

Beyond Equilibrium: Revisiting Two-Sided Markets from an Agent-Based Modeling Perspective

Torsten Heinrich^{*} and Claudius Gräbner[†]

Abstract

Two-sided markets are an important aspect of today's economies. Yet, the attention they have received in economic theory is limited, mainly due to methodological constraints of conventional approaches: two-sided markets quickly lead to non-trivial dynamics that would require a computational approach, as analytical models quickly become intractable.

One approach to this problem is to opt for models that operate on an aggregated level, abstracting from most of the (micro-level) causes of these non-trivial dynamics. Here we revisit a well known equilibrium model by Rochet and Tirole of two-sided markets that has taken this approach. Analyzing the model from an agent-based perspective, however, reveals several inconsistencies and implicit assumptions of the original model. This, together with the highly implausible assumptions that are required to make the model analytically tractable, limits its explanatory power significantly and motivates an alternative approach.

^{*}Mathematical Institute and Institute for New Economic Thinking, University of Oxford, Oxford, UK.
Email: torsten.heinrich@maths.ox.ac.uk.

[†]Institute for Institutional and Innovation Economics, University of Bremen, Germany. Email: graebner@uni-bremen.de.

The agent-based model we propose allows us to study the phenomenon of two-sided markets in a more realistic and adequate manner: Not only are we able to compare different decision making rules for the providers, we are also able to study situations with more than two providers. Thus, our model represents a first step towards a more realistic and policy-relevant study of two-sided markets.

Keywords: Two-sided markets; Network externalities; Agent-based modeling; Simulation; Heuristic decision making; Reinforcement learning; Satisficing; Differential evolution; Evolutionary economics; Market structure; IT economics; Equilibrium dynamics

1. Introduction

Two-sided-markets (TSM) are of tremendous significance in the economy: a market is called two-sided if there are two groups of customers, usually buyers and sellers, that require the service of a third party to conduct transactions with each other. This third party is called *provider*. Examples for TSM include credit card systems (where the provider is the credit card company), video games (where the provider is the producer of video game consoles) or personal computers (where the provider is the producer of an operating system).¹

TSM pose huge difficulties for conventional economic analysis: the complex interaction among the three types of agents quickly yields nonlinear dynamics that are very difficult to capture with standard equilibrium models. One reason is that the provider's service exhibits network externalities in the sense that her service becomes more valuable to the customers the more customers are already signed up. Usually, network externalities for users of any one side are generated by the network of users on the other side. In such cases, customers of both sides must be present for the network externalities to materialize either directly or indirectly. On top of this, the strategy space for the provider is *at least* four dimensional, as she can charge customers for general membership in her network, and per executed transaction through his network. She can also set different prices for

the two sides. In extreme cases one side might even be subsidized (i.e. ‘charged’ negative prices). This is a common business practice : temporary premium benefits or even monetary incentives, for example, are frequently offered for new users in online market platforms. Credit card companies for instance often maintain bonus point systems and may give cash rewards to frequent users.

In this paper we contribute to a better understanding of TSM on two levels: On the methodological level we illustrate the inadequacy of an equilibrium approach to the study of TSM. We do this by focusing on one of the canonical equilibrium models, the one by Rochet and Tirole (2003, 2006). Analytically we show that the model remains intransparent regarding the strong assumptions required to ensure a unique and stable equilibrium. We also build a computational agent-based version of the model that helps to make transparent many of the other implicit assumptions made by Rochet and Tirole – a strategy that has proven successful since at least Gode and Sunder (1993). By doing so we illustrate the limitations of the analytical equilibrium-approach: in order to keep the model tractable one needs to reduce the possibility space for the providers to a minimum, and must assume away most of the non-linearities present in the system. The resulting models are unable to capture either the price differentiation between different user groups, business strategies that depend on disentangling entry and transaction fees, or the true uncertainty faced by the providers. This would be less severe if the consequences of these implausible assumptions were negligible. But thanks to its modularity, our computational model shows that taking even a small step towards a more realistic description of the TSM system has significant consequences. This raises the methodological concern of whether equilibrium models are credible means to study TSM at all.

Our second contribution is on the substantial level and consists of the ABM as such: our model is able to capture the most important aspects of TSM and provides some first key insights into the functioning of TSM that already go beyond what can be achieved with current equilibrium models. In particular we are able to show how the number of providers affects the market outcome, what kinds of strategies providers can develop to sustain their profits (and what kinds of strategies are likely to fail), and what regularities

or characteristics are to be expected in the development of these systems. Our model thus represents a first step in studying TSM from an ABM perspective and serves as an illustration of both the usefulness and the necessity of such an approach.

