

Agents playing Hotelling's game: an agent-based approach to a game theoretic model

Eveline van Leeuwen¹ · Mark Lijesen¹

Received: 11 February 2015 / Accepted: 24 September 2015 / Published online: 19 October 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract This paper combines game theory and agent-based modelling, two powerful tools that economists use to understand the behavior of economic agents. We construct an agent-based version of Hotelling's two-stage game of spatial competition and explore the possibilities of creating synergies between the two approaches. Game theoretic insights into strategic behavior and equilibrium states can provide useful theoretic underpinnings for agent-based approaches in regional science. By combining the two, we can model micro-based social order as it emerges out of local interactions. The use of agent-based modelling in the context of a multistage game is new and hence provides a valuable contribution to both streams of the literature. We show that combining the two approaches is feasible, also in the context of a more complex two-stage game. The model correctly reproduces the analytical results and also allows for more complex situations. As an example, we show the effect of different levels of consumer tastes for variety in Main Street. The reconstruction of Hotelling's model of spatial competition opens up a wide variety of possibilities for further extensions that can lead to a better understanding of the variations we observe in reality. For some extensions, the use of a single-stage model would probably be more feasible though.

JEL Classification C63 · C72 · D21 · L13

1 Introduction

Hotelling's metaphor of spatial competition on Main Street ([Hotelling 1929](#)) is widely accepted as one of the most important models in understanding strategic product

✉ Eveline van Leeuwen
e.s.van.leeuwen@vu.nl

¹ Department of Spatial Economics, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

differentiation. The known standard outcomes allow us to evaluate the results of the basic model. A vast literature has evolved, discussing extensions of the model with respect to the nature of transport costs (Economides 1986), distributions of consumers (or preferences) (Tabuchi and Thisse 1995; Meagher and Zauner 2005), dimensions (Irmen and Thisse 1998; Chen and Riordan 2007), and so on.

Although being widely used, game theory has to deal with several challenges, both from a theoretic perspective and from the perspective of translating its outcomes to the real world. Many economic games assume all actors to be fully rational and perfectly informed, which may be practical in finding a solution, but is unlikely to hold in real life. Although game theory as such allows for relaxing these assumptions, doing so often leads to intractable results or the nonexistence of equilibria (Halpern and Pass 2015). Likewise, allowing for alternative underlying distributions (e.g., of consumers over space in the Hotelling model) might yield a model without equilibria in pure strategies (Caplin and Nalebuff 1991). Studying states and processes evolving from the behavior of non-rational agents is at the core of agent-based modelling.

The number of researchers involved in modelling agents' behavior has increased significantly over the last years (Crooks and Heppenstall, 2012). The so-called agent-based models are developed to better understand the theories of political identity and stability (Lustick 2002); economic processes as dynamic systems of interacting agents (Tesfatsion 2006); company size and growth rate distributions (citealtA99); burglary activities (Malleon et al. 2010); size-frequency distributions for traffic jams (Nagel and Rasmussen 1994); spatial patterns of unemployment (Topa 2001); size distributions of cities (Mansury and Gulyás 2007); land system dynamics and well-being (Robinson et al. 2012); and so on. However, until recently, agent-based models are not well established in spatial economic research mainly due to the lack of micro-foundations underlying individual motivation and behavior. In general, it is not easy to translate micro-foundations into agent behavior in complex (spatial) situations.

Our contribution to the literature is to combine game theory and agent-based modelling in such a way that both tools strengthen each other. It is striking to see that the strengths of agent-based modelling correspond with the weaknesses of game theory and vice versa. We develop an agent-based version of Hotelling's model of spatial competition to explore the potential for synergy of using these tools simultaneously.

Finding the optimal location of a shop essentially depends on the trade-off between market power (shops far away from each other) and market share (shops close to the center of the market). The players of this game are thus the shops and their payoff of the profit based on the number of customers they attract at a certain price. In an earlier approach, Ottino et al. (2009) developed an agent-based model that uses the concepts of Hotelling in both a one- and a two-dimensional world. In this model, the shops change their location to gain market share and then change their price to see whether they can also improve their profit at that specific location. By construction, this puts the emphasis on the market-stealing effect, which, as can be seen in the model, always results in shops ending up being located close to each other with very low revenues. Furthermore, the shops do not interact with each other; after they have moved simultaneously, they respond to the new situation. Since the shops do not search for best responses (to the other's shop strategy), this agent-based version lacks the strategic interaction that characterizes game theoretic models.

Our agent-based model goes a step further than the model of [Ottino et al. \(2009\)](#), as it aims to capture the full game theoretic nature of Hotelling's model. Since this model is a multistage game, an important challenge is to find a computational alternative for the analytical tools used. Alternatives for backward induction have to our knowledge not been discussed in the literature on economic games. The use of computational models in multistage games is well developed in a broader context however, the most obvious example being chess computers (or software). Chess computers calculate the payoff of each possible sequence of moves to select the move with the highest expected payoff. We use the same basic mechanism. First of all, we incorporate strategic behavior by asking the shops to respond to each other's moves in order to maximize their profit. Furthermore, the model maintains the two-stage nature of Hotelling's model in which the interaction between the location and the price (i.e., the two-staged nature) is essential.¹

The rest of this article is organized as follows: Sect. 2 briefly discusses game theory and agent-based models. Section 3 outlines Hotelling's model of spatial competition, and Sect. 4 shows how it can be translated into an agent-based model. The model validation is shown in Sect. 5, followed by an application of the model to consumer-specific heterogeneity. Finally, in Sect. 7 conclusions are drawn.

