

# Competition and spatial efficiency

Shane Auerbach and Rebekah Dix  
University of Wisconsin-Madison

November 4, 2017

## Abstract

Does competition yield efficient spatial allocations? In this paper, we show evidence to the contrary by proposing a measure of spatial inefficiency and applying it to several empirical case studies. We then turn to the question of how to build predictive models of spatial competition. Literature on location theory reveals the complexity in static models of spatial competition. In one-dimensional spaces, prediction is impeded both by equilibrium multiplicity and non-existence. In two-dimensional spaces, we lose tractability. We extend location theory to propose dynamic, agent-based models with myopic agents. The behavioral assumption of myopia allows for the modeling of a dynamic game as a series of static optimizations, yielding tractability even in rich, two-dimensional environments. We present our own experiment results that support the assumption of myopia in dynamic spatial games.<sup>1</sup>

## 1 Introduction

When 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 Uber drivers are, and Uber shows nearby drivers only to passengers, not drivers.<sup>2</sup> If in the passenger app an Uber 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 on Uber, they are matched 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 Uber 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 Uber 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

---

<sup>1</sup>We thank Antonio Penta, Daniel Quint, Kenneth Hendricks, Lones Smith, Justin Sydnor, Kathryn Carroll, UW-Madison’s BRITE Laboratory, and seminar participants at the University of Wisconsin and Carleton College for their feedback and support.

<sup>2</sup>Lyft does essentially the same, but with driver and passenger modes combined in the same app.

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.




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

Why do ride-sharing apps 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 a group of customers desirous of a product or service and no firm within the proximity, an entrant should profit from moving in. This intuition has been both supported and questioned in theoretical literature dating back to Hotelling (1929). In this paper, we show empirical evidence that contradicts this intuition and discuss how best to model spatial competition.

To test the intuition, we propose a measure of spatial inefficiency and apply it to data on allocations of gas stations and supermarkets. Our measure of spatial inefficiency requires us to compute optimal spatial allocations, and we develop algorithms for this purpose. Our algorithm for optimizing two-dimensional allocations works on the logic of Lloyd’s algorithm (or Voronoi relaxation), which is commonly used in computer science and electrical engineering (Lloyd, 1982). A modified version of our algorithm could allow a ride-sharing service to reduce wait-times for passengers by giving drivers guidance on where to locate.

In comparing actual allocations to optimized ones, we find significant inefficiencies in the allocations of gas stations and supermarkets. We also find that spatial inefficiency is significantly less acute for allocations of hospital and fire stations, which result from mechanisms that involve more central planning and less competition. We estimate, conservatively, that spatial inefficiency in the allocation of supermarkets in the Chicago Metropolitan Statistical Area costs consumers over 100 million dollars annually in additional transportation costs.

To evaluate mechanisms that yield spatial allocations and test related policies such as exclusive territories and zoning, we need a model of spatial competition that is both tractable and rich enough to be predictive in applied settings. Static analyses in location theory reveal many of the challenges in developing such a model: static spatial games may have multiple Nash equilibria, none at all, or be intractable even in relatively simple environments—we discuss this in Section 3.1. Therefore, we propose modeling 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. Microeconomic theory is usually deductive in that results are derived directly from assumptions. Agent-based modeling is inductive in that one makes assumptions on agents and then watches phenomena emerge through agent interaction. We discuss this further in Section 3.2.

To support our behavioral assumption that agents optimize myopically, we present results from an experiment that we designed to test for myopia in a dynamic spatial game. In the experiment, we provided players with calculator software to compute flow payoffs for any possible spatial allocation. As players used the calculator to consider choices before making them, we observe not their choices but also, in the calculator data, indications of the allocations that they considered in making those choices. Looking at both, we find very little evidence that contradicts our assumption of agent myopia. Most of the choices made were myopically optimal in that they maximized the payoff flow at the moment the choice was made. Of the choices that were not, the calculator data suggests that most are attributable to error. Furthermore, the number of times a player chose a myopically optimal move had 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.

The spatial game in our experiment was designed to roughly match the ride-sharing motivation. It is a dynamic game with reversible decisions, no price competition, and fairly low stakes. Assuming myopia in the location choices of gas stations and supermarkets in spatial games with irreversible decisions, price competition, and higher stakes is stronger—businesses do extensive market research in location choices, but this is consistent with myopic optimization unless reasoning on potential future entry, exit, or competitor relocation induces choices that do not maximize myopic profits. What the two environments share is complexity. That complexity makes traditional game-theoretic analysis difficult. But it also induces agents to make decisions through heuristics or rules of thumbs. Insofar as we can identify these, we can build distinct agent-based models for differing environments.

Given the complexity of the environment and uncertainty on competitors, we think it is reasonable to assume that firms generally optimize myopically.

The rest of this paper is organized as follows: Section 2 presents our empirical work. In Section 3, we turn our attention to modeling by reviewing insights from static spatial games and introducing agent-based models with myopic agents. In Section 4, we present experiment results that support the assumption of myopia in dynamic spatial environments. We conclude in Section 5.

## 2 Evaluating spatial allocations empirically

In Section 2.1, we define spatial inefficiency formally. In Section 2.2, we evaluate the spatial inefficiency of allocations of gas stations on segments of Interstate 90. In Section 2.3, we evaluate it for allocations of supermarkets, fire stations, and hospitals in three cities.

## 2.1 Defining spatial inefficiency

To measure spatial inefficiency, we must define it precisely. Let  $(X, d)$  denote a normed vector space with a density function  $f : X \rightarrow \mathbb{R}_+$  representing the distribution of customers across the space and  $d$  a distance metric on it. In our Uber example,  $X \subset \mathbb{R}^2$  represents the city in latitude-longitude space. A spatial allocation,  $s = (s_1, \dots, s_n)$ , is a vector of positions for the  $n$  drivers. For each  $i$ ,  $s_i \in X$  is a latitude-longitude position. Let  $S_n(X)$  denote the set of all possible spatial allocations of  $n$  drivers in  $X$ . Supposing Uber prefers shorter wait-times for passengers, we're interested in the average drive time to a passenger's location from her nearest driver as defined in (1)—the lower the better.<sup>3</sup>

$$\bar{d}(s) = \frac{\int_X f(x) \cdot \min_i (d(x, s_i)) dx}{\int_X f(x) dx} \quad (1)$$

To compare across different applications, we look not just at the average distance but also the difference between the average distance and what the average distance would be in an optimized allocation, holding  $n$  fixed. (2) defines an  $n$ -optimal spatial allocation,  $s^{*n}$ , as one which minimizes the average distance.

$$s^{*n} \in \arg \min_{w \in S_n(X)} \bar{d}(w) \quad (2)$$

Finally, (3) defines a measure of spatial inefficiency,  $\xi(s)$ , as the difference between the average distance in  $s$  and that in  $s^{*|s|}$ , dividing by the latter to get a percentage difference that abstracts from units.

$$\xi(s) = \frac{\bar{d}(s) - \bar{d}(s^{*|s|})}{\bar{d}(s^{*|s|})} \quad (3)$$

Note that our measure of spatial inefficiency does not account for prices. This is by design as it allows us to apply it to a variety of spatial applications, only some of which will involve price competition. However, it also means that a social planner may prefer a spatially inefficient allocation over one that is spatially efficient if the former yields advantageous consequences in terms of prices.

## 2.2 Gas stations on Interstate 90 (1D)

The simplest environment to study spatial competition is a bounded line with uniform customer density. This is the environment of Hotelling (1929), which we review in Section 3.1. We believe that the spatial competition of gas stations choosing locations on rural segments of an interstate highway roughly matches this. If most of the traffic on a rural segment of an interstate is driving from one end of the segment to the other, and we average over time, the customer density is effectively uniform across the segment.

### 2.2.1 Data

We take Interstate 90 (I90) between Seattle, Washington and Madison, Wisconsin as the focus for this analysis because it has many long, rural segments. There are still significant

---

<sup>3</sup>Our methods are easily extended to quadratic transport costs and average squared distances.

urban centers along this half of I90, and we would expect customer density to be significantly higher in these centers. Therefore we divide I90 into seven segments, cutting it at each point at which I90 passes through an urban area.<sup>4</sup> The seven segments, between eight urban centers, are detailed in Table 1. With each segment, we remove portions at each end that lie in the urban areas. We then optimize allocations and evaluate efficiency on each segment separately. In doing so, we are assuming uniform customer density only within each segment—two points on I90 between Missoula and Billings are assumed to have the same customer density, but a point in one segment is not assumed to have the same density as a point in another.




Figure 2: Interstate 90 from Seattle to Madison

Table 1: Summary statistics on segments of Interstate 90

Segment	Miles	Exits	Gas Stations
(1) Seattle to Spokane	252	58	46
(2) Coeur d'Alene to Missoula	153	41	22
(3) Missoula to Billings	333	66	47
(4) Billings to Rapid City	354	73	50
(5) Rapid City to Sioux Falls	335	63	50
(6) Sioux Falls to La Crosse	288	61	43
(7) La Crosse to Madison	120	23	41

Our gas station data was compiled in two steps. First, we downloaded all gas stations in the relevant states from the OpenStreetMap (OSM) project (OpenStreetMap contributors, 2017). We selected gas stations within a 3-minute drive of I90 using Geographic Information

<sup>4</sup>For urban areas, we use the 2010 Census Urbanized Areas (UA), densely settled cores of census tracts/blocks with a total of at least 50000 people. We ignore Urban Clusters (UC), which are defined similarly but with populations between 2500 and 50000.

System (GIS) software. Unfortunately, because OSM maps are maintained by volunteers, they tend to be least accurate/complete in rural areas, where there are few contributors. Therefore, we also manually scrolled Google Maps over the entirety of I90 with “gas stations” in the search bar. At the appropriate level of zoom, this reveals all features coded as gas stations in Google’s database (Google, 2017). Where these were missing from the OSM data, we added them.

Gas stations are, of course, not located along I90 itself, but rather at or near exits from the highway. In this sense, the game of spatial competition is not played along a line, but really across a set of discrete points positioned on the line. While we have the geocoded positions of the gas stations, for the purposes of our analysis a gas station’s location is its exit number. Where two or more gas stations are closest to the same I90 exit, we treat them as colocated. Exit numbers on I90 are mile markers, specifying highway mileage from the westernmost point of I90 in the current state. Since most segments span multiple states, we convert the state exit numbers to national ones, i.e. I90 distances from Seattle, for the analysis.<sup>5</sup>

The 3-minute drive time was used more as a rule of thumb than a strictly applied criterion for inclusion. In some urban clusters, there are several gas stations located right at the highway exit and then additional gas stations near the center of the town. In these cases, we reason that highway traffic is served by the stations placed nearest to the exit and the stations in the center of the town are primarily serving residents. We exclude the latter even if they are less than 3 minutes from the exit. In other cases, stations may be slightly further than 3 minutes from an exit but are likely serving I90 drivers given the lack of nearby alternatives—we include these.

### 2.2.2 Computing optimal allocations

In a world with full information and perfectly rational agents that plan their stops, the locations of gas stations would be unimportant—as long as there are no absurdly long stretches of highway without a station, a driver could plan her stops to avoid running out of gas. In a more realistic setting, however, drivers do not plan their stops and gas stations also offer food, coffee, and restrooms. Insofar as hunger, fatigue, and one’s bladder are less predictable than a car’s gas consumption, there is a benefit to having gas stations spread out so that a driver in need of these services is never too far from the next gas station.

Our primary measure of spatial inefficiency,  $\xi(s)$ , compares the actual allocation of gas stations with an optimized one with the same number of features. If gas stations could be placed along I90 itself, rather than at its exits, the optimal allocation on each segment would evenly divide the segment to minimize the average distance between a driver and her nearest gas station. But using this as the optimized allocation would mean that the resulting  $\xi(s)$  would conflate the inefficiency resulting from the allocation of the actual gas stations across the exits with that resulting from the locations of exits. To isolate the former, we compute the optimal allocation of gas stations constraining them to the actual exits, an exercise in combinatorial optimization.

Exhaustive search, which would require 58 choose 46 (almost a trillion) computations on

---

<sup>5</sup>In our figures, however, we include the actual, state-specific exit numbers.

just the I90 segment from Seattle to Spokane, is not feasible. Instead, we first distribute the actual number of gas stations of the segment arbitrarily, but without collocation, across the exits in the segment. Then, we iterate over the gas stations, checking for each if the average distance to gas station on the segment could be reduced by moving the gas station to an exit that does not have a gas station in the current iteration. When we find such a beneficial move, we implement it and repeat the process on the new allocation. The optimal allocation is a fixed point in this algorithm and the algorithm must converge to a fixed point. We are unsure as to the uniqueness of the fixed point—we have not observed multiple fixed points on any of the segments after some experimentation with differing initializations. If there are multiple fixed points, they need not be equally efficient, and therefore our “optimal” allocations may be suboptimal. In this case, our reported inefficiencies would be underestimates.

