subsequent behavior, we would expect them to run calculations that considered potential subsequent movements from opponents. We saw this only rarely. Players tended to keep their opponents in place relative to the current allocation of players in the main console and tested allocations with only their own locations adjusted. In fact, our second result of note is that a majority of calculations made were on move options. Table 2 summarizes the positions of opponents in calculations.

Result 2 In 82% of calculations, all opponents were positioned as they were in the current iteration. 76% of calculations were of move options.

Table 2: Positioning of opponents in calculations.

	Freq.	**%**
Total	24 765	
Opponents in place	20 341	82
Move option	18 869	76

We can also consider whether players made forward-looking calculations with respect to their own locations. As players could only move one square per turn, a calculation with the player moved more than one square tested a non-feasible move. A player positioning herself more than one unit away from her position in the current iteration could suggest some level of higher-order reasoning. Table 3 shows the distances that the player who made the calculation was from her current position in the iteration. Very few forward-looking calculations were made. Instead, a majority of the calculations were either updating to the current iteration or testing a move option of the current iteration. Again, we see that opponents were often in place relative to the current iteration.

Table 3: Distribution of player distances from their iteration position in calculations.

Distance	**Freq.**	**%**	**Of which opponents in place**	**%**
0	12 046	49	9546	79
1	10 733	43	9323	87
2	1205	5	957	79
3	360	2	253	70
$\geq$4	421	2	262	62

Result 3 In 92% of calculations, the calculating player was within one square of her position at the the time of calculation.

We now turn to an analysis of non-FP1 moves to address the question of how many of them are suggestive of higher-order reasoning on opponent responses. Here, we define the relevant calculation interval for a move to be the timespan from the last change in the allocation up until the move. The set of relevant calculations, then, is the set of all calculations made during a move's relevant interval. To partition the set of non-FP1 moves, we first determine whether the FP1 move was calculated during the relevant interval. In a majority of non-FP1 moves, the FP1 move was not calculated, suggesting to us that these choices are likely attributable to a failure to consider or calculate the FP1 move rather than higher-order reasoning.

To further refine this partitioning, we then determine whether the move that was chosen was actually calculated. For moves where the FP1 move was not calculated but the move chosen was calculated, we can ask whether the player chose the highest-scoring calculated move (HSCM) of all moves calculated in the relevant interval. Although the FP1 move was not chosen, the moves in this subset suggest MBR-like behavior. See Table 4 for the full partitioning of non-FP1 moves.

Table 4: Partitioning the 863 non-FP1 moves.

	FP1 not calculated	**FP1 calculated**
	535	328
Choice not calculated	325	21
No calculations	(205)	
Some calculations	(120)	
Choice calculated	210	307
Chose HSCM	(126)	
Did not choose HSCM	(84)	

Within the set of non-FP1 moves, we define a move to be suggestive of higher-order reasoning (SHO) if the player calculated both the FP1 move and the move she ultimately chose, but did not choose the FP1 move. For these 307 moves (out of 2178 total moves), the player was accurately using the calculator and deviating from the FP1 move. Approximately 24% of SHO moves were made by less than 7% percent of players.

Another thought is that players may be following the MBR assumption roughly but then deciding by intuition when flow payments of two or more moves are very close. To get at this, we look at the flow payment differences between the FP1 move and that of the chosen moves. The average difference for non-FP1, non-SHO moves, between the move selected and the FP1 move, was 1.74 grid squares. The average score difference for non-FP1, SHO moves was only .82 grid squares, suggesting that players are relying on intuition and higher-order reasoning mostly when the flow payments are close.

Result 4 Differences in flow payments between SHO moves and their FP1 alternatives were significantly smaller than those between non-SHO, non-FP1 moves and their FP1 alternatives.

Finally, we can examine the determinants of success in the experiment through regression, which gives us our final key result.

Result 5 The number of SHO moves a player chose had no statistically significant impact on her performance. The number of FP1 moves had a significant positive impact.

5 SIMULATIONS WITH MYOPIC AGENTS

5.1 In the Environment of the Behavioral Experiment

Recall that we designed the experiment to represent the incentives of idle drivers in a ridesharing setting. Because we use a grid and the ℓ_1 norm for distance, the experimental environment represents competition on an urban road grid with uniformly distributed expected demand. As we described in Section 1, a ridesharing platform may wish for idle drivers to position themselves so as to minimize expected wait-times for passengers. An obvious question then is whether the myopic best responding that we find in the experiment yields spatial allocations that do well by this metric, measured by $\xi(s_t)$ as defined in Section 3.