2 Game theory and agent-based modelling: never the twain shall meet?

To many economists, game theory and agent-based modelling may seem to be two completely different things, explaining the lack of economic literature on the joint application of the two. The differences between the two approaches are indeed apparent, but we claim that they are partly superficial and partly result in the two approaches being complementary, creating ample room for synergy. In this section, we systematically explore the strengths and weaknesses of both approaches, in order to identify the scope for synergy between the two.

2.1 Game theory

Game theory is defined by [Shoham and Leyton-Brown \(2009\)](#) as '...the mathematical study of interaction among independent, self-interested agents' (op. cit, p. 47). The basic principle behind economic applications of game theory is that two or more rational agents try to maximize their target function, while being aware that their actions impact each other's outcomes. This awareness leads to strategic behavior, where the behavior of each agent takes into account the anticipated response of the other agent(s). Based on the target the agent is trying to achieve, we can determine the agent's best response to any strategy of the other agent(s).

In mathematical game theory models, the best response is expressed in an equation where the optimal value of one player's decision value is a function of the decision

¹ [Ottino et al. \(2009\)](#) also suggest that making the shops 'more intelligent' by 'considering their moving and pricing options in conjunction' would be an important extension of the model.

value of the other(s). If two firms compete in prices for instance, the best response in prices of firm 1 is expressed as an equation (referred to as the best response function) where the price of firm 1 is a function of the price of firm 2. Similarly, firm 2 has a best response function where the price of firm 2 is a function of the price of firm 1.

If each agent plays its best response to the other agents' best responses, the so-called Nash equilibrium will be reached, that is an equilibrium where no agent feels the desire to unilaterally change its behavior. In a mathematical model, this boils down to solving the set of equations that reflect the best responses of the players. Since the number of variables as well as the number of equations is equal to the number of players, such a system can be solved by substitution and the Nash equilibrium will be found, provided that it exists.² These Nash equilibria are the core concept of most economic applications of game theory, including Hotelling's model of spatial competition.

2.2 (Bounded) Rationality?

Mainstream economists tend to agree that economics is about agents behaving rationally in order to maximize their individual utility. It often goes along with fully rational agents that use perfect information to make their choices. These assumptions are rigorous but important for the purpose of analytical convenience (Becker and Mulligan 1997).

Samuelson already stated in 1937 that it is doubtful whether we can learn much from considering the 'homo economicus' with its unchanged tastes, seeking for maximum functional consumption, in a perfect world where all things are certain, clear, and synchronized. Most important critiques against the 'homo economicus' concern the descriptive inaccuracy of the underlying assumptions of rationality and the reliance on context-free individual utility maximization. Clearly, one can argue that at the individual level, the 'homo economicus' does not exist. However, when looking at a more aggregated level, these assumptions might cause less distortions. What is necessary, is a better understanding of how individual behavior sums up to meso- and macroprocesses.

Over the years, many researchers have focussed on implementing alternative assumptions and approaches to individual behavior within economics. Important concepts that are used are bounded rationality, judgment of probabilities, prospect theory, interdependent utility, adaptive and relative preferences, and intertemporal choice. There are now a variety of approaches in economics that attempt, each in its own way, to move beyond the rational agent. These include experimental economics, evolutionary economics, and certain computational approaches (Axtell 2000).

In this paper, we do not drop the assumptions of rationality and perfect information, but rather develop a flexible version of the model of spatial competition that can both adopt the assumptions as they are commonly used and relax them to evaluate their impact.

² At least one Nash equilibrium always exists, but it might be a mixed equilibrium. In this paper, the existence of an equilibrium refers Nash equilibria in pure strategies.

2.3 Agent-based modelling

Agent-based models are usually based on a set of autonomous agents, capable of interacting with each other, as well as with the environment according to rules of behavior. An important advantage of agent-based models is that they allow for 'real world' characteristics, such as limited rationality and heterogeneous agents. Another advantage is that any model run provides an entire dynamical history of the process under study, since the model is 'solved' merely by executing it (Axtell 2000; Webster and Wu 2001).

An agent-based model is a computerized simulation of a number of decision-makers, the so-called agents, that interact through prescribed rules. These models do not rely on the assumption that the economy will move toward a predetermined equilibrium state, as for example game theoretic models do. Instead, at any given time, each agent acts according to its current situation, the state of the world around it, and the rules governing its behavior (Farmer and Foley 2009).

Thus, agent-based models consist of three main components:

- Individual agents with individual characteristics.
- The environment or spatial context the agents are located in.
- Rules of behavior and interaction between agents and between an agent and its perceived environment.

Although those three components are part of the basic design of all agent-based models, the implementation can be more or less complex. A well-known example of a simple but very powerful approach is Schelling's model of segregation, in which two types of agents move (randomly) through space, looking for a place to live where the number of similar neighbors is above a certain a priori set threshold value. Depending on the threshold value set and the density of agents on the grid (i.e., the number of empty locations), different spatial patterns of segregation appear. Examples of far more complex models are the SIMpop model in which cities are seen as agents and their development patterns in Europe between 1300 and 2000 are modelled (Bretagnolle and Pumain 2010).