### 2.2.3 Results

To analyze the efficiency of the actual allocations on each segment, we report in Table 2 several statistics for each of the seven segments. First, we include the collocation percentage of gas stations in the actual allocation:  $CP(s^{act}) = x\%$  implies that  $x$  percent of the gas stations in the actual allocation are located at exits which have two or more gas stations. Two columns,  $\hat{d}(s^{act})$  and  $\hat{d}(s^*)$ , report the longest stretch without a gas station for the actual and optimized allocations respectively.<sup>6</sup> We also include three relevant measures of average distance. The second and third displayed,  $\bar{d}(s^{act})$  and  $\bar{d}(s^*)$ , are the average distances (in miles) to the nearest gas station, as defined in (1) and (2) and used in (3), for the actual and optimal allocations respectively.<sup>7</sup> The first,  $\bar{d}(s^{ran})$ , is an estimate of the average distance that would result from a random allocation. To calculate these, we randomly assign gas stations to exits (with replacement of exits, allowing collocation) a thousand times, calculate the average distance of each, and then average these averages. We think these provide an additional interesting point of comparison with the actual allocations. Finally, we include our measure of spatial inefficiency,  $\xi(s^{act})$ .

The results are surprising. The last segment, in Wisconsin, is an apparent outlier. Along this segment, the most populated, there is a gas station at all but one exit. Demand is sufficiently high such that spatial scarcity is limited and its binding constraint is the number of exits rather than the number of gas stations—we would expect to see similar results for segments of I90 east of Madison.

For all other segments, the our measurements of spatial inefficiency are huge. The 4th segment, from Billings to Rapid City, is both the longest segment and the least efficient. Its  $\xi(s^{act})$  of 279% implies that the average distance with the actual allocation of gas stations is almost four times that of the optimal allocation with the same number of gas stations. While this is the largest, the spatial inefficiency is above 100% for all of the non-outlier segments, suggesting average distances over twice as long as they would be in the optimal allocations. The average distances for the actual allocation are even worse than those from an average

---

<sup>6</sup>The  $s^*$  here are the allocations optimized as described above, i.e. to minimize average distance. They do not necessarily minimize the maximum distance, though they usually coincide fairly closely.

<sup>7</sup>Note that the average distance to nearest gas station ignores directional considerations. For a driver, the more relevant number would be the expected distance to the next gas station (in the right direction), for which the number would be roughly double the average distance to nearest gas station.

Table 2: Evaluating efficiency on stretches of Interstate 90

Segment	Colocation	Max. dist.		Average distance			Inefficiency
	CP%	$\hat{d}(s^{\text{act}})$	$\hat{d}(s^*)$	$\bar{d}(s^{\text{ran}})$	$\bar{d}(s^{\text{act}})$	$\bar{d}(s^*)$	$\xi(s^{\text{act}})\%$
(1)	74	41	21	3.6	5.0	1.9	160
(2)	59	28	10	4.0	4.1	1.8	126
(3)	83	39	17	4.3	6.2	2.0	206
(4)	74	66	16	4.2	7.5	2.0	279
(5)	64	31	16	3.8	4.1	1.8	130
(6)	72	33	12	4.0	4.5	1.7	158
(7)	83	14	14	2.5	2.0	2.0	2

$s^{\text{ran}}$ ,  $s^{\text{act}}$ ,  $s^*$  are the average randomly generated, the actual, and the optimal allocations, respectively.


random allocation in all non-outlier segments.

The first hint at the explanation of these results is in the colocation percentages—in each segment, a large majority of the gas stations are colocated with other gas stations despite the fact that there are also many exits without gas stations. But the narrative driving these results is particularly apparent when one looks at the allocations on a particular segment. In Figure 3d, we show  $s^{\text{act}}$  and  $s^*$  for the fourth segment. The segment has been cut into five horizontal lines. Each tick mark is an exit, with exit numbers resetting to zero in the middle of the second and fourth rows as I90 crosses state lines from Montana into Wyoming and then into South Dakota. Red circles above each bar denote actual gas stations at each exit while the blue circles below represent gas stations in the optimal allocation. Finally, shading has been included above and below the bar to denote distance from the nearest gas station for that point on the highway in the actual and optimal allocations respectively. The analogous figures for other segments, 3a–3g, are included in the appendix.

Figure 3d shows that gas stations along this rural segment are clustered around population masses (urban clusters by census definitions). Exit 495 on the first row, with five stations, is the small town of Hardin, Montana. Exits 20, 23, and 25 in the second row are all exits for Sheridan, Wyoming. Exits 124, 126, and 128 on the third are exits for Gillette, Wyoming. Gas station clusters in the fourth and fifth rows similarly identify other small towns. Even the longest stretch of highway without a gas station, 66 miles in Wyoming from Buffalo (exit 58) to Gillette (exits 124–128) coincides with a long stretch of highway without a population mass. Essentially, the actual allocation of gas stations mirrors the locations of residents along I90, seemingly in contradiction with our prior that most of the demand for gas came from uniformly distributed through-drivers.

We see a few possible explanations for this. One possible explanation is that our assumption of uniform customer density is incorrect. We assume uniform customer density because we believe that most drivers on rural segments of I90 are driving through that segment, not residents in towns along the segments. The fact that there is significant traffic flow even on remote sections of I90 would seem to support this. Alternatively, perhaps significant demand comes from both through-drivers and residents but the residents are better served because of





Numbers/ticks are highway exits. Red dots (above the line) and blue dots (below) are gas stations in the actual and optimal allocations, respectively. Shading represents distance from nearest gas station.

Figure 3d: Gas stations on I90 between Billings and Rapid City

an informational story akin to a *tourists and natives* model (Salop and Stiglitz, 1977). There could also be complementarities at play if through-drivers prefer to get food or lodging in the same location as gas. Finally, there could be reasons on the supply side—station owners could prefer to have their businesses near their home or those of their employees, or perhaps distribution is cheaper when gas stations are colocated. What we can say without hesitation is that the spatial allocation of gas stations on I90 is highly inefficient for the purposes of through-drivers. In the following section, we do a similar exercise in two dimensions, taking into account customer density rather than assuming it uniform.

### 2.3 City supermarkets, hospitals, and fire stations (2D)

Most firms engaged in spatial competition face a two-dimensional game with non-uniform customer distributions. To consider efficiency in these settings, we need data on the actual allocations of firms and the customer distribution. We also need to be able to compute optimal allocations to form the reference points for our calculations of spatial inefficiency.

In this section, we evaluate allocations of supermarkets, hospitals, and fire stations in three major cities. We refer to these generally as features, rather than firms, to accommodate the inclusion of fire stations as well as non-profit hospitals, which may or may not behave as firms. We chose these three classes of features because we suspect that the mech-

anisms generating their allocations span the spectrum from regulated market competition (supermarkets) to central planning (fire stations), with hospitals somewhere in between.<sup>8</sup> By comparing  $\xi(s^{\text{act}})$  across these classes of features, we can get a rough empirical indication of how spatially efficient competitive mechanisms are relative to central planning. That is, we look to identify the spatial inefficiency resulting from the competitive mechanism by looking at a difference in differences—our evidence is not that  $\bar{d}(s^{\text{act}})$  is greater than  $\bar{d}(s^*)$  for supermarkets, though this is true and the difference is large, but rather that the percentage difference  $\xi(s^{\text{act}})$  is much larger for supermarkets than for hospitals and fire stations.

Analysis with (1)-(3) in this context requires strong assumptions. First, we are viewing our customers as static components of the environment, and identical but for location. In the ride-sharing application, the idea that customers are static is natural. However where we look at spatial allocations of supermarkets, that assumption is stronger—the customers, in their choices of residence, are perhaps as mobile as the firms.<sup>9</sup> Second, we are assuming that the firms are identical and can serve arbitrarily many customers, akin to Bertrand (1883). Third, we are treating the customer as uniquely located, i.e. ignoring the possibility that a customer could frequent a distant supermarket with little inconvenience due to it being located on a commute. In using Euclidean distances, we ignore transportation networks. We also ignore important empirical infeasibilities, both physical and regulatory. These include infeasibilities of traveling in a straight line between pairs of locations as well as infeasibilities in feature placement due to lakes, zoning, etc.

We make these strong assumptions because they allow us to compute optimal spatial allocations, a challenging exercise, and therefore also to meaningfully evaluate  $\xi(s^{\text{act}})$ . With all of these assumptions, even if an empirically observed spatial allocation were actually optimal for our measure given real-world constraints, we would still likely calculate a non-negligible  $\xi(s^{\text{act}})$  for it given that our optimal  $s^{*n}$  is computed without these real-world constraints. However, to argue that the spatial allocations of supermarkets are inefficient, we show not just that their inefficiencies are large, but more importantly that they are much larger than those for allocations of hospitals and fire stations. Unless the assumptions are substantially more problematic for supermarkets than they are for hospitals and fire stations, differences in  $\xi(s^{\text{act}})$  can still be attributed to the difference in mechanisms generating the allocations.

### 2.3.1 Data

Our data on supermarket locations comes from OSM.<sup>10</sup> OSM’s definition of a supermarket is a large store for groceries and other goods. It includes only full-service grocery stores, meaning most specialty and ethnic grocers are not included. Data for hospitals comes from two sources. Our primary source is the *Hospital General Information* dataset in Medicare’s *Hospital Compare* data (Medicare, 2017). From this, we select all acute care and critical

---

<sup>8</sup>While this vague conjecture is as deep as we go into the mechanisms generating the allocations in this section, a goal of the broader research agenda is to design or improve mechanisms to yield spatially efficient outcomes.

<sup>9</sup>While people may move frequently, housing stocks, and therefore population distributions, change slowly.

<sup>10</sup>See OpenStreetMap contributors (2017). Fortunately, this dataset appears to be far more accurate and complete in urban areas than it is for gas stations in rural areas.

access hospitals that have emergency services. Because the geocoding in that database is incomplete, we cross-reference each hospital with its entry in the *Hospitals* database of the Department of Homeland Security’s Homeland Infrastructure Foundation-Level Data (HIFLD) to get its location. Our fire station data also comes from HIFLD.<sup>11</sup>

We conduct our analyses on three cities: Atlanta, Chicago, and Los Angeles. We selected these cities both for certain desirable characteristics as well as technical reasons. For characteristics, we wanted cities of different sizes and in different regions to argue the external validity of our results. Substantial racial diversity in the three cities allows us to investigate the role of race in feature allocations through spatial regressions in another project. As for technical reasons, we wanted cities with statistical areas that were surrounded by areas of low population density—this allows us to analyze allocations on the city’s statistical area and ignore users of the features outside of that area without introducing much error. Additionally, our optimization algorithm, described below, may struggle on cities that have sizable interior areas with zero customers, so we avoided cities with significant interior areas of water.<sup>12</sup> For each city, we conduct our analyses on both that city’s metropolitan statistical area (MSA) and on a much smaller area roughly corresponding to the city lines.<sup>13</sup>

Finally, we take our population data from two sources. We account for population and population density at the level of census tracts with TIGER data from the 2010 US Census (U.S. Census Bureau, 2016). We then augment that with updated 2015 population estimates from Esri et al. (2017). TIGER data includes coordinates for each vertex of each census tract. We cannot use latitude and longitude coordinates directly for Euclidean distances because the distance of moving a degree North/South does not equate with that of moving a degree East/West. Instead, for each area, we convert the census tract vertices and feature locations to a metric stereographic projection centered in that area so that Euclidean distances between the points can be interpreted as distances in meters with only minimal error resulting from the earth’s curvature.<sup>14</sup>

### 2.3.2 Computing optimal allocations

To compute optimal feature allocations, we use a numerical algorithm related to Lloyd’s algorithm, which finds evenly spaced sets of points in subsets of Euclidean spaces (Lloyd, 1982). While more commonly used in computer science and electrical engineering, the al-

---

<sup>11</sup>See Oak Ridge National Laboratory (2017) for HIFLD hospital data and TechniGraphics, Inc. (2010) for HIFLD fire station data.

<sup>12</sup>Having water form a boundary, as is the case with Chicago and Los Angeles, is not an issue. Small rivers and lakes are also fine. But cities like Boston, New York, and San Francisco would yield additional challenges.

<sup>13</sup>The boundaries of American cities are complicated—many cities contain non-city enclaves and include disconnected exclaves. We take convexifications of the actual cities. Chicago City includes all census tracts that intersect city boundaries with the exclusion of those around O’Hare airport (a city exclave) and with the inclusion of tracts in Norridge (a non-city enclave). Atlanta City includes all tracts in Fulton and DeKalb counties that intersect city boundaries, except one (FIPS: 13089020802) that juts out and includes a non-city exclave. Los Angeles City includes all tracts that intersect city boundaries, adding in several non-city enclaves (e.g. Beverly Hills and Santa Monica), and truncating the corridor leading down to the Port of Los Angeles by excluding all tracts south of Interstate 105.