We evaluate the efficiency of the dynamic paths of our experimental sessions and compare those with the paths of MBR simulations starting from various initializations. In our MBR simulations, each player chooses the move that maximizes her flow playoff (FP1) at the time of her decision.

Figure 3 shows the results. Avg. Exp. takes the spatial inefficiency from each experimental session at iteration t and averages them. The initialized allocation in the experiment was quite efficient, but players' early choices increased efficiency further. Avg. Sim1 mimics the experiment with 18 sessions but replaces the observed behavior with that of simulated MBR agents. In this setting, MBR usually led players to optimal or near-optimal spatial allocations. Avg. Sim2 and Avg. Sim3 are equivalent but were run from different, less efficient, initial allocations.

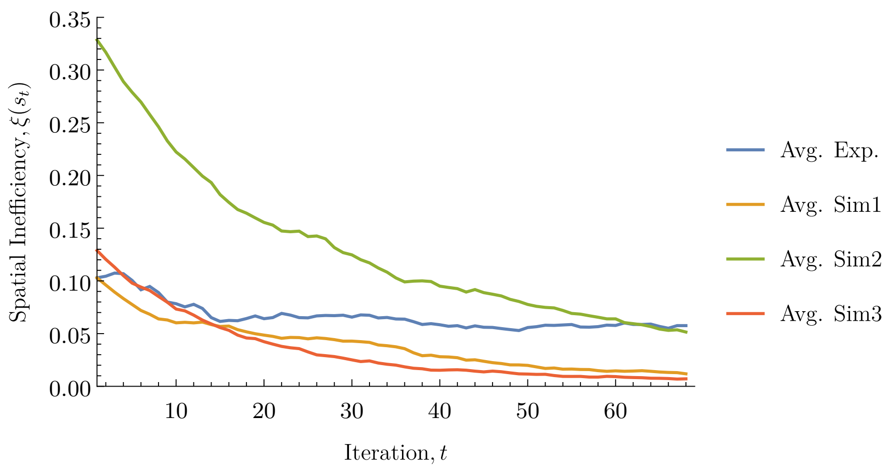

Figure 3: Spatial inefficiency in experiment and simulations.

The players in our experiment did choose the FP1 move a majority of time. However, the average spatial inefficiency in the experiments remained substantially higher than that of allocations achieved by agents who perfectly implement MBR in simulations. Given that the interests of myopic best responding drivers coincide well with the interests of the ridesharing platform, we believe that ridesharing platforms may wish to allow drivers to see other nearby drivers and assist them in best responding to their neighbors.

5.2 In a Ridesharing Environment with a Realistic Transportation Network

We now explore whether the myopic best responding that we find in the experiment yields efficient spatial allocations on a more realistic transportation network. Suppose $\mathscr{I}$ is the set of drivers for a ridesharing platform. A dynamic game of spatial competition represents drivers competing for passengers on a transportation network $G = (\mathscr{N}, A)$. The nodes of this network represent intersections, and the edges represent roads. The length of the edges are proportional to the length of the roads in the city.

We recreate the road network of the City of Oldenburg, Germany using data from Brinkhoff (2002). There are 6,105 nodes and 7,029 edges, and the average degree of the nodes is 2.3. We compute an approximately optimal allocation of 60 drivers in Oldenburg using a myopic (greedy) heuristic, as in Kuehn and Hamburger (1963). Call this optimal allocation s^*—we find that $\bar{d}(s^*) = 480.18$ meters.

To simulate the game played between MBR drivers, we generate an allocation of the set $\mathscr{I}$ of drivers on the nodes of the transportation network and run the MBR algorithm for $T = 5000$ iterations. Figures 4 and 5 show the initial and final allocations, respectively. The spatial inefficiency of the initial allocation is $\xi(s_1) = 2.02$ and the spatial inefficiency of the final allocation is $\xi(s_{5000}) = 0.55$. Thus, the decisions made by MBR drivers resulted in a large decrease in spatial inefficiency, which implies a large decrease in expected consumer wait-times. Figure 6 shows spatial inefficiency along the dynamic path of one simulation. We observe this same tendency in all simulation runs. We note that spatial inefficiency did not decrease monotonically along the dynamic path.

However, in the simulation discussed above, MBR agents did not converge to an optimal spatial allocation. Within 5000 iterations, drivers converged to a cycle of spatial allocations in which spatial inefficiency was near .55 for all allocations in the cycle. This result suggests that allowing agents to myopically best respond may reduce consumer wait-times but it need not converge to an optimal outcome.