The above example implies that agents are not necessary persons, and they can also be households, firms, or other (hypothetical) decision-making units, such as cities. Thus, what qualifies something as an agent? First of all, it needs to be autonomous, making its own decision(s) and thus being goal-directed. Furthermore, in order to be able to interact with other agents and with the environment it is in, an agent needs to be able to sense and to communicate; in other words, it needs to be able to retrieve and to send information. Additionally, certain agents are able to learn and to adapt their behavior or even their personal preferences (O'Sullivan and Haklay 2000). In doing so, agent-based models can handle a far wider range of nonlinear behavior than conventional equilibrium models. Policy-makers can use them, for example, as a living laboratory to simulate an artificial economy under different policy scenarios and quantitatively explore their consequences (Farmer and Foley 2009).

2.4 Agent-based modelling of games

A well-known game theoretic model in regional science is Hotelling's model of spatial competition. It is the so-called two-stage model, where agents make subsequent choices; they choose location first, and once located, they set prices. In analytical game theoretic models, multistage games are solved through backward induction; the last stage is solved first, then the next-to-last, and so on.

The concept of backward induction is directly related to the assumptions of perfect rationality and perfect information. When applying backward induction, one implicitly assumes that the actors fully understand the impact of their first-stage actions on the equilibrium outcome of subsequent stages. Like a chess player predicts his opponent's best responses several turns (stages) ahead, actors in a multistage game are aware of the impact of their first move on subsequent moves by other actors. Although it is hard to imagine that these assumptions hold in complex situations, the concept of backward induction is very convenient for analytical purposes and conceivable for simpler settings.

In agent-based models, agents can be backward *looking* when they base their expectations on past behavior. This past behavior is often related to rewards (i.e., realized profits) or punishments and is fundamentally different from backward induction. In essence, they look forward, by rehearsing the lessons of the past, and by doing so the final outcome (the future) is based on repetition, instead of prediction (Macy and Flache 2009). When agents are forward looking, they aim to move toward a (pre-defined) goal. Their strategic behavior is generally based upon expected profits or rewards, implying the ability of agents to (rationally) predict outcomes. Often, the expectation explains the action, not the actual consequences, which as such might not even occur (Scott 2000). In the game theoretic approach of reasoning backward (in time), it means that a decision-maker anticipates on what the other player is expected to do. As such, those agents can be referred to as forward looking as well. However, in particular in the case of multistage games, like the Hotelling game where both prices and locations are important, this cannot be solved by an individual agent due to its complexity.

We think that game theoretic insights into strategic behavior and equilibrium states can be a very useful theoretic underpinning for agent-based approaches in regional science. In addition, agent-based models offer the possibility to relax the restrictive behavioral assumptions of (classical) game theory and to add dynamics and heterogeneous agents into the framework (Macy and Flache 2009). By combining the two, we can model a micro-based social order as it emerges out of local interactions.

3 Hotelling's model of spatial competition

Hotelling's model of spatial competition is one of the many game theoretic applications in economics. In his original paper, Hotelling (1929) used the analogy of two stores locating on Main Street to analyze the phenomenon of strategic product differentiation. However elegant the analogy, Hotelling's original model does not result in a Nash equilibrium in pure strategies. Both firms locate in the middle of Main Street with

zero profits. It is straightforward to see (although addressed 50 years later) that either firm could increase its profits by moving out of the middle. D'Aspremont et al. (1979) formally prove this nonexistence problem and propose an alternative specification, with quadratic transportation costs. Salop (1979) uses a slightly different approach by assuming a reservation price, above which consumers will buy an outside good. Salop's model is fundamentally different from Hotelling's as it assumes equidistance between the shops and hence imposes locations rather than analyzing them.³ Economides (1984) explores the concept of a reservation price in the context of the original Hotelling model and finds that, for a sufficiently low reservation price, firms have an incentive to move out of the center to obtain a local monopoly position. Hinloopen and Van Marrewijk (1999) conclude that an equilibrium in pure strategies does not exist if the reservation price is 'high' (one can think of Hotelling's original model as a model with an infinite reservation price); that a continuum of local monopoly equilibria exists if the reservation price is 'low'; and that a unique Nash equilibrium in pure strategies exists if the reservation price is 'intermediate.' Lijesen (2013) drops the assumption that the reservation price is exogenous and introduces a Web shop in the location game. If the delivery costs of the Web shop are sufficiently low (i.e., the outside good is not too different from the goods in the analysis), a Nash equilibrium in prices exists, as well as a continuum of local monopoly equilibria.

Anderson et al. (1992) approach Hotelling's model from a slightly different perspective. Rather than consumers choosing shops purely on a deterministic basis, they use a logit framework for consumers choosing a shop, allowing for customer heterogeneity. This adds an idiosyncratic preference on top of the structural choice component. The logit framework can handle both linear and quadratic transport costs, and the model with linear transport costs can be solved, provided that consumer heterogeneity is sufficiently large. We base our model on the framework of Anderson et al. (1992), as the concept of customer heterogeneity nicely fits the agent-based approach. Furthermore, we assume that transport costs are quadratic, as this guarantees the existence of a price equilibrium at every location pair without having to impose further assumptions.

3.1 Model setup

For the development of the model, we used the NetLogo platform (Wilensky, 1999). A two-dimensional world has been created, inhabited (partly) by consumers, as well as by two shops. Consumers are located on the interval 0–1, following distribution $f(x)$, shop A locates at a and shop B locates at $1 - b$ (Fig. 1).⁴

Consumers buy exactly one unit and choose stores based on price and quadratic transport costs. We adopt the approach in Anderson et al. (1992), where the choice process is probabilistic.⁵ Apart from the fact that the concept of taste heterogeneity fits

³ Salop (1979) presents outcomes in terms of number of firms (or varieties) and price equilibria.