<sup>14</sup>With our projection on Chicago, for instance, the calculated distance between two points 200 kilometers apart will have an error of about 76 meters.

gorithm’s underlying logic of reaching a global optimum through iterated local optimization serves our purposes equally well.




Figure 4: Local adjustments in Lloyd’s algorithm

Consider the problem of trying to place features to minimize average *squared* distances between customers and their nearest features on a Euclidean plane with uniform customer density. Start from the arbitrary allocation represented in Figure 4, in which the original positions of seven features are represented by seven circles. The irregular hexagon with a solid perimeter is the Voronoi cell of the central feature.<sup>15</sup> Lloyd’s algorithm is the iterative movement of features to the centroids of their Voronoi cells.<sup>16</sup> In the figure, the centroid of the central feature’s Voronoi cell is shown with a diamond.

The argument for why moving the central feature to the centroid of its cell must reduce the overall average squared distance is simple. First, note that the centroid of a shape is, by definition, the point that minimizes the average squared distance between a single point and all others in the shape. So if Voronoi cells were unaffected by a movement, the adjustment would have to yield an overall improvement—we would be reducing average squared distances within a particular cell, the solid hexagon, without affecting average squared distances elsewhere.

Of course, any feature movement will shift and contort its Voronoi cell. After moving to the diamond, the central feature’s new Voronoi cell is the irregular hexagon with a dashed perimeter. The central feature has lost customers in the red area and gained those in the green area. In the previous paragraph, we argued that the adjustment yielded a net improvement in the average squared distance for those customers within the solid hexagon. In that assessment, we ignored the welfare gains of customers in the green area. Because

<sup>15</sup>A generator point’s Voronoi cell is the set of points that are closer to that generator point than any of generator point. In this context, the generator points are the locations of firms.

<sup>16</sup>The centroid of a shape is the arithmetic mean position of all points in the shape.

they are in the Voronoi cell of the central feature after the adjustment, they must be closer to the central feature after the adjustment than they were to their original nearest feature. As for the red area, because it lies inside of the solid hexagon, their welfare loss from the adjustment was included in the analysis of the previous paragraph that suggested that the adjustment was an improvement. However that analysis actually overstates the welfare loss of those in the red area by not allowing for the fact that these customers would visit other closer features instead of the central feature following the adjustment. So the previous paragraph undercounted welfare gains from the adjustment and overcounted welfare losses, and it still showed that the adjustment was advantageous—it follows that the adjustment is advantageous.

Since each adjustment of a feature to the centroid of its Voronoi cell reduces the average squared distance, adjusting all features iteratively (Lloyd’s algorithm) converges to a fixed point. And the optimal allocation must be a fixed point. Unfortunately, Lloyd’s algorithm may have multiple fixed points. As an example, the allocation in the center and on the right of Figure 1 are each fixed points in Lloyd’s algorithm on a disk with six features.

In optimizing feature allocations on cities, we face two challenges that go beyond standard implementations of Lloyd’s algorithm. First, our goal is to minimize the average customer distance to nearest feature, not the average squared distance. Second, our customer distribution function, drawn from census data, is non-uniform.

In minimizing average distance, we can still rely on the same Lloyd’s algorithm logic that we discuss in Figure 4. The difficulty is that instead of adjusting a feature to the centroid of its Voronoi cell, we instead want to adjust it to the point within its Voronoi cell that minimizes the average distance to all other points in the cell, i.e the geometric median. Finding the geometric median for a discrete set of points is well-studied as an important problem in facility location.<sup>17</sup> Of course, ideally we would want to find the geometric median not of a discrete set of points but that of a region. There are algorithmic solutions for this continuous Fermat-Weber problem,<sup>18</sup> but none that we are able to implement quickly enough (computationally) to be feasible within our version of Lloyd’s algorithm. Instead, we approximate the region of the Voronoi cell by a large set of discrete points drawn randomly from the region, and then compute the geometric median of those points. This can be done almost instantaneously even with tens of thousands of points.

Our implementation of the geometric median also facilitates our extension to non-uniform customer density. Figure 5 shows the Voronoi cell of a supermarket in the Chicago MSA. The figure’s interior lines represent boundaries of census tracts—the cell intersects roughly 30 tracts. The census tracts are shaded red with darker reds representing higher population densities. The black region to the right of the figure is Lake Michigan. The black circle is the current location of the supermarket. To calculate the geometric median, the gray circle, we draw points randomly from each tract-cell intersection, with the number drawn from each determined by the area of the tract-cell intersection multiplied by the population density of the tract, all divided by the sum of these products across all tract-cell intersections. This

---

<sup>17</sup>A firm wanting to minimize the total distance between itself and its suppliers would locate at the geometric median of its suppliers. Finding that point is often called the *Weber problem* after Weber (1929) or the *Fermat-Weber problem*, recognizing Fermat for having originally posed the problem with three suppliers. See Drezner and Hamacher (2001) for a review.

<sup>18</sup>See Fekete et al. (2005), for example.

gives us a percentage of our sample size to draw from each tract-cell intersection. We then construct the sample of discrete points and find the geometric median of the sample.




Figure 5: Voronoi cell of Super Fresh Market in Waukegan, IL

Fixed point multiplicity is a major issue in our algorithm. With the gas station analysis, we ran that algorithm from several initializations and always converged to the same fixed point. When we run our algorithm here from different initializations, we converge to many different fixed points. The reason for this lies in the non-uniformity of the customer distribution and the fact that our algorithm involves only local adjustments. The algorithm essentially pulls each feature towards its nearest population mass, as if the model were gravitational, and then optimizes the allocation of features over each population mass. What it does not do as well is distribute features efficiently across the population masses. As an example, in the Chicago MSA, there may be one fixed point with 4 supermarkets serving Aurora and 2 serving Joliet, another with the numbers reversed, and another with 3 serving each. Because the population density is low between the two suburbs, gas stations are not necessarily pulled across from one suburb to another even if it would be optimal to do so.

Our approach to dealing with fixed point multiplicity is to run the algorithm from fifty randomly generated initial allocations as well as the actual one for each optimization, selecting the best. For about half of our optimizations, the actual allocation proved a better initialization than any of the fifty randomly generated initializations. This is not surprising—we show below, in Figure 6, that actual feature allocations match population densities fairly well. Insofar as they have the right number of features in each community, the actual allocations serve as good initializations for the algorithm.

Lloyd’s algorithm technically never requires the computation of average (squared) distances. However, we need to be able to evaluate this for two purposes. First, it’s an obvious metric over which to define a tolerance and determine convergence of the algorithm. Second, we need to evaluate average distance for both the actual and the optimal allocations to measure  $\xi(s^{\text{act}})$ . To this end, we compute the average distance over the entire city or MSA by computing the average distance within each census tract through numerical integration and then taking a weighted average across the census tracts based on their populations. Importantly, the approximation of a region by drawing points randomly that we describe above and use in our implementation of the algorithm is not relevant to the final evaluation of average distances for both the observed and optimal allocations. So while it is true that the

allocation we call optimal is only approximately so, the precision of our evaluations of the average distances is only limited by the minor potential imprecision of numerical integration.

### 2.3.3 Results

In Table 3, we present summary data on the city and MSA of each of our three cities. We include their areas in square miles, their populations, and the number of supermarkets, hospitals, and fire stations contained in each.

In Tables 4 and 5, we replicate the efficiency analysis that we did for gas stations in Table 2. Table 4 shows the average mileages to features for each feature, each region, and each of three allocations. Just as in the gas station analysis,  $\bar{d}(s^{\text{ran}})$  is the average mileage across a thousand randomly generated allocations with feature numbers equal to those in the corresponding actual/observed allocations. We then show the average distances in the actual allocations,  $\bar{d}(s^{\text{act}})$ , and those in the optimal allocations,  $\bar{d}(s^*)$ , which were generated by our algorithm. In Table 5, we report the implied spatial inefficiencies for each feature in each city.

For hospitals and fire stations, the actual allocations are better than the average of the randomly generated allocations for all six regions. For supermarkets, the actual allocations are worse than the randomly generated allocations for the cities but better for the MSA's, with the exception of Atlanta MSA, for which the two are very close. This is a first indication of our general point that supermarkets are inefficiently allocated.

The inefficiency of the supermarket allocations is yet more stark in Table 5. In each of the six regions, the spatial inefficiency associated with the supermarket allocations is significantly greater than that of hospital and fire station allocations. While the magnitudes are not as large as those that we found for gas stations,<sup>19</sup> they still represent average distances between 56% and 96% longer than the optimal configurations. The difference between being 1.09 miles and 0.6 miles from a supermarket in the city of Chicago, for instance, is fairly significant for customers without a car.<sup>20</sup>

One might question whether our measure of inefficiency would correlate with the number of features, and thus a comparison between two allocations of different  $n$  would not be meaningful. In this case, however, there are more supermarkets in each region than there are hospitals and less supermarkets than fire stations, and both hospital and fire station allocations are more efficient. Another contention would be that our assumption that supermarkets can serve arbitrarily many customers is influencing the results. But this is also true of hospitals and fire stations,<sup>21</sup> for which we find less inefficiency.

We include figures to directly compare the actual allocation of supermarkets with the optimal allocation. Figures 6a and 6b show the population densities of census tracts in

---

<sup>19</sup>The inefficiency in the supermarket allocations appears to come from their tendency to cluster near areas of higher population density. In this sense, they are at least clustered where the people are. For the gas stations, given our uniform customer distribution assumption, the gas stations were clustered in essentially arbitrary locations. This is the intuition of why the inefficiency magnitudes were much larger in that analysis.

<sup>20</sup>While we are using supermarkets as an example to illustrate a broader point about spatial allocations resulting from competition, our analysis overlaps here with literature on food deserts and access. Distances affect food choice, not just convenience. For a review of this literature, see Walker et al. (2010). Also see ERS-USDA (2017) for a related spatial atlas.

<sup>21</sup>Fire stations' duties may scale more with area than population.

Table 3: Summary statistics on regions and features therein

<b>Region</b>	<b>Sqmi</b>	<b>Pop. (mil.)</b>	<b>Tracts</b>	<b>S</b>	<b>H</b>	<b>F</b>
(1) Atlanta City (AC)	163	0.47	130	19	4	38
(2) Atlanta MSA (AM)	8835	5.53	951	174	37	502
(3) Chicago City (CC)	227	2.78	803	76	25	96
(4) Chicago MSA (CM)	6304	9.56	2210	259	80	759
(5) LA City (LC)	536	4.44	1109	100	29	137
(6) LA MSA (LM)	4754	13.14	2923	277	86	535

Table 4: Average mileages to features for random, actual, and optimal allocations

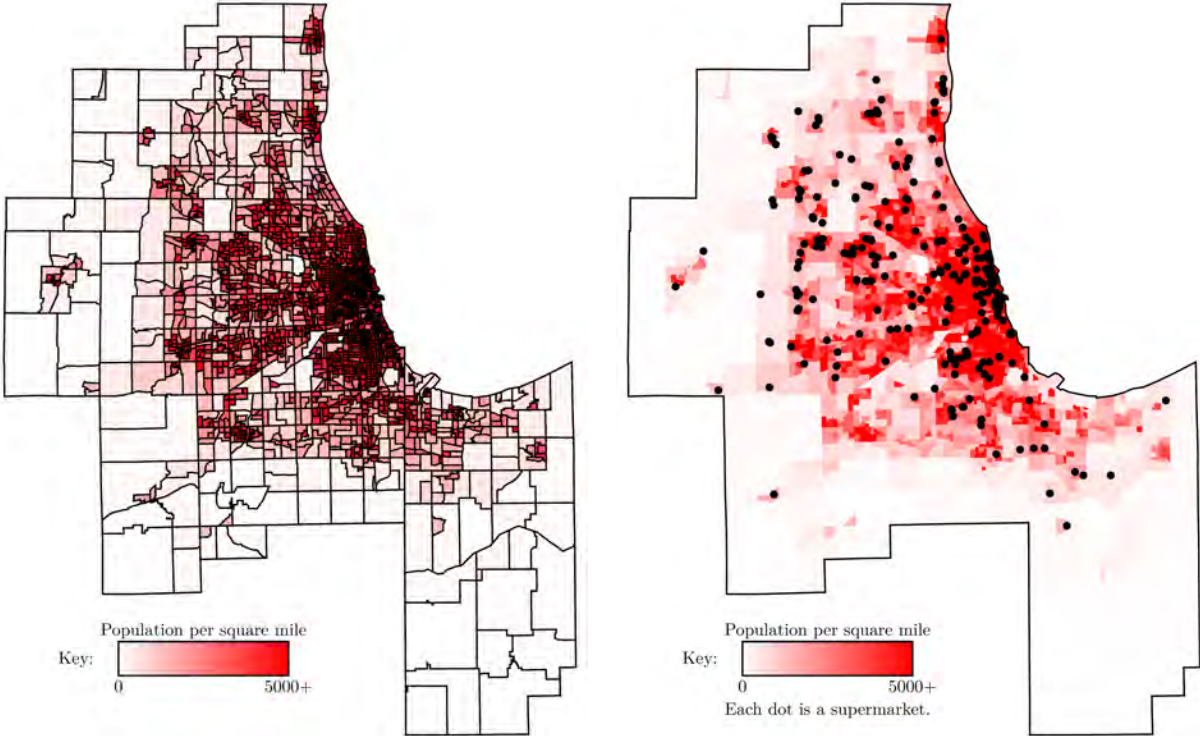
<b>Region</b>	<b>Supermarkets</b>			<b>Hospitals</b>			<b>Fire Stations</b>		
	$\bar{d}(s^{\text{ran}})$	$\bar{d}(s^{\text{act}})$	$\bar{d}(s^*)$	$\bar{d}(s^{\text{ran}})$	$\bar{d}(s^{\text{act}})$	$\bar{d}(s^*)$	$\bar{d}(s^{\text{ran}})$	$\bar{d}(s^{\text{act}})$	$\bar{d}(s^*)$
(1) AC	1.58	1.90	1.01	3.58	3.47	2.24	1.10	0.86	0.71
(2) AM	3.58	3.59	2.01	7.83	5.13	4.24	2.11	1.47	1.18
(3) CC	0.93	1.09	0.60	1.67	1.49	1.06	0.82	0.61	0.54
(4) CM	2.74	2.40	1.39	5.01	3.06	2.34	1.58	0.90	0.75
(5) LC	1.20	1.43	0.73	2.33	2.06	1.36	1.02	0.75	0.63
(6) LM	2.12	1.77	1.00	3.83	2.30	1.77	1.52	0.82	0.68