We believe one contributing factor to inefficiency in fixed points (or cycles) under the MBR dynamic is boundary behavior, where individuals near the periphery have incentives to move inwards as they can do so without sacrificing market share on the periphery. To isolate this, consider the approximately optimal allocation of drivers s^* on Oldenburg's transportation network, and suppose we fix the positions of the drivers that are located on the outer periphery of the network. Let $\mathscr{I}_{outer}$ be the set of 14 drivers on the outer periphery. We create a new spatial allocation in which the drivers in $\mathscr{I}_{outer}$ are at their locations in s^*, then we generate a random allocation of the remaining drivers. Finally, we simulate another game played between MBR drivers on Oldenburg's transportation network, but we only allow drivers in the set $\mathscr{I} \setminus \mathscr{I}_{outer}$ to move in the simulation. In this setting, the spatial inefficiency of the initial allocation is $\xi(s_1) = .34$ and that of the final allocation is $\xi(s_{1500}) = 0.2$. Fixing drivers on the periphery and allowing for MBR movement in the center of the transportation network resulted in significantly lower final spatial inefficiency, suggesting that boundary behavior is indeed a contributing factor to inefficiency in MBR simulations.

6 DISCUSSION

We cannot argue affirmatively that players myopically best respond, but we do fail, in our experiment, to find significant evidence that they violate the assumption. Our regression results suggest further that learning may push players toward, not away from, MBR behavior. Of course, this negative result is good news for those seeking to do agent-based modeling with an MBR assumption. The irony is that the underlying complexity of spatial competition that makes equilibrium analysis difficult may also make players behave quite predictably, thereby facilitating agent-based modeling. Given the results of our experiment, we think it best to model agents in a spatial agent-based model as noisy MBR agents. These agents would usually choose FP1 moves, but would randomly choose non-FP1 moves, doing so more often when non-FP1 moves are closer in flow payments and in these cases usually selecting moves with relatively high flow payments.

Given that the interests of myopic best responding drivers appear to coincide well with the interests of the ridesharing platform even in a setting with a more realistic transportation network, as we show through simulation, we believe that ridesharing platforms may wish to allow drivers to see the locations of nearby drivers and assist them in best responding to their neighbors. Furthermore, because of inefficiency due to boundary behavior, as ridesharing platforms start to incorporate automated vehicles alongside regular vehicles, they may want to have these automated vehicles target the periphery, allowing spatial competition between drivers in the interior.

Deriving behavioral insight in a laboratory environment and then investigating the implications of that behavior through simulation is a workflow that could be applied to other environments in which complexity impedes traditional theoretical analysis.

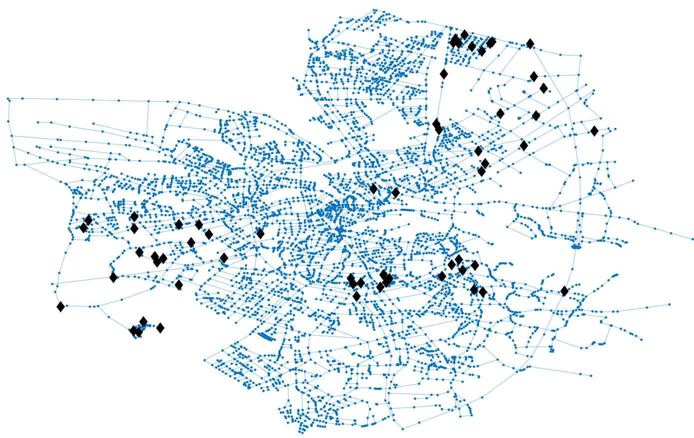

Figure 4: Initial allocation of 60 drivers in Oldenburg; $\xi(s_1) = 2.02$.

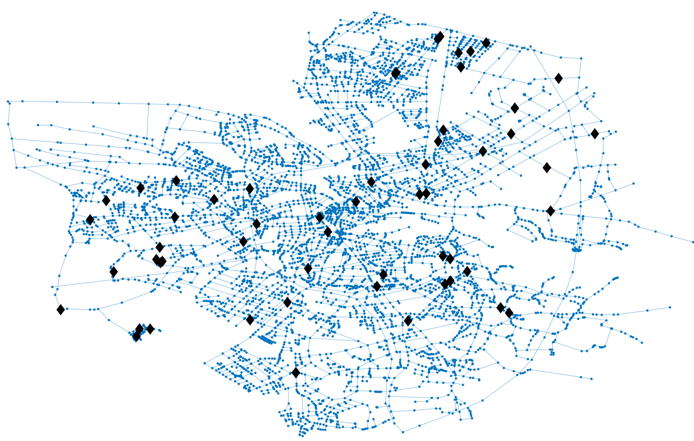

Figure 5: Final allocation of 60 drivers in Oldenburg; $\xi(s_{5000}) = 0.55$.