⁴ In most applications of the model, a uniform distribution is assumed, which simplifies the model considerably.

⁵ The vast majority of the literature use a deterministic version of the model, where an indifferent consumer is identified and every consumer to the left of that consumer is assumed to go to shop A.




Fig. 1 Hotelling’s analogy of spatial competition

the agent-based approach, it also helps us to avoid that the demand function becomes discontinuous as a result of the discrete locations shops and consumers.⁶

The probability (or in the case of repeated buys, the share) that consumer i located at x_i buys at shop A defined as:⁷

$$s_{A,i} = \frac{\exp\left(\frac{-p_A - (x_i - a)^2}{\mu}\right)}{\exp\left(\frac{-p_A - (x_i - a)^2}{\mu}\right) + \exp\left(\frac{-p_B - (1 - b - x_i)^2}{\mu}\right)}. \tag{1}$$

where p_A and p_B are the mill prices of shop A and B, respectively, and parameter μ reflects the level of taste heterogeneity of consumers. This may be viewed as a taste for variety. For higher levels of μ , consumers attach a higher weight to variety, implying that the importance of the mill price and transport costs becomes smaller. The expected demand for firm A is defined as the sum (or integral in the case of continuous consumers) of the probabilities for all consumers:

$$q_A(p_A, p_B, x_A, x_B) = \sum_i s_{A,i} f(x_i). \tag{2}$$

The model is solved through backward induction. Since shops set locations and then prices, we must solve the price stage first. This entails finding a Nash equilibrium in prices given the locations of both shops. Under the (in spatial competition models) commonly used assumption that firms produce at zero costs, firms’ profits are defined as:

$$\pi_A = p_A q_A(p_A, p_B, x_A, x_B). \tag{3}$$

Maximizing profits by setting prices yields a first-order condition, which can be rewritten to a best response in prices:

$$p_A = \text{BR}(p_B, x_A, x_B). \tag{4}$$

⁶ Customers in the model are located at discrete locations. This feature of the model enables us to model their behavior in an agent-based model and has the advantage of adding to the realism of the model. A less fundamental problem is that of rounding. If the equilibrium location of a shop is not an integer, the model recognizes one of the nearby locations as the equilibrium. Rounding does not lead to incorrect outcomes, but outcomes might be imprecise. This problem can be countered by increasing the size of Main Street. Obviously, there is a trade-off between the level of precision and computing times. The computing time is a function of the level of precision to the power 4. Exact times obviously depend on the type of processor used.

⁷ Wlog, we scale the transport cost parameter to unity, allowing us to simplify notation.

Formulating a best response function for each shop yields a system of two equations and two unknowns, which can be solved through substitution to find the Nash equilibrium in prices, $p_A^N(x_A, x_B)$, $p_B^N(x_A, x_B)$.

Solving a system of n equations to find n unknowns is common practice in the analytical treatment of economic models. It has become so common that one tends to forget the implicit assumptions underlying it. The case of finding a Nash equilibrium from two best response functions requires the assumption that each player can perfectly predict the other player's response, since all parameters are known to both players and both players are assumed to be perfectly rational.

To obtain optimal locations, the Nash equilibria of the pricing stage are then substituted back into the model, yielding a profit maximization problem that depends on locations alone:

$$\pi_A = p_A^N(x_A, x_B) q_A \left(p_A^N(x_A, x_B), p_B^N(x_A, x_B), x_A, x_B \right). \quad (5)$$

Solving this problem for both shops then leads to a Nash equilibrium in locations. Like in the pricing stage, optimal solutions are found by obtaining best responses and then substituting them, again requiring perfect information and rationality.

4 An agent-based model of spatial competition

This section provides the basis for developing an agent-based model, based on decision rules that are consistent with the concept of Nash equilibrium. The purpose of the model is to simulate the optimal location of two (or more) shops in an area where they compete for consumers. The world we are looking at is one-dimensional (Hotelling's Main Street) and consists of two different types of agents: consumers and shops. Consumers choose a shop based on total costs, as defined by Eq. (1). The level of demand is fixed; consumers only choose where to buy, not whether they will buy. Shops maximize their revenues, that is a shop's price times the number of consumers that choose that particular shop.⁸ Consumers are distributed uniformly over Main Street. Since the theoretic model predicts optimal shop locations outside the inhabited area of Main Street, we expand the size of the world to twice the length of the inhabited part.

As mentioned in the introduction, we use a similar approach as chess computers do, providing a decision tree as illustrated in Fig. 2. For any given location pair, the model calculates the payoff of any given pair of prices. The payoff of the price pair that provides a Nash equilibrium is substituted into the location stage.⁹ These are then evaluated, and the location pair that provides a Nash equilibrium in locations is selected. This basic principle of coupling the stages is consistent with the analytical approach of backward induction.

Using this mechanism in its pure form for a problem with m location pairs and n price pairs would require calculating $m \times n$ payoff pairs. The potential number of

⁸ Costs for the shops are assumed to be zero for simplicity, but could easily be added to the model.

⁹ Throughout this section, the term equilibrium may also be interpreted as multiple equilibria.




Fig. 2 Payoffs of sequential moves in a multistage game

location pairs could be several hundreds, and the potential number of price pairs could be several thousands or even millions (both depending on the level of precision). This implies that we would have to leave out some options to obtain reasonable computation times, similar to the practice of ‘pruning’ in the case of chess computers.