Table 5: Spatial inefficiency of actual feature allocations

<b>Region</b>	$\xi(s^{\text{act}})\%$		
	<b>S</b>	<b>H</b>	<b>F</b>
(1) Atlanta City	87	55	22
(2) Atlanta MSA	79	21	25
(3) Chicago City	82	41	14
(4) Chicago MSA	73	30	21
(5) LA City	96	51	19
(6) LA MSA	78	30	20



Chicago MSA as well as the actual spatial allocation of supermarkets. The latter shows that supermarkets do tend to locate in high-density areas, which is efficient, but they also tend to cluster, which is not efficient by our measure. Figures 6c and 6d show distances of all points in Chicago MSA for the actual and optimal allocations respectively. It is immediately apparent that most of the MSA is much closer to their nearest supermarket with the optimal allocation. Of course, what matters is not how close points in the region are to their nearest supermarket, but rather how close people are. Therefore, we also provide *scarcity plots* that reveal where there is significant population that is distant from a supermarket. To this end, in Figures 6e and 6f we map over the region a density plot in which the shade of coloring at a point is determined by the product of the population density of the containing tract and the distance from that point to its nearest supermarket.<sup>22</sup>




(a) Census tracts and pop. density

(b) Supermarkets and pop. density


Figure 6: Chicago MSA

Figures 6c-6f also show that a consequential difference between the actual and optimal allocations is that a significant number of supermarkets have been pushed out from the city center to Chicago’s suburbs and exurbs in the optimal allocation. This is, indeed, optimal in terms of minimizing average distances. Figure 6a is a little misleading in its portrayal of population density as it seems to suggest that there is very little population at the periphery of the MSA. But the almost-white census tracts on the periphery actually have roughly the same population as census tracts closer into the city—census tracts are


<sup>22</sup>We subtract one mile from the distance in the figure coloring as otherwise high-density areas suggest scarcity even quite close to a supermarket.




(c) Distance to supermarket in  $s^{\text{act}}$



(d) Distance to supermarket in  $s^{\text{opt}}$



(e) Supermarket scarcity in  $s^{\text{act}}$



(f) Supermarket scarcity in  $s^{\text{opt}}$

Figure 6: Chicago MSA (continued)

defined to have similar populations. There are even some significant towns of concentrated population density in some of these peripheral census tracts. But the population densities for these tracts, which determines the shading, is still almost negligible because of their size. In any case, it should come as little surprise that the optimal allocation attempts to serve these communities.<sup>23</sup> The fact that the actual allocation does not hints at the sort of principle of minimum differentiation clustering that was posited in Hotelling (1929).<sup>24</sup> All of this notwithstanding, we did the analyses on the city boundaries also to show that not all of our efficiency gains in the optimal allocations come from movements of features from the inner city to suburbs and exurbs. Figures for other features in Chicago MSA and for supermarkets in Chicago city are provided in Appendix B.

We can also estimate the cost of the spatial inefficiency in supermarket allocations in monetary terms. Our optimal allocation reduces the average distance to nearest supermarket from 2.4 miles to 1.4 miles. Since we have no evidence as to the feasibility of the optimal allocation, let us compare the actual with an allocation of supermarkets that has the same spatial inefficiency of 30% as Chicago hospitals, the less efficient of the other two feature allocations. Matching that 30% inefficiency means bringing the distance from 2.4 miles down to about 1.8 miles. For a single round-trip, this represents an average savings of 1.2 miles. If we multiply this by the roughly 3.4 million households in the Chicago MSA and suppose that a member of each household visits a supermarket weekly, we get a reduction of about 214 million miles traveled annually. If we multiply this by the IRS standard mileage rates used to calculate the deductible costs of operating an automobile for business, 53.5 cents per mile, we would value these miles at about 114 million dollars.<sup>25</sup> This estimate includes only direct transportation costs, not time costs. Even for the transportation costs, we think it is highly conservative—reducing the Euclidean distance to supermarkets by 0.6 miles usually reduces the travel distance by more than 0.6 miles, and we suspect per-mile travel costs to be greater than 53.5 cents per mile for those without cars.<sup>26</sup>

Further research is required to determine the degree to which inefficiencies we find for gas station and supermarket allocations can be generalized to other allocations resulting from competition. We also do not isolate to what degree current regulation may either limit or exacerbate these inefficiencies. But we believe that the tendency of similar firms to cluster is fairly general. In retail and dining with limited consumer information, this may be efficient—consumers enjoy shopping at multiple retail stores at one location and diners can choose a location and then compare dining options. But when firms are homogeneous, as is the case with gas stations and supermarkets, we view clustering as an inefficient phenomenon.

Insofar as the clustering of homogeneous firms is inefficient, we are interested in policies that could limit clustering. But our empirical analysis has no direct policy conclusions. We do not propose optimizing the spatial allocations of supermarkets or gas stations because this

---

<sup>23</sup>We do not consider different rates of car ownership between urban and sub/exurban tracts—if our goal was to generate truly optimal allocations, considering this would be necessary.

<sup>24</sup>Remember that our optimal allocation is for the  $n$  from the actual allocation. We do not suggest that this  $n$  is optimal, nor that policymakers should seek to implement  $s^*$ . Our focus is the inefficiency of  $s^{\text{act}}$ .

<sup>25</sup>Analogous calculations for Atlanta MSA and Los Angeles MSA yield 127 million dollars and 114 million dollars, respectively.

<sup>26</sup>We grant that some households get groceries from vendors other than full-service supermarkets. While this may reduce their transportation costs, it has other costs in terms of food choice and public health.

could yield additional deadweight loss from price competition—we suspect that the clustering of homogeneous firms has advantageous effects on prices. Our optimized allocations also take the number of firms as given, and we have not considered the different sizes of our features and the value of land.<sup>27</sup> To consider policy, we need a predictive model of spatial competition, the development of which we discuss in the next section.

### 3 Modeling spatial competition

In this section, we build towards agent-based models of spatial competition. In Section 3.1, we review insights from static games. In Section 3.2, we describe how location theory can be extended through dynamic, agent-based models with myopic agents.

#### 3.1 Insight from static spatial games

Spatial games are typically modeled as static games despite the fact that mechanisms generating empirical spatial allocations are likely continuous-time dynamic games.<sup>28</sup> Analyses vary in the number of firms, the space, and the customer distribution.<sup>29</sup> Each firm’s profit is proportional to the mass of customers that is closer to the firm than any other. If two or more firms are colocated, they split equally their joint mass of customers. Each firm simultaneously chooses its location,  $s_i$ , to maximize its profits. A Nash equilibrium spatial allocation,  $s$ , is such that no firm could increase its profits by unilaterally deviating to another location. In discussing equilibrium analysis on static spatial games, we illustrate three key impediments to using these models for applied prediction: i) multiplicity, ii) non-existence, and iii) intractability.<sup>30</sup>

Analysis of spatial competition dates back to Hotelling (1929).<sup>31</sup> Hotelling’s canonical

---

<sup>27</sup>Our optimized allocation has supermarkets further from the city center, on average, than the actual allocation, which likely means that the optimized allocation would occupy less valuable land than the actual one.

<sup>28</sup>Sequential games with irreversible decisions are commonly modeled as static games. In our setting, there may be uncertainty about timing in the game, and beliefs on that timing. See Penta and Zuazo-Garin (2017) for analysis on rationalizability in this context.

<sup>29</sup>For reviews of this literature, see Graitson (1982) and Gabszewicz and Thisse (1992).

<sup>30</sup>The challenges of theoretical analysis also motivate the Structure-Conduct-Performance (SCP) paradigm developed in Mason (1939, 1948), which looks for empirical evidence of relationships between industry structure and outcomes. See Bain (1951, 1956) for across-industry analyses and Stigler et al. (1983) for a critique that favored price theory models. In some sense, Von Neumann and Morgenstern (1944) developed game theory as an alternative to the SCP paradigm.

<sup>31</sup>Hotelling (1929) was a response to Bertrand (1883) and extensions in Edgeworth (1897). In turn, Bertrand (1883) was a paradox proposed in critique of Cournot (1838) and Walras (1883). Cournot’s duopoly with firms choosing quantities yields lower quantities and higher prices than the social optimum. Bertrand’s model is similar but with firms choosing prices, not quantities. Intuitively this should not matter given that prices determine quantities and vice versa, yet Bertrand’s model predicts no deadweight loss in the duopoly, thus the paradox. Hotelling’s critique of Bertrand (1883) focused on what he viewed as the unrealistic discontinuity in the Bertrand model where one seller goes from serving no customers to serving all of them as she moves her price from minimally above her rival’s to below: “... a discontinuity, like a vacuum, is abhorred by nature” (Hotelling, 1929, p.44). Ironically, Hotelling’s analysis, which had firms competing on both price and location, was incorrect because he failed to take account of discontinuities in his firms’

main-street, or linear-city, model has two firms competing on location. His key result, which came to be known both as Hotelling’s Law and the principle of minimum differentiation, is that firms may be incentivized to make their products as similar as possible. In a spatial model, this manifests as colocation. Spatial competition also relates to monopolistic competition, à la Chamberlin (1933), except that the product differentiation comes from the location of the sellers. Many real markets involve sellers in different locations selling differentiated products at different prices—high dimensionality makes it very difficult to make theoretical predictions in such a rich model.

*Model 1* of Eaton and Lipsey (1975) has a bounded line as its space and uniform customer density—it is a fairly straight-forward extension of Hotelling (1929), without prices, to cover  $n \geq 2$  firms. For  $n = 3$ , there is no equilibrium. For  $n = 4$  and  $n = 5$ , there exists a unique equilibrium with two firms collocated near each boundary, and one additional firm in the middle in the  $n = 5$  case. For  $n \geq 6$ , there is multiplicity—while the two firms nearest each boundary must be collocated, each interior firm may be collocated or uniquely located. Flexibility in the equilibrium positioning of the interior, individually-located firms means that there are infinitely many equilibria for any  $n \geq 6$ . For exposition, assume here that  $n$  is even. Again we take spatial efficiency as the inverse of the expected distance from a customer to her nearest firm. Then, in the optimal allocation, which is not an equilibrium, the  $n$  firms evenly divide the line. The most efficient equilibrium has all but the boundary firms uniquely located. The least efficient has all firms collocated in pairs. We show the three for  $n = 10$  in Figure 7.<sup>32</sup>




Figure 7: Equilibrium spatial allocations with ten firms

As we abstract from the boundary behavior by increasing  $n$ , the expected distance in the best equilibrium allocation converges to that of the optimal allocation, while that of

---

best response functions. This error was mentioned in Shubik (1959) and Vickrey (1964) but not corrected until d’Aspremont et al. (1979), which showed that no equilibrium existed in Hotelling’s original model. While Hotelling’s predictions were incorrect for his price-location model, they are correct for the same model without prices, and most now remember Hotelling’s contribution as one of pure spatial competition.