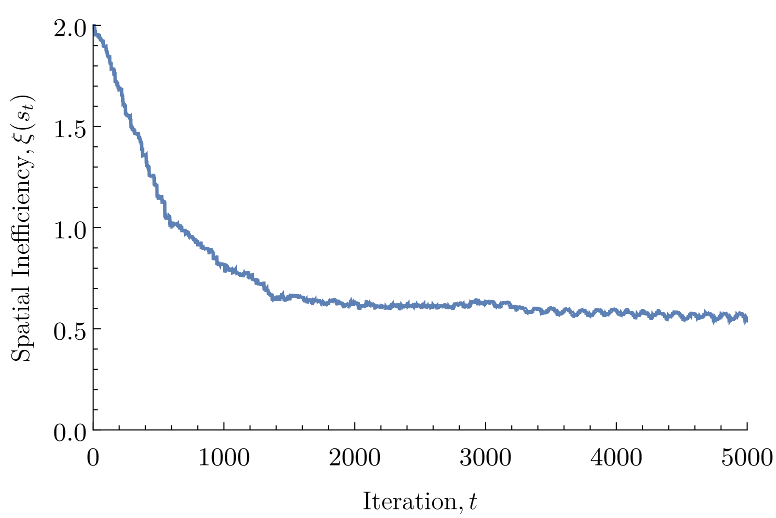

Figure 6: Spatial inefficiency along dynamic path of simulation with $T = 5000$.

REFERENCES

Alaoui, L., and A. Penta. 2016. "Endogenous Depth of Reasoning". *The Review of Economic Studies* 83 (4): 1297–1333.

Blume, L. E. 1993. "The Statistical Mechanics of Strategic Interaction". *Games and economic behavior* 5 (3): 387–424.

Blume, L. E. 1995. "The Statistical Mechanics of Best-Response Strategy Revision". *Games and Economic Behavior* 11 (2): 111–145.

Brinkhoff, T. 2002. "A Framework for Generating Network-Based Moving objects". *GeoInformatica* 6 (2): 153–180.

Brown-Kruse, J., M. B. Cronshaw, and D. J. Schenk. 1993. "Theory and Experiments on Spatial Competition". *Economic Inquiry* 31 (1): 139–165.

Collins, R., and K. Sherstyuk. 2000. "Spatial Competition with Three Firms: an Experimental Study". *Economic Inquiry* 38 (1): 73–94.

Crooks, A., C. Castle, and M. Batty. 2008. "Key Challenges in Agent-Based Modelling for Geo-Spatial Simulation". *Computers, Environment and Urban Systems* 32 (6): 417–430.

Crooks, A. T., and A. J. Heppenstall. 2012. "Introduction to Agent-Based Modelling". In *Agent-based models of geographical systems*, 85–105. Springer.

Halpern, J. Y., and R. Pass. 2015. "Algorithmic Rationality: Game Theory with Costly Computation". *Journal of Economic Theory* 156:246–268.

Hotelling, H. 1929. "Stability in Competition". *The Economic Journal* 39 (153): 41–57.

Keynes, J. 1936. *The General Theory of Employment Interest and Money*. London: Macmillan.

Kruse, J. B., and D. J. Schenk. 2000. "Location, Cooperation and Communication: An Experimental Examination". *International Journal of Industrial Organization* 18 (1): 59–80.

Kuehn, A. A., and M. J. Hamburger. 1963. "A Heuristic Program for Locating Warehouses". *Management science* 9 (4): 643–666.

Nagel, R. 1995. "Unraveling in Guessing Games: An Experimental Study". *The American Economic Review* 85 (5): 1313–1326.

Rampal, J. 2017. "Limited Foresight Equilibrium". Technical report, Indian Institute of Management, Ahmedabad, India.

Van Leeuwen, E., and M. Lijesen. 2016. "Agents playing Hotellings game: an Agent-Based Approach to a Game Theoretic Model". *The Annals of Regional Science* 57 (2-3): 393–411.

AUTHOR BIOGRAPHIES

SHANE AUERBACH is a Research Scientist in Marketplace Labs at Lyft Inc. Before joining Lyft, he taught at Carleton College. He completed his Ph.D. in Economics at the University of Wisconsin-Madison, where his research focused on microeconomic theory, industrial organization, and behavioral economics. His website is www.shaneauerbach.com and his email address is shaneauerbach@gmail.com.

REBEKAH DIX is an advanced undergraduate at the University of Wisconsin-Madison majoring in mathematics and economics with a minor in computer science. Her primary research interests are in applied microeconomic theory, computational economics, and randomized algorithms. She will be applying for Ph.D. programs in economics to start in Fall 2019. Her email address is rebekahanned@gmail.com.