4.1 Pricing stage

In the price optimization part, the shops look for the optimal prize to ask for their product. This module of the model is repeated for every location pair taken into consideration, so any efficiency gain made here has a large impact on total computing time. There is also ample room for efficiency increases in this stage, as the number of possible price pairs is very high. To illustrate this point, we note that the optimal price of a symmetric model precision level 20 equals 600 units, whereas it would nearly equal double that amount for a precision level of 28.¹⁰ Evaluating all price pairs between 0 and 600 would imply 360 thousand computations for every location pair at a precision level of 20 and 1.4 million computations for a precision level of 28.

Rather than using brute computational power, the shops first calculate the so-called reservation price for each consumer (or location). This is the price where, given the price of the other shop, that consumer would just buy their product at the shop. For shop A to set the reservation price, it calculates the difference between the costs the consumers at each location make to visit shop B and the travel costs to shop A. This difference is the highest price shop A can ask. After having calculated this for each location, a list is made of all reservation prices. Shop A then evaluates its profit for each of these reservation prices and picks the one that yields the highest profit. After that, the other shop will look for its best reservation price. This process is repeated up to the point where neither shop can increase their profits by selecting a different reservation price from the list.

However efficient this procedure may be, it does not provide a sufficiently precise price equilibrium. Therefore, shops use the best response in reservation prices as a starting part for fine-tuning their best response in prices. Shops then increase or

¹⁰ The precision level is defined as the number of patches that is inhabited, that is the length of Main Street.

$x_A \backslash x_B$	1	2	3	4	5	6	7	8	9
1									
2								Nash eq	
3									
4									
5									
6									
7									
8									
9									

Fig. 3 Evaluating best responses on a turn-by-turn basis

decrease their price by one unit at a time. Consumers will respond, and the shop estimates its profit. By comparing it with the old profit, the shop decides to add or to subtract another unit to the price. Every time one shop finds an optimal price, the other shop can adapt its price as well, until both shops found the same profit twice. When this happens, a precise Nash equilibrium in prices is found.

4.2 Location stage

The approach to calculating payoffs for every location pair is fairly similar to the approach used by chess computers, although some differences apply. The game of chess is turn-based, whereas the game we are looking at is characterized by simultaneous actions of players. This difference is, however, not fundamental. In fact, the algorithm that we use to identify the Nash equilibrium is turn-based, as this approach intuitively reflects the concept of a Nash equilibrium. In an analytical model, best responses take the form of equations that can be substituted to find a solution. In the computational model, this is not an option. Instead, a shop plays its best response to the current strategy of its opponent. The opponent then plays its best response to that outcome, and shops take turns up to the point where neither shop wishes to change its response, in which case a Nash equilibrium is reached. Investigating best response on a turn-by-turn basis greatly improves the efficiency of the model compared to calculating all possible payoffs.¹¹

Suppose, as illustrated in Fig. 3, shop B starts at position 7. Shop A evaluates all possible positions given that shop B locates at position 7 and finds $X_A = 3$ to be optimal, so he moves there. Now, shop B evaluates all positions given that shop A is located at position 3 and chooses $X_B = 8$. Again, shop A evaluates its options and chooses $X_A = 2$. Again, shop B evaluates its options and finds that it is already in the

¹¹ A similar mechanism is used in finding the best responses for the reservation price and the precise price.

optimal position ($X_B = 8$). Since neither shop now has an incentive to move, a Nash equilibrium in locations is reached.

The efficiency of the solving algorithm is clear from Fig. 3. Rather than evaluating all possible (81) combinations of positions, as a chess computer would do, each shop evaluates 18 positions. The efficiency gain obviously depends on the number of moves needed to reach a Nash equilibrium. A typical model run for precision level 20 would take shop A 4 turns and shop B 3 turns, leading to a total of $(3 + 4) \times 2 \times 20 = 280$ payoffs, instead of the maximum number of $2 \times 40 \times 40 = 3200$. The efficiency gain, however, also has a serious drawback. Since not all position pairs are evaluated, the model may only find one equilibrium if multiple equilibria exist. This problem can be overcome by running the model from different (pairs of) starting positions, but this obviously reduces the efficiency gains described above.

Figure 4 provides a flowchart of the resulting model structure, based on the main elements described above.

5 Model validation

An important step in developing agent-based models is to evaluate and validate the model very carefully. Most models have several types of behavior going on at the same time, and it is important to understand what drives what and what affects what (Crooks et al. 2008). Although our basic model is fairly simple, it still needs to be validated and evaluated.

We validate the model by comparing the outcomes to the analytical outcomes provided by Anderson et al. (1992). We run simulations with the model for different precision levels and a range of values for heterogeneity parameter μ . Figure 5 shows the plot of the outcomes alongside the outcome presented by Anderson et al. (1992, p. 373). The solid line represents figure 9.11 in Anderson et al. (1992). It has to be noted that it is not clear from the text in Anderson et al. (1992), whether their line is an exact representation of equilibria or an approximation based on outcomes for $\mu = 0$ and $\mu = 0.5$.

For very small values of μ below 0.1, our model does not solve at the precision levels presented here. This is a purely numerical problem, because (very) small values of μ require small steps in prices and locations to yield a valid solution.¹²

Our outcomes follow the same pattern as the outcome in Anderson et al. (1992), but reach the centered equilibrium at a higher level of μ . The resemblance in the overall pattern suggests that our model represents the outcome of Anderson et al. (1992) well enough to be able to draw general conclusions.

6 Consumer-specific heterogeneity

In the analytical version of the model by Anderson et al. (1992), the heterogeneity parameter μ is considered to be the same for all consumers. In practice, different

¹² At a precision level of 52, the model solves for $\mu \geq 0.03$.