<sup>32</sup>We thank Liyang Liu for his assistance in extending the Eaton and Lipsey (1975) characterization to arbitrary  $n$ .

the worst equilibrium converges to twice that of the optimal allocation. The multiplicity of equilibria means that equilibrium analysis gives no conclusive answer to the question of spatial efficiency even in this simple setting of uniform customer density of a line.<sup>33</sup>

*Model 3* in Eaton and Lipsey (1975) departs from the assumption of a uniform customer density and finds a non-trivial necessary condition for the existence of equilibrium: the number of firms cannot exceed twice the number of modes of the (assumed continuous) customer density function. In particular, this means that for any  $n \geq 3$ , there exists no equilibrium on if customer density is single-peaked—imagine a linear city with population density highest at the center and tapering toward the boundary in each direction. Where multiplicity impeded prediction with uniform customer density functions, here non-existence is the challenge.

Spatial competition in applied settings usually takes place in two-dimensional spaces. Lösch (1954, p.94-97) considers the unbounded Euclidean plane with a uniform customer distribution and suggests the optimality and stability of an offset grid configuration of firms, the Voronoi diagram<sup>34</sup> of which is a hexagonal covering as shown in Figure 8a—both the letters and points represent firms and the hexagons are Voronoi cells.<sup>35</sup>

Lösch’s equilibrium is supported numerically in Eaton and Lipsey (1975) and proven analytically in Okabe and Aoyagi (1991).<sup>36</sup> Its optimality is proven in Bollobas and Stern (1972). In this setting, equilibrium and optimality coincide.

But Eaton and Lipsey (1972, 1976, 1978) find many flaws in the Löschian model and argue that competition need not yield that hexagonal covering. Further, even if one were to accept the prediction of the hexagonal covering in this setting, there is no evidence that this prediction generalizes beyond an unbounded Euclidean space with uniform customer density. Boundaries in the linear model implied inefficient colocation near them in equilibrium, while departing from uniform customer density raised questions of existence. The former issue extends fairly intuitively to two dimensions—a uniquely-located firm near a boundary is incentivized to move away from the boundary as it gains customers towards the interior without losing its essentially captive customers between it and the boundary.<sup>37</sup> As for the latter, little is known.

In fact, beyond the equilibrium characterization, very little is known even about the unbounded, uniform case. In equilibrium, each firm is located optimally given the locations of all other firms. Yet for a particular firm,  $A$ , we have no analytical solution for the optimal

<sup>33</sup>One might favor the worst equilibrium as a prediction given that it is the only strict equilibrium.

<sup>34</sup>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. It is a formalization of what we called *catchment regions* in describing Figure 1.

<sup>35</sup>Lösch (1954) attempts game-theoretic equilibrium analysis before the tools were well understood—his equilibrium conditions are a mix of behavioral postulates and conditions that purportedly follow from them. One of the equilibrium conditions, incorporating prices, was a zero-profit condition justified by free entry. But Eaton and Lipsey (1978) shows that the neoclassical result of free entry yielding zero profits does not survive an extension to space with scale effects, calling into question many spatial analyses that used that assumption—Eaton and Lipsey (1978, p.455) offers a list of such analyses in a footnote.

<sup>36</sup>Okabe and Aoyagi (1991) also proves the existence of a square-covering equilibrium. See Knoblauch (2002) for a particularly elegant proof.

<sup>37</sup>This is more nuanced in 2D where Voronoi cell walls pivot with any movement such that only those customers on the line of the direction of movement and between the firm and boundary are truly captive.



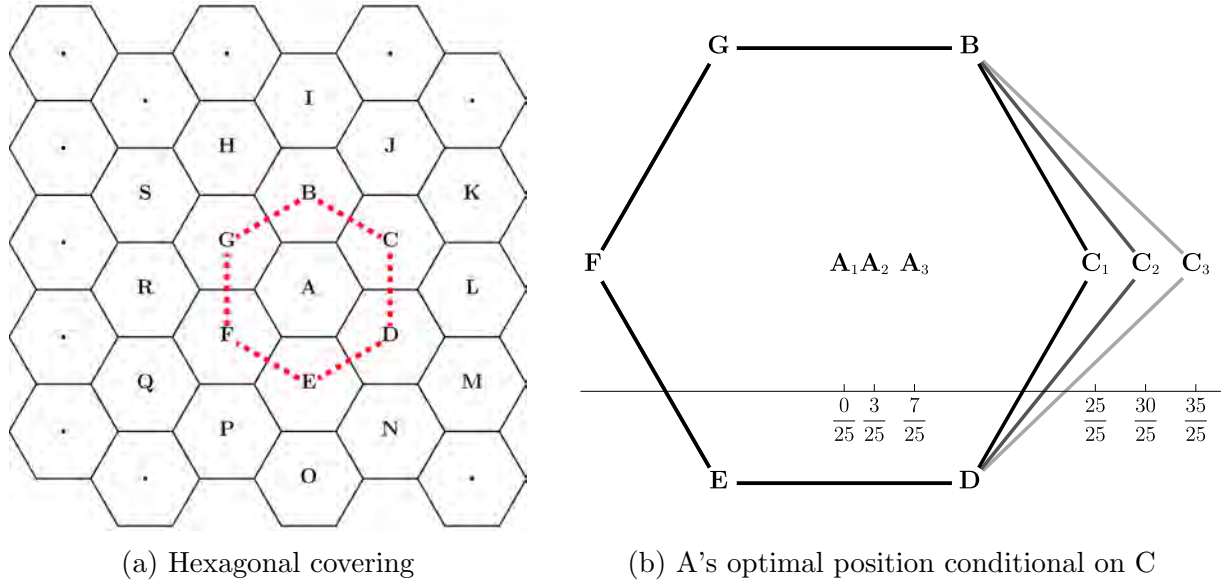


Figure 8: Löschian equilibrium hexagonal covering and individual incentives

location choice given the locations of all other firms except in a few special cases. This problem, maximizing a Voronoi region, is an open problem in computational geometry.<sup>38</sup> If we constrain firm  $A$  to select within a particular region<sup>39</sup> meeting minor technical assumptions, Dehne et al. (2005) proves that there exists a unique location that maximizes the area of  $A$ 's region. The same paper shows that if the convex hull of  $A$ 's neighbors happens to form a regular polygon with  $n > 4$  sides, then  $A$ 's optimal location is at the center of the polygon. Where the convex hull of opponents is an irregular polygon, we can use area formulas in that paper to find the optimal location for  $A$  numerically. As we show in Figure 8b, which looks at how  $A$ 's optimal position changes as one of its neighbors,  $C$ , is moved, a firm's optimal location appears to be *at the center* of the convex hull of its neighbors.<sup>40</sup> But it does not coincide with any of the myriad notions of polygon centrality that we have considered. This leaves us without an analytical solution to the problem beyond the very specific case in which the convex hull of  $A$ 's neighbors forms a regular polygon.

There is also an obvious heuristic argument against the presumption of equilibrium existence in two-dimensional spaces. Consider Figure 8a. Firm  $C$  has six neighbors, with the convex hull of their positions forming a regular hexagon, as is the case for Firm  $A$  in Figure 8b. Consideration of the latter figure suggest that if we assume that i)  $C$  is located optimally and ii) the locations of  $C$  and her other neighbors ( $B, J, K, L$ , and  $D$ ) are given, then we know  $A$ 's position. But the same argument is true for each of  $A$ 's neighbors:  $B$  along with her other neighbors ( $G, H, I, J$ , and  $C$ ) pins down  $A$ , as does  $G, F, E$ , and  $D$ , each with

<sup>38</sup>For examples in that literature, see Cheong et al. (2004, 2007) and Fekete and Meijer (2005).

<sup>39</sup>A neighborhood cell, which is a set of points such that  $A$  would have the same set of Voronoi neighbors locating at any one of those points.

<sup>40</sup>Uber drivers appear to have worked this out intuitively. There are several how-to videos on YouTube in which an experienced Uber driver shows the process of switching into the passenger app and then advises the audience that they will get a ride sooner if they move to the center of an unoccupied region, essentially mimicking our numerical analysis in Figure 8b.

their neighbors. And this is true for all firms: if a firm has  $k$  Voronoi neighbors, then its position must, in equilibrium, solve  $k$  equations.<sup>41</sup> As such, the existence of equilibrium requires a solution to a significantly overdetermined system. Such systems can have solutions, of course. Lösch’s hexagonal covering on the unbounded Euclidean plane with uniform customer density is a perfect example—each of the six equations pinning down  $A$  yields the same location. But we suspect that equilibrium existence is non-generic, a pleasant quirk of unboundedness and uniform customer density.<sup>42</sup> Unfortunately, we cannot speak to the rank of the system without an analytical solution to the problem of maximizing a Voronoi region.

Finally, we also note a few related models that combine competition on price with that on location or product characteristic: Capozza and Van Order (1978) offers a generalization of some earlier Löschian models. Salop (1979) adds an outside good/industry to the Hotelling model. Novshek (1980) considers alternatives to Nash equilibrium. Economides (1984, 1986a,b) experiment with the addition of reservation prices, the adjustment of the convexity of utility functions, and an expansion to two dimensions. There are also parallels with models of non-spatial product differentiation: Rosen (1974) offers a perfect competition model, with a continuum of firms, of hedonic pricing where products are differentiated and priced based on attributes—note that Rosen’s proof of equilibrium existence is exclusive to the one-dimensional case. Gabszewicz and Thisse (1979) models quality where customers have the same tastes but varying incomes.

In summary, there exists a rich theoretical literature on spatial competition, but using it to generate predictions in an applied setting is impeded by multiplicity of equilibria, potential equilibrium non-existence, and tractability issues. This motivates our agenda to pursue agent-based models.

### 3.2 Agent-based models with myopic best responses

An agent-based model (ABM) is a computational model for simulating the interactions of autonomous agents to assess their effects on the system. While this inductive approach to modeling comes largely from computer science, it is becoming increasingly prevalent in the social sciences. Axtell (2000) describes three distinct uses of ABM in the social sciences. Tesfatsion (2006) argues that it is a new, constructive approach to theory. Farmer and Foley (2009) argue for its use in macroeconomics given the complexity of macroeconomic systems, particularly in light of the most recent financial crisis. Hommes (2008) surveys their use in finance, particularly asset pricing. More related to our question, 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.<sup>43</sup>

---

<sup>41</sup>In the linear case, each firm has at most two Voronoi neighbors. And even in this *less overdetermined* system, we know that equilibrium existence is not guaranteed, as shown in Eaton and Lipsey (1975).

<sup>42</sup>We are not the first to make such a conjecture. In Eaton and Lipsey (1975), the authors conjecture that there exists no equilibrium on a disk with uniform density for  $n > 2$  firms. Shaked (1975) proves that conjecture for  $n = 3$ . Dasgupta and Maskin (1986) connects existence issues in location games to those in other games with discontinuities in payoff functions.

<sup>43</sup>See van Leeuwen and Lijesen (2016). NetLogo, software for ABM, even includes a Hotelling model in its model library (Ottino et al., 2009).



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.<sup>44</sup> To define agent myopia, consider a dynamic continuous-time<sup>45</sup> spatial game with payoff flows 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.

In motivating spatial competition with the ride-sharing application and designing a related experiment, we focus on a very specific environment without price competition.<sup>46</sup> Most applications of spatial competition, including those in our empirical work, also involve competition on prices. We can include prices in spatial ABM’s with the MBR dynamic in two ways. The simplest would be to extend the behavioral assumption to also cover prices—when an agent chooses a location and price, she takes her opponents’ locations and prices as given and maximizes her instantaneous payoff flow at the moment of the decision. Alternatively, we could construct hybrid models where agents take their opponents’ locations as fixed but anticipate prices resulting from a Nash equilibrium in a static price game, given the locations, immediately following their location choice.<sup>47</sup>

The MBR dynamic has important connections to Nash equilibrium. One interpretation of Nash equilibrium as a prediction is that it could result from an evolutive process in an environment in which agents know little about the structure of the game but best respond given their limited information and learn from outcomes.<sup>48</sup> In this sense, Nash equilibrium in a static game represents a fixed point under the MBR dynamic in a dynamic game. Brown (1951) computes equilibria in static games algorithmically by applying the MBR dynamic iteratively in fictitious dynamic play. Similarly, fixed points in the MBR dynamic on our dynamic games represent Nash equilibria in their static analogs.<sup>49</sup>

The consequences of MBR dynamics on a lattice with local interactions are studied in Blume (1993, 1995). The MBR dynamic is also common in evolutionary game theory,<sup>50</sup> although that literature focuses on population games with many players. Agent myopia

---

<sup>44</sup>Myopia is a dynamic-game analog of zero conjectural variance (ZCV), which historically has meant simply that Nash equilibrium is being applied as a solution concept to a static game. In that setting, when one checks a potential equilibrium by looking for a profitable unilateral deviation, the possibility of an opponent response is ruled out by the fact that the game is static. The first use of ZCV that we find is in Eaton and Lipsey (1972). The authors employ the term in critique of Löschian models in which, they argue, equilibrium is not well defined due to the absence of an explicit assumption on conjectural variance. While ZCV comes automatically in Nash equilibrium analysis on a static games, Lösch does not appeal to Nash equilibrium—his work was contemporaneous with that of Nash. Therefore, equilibrium in the Löschian model is not well defined without an explicit assumption on conjectural variance.