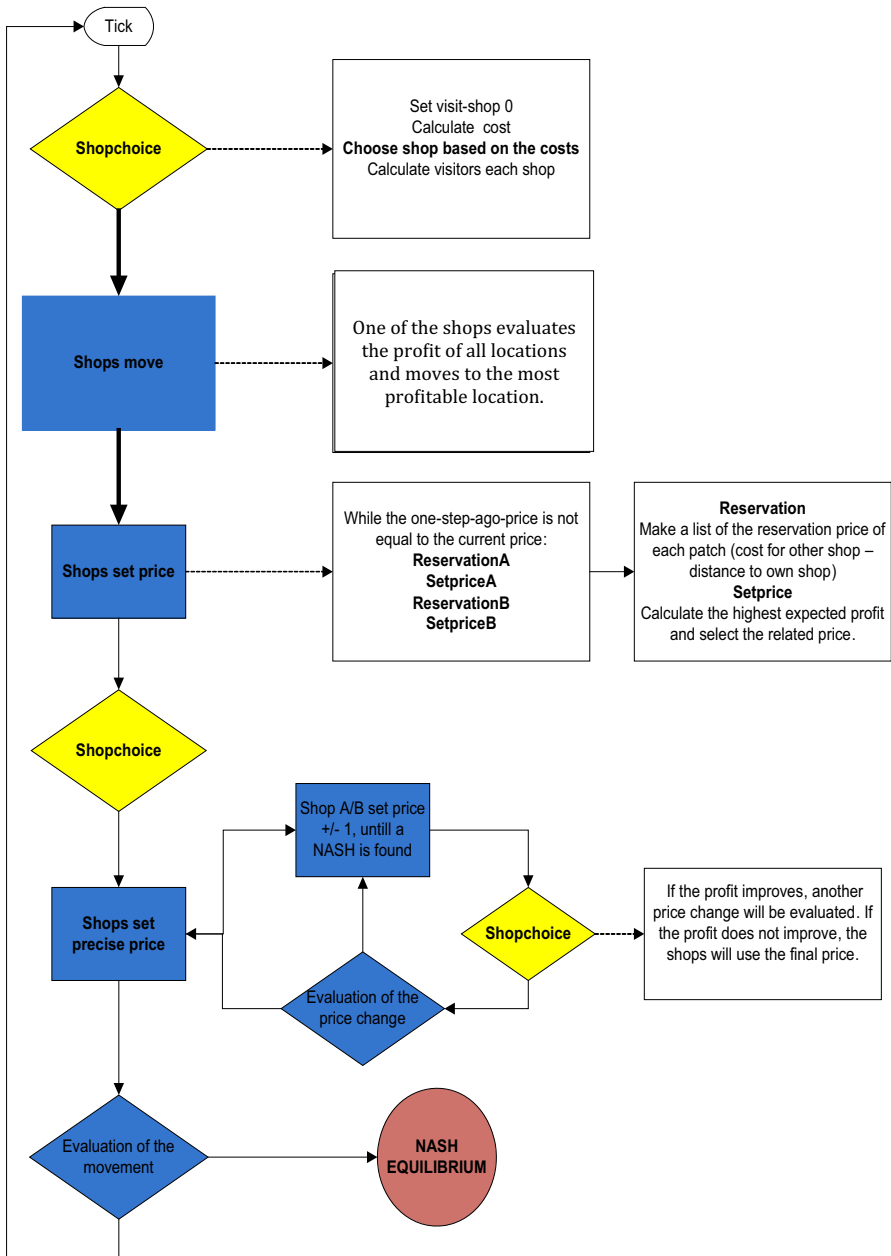


Fig. 4 Structure of the model

consumers could well have different levels of μ . Parameter μ reflects the (inverse) weight of deterministic (or observed) utility in the decision process and therefore depends on the weight that other (stochastic or unobserved) factors have. In a more down-to-earth sense, we could think of this parameter as a reflection of the taste for




Fig. 5 Location equilibria versus the level of consumer heterogeneity

variety. Some consumers stick firmly to choosing based on price and transport costs, whereas others like to deviate from their ‘optimal’ choice from time to time for the sake of variation. This taste for variety is likely to differ substantially between individuals. This can be taken into account by allowing the heterogeneity parameter to vary over individuals.

Although easy to impose, allowing the heterogeneity parameter to vary over individuals stands in the way of finding an analytical solution to the model, as it renders the demand curve for each of the firms discontinuous. With discontinuous demand functions, the existence of a Nash equilibrium is no longer guaranteed, although it may still exist. A computational model is capable of overcoming the problem of a discontinuous demand curve, although it obviously cannot find the Nash equilibrium if it does not exist.

In the original version of the model, all consumers have the same preference for heterogeneity μ , and we call this level μ_W . We allow part of the consumers to have a different level of preference, which is defined by a constant μ_X added to μ_W . Higher levels of μ imply that consumers are less sensitive to (small) price differences.

We look at two scenarios: Firstly, consumers in the third quartile of Main Street are assigned a positive level of μ_X (scenario 1); secondly, consumers in the fourth quartile are assigned a positive level of μ_X (scenario 2). The rationale for a consumer in a certain section of Main Street having a greater lover for variety can be found in an example of strong tastes, such as spicy foods. Those with a preference for spicy foods are more likely to also like less spicy foods than the other way around and would hence be viewed by an observer as having a stronger preference for variety (scenario 2). To analyze the effect, different levels of μ_W (between 0.1 and 0.8), together with different levels of μ_X (between 0 and 0.5) are tested. The results are shown in Figs. 6 and 7.