<sup>45</sup>This argument also applies to discrete-time games so long as at most one agent has a choice at any time.

<sup>46</sup>Since Uber and Lyft set prices centrally, drivers cannot compete on price.

<sup>47</sup>We can only do this if pure-strategy price equilibria exist and are computable. Relevant conditions for existence are given in Caplin and Nalebuff (1991).

<sup>48</sup>See Binmore (1987, 1988) and Gilboa and Matsui (1991).

<sup>49</sup>In this sense, our dynamic agent-based models may yield further insight on static games.

<sup>50</sup>See Sandholm (2010, Chapter 6) for best response dynamics in that setting.

and the MBR dynamic are tremendously powerful in that they allow us to model complex continuous-time games as sequences of static individual optimizations. Agent-based models in the social sciences commonly exploit behavioral assumptions to allow the representation of individual agents as automata. Because of the assumed relative simplicity of agent behavior, we can work in much richer environments than those used in game-theoretic analyses without losing tractability, and this may allow us to better compare mechanisms that generate spatial allocations and to evaluate relevant policies such as zoning and exclusive territories. Of course, the accuracy of our predictions is likely to depend on the validity of the behavioral assumption, which we support with an experiment in Section 4.

## 4 Experimental support for MBR in spatial games

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 payoffs at the moment of her decision, ignoring potential future opponent movements, we say that her choices satisfy the MBR assumption.

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.

There are two experimental literatures that are relevant to our work. One is a small literature that seeks to offer an experimental answer to the question of whether competition yields spatially efficient outcomes. 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, 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. 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 role of risk aversion as agents choose low-risk locations instead of the risk-neutral equilibrium predictions.

The second relevant literature is that on 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, 2017a,b) offer a model of, as well as experimental support for, 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.

Insofar as 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.

## 4.1 Experiment design

In the experiment, five participants (players, henceforth) play a location game on a  $21 \times 21$  grid. Players are given the opportunity, one-by-one, to move one square in any cardinal direction. Colocation is not allowed. We used the same initial allocation of players, shown in Figure 9, 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  $\ell_1$  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 player, 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, shown in Figure 10, 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 11, 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




Figure 9: Initial allocation of players on grid

many cases, we have data on not only the choices they make 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.<sup>51</sup> 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 1 unit of the current player's position, she has five squares to choose from. For example, in Figure 10, we see the squares Player 3 can choose to move to are highlighted in a dark shade of purple. 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

<sup>51</sup>The time limit was only to prevent a potential misanthropic participant from stalling the experiment indefinitely, not to put any time pressure on play. The limit was reached only a handful of times. In each of those occasions, the experiment leader then asked the player where she wanted to move and the player responded immediately.

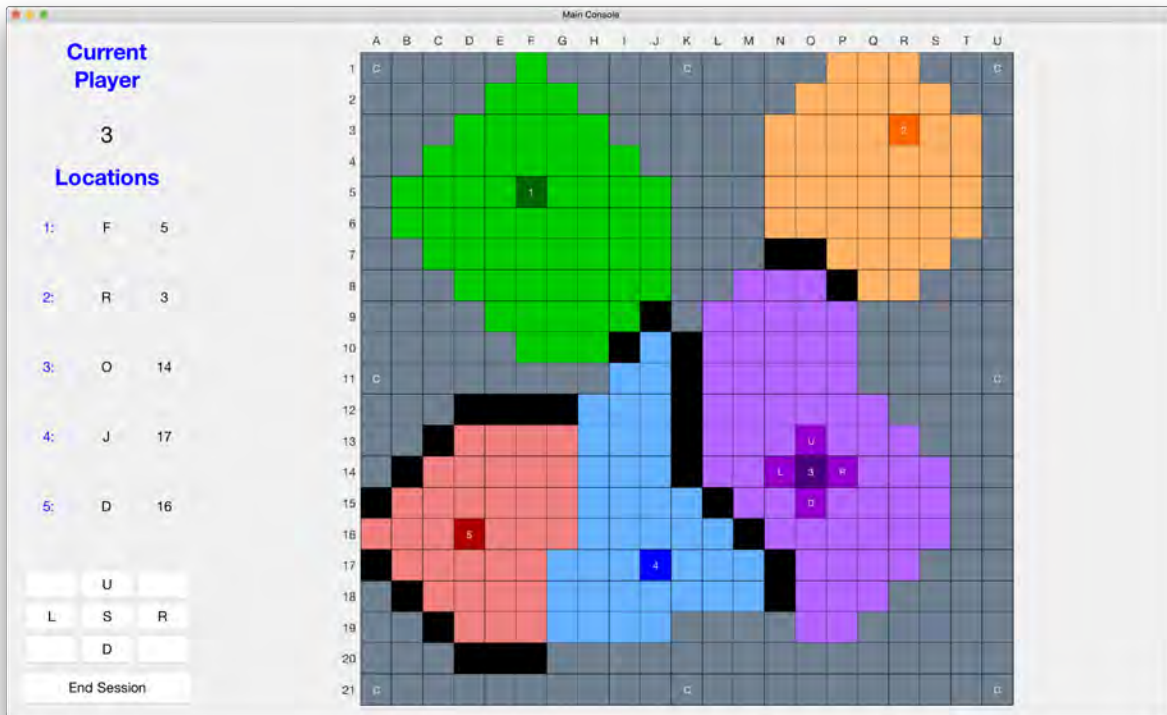


Figure 10: Main console software

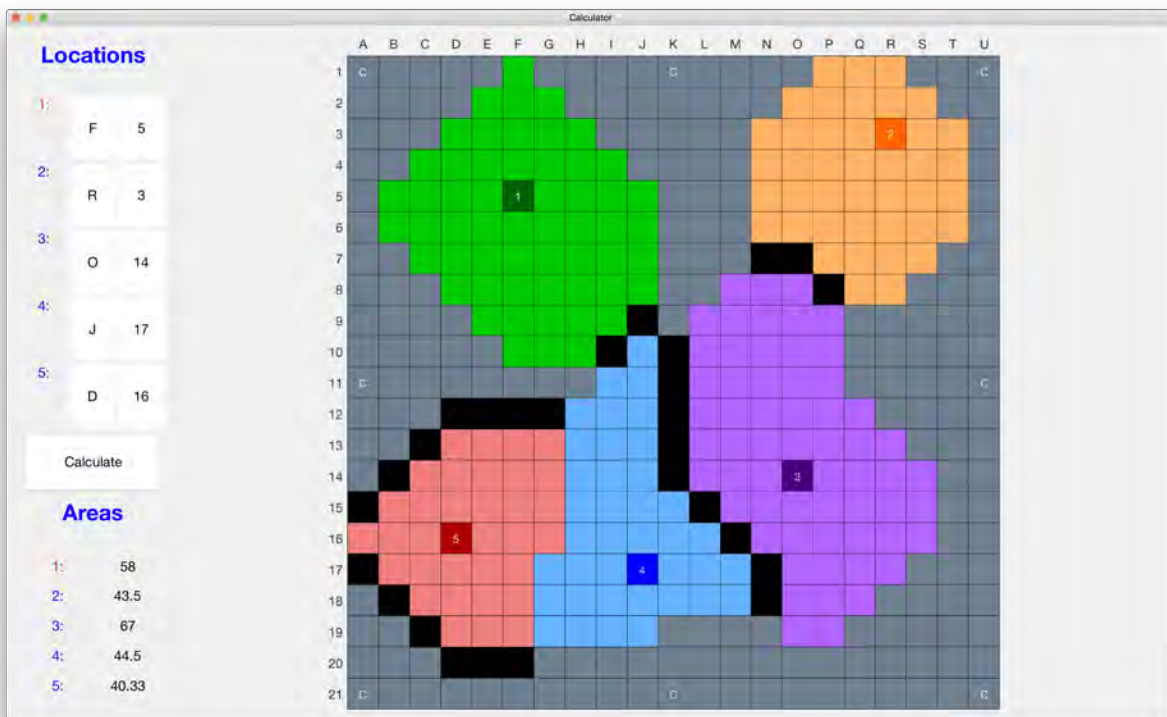


Figure 11: Calculator software

all move options, we rank the move options by their flow payments: the FP1 move is the move 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. A transcript of the instructions and a link to the video are included in Appendix C.

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 positioning of all players on the grid 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.<sup>52</sup>

## 4.2 Experiment results

There were a total of 2178 main console moves across the 18 sessions. Table 6 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.

Table 6: Distribution of move FP rankings

FP ranking	All moves		After 10	
	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

<sup>52</sup>Each iteration is scored by calculating the size of each player’s Voronoi region in the Voronoi mesh of the players over the grid. Let  $m_{i,t}$  denote player  $i$ ’s score for iteration  $t$ . It is calculated as the number of squares in the grid that are closest to player  $i$  in iteration  $t$ , i.e., the size of player  $i$ ’s Voronoi region. A square that is equidistant from multiple players is divided evenly amongst those players for scoring purposes. Player  $i$ ’s current session score in iteration  $\tau$  is  $M_{i,\tau} = \frac{1}{\tau} \cdot \frac{1}{21^2} \cdot \sum_{t=1}^{\tau} m_{i,t}$ . The session score calculates the average percentage of the  $21 \times 21$  grid controlled by player  $i$  over the  $\tau$  iterations. After the last iteration ( $t = T$ ), player  $i$  was paid  $\$180 \cdot M_{i,T}$ . While \$180 was the technical session prize pool, the actual amount paid out to players was close to \$100 given that the computer agents win a significant portion of the pool.

**Result 1.** *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 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.<sup>53</sup> 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—where their own location in the calculation is within one unit of their current location, we call this a move option. In fact, our second result of note is that a majority of calculations made were on move options. Table 7 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 7: Positioning of opponents in calculations

	<b>Freq.</b>	<b>%</b>
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 8 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.

<sup>53</sup>In some sense the modeler may not care how agents arrive at FP1 moves—if players are systematically making these choices, they are predictable regardless of how they reached them. On the other hand, such a coincidence of reasoning and predictable choice is less transportable to other models than a behavioral assumption that explains the coincidence.

Table 8: Distribution of player distances from their iteration position in calculations

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

**Result 3.** *In 92% of calculations, the calculating player was within one square of her position at the the time of calculation.*<sup>54</sup>

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.<sup>55</sup> 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 theses 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 9 for the full partitioning of non-FP1 moves.<sup>56</sup>

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 actively deviating from the FP1 move.<sup>57</sup> Table 10 shows how many players, and in what quantities, were responsible for the SHO moves. Approximately 24% of SHO moves were made by less than 7% percent of players.

<sup>54</sup>The frequency of zero-distance calculations is high partially because players’ calculators do not automatically update to the iteration positions after a move is made—many of these calculations should be attributed to players updating their calculators to a new iteration.

<sup>55</sup>This is to account for instances where players before the player in question choose to remain in place.

<sup>56</sup>We also saw a slight bias for non-FP1 moves towards the center—of the 863 non-FP1 moves, 362 were towards the center, 249 were towards the perimeter, and 252 remained in place.

<sup>56</sup>Although we made it clear in each session that the calculators did not automatically update for each new iteration (by design), we saw quite a few instances where players were doing MBR-like calculations but with the opponents in the position of a previous (outdated) iteration. This accounts for many of the 120 cases where some calculations were made but neither the choice nor the FP1 move was calculated.

<sup>57</sup>It is also possible that they simply forgot the results of some of their calculations. They were not given pen and paper, or any other means to record their calculations.



Table 9: Partitioning the 863 non-FP1 moves

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

Table 10: Number of SHO moves by player

# SHO Moves	# Players	% Players	% Total SHO
0	19	21	0
1	15	17	5
2	11	12	7
3	9	10	9
4	11	12	14
5	6	7	10
6	4	4	8
7	4	4	9
8	3	3	8
9	2	2	6
$\geq 10$	6	7	24

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.<sup>58</sup> Table 11 compares the score differences.

**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. Table 12 shows the impact of several behaviors on a player’s session score. Since there is

<sup>58</sup>SHO moves with very small flow payment differences could also potentially be attributed to errors in noticing the small differences.

Table 11: Score differences for non-FP1 moves

	SHO %	non-SHO %
Score Diff > 1	22	49
Score Diff ≤ 1	78	51
Score Diff < .5	33	14
Score Diff < .25	17	6

expected variation in the session scores between player numbers because of the unequal areas in the initial allocation, we calculate the mean session score for each player number and determine the difference from this mean for each player. We use this as our measure of player performance and the dependent variable in the regression analysis. Moreover, because turn order was randomly determined and sessions were limited by total amount of time playing the game, rather than total number of iterations, players had different numbers of turns. We control for this in the regressions.

Table 12: Difference (in thousandths) from mean score (by player number)

	(1) All	(2) All	(3) SHO ≥ 3	(4) All	(5) All	(6) All	(7) SHO ≥ 3
# Turns	-0.171 (0.104)	0.0218 (0.0871)	-0.0375 (0.137)	-0.0288 (0.0851)	-0.179 (0.103)	-0.00974 (0.0852)	-0.0730 (0.135)
# FP1	0.441** (0.158)				0.391* (0.158)		
# SHO		-0.221 (0.234)	-0.392 (0.409)			-0.375 (0.234)	-0.381 (0.400)
# Calcs.				0.0137* (0.00605)	0.0112 (0.00596)	0.0162* (0.00619)	0.0147 (0.00864)
<b>N</b>	90	90	45	90	90	90	45
<b>adj. R<sup>2</sup></b>	0.062	-0.013	-0.023	0.034	0.088	0.051	0.021

Standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$

In the first column, we find that players who chose more FP1 moves achieved higher session scores than those choosing fewer. In the second and third columns, we find no statistically significant relationship between the number of SHO moves chosen and a player's session score. The fourth column shows that players who made more calculations earned higher session scores. But when we include both the number of calculations made, and the number of FP1 moves, it is the latter that maintains its significance. In the third and seventh columns we run similar regressions but restrict our sample to the moves of players who make

at least three SHO moves.

**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.*

A potential critique of this experiment involves its external validity: perhaps undergraduate students participating in low-stakes experiments behave in accordance with a simplistic behavioral assumption, willfully ignoring inevitable opponent responses to their choices. Perhaps this could also apply to Uber driver given the relatively low stakes. But it may seem less intuitive that gas station and supermarket owners would behave so simplistically in their higher-stakes location choices. First, we think the regression results push against this. The more players behaved in this simplistic MBR manner, the better they did. If players are to learn from their performance, it would lead them to select FP1 moves more often, not less.

We also think that the MBR assumption may be realistic in an applied setting. For Uber drivers, given the reversibility of their decisions and the low stakes, sophisticated reasoning on the movements of other drivers is unlikely to be worthwhile. It is a much stronger assumption to make in games with irreversible decisions and large stakes. Businesses do extensive market research before deciding where to place new facilities, but this due diligence may still be analogous to MBR agents considering move options in our experiment unless businesses explicitly reason on potential future entry, exit, or firm relocation.

We cannot argue affirmatively that players behave in accordance with the MBR assumption. But we do fail to find significant evidence that they violate it. And our regression results suggest 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.<sup>59</sup>

Given the results of our experiment, we think it best to model agents in a spatial agent-based model as noisy MBR agents.<sup>60</sup> 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. This extends easily to agent-based models with spatial and price competition as we described in Section 3.2. Our finding that agent behavior is quite predictable in a complex dynamic spatial game might also be relevant to other complex dynamic games.

## 5 Conclusion

In this paper, we have shown empirical evidence that competition can yield spatially inefficient outcomes. In doing so, we have defined spatial inefficiency and developed algorithms to optimize spatial allocations. We have argued for the use of agent-based models with myopic best response dynamics for predictions and policy analysis in complex, dynamic spatial

---

<sup>59</sup>Incidentally, the average spatial inefficiency across iteration allocations in this experiment was only about 5%—FP1 moves tended to push players towards the optimal allocation in this setting.

<sup>60</sup>Models in which all agents myopically best respond are also useful both as a benchmark and as a way of computing Nash equilibria in analogous static games.

games. We have found justification for the MBR assumption in our experiment. Future work will involve the implementation and evaluation of the proposed agent-based models as well as analysis of whether ABM with MBR (or other) dynamics may be appropriate and useful in other environments relating to industrial organization.

## References

- Alaoui, L. and Penta, A. (2016). Endogenous depth of reasoning. *The Review of Economic Studies*, 83(4):1297–1333.
- Alaoui, L. and Penta, A. (2017a). Cost-benefit analysis in reasoning. Technical report.
- Alaoui, L. and Penta, A. (2017b). Reasoning about others' reasoning. Technical report.
- Axtell, R. (2000). Why agents? On the varied motivations for agent computing in the social sciences. Working Paper 17, Center on Social and Economic Dynamics, Brookings Institution, Washington, DC.
- Bain, J. S. (1951). Relation of profit rate to industry concentration: American manufacturing, 1936–1940. *The Quarterly Journal of Economics*, 65(3):293–324.
- Bain, J. S. (1956). Barriers to new competition, their character and consequences in manufacturing industries. Technical report.
- Bertrand, J. (1883). Review of theorie mathematique de la richesse sociale and of recherches sur les principes mathematiques de la theorie des richesses. *Journal des savants*, 67:499–508.
- Binmore, K. (1987). Modeling rational players: Part i. *Economics & Philosophy*, 3(2):179–214.
- Binmore, K. (1988). Modeling rational players: Part ii. *Economics & Philosophy*, 4(1):9–55.
- 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.
- Bollobas, B. and Stern, N. (1972). The optimal structure of market areas. *Journal of Economic Theory*, 4(2):174–179.
- Brown, G. W. (1951). Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1):374–376.
- Brown-Kruse, J., Cronshaw, M. B., and Schenk, D. J. (1993). Theory and experiments on spatial competition. *Economic Inquiry*, 31(1):139–165.
- Caplin, A. and Nalebuff, B. (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica: Journal of the Econometric Society*, pages 25–59.
- Capozza, D. R. and Van Order, R. (1978). A generalized model of spatial competition. *The American Economic Review*, 68(5):896–908.
- Chamberlin, E. (1933). *Theory of monopolistic competition*. Harvard University Press.

- Cheong, O., Efrat, A., and Har-Peled, S. (2007). Finding a guard that sees most and a shop that sells most. *Discrete & Computational Geometry*, 37(4):545–563.
- Cheong, O., Har-Peled, S., Linial, N., and Matousek, J. (2004). The one-round Voronoi game. *Discrete & Computational Geometry*, 31(1):125–138.
- Collins, R. and Sherstyuk, K. (2000). Spatial competition with three firms: an experimental study. *Economic Inquiry*, 38(1):73–94.
- Cournot, A.-A. (1838). *Recherches sur les principes mathématiques de la théorie des richesses*. Hachette.
- Crooks, A., Castle, C., and Batty, M. (2008). Key challenges in agent-based modelling for geo-spatial simulation. *Computers, Environment and Urban Systems*, 32(6):417–430.
- Crooks, A. T. and Heppenstall, A. J. (2012). Introduction to agent-based modelling. In *Agent-based models of geographical systems*, pages 85–105. Springer.
- Dasgupta, P. and Maskin, E. (1986). The existence of equilibrium in discontinuous economic games, I : Theory. *The Review of economic studies*, 53(1):1–26.
- d’Aspremont, C., Gabszewicz, J. J., and Thisse, J.-F. (1979). On Hotelling’s “Stability in competition”. *Econometrica: Journal of the Econometric Society*, 47(5):1145–1150.
- Dehne, F., Klein, R., and Seidel, R. (2005). Maximizing a Voronoi region: the convex case. *International Journal of Computational Geometry & Applications*, 15(05):463–475.
- Drezner, Z. and Hamacher, H. W. (2001). *Facility location: applications and theory*. Springer Science & Business Media.
- Eaton, B. C. and Lipsey, R. G. (1972). Unsuspected perversities in the theory of location. Technical Report 88, Queen’s Economics Department.
- Eaton, B. C. and Lipsey, R. G. (1975). The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *The Review of Economic Studies*, 42(1):27–49.
- Eaton, B. C. and Lipsey, R. G. (1976). The non-uniqueness of equilibrium in the loschian location model. *The American Economic Review*, 66(1):77–93.
- Eaton, B. C. and Lipsey, R. G. (1978). Freedom of entry and the existence of pure profit. *The Economic Journal*, 88(351):455–469.
- Economides, N. (1984). The principle of minimum differentiation revisited. *European Economic Review*, 24(3):345–368.
- Economides, N. (1986a). Minimal and maximal product differentiation in hotelling’s duopoly. *Economics Letters*, 21(1):67–71.

- Economides, N. (1986b). Nash equilibrium in duopoly with products defined by two characteristics. *The RAND Journal of Economics*, pages 431–439.
- Edgeworth, F. Y. (1897). La teoria pura del monopolio. *Giornale degli economisti*, pages 13–31.
- ERS-USDA (2017). Food access research atlas. <https://www.ers.usda.gov/data-products/food-access-research-atlas/>.
- Esri, TomTom North America, Inc., and US Census Bureau (2017). USA census tract boundaries (shapefile). Retrieved from <https://www.arcgis.com/home/item.html?id=ca1316dba1b442d99cb76bc2436b9fdb>.
- Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256):685–686.
- Fekete, S. P. and Meijer, H. (2005). The one-round Voronoi game replayed. *Computational Geometry*, 30(2):81–94.
- Fekete, S. P., Mitchell, J. S., and Beurer, K. (2005). On the continuous fermat-weber problem. *Operations Research*, 53(1):61–76.
- Gabszewicz, J. J. and Thisse, J.-F. (1979). Price competition, quality and income disparities. *Journal of Economic Theory*, 20(3):340–359.
- Gabszewicz, J. J. and Thisse, J.-F. (1992). Location. *Handbook of game theory with economic applications*, 1:281–304.
- Gilboa, I. and Matsui, A. (1991). Social stability and equilibrium. *Econometrica: Journal of the Econometric Society*, pages 859–867.
- Google (2017). Google maps. Retrieved, for example, from <https://www.google.com/maps/search/gas+stations/@44.265134,-106.5989458,11.25z>.
- Graitson, D. (1982). Spatial competition a la hotelling: A selective survey. *The Journal of Industrial Economics*, pages 11–25.
- Halpern, J. Y. and Pass, R. (2015). Algorithmic rationality: Game theory with costly computation. *Journal of Economic Theory*, 156:246–268.
- Hommes, C. (2008). Interacting agents in finance. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, Basingstoke.
- 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.
- Knoblauch, V. (2002). An easy proof that a square lattice is an equilibrium for spatial competition in the plane. *Journal of Urban Economics*, 51(1):46–53.

- Kruse, J. B. and Schenk, D. J. (2000). Location, cooperation and communication: An experimental examination. *International Journal of Industrial Organization*, 18(1):59–80.
- Lloyd, S. (1982). Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137.
- Lösch, A. (1954). *Economics of location*. Yale University Press.
- Mason, E. S. (1939). Price and production policies of large-scale enterprise. *The American Economic Review*, 29(1):61–74.
- Mason, E. S. (1948). The current status of the monopoly problem in the united states. *Harv. L. Rev.*, 62:1265.
- Medicare (2017). Hospital compare: Hospital general information. Retrieved from <https://data.medicare.gov/Hospital-Compare/Hospital-General-Information/xubh-q36u>.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326.
- Novshek, W. (1980). Equilibrium in simple spatial (or differentiated product) models. *Journal of Economic Theory*, 22(2):313–326.
- Oak Ridge National Laboratory (2017). Hospitals. Data retrieved from Homeland Infrastructure Foundation-Level Data (HIFLD), <https://hifld-geoplatform.opendata.arcgis.com/datasets/hospitals>.
- Okabe, A. and Aoyagi, M. (1991). Existence of equilibrium configurations of competitive firms on an infinite two-dimensional space. *Journal of Urban Economics*, 29(3):349–370.
- OpenStreetMap contributors (2017). <https://www.openstreetmap.org>. Feature shapefiles retrieved from <http://osm2shp.ru>.
- Ottino, B., Stonedahl, F., and Wilensky, U. (2009). Netlogo hotellings law model. *Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL*. Retrieved from <http://ccl.northwestern.edu/netlogo/models/HotellingsLaw>.
- Penta, A. and Zuazo-Garin, P. (2017). Rationalizability and observability. Technical report.
- Rampal, J. (2017). Limited foresight equilibrium. Technical report.
- Rosen, S. (1974). Hedonic prices and implicit markets: product differentiation in pure competition. *Journal of political economy*, 82(1):34–55.
- Salop, S. and Stiglitz, J. (1977). Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies*, pages 493–510.
- Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, pages 141–156.