When assigning μ_X to the consumers in the third quartile (scenario 1), it appears that not all configurations result in a solution. In particular, low levels of μ_W (<0.3),




Fig. 6 Location equilibria for different levels of μ_W and μ_X assigned to consumers in the third quartile (scenario 1)




Fig. 7 Location equilibria for different levels of μ_W and μ_X assigned to consumers in the fourth quartile (scenario 2)

together with (relative) high levels of $\mu_X(\mu_W + 0.5)$ do not solve. For scenario 2, only the higher levels of μ_X together with μ_W 0.1 do not solve. Apparently, there are two different stories to be told.

In scenario 1, we increase the love for variety of the consumers that originally had a slight preference for shop B (i.e., the third quartile of the line segment). In the original situation (with $\mu_X = 0$), shop B did not move too far away from the center, as this would cause the consumers in the third quartile to switch to shop A (the market-stealing effect). As these consumers are now less sensitive to costs, this incentive becomes smaller, allowing shop B to move to the right. In terms of the spicy food example, if consumers with a slight preference for spicy food (the third quartile) have a stronger love for variety, they are less sensitive to how spicy exactly the food is, and the supplier of spicy food (shop B) can increase the taste difference between the spicy and non-spicy food (as supplied by shop A), thus increasing its price (market power effect) without losing too many consumers.

It is clear from Fig. 6 that in certain cases shop B ends up at the right edge of the (already expanded) market, especially for high levels of μ_X . Shop A follows shop B's move to the right, but to a smaller extent. However, when all consumers are less sensitive to costs (high levels of μ_W), the effect diminishes gradually and as in the original situation, they both locate in the middle. Apparently, it is in particular the relative difference between μ_W and μ_X that causes shops to behave differently.

In the second scenario, the consumers that are located in the far right of the world are assigned higher levels of μ_X . Compared to the original situation, shop B loses (some of) its captive consumers (who love very spicy food). Therefore, it moves to the center of the world (makes its food less spicy) to try to steal part of the market of shop A. Shop A then also moves to the center, responding to shop B's aim to steal market share; although for low levels of μ_W and μ_X , the direction of movement small and is ambiguous.

Again, with higher levels of μ_W the effect diminishes; however, both shops locate one position more to the left than in the original situation and in the situation of scenario 1.

In both scenarios, we changed the settings such that a subset of consumers becomes less sensitive to costs (price plus transport costs). When comparing both scenarios, we note that it makes a difference which group of consumers has a lowered cost sensitivity. If those consumers had a slight preference toward one of the shops, the market-stealing incentive is lowered and one firm moves outward (i.e., produces a stronger differentiated variant of the product). The other firm follows, but not fully, so the distance between them (and hence the price) is increased. If on the other hand the consumers that now have a lower cost sensitivity, had a *strong* preference toward one of the shops, the market power effects are weakened and both firms move toward the center in order to gain more market share.

7 Conclusions

In this paper, we demonstrate that a multistage game theoretic model, i.e., Hotelling's model of spatial competition, can be adequately shaped into a computational agent-based model. The biggest challenge lies in translating the mechanism of backward

induction from the analytical models to a method that is more feasible for an agent-based model. The basic approach here follows the same principle as chess computers do, which requires some efficiency steps to run the model within reasonable computing times. For any given location pair, the model calculates the payoff of any given pair of prices. The payoff of the price pair that provides a Nash equilibrium is then substituted into the location stage. These are then evaluated and the location pair that provides a Nash equilibrium in locations is selected.

The resulting model does not abandon the assumptions of rationality and perfect information. Instead, it provides the researcher with a choice to either strictly follow, adjust, or relax these assumptions. This paper provides a first exploration of the possibilities in this respect by including heterogenous consumers into the model. By doing so, it opens up a wide variety of possibilities for further exploration that can lead to a better understanding of the variations we observe in reality.

The application of the model to different distributions of consumer heterogeneity (i.e., less costs sensitive consumers in the third or fourth quartile of Main Street) shows how the results are not so straightforward as might be expected. First of all, location behavior is clearly different between the two scenarios: In scenario 1, shop B moves as far away from the center as possible, and in scenario 2 both shops move to the center. Configurations with relatively large differences between the consumers show the clearest distinct patterns. In both scenarios, when the whole world becomes less sensitive (i.e., higher levels of μ_W), both shops locate in the center, similar to the original situation with $\mu_X = 0$.

In its current form, the model strictly adheres to the Nash equilibrium conditions. Venues for further research include both deviations from the assumptions of rationality and perfect information and regular extensions of Hotelling's model. This is, however, difficult to accomplish with the current version of the model, since it still searches for (rather strict) equilibria, for example in prices. Loosening these conditions and still being able to find (a set of) optimal locations is a necessary next step.

Moreover, the distribution of consumers (or preferences) is unlikely to be uniform in real life, as most theoretic applications of the model assume. A promising direction of further research is applying a variety of different consumer distributions within the model, as this is typically something where the analytical version of the model would run into trouble.