- Sandholm, W. H. (2010). *Population games and evolutionary dynamics*. MIT press.
- Shaked, A. (1975). Non-existence of equilibrium for the two-dimensional three-firms location problem. *The Review of Economic Studies*, 42(1):51–56.
- Shubik, M. (1959). *Strategy and market structure*. John Wiley.
- Stigler, G. J. et al. (1983). The organization of industry. *University of Chicago Press Economics Books*.
- TechniGraphics, Inc. (2010). Firestations. Data retrieved from Homeland Infrastructure Foundation-Level Data (HIFLD), <https://hifld-geoplatform.opendata.arcgis.com/datasets/fire-stations>.
- Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. *Handbook of computational economics*, 2:831–880.
- U.S. Census Bureau (2016). 2016 tiger/line shapefiles (machine-readable data files). Retrieved from <https://www.census.gov/geo/maps-data/data/tiger-line.html>.
- van Leeuwen, E. and Lijesen, M. (2016). Agents playing hotellings game: an agent-based approach to a game theoretic model. *The Annals of Regional Science*, 57(2-3):393–411.
- Vickrey, W. S. (1964). *Microstatics*. Harcourt, Brace & World.
- Von Neumann, J. and Morgenstern, O. (1944). *Theory of games and economic behavior*.
- Walker, R. E., Keane, C. R., and Burke, J. G. (2010). Disparities and access to healthy food in the united states: A review of food deserts literature. *Health & place*, 16(5):876–884.
- Walras, L. (1883). *Théorie mathématique de la richesse sociale*. Guillaumin.
- Weber, A. (1929). *Theory of the Location of Industries*. University of Chicago Press.

# A Additional figures from 2.2




Figure 3a: Gas stations on I90 between Seattle and Spokane




Figure 3b: Gas stations on I90 between Couer d'Alene and Missoula




Figure 3c: Gas stations on I90 between Missoula and Billings




Figure 3d: Gas stations on I90 between Billings and Rapid City




Figure 3e: Gas stations on I90 between Rapid City and Sioux Falls




Figure 3f: Gas stations on I90 between Sioux Falls and La Crosse





Figure 3g: Gas stations on I90 between La Crosse and Madison

## B Additional figures from 2.3


In this appendix, we provide a few extra figures for comparison with those presented in the body of the paper. First, we continue Figure 6 to show allocations of hospitals and fire stations in Chicago MSA. Then, in Figure 9, we show the supermarket allocations in Chicago city as an additional comparison to those in Chicago MSA. Figures for other regions and features are available upon request.




(g) Hospitals and pop. density

Figure 6: Chicago MSA (continued)







(h) Distance to hospital in  $s^{\text{act}}$



(i) Distance to hospital in  $s^{\text{opt}}$

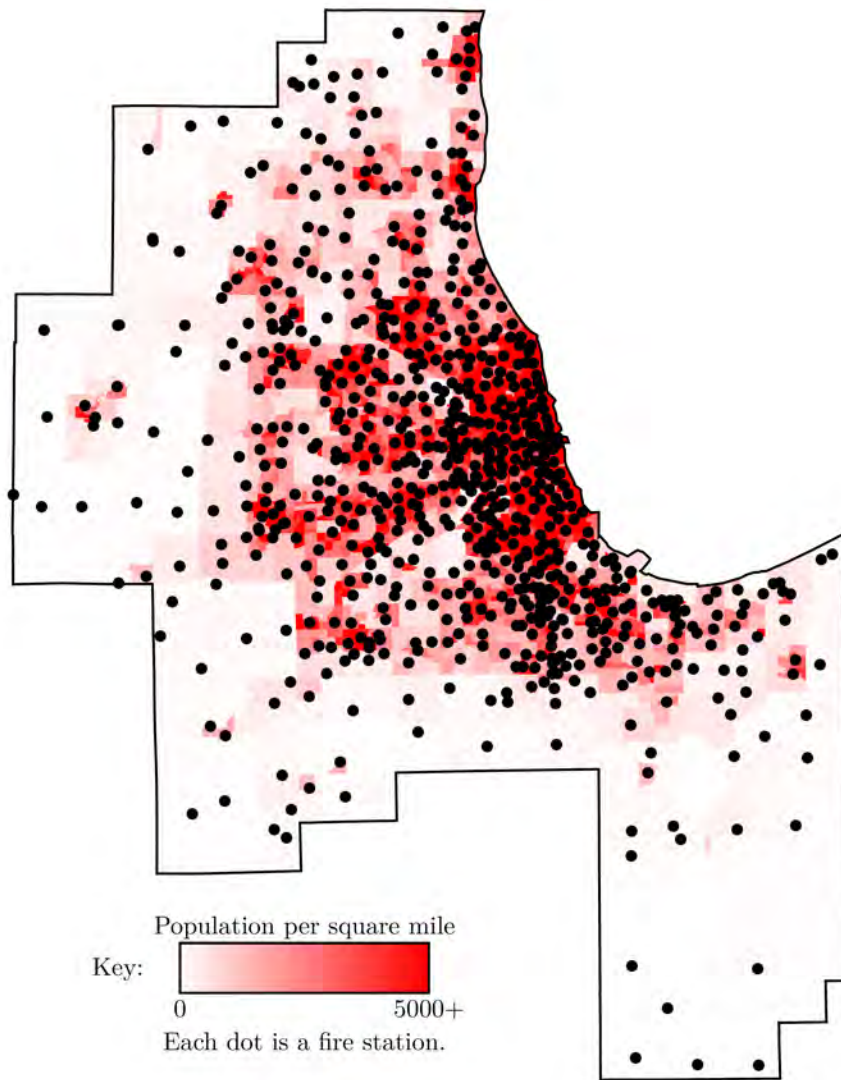


(j) Hospital scarcity in  $s^{\text{act}}$




(k) Hospital scarcity in  $s^{\text{opt}}$

Figure 6: Chicago MSA (continued)




(1) Fire stations and pop. density


Figure 6: Chicago MSA (continued)



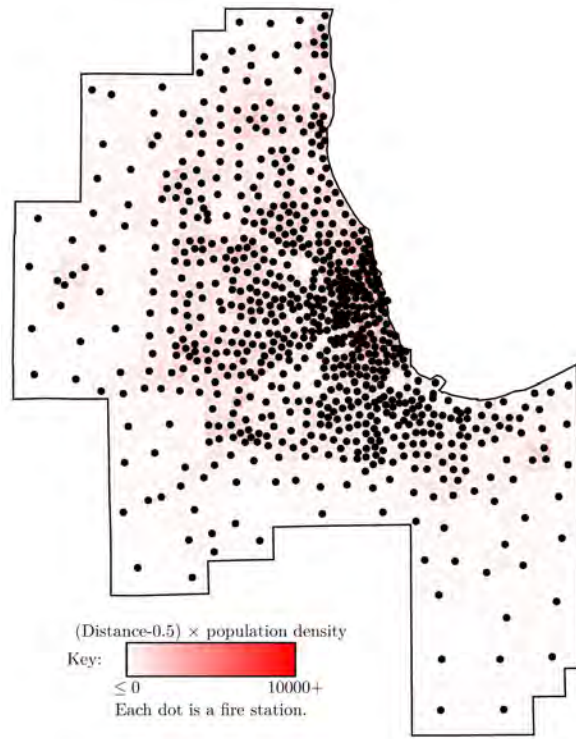
(m) Distance to fire station in  $s^{\text{act}}$



(n) Distance to fire station in  $s^{\text{opt}}$




(o) Fire station scarcity in  $s^{\text{act}}$




(p) Fire station scarcity in  $s^{\text{opt}}$

Figure 6: Chicago MSA (continued)






(a) Census tracts and pop. density




(b) Supermarkets and pop. density


Figure 12: Chicago city




(c) Distance to supermarket in  $s^{\text{act}}$



(d) Distance to supermarket in  $s^{\text{opt}}$



(e) Supermarket scarcity in  $s^{\text{act}}$



(f) Supermarket scarcity in  $s^{\text{opt}}$

Figure 12: Chicago city (continued)

## C Experiment instructions

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 were then given the chance to ask questions before the game began. See the instructional video at <http://youtu.be/7hcN24RFI3M>. The following is a transcript of the instructional video:

Welcome to the BRITE Lab and thank you in advance for participating in our experiment on spatial competition. You are about to play an experimental game of spatial competition. The duration of the experiment will be about 90 minutes. Please do not talk to other participants during the experiment or use the computers in ways other than those described here. After participating in this experiment, please do not discuss it with others who may also participate in the future. After watching these video instructions, you may ask questions.

This is the game board that you will see projected. It may look slightly different, depending on the computer. At the top left, we see that it is currently player 3's turn. You have already been assigned player numbers. The turn order is randomly determined after each turn, it is as if the next player's turn is decided by rolling a five-sided die. This can result in you having multiple turns in a row or going for long stretches of time without having a turn. Below the turn indicator, you can see a list of the indexed locations of each of the five human players. These indices will be important for operating the calculator on your computer. Reading the top line, we see that player 1 is located at grid coordinate  $F5$ .  $F$  represents the column and 5 represents the row. Sure enough, we can see player 1 at coordinate  $F5$ .

At the bottom-left, you can see the controls that the experiment leader will use to move you and the other human players around the board. When it is your turn, you will select to move either up, down, left, or right. You may also stay in your current location. Once you have made your decision, you will communicate it to the experiment leader who will then input the choice using these buttons. Suppose player 3 chooses to move up and communicates this to the leader. Now the leader hits the  $U$  button for up, and we see that the gameboard is redrawn with player 3 one row higher than before. We also see that player 4 has been randomly selected for the next turn.

So far, we have not discussed how you should decide where to move. We will show you precisely how your cash payment is determined in a moment, but let's first focus on the grid displayed on the game board. The game is played on a 21 by 21 grid. There are actually 13 players in this game, including 1-2-3-4-5 human players and 1-2-3-4-5-6-7-8 computer players surrounding the grid. The computer players will remain in their current positions throughout the game.

Each human player is given a specific color. Player 4 is blue. Her current location is marked in dark blue. Then, all squares in the grid that are closer to player 4 than any other player (including the computer players) are marked with light

blue. For instance, look at grid coordinate  $J10$ , shaded light blue. It is 1-2-3-4-5-6-7 squares from player 4's current location. This is less than the distance to any other player. Player 3 is 1-2-3-4-5-6-7-8 squares away. Player 1 is nine squares away: 1-2-3-4-5-6-7-8-9. All of the light-blue shaded squares are currently in player 4's area.

Some squares are equally far from two or more players. For instance,  $I10$  is 1-2-3-4-5-6-7-8 squares from player 4. It is also 1-2-3-4-5-6-7-8 from player 1. Squares equally far from multiple players (of which at least one of whom is a human player) are shaded in black. For scoring purposes, such a square is split evenly among the players who are equidistant from it. So  $I10$  contributes half of a unit to player 4's area, and half of a unit to player 1's area. Sometimes squares may also be split between 3 or more players. Finally, squares shaded grey are closer to the computer players than to any human player.

The basic objective of the game is to have as large of an area as possible since your payment will depend on the average size of your area over the duration of the experiment. Therefore, to maximize your cash payment, it is in your interests to move strategically to have as large of an average area over the course of the experiment as possible.

When it is your turn, you make take up to two minutes to decide your next move. Note that you are not able to move into a square currently occupied by another player (including human and computer players). If player 5 was in square  $I17$ , for instance, player 4 would not be allowed to move left. Nor would player 5 be allowed to move right into  $J17$ . Now lets look at the calculator program that is on your computer screen.

We provide you with a calculator to facilitate your decision-making. The calculator can be used to calculate (and show) the areas for any possible allocation of the five human players. Each of you has the calculator in front of you on your computer.

At the top left of the calculator, you may input player locations. Recall that the indices of the current locations are provided on the gameboard that's projected for all to see. Once you have inputted locations, hit calculate to redraw the simulated gameboard for that allocation of players. As an example, let's try moving player 4 down one position from her current location, from  $J17$  to  $J18$ . If we wanted to, we could move multiple players at once by editing all of the locations. You must hit calculate after editing the locations to regenerate the gameboard.

At the bottom left of the calculator, you can see the size of the areas resulting from the locations that you have inputted. These areas include the appropriate portions of squares that are equally far from multiple players. On the grid itself, you can see what the gameboard would look like if players were in the locations that you inputted.

Please note that the calculator does not update itself to the current locations when players move. If you want to reset your calculator to the current locations,

you have to do it manually by inputting the location indices listed for each player on the game board. You may use the calculator at any point in the experiment, whether or not it is your turn. You may use it as much or as little as you want. We do collect data from the calculators. But nothing you do with the calculator will directly affect your payment. Only the selected moves by all participants affect payoffs.

Now let's look at how your cash payment is determined.

To calculate payments, call each player's turn one iteration. The size of your area is calculated in each iteration. Then, we calculate your average area size, averaging over all iterations. Finally, we multiply your average by a number,  $X$ , to calculate your cash payment, rounding up to the nearest dollar.  $X$  is the same for all players.

The multiplier,  $X$ , has been selected to target total payments to human players in the game at around \$100. This implies an anticipated average participant payment of \$20. Your actual payment depends on the choices of all participants, so we can make no minimum payment guarantees. Also, your final payment is not necessarily a good measure of your performance in the experiment, as some players start with more favorable positions.

The experiment leader will terminate the game shortly before 90 minutes has elapsed from the experiment start time. We will have as many turns as time allows. This is the end of the instructional video. You may now ask questions to the experiment leader.