An interesting, and more challenging, venue for further research would be to add dynamic adaptation processes to the model. Given the 'turn-based' approach in the treatment of the multistage character of Hotelling's model, we conjecture that an attempt in this direction may lead to excessive computing times. Adding dynamics to a single-stage game would probably yield similar insights into a much lower level of complexity and computing times.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Anderson SP, De Palma A, Thisse JF (1992) *Discrete choice theory of product differentiation*. MIT Press, Cambridge
- D'Aspremont C, Jaskold Gabszewicz J, Thisse JF (1979) On Hotelling's 'Stability in competition'. *Econometrica* 47:1145–1150
- Axtell R (1999) The emergence of firms in a population of agents: local increasing returns, unstable Nash equilibria, and power law size distributions. Center on Social and Economic Dynamics, Brookings Institution, Washington, DC
- 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
- Becker GS, Mulligan CB (1997) The endogenous determination of time preference. *Q J Econ* 112(3):729–758
- Bretagnolle A, Pumain D (2010) Simulating urban networks through multiscale space–time dynamics (Europe and United States, 17th–20th centuries). *Urban Stud* 47(13):2819–2839
- Caplin A, Nalebuff B (1991) Aggregation and imperfect competition: on the existence of equilibrium. *Econometrica* 59(1):25–59
- Chen Y, Riordan MH (2007) Price and variety in the spokes model. *Econ J* 117(522):897–921
- Crooks AT, Castle CJE, Batty M (2008) Key challenges in agent-based modelling for geo-spatial simulation. *Comp Environ Urban Syst* 32(6):417–430
- Crooks AT, Heppenstall AJ (2012) Introduction to agent-based modelling. In: Heppenstall AJ, Crooks AT, See LM, Batty M (eds) *Agent-based models of geographical systems*. Springer, Dordrecht, pp 85–105
- Economides N (1984) The principle of minimum differentiation revisited. *Eur Econ Rev* 24(1984):345–368
- Economides N (1986) Minimal and maximal product differentiation in Hotelling's duopoly. *Econ Lett* 21(1):67–71
- Farmer JD, Foley D (2009) The economy needs agent-based modelling. *Nature* 460(7256):685–686
- Halpern JY, Pass R (2015) Algorithmic rationality: game theory with costly computation. *J Econ Theory* 156:246–268
- Hinloopen J, Van Marrewijk C (1999) On the limits and possibilities of the principle of minimum differentiation. *Int J Ind Organ* 17(5):735–750
- Hotelling H (1929) Stability in competition. *Econ J* 39(153):41–57
- Irmen A, Thisse J-F (1998) Competition in multi-characteristics spaces: Hotelling was almost right. *J Econ Theory* 78(1):76–102
- Lijesen MG (2013) Hotelling's webshop. *J Econ* 109(2):193–200
- Lustick I (2002) PS-I: A user-friendly agent-based modelling platform for testing theories of political identity and political stability. *J Artif Soc Soc Simul* 5(3). <http://jasss.soc.surrey.ac.uk/5/3/7.html>
- Macy MW, Flache A (2009) Social dynamics from the bottom up: agent-based models of social interaction. In: Hedström P, Bearman P (eds) *The Oxford handbook of analytical sociology*. Oxford University Press, Oxford, pp 245–268
- Mallesen NS, Heppenstall AJ, See LM (2010) Simulating burglary with an agent-based model. *Comp Environ Urban Syst* 34(3):236–250
- Mansury Y, Gulyás L (2007) The emergence of Zipf's Law in a system of cities: an agent-based simulation approach. *J Econ Dyn Control* 31(7):2438–2460
- Meagher KJ, Zauner KG (2005) Location-then-price competition with uncertain consumer tastes. *Econ Theory* 25:799–818
- Nagel K, Rasmussen S (1994) Traffic at the edge of chaos. In: Brooks R (ed) *Artificial life IV*. MIT Pr, Cambridge, MA, pp 222–236
- O'Sullivan D, Haklay M (2000) Agent-based models and individualism: is the world agent-based? *Environ Plan A* 32(8):1409–1425
- Ottino B, Stonedahl F, Wilensky U (2009) NetLogo Hotelling's Law model. <http://ccl.northwestern.edu/netlogo/models/Hotelling'sLaw>
- Robinson DT, Murray-Rust D, Rieser V, Milicic V, Rounsevell M (2012) Modelling the impacts of land system dynamics on human well-being: using an agent-based approach to cope with data limitations in Koper, Slovenia. *Comput Environ Urban Syst* 36(2):164–176
- Salop SC (1979) Monopolistic competition with outside goods. *Bell J Econ* 10(1):141–156
- Samuelson PA (1937) A note on measurement of utility. *Rev Econ Stud* 4(2):155–161

- Scott J (2000) Rational choice theory. In: Browning G, Halcli A, Webster F (eds) *Understanding contemporary society, theories of the present*. Sage, London, p 129
- Shoham Y, Leyton-Brown K (2009) *Multi-agent systems, algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, Cambridge
- Tabuchi T, Thisse JF (1995) Asymmetric equilibria in spatial competition. *Int J Ind Organ* 13:213–227
- Tesfatsion L (2006) Agent-based computational economics: a constructive approach to economic theory. In: Tesfatsion L, Judd KL (eds) *Handbook of computational economics: agent-based computational economics, vol 2*. North-Holland, Amsterdam, pp 831–880
- Topa G (2001) Social interactions, local spillovers and unemployment. *Rev Econ Stud* 68(2):261–295
- Webster C, Wu F (2001) Coase, spatial pricing and self-organizing cities. *Urban Stud* 38:2037–2054
- Wilensky U (1999) NetLogo. Center for connected learning and computer-based modeling, Northwestern University, Evanston, IL. <http://ccl.northwestern.edu/netlogo/>