The rest of the paper is structured as follows: After giving a short overview over the existing literature in section 2, section 3 explains the fundamental problems of the equilibrium approach to TSM as taken by Rochet and Tirole (2006). A rigorous discussion of the model – which provides analytical insights that go beyond what Rochet and Tirole show in their papers – can be found in the appendix. Section 4 introduces our ABM and illustrates how sensitive the results of the equilibrium model depend on several crucial assumptions. In section 5 we discuss the additional insights that can be derived from our computational model before we conclude in section 6.

2. Literature review

The concept of two-sided markets was originally developed in the literature on banking and credit card systems to study the coordination of prices and fees between the involved financial institutions.² The term was first used in Rochet and Tirole (2003) which was previously published as a working paper in 2000. It was quickly realized that network externalities are an important feature of the systems under investigation without which those systems could not be properly understood. Subsequent models (Rochet and Tirole, 2003, 2006; Armstrong, 2006; Choi, 2010; Peitz et al., 2011) therefore generally referred to the game-theory-based equilibrium and welfare analysis of network externalities following Katz and Shapiro (1985). They did, however, ignore the literature on path-dependence, lock-in, and non-equilibrium systems (David, 1985; Arthur et al., 1987) that are likely to arise from network externalities. From this perspective, and considering the argument that real business strategies are too complex to be captured by simple optimization behavior, computational non-equilibrium methods such as agent-based simulations seem to be an intuitive modeling choice. Indeed, in recent years, this approach to TSM is reluctantly taken: Peitz et al. (2011) perform Monte Carlo simulations to find optimal strategies in their model. Meyer (2012) proposes an agent-based model of two-sided markets but

mainly to investigate the network effects. He keeps pricing exogenous and focuses on computationally establishing the viability of his theoretical results. We now try to close the existing research gap by proposing a full-fledged agent-based model that can be used to study the pricing behavior in TSM.

3. Formal problems in the canonical Rochet-Tirole model

Analytical equilibrium models are currently the standard way to analyze TSM. It is one of the main messages of the present article that such an approach can come with some serious drawbacks. We chose the model of Rochet and Tirole (2003, 2006) to illustrate this drawbacks since it is often considered a canonical model for TSM and is therefore representative for the conventional treatment of TSM.

The Rochet/Tirole model – as most equilibrium models – consists of 3 types of agents, buyers, sellers, and one single provider. Every buyer potentially interacts with every seller and the provider is able to facilitate all resulting transactions which are a constant share of all potential transactions - a strong assumption justified by the authors by referring to it as being common practice (Rochet and Tirole, 2006, p. 652).

Buyers and sellers receive a total benefit from being part of the provider's network and a per-transaction benefit. The provider charges them a membership fee and a transaction fee while incurring a per-customer and a per-transaction cost. The authors then postulate the existence of a unique equilibrium for the respective sizes of the network and derive some comparative statics results for the resulting prices.

Such an approach is, however, not without difficulties. One problem of the model is its lack of transparency concerning the mathematical derivations and assumptions. We therefore provide a detailed and rigorous description of the model that goes beyond what Rochet and Tirole offer in their papers in the appendix. Here we focus on the fundamental problems of their approach to justify an alternative treatment of TSM that encompasses these limitations.

3.1. Incapacity to consider asymmetry among buyers and sellers

Rochet and Tirole use the product of the total number of buyers as sellers as a proxy of the actual number of transactions. This implies the probability of interaction between any buyer and seller agent as equal. This is not the case in real-life interaction networks. The assumption not only cuts away a large part of the micro-layer of the system but also renders any distinction between entry fee and per-transaction fee irrelevant. For heterogeneous interaction frequencies, high entry fees would encourage membership among more well-connected, more central agents while high per-transaction fees would discourage this (and in relative terms would encourage membership among less well-connected agents instead). Credit card systems serve as an illustrative example: it must be assumed that geographical distance matters for the interaction patterns and that some buyers travel more widely than others, and that some sellers tend to cater to larger or geographically more dispersed groups of buyers. It is likely that the providers' strategies take such things into account.

3.2. Necessity to assume perfect information of customers

Rochet and Tirole assume that agents can calculate their utility from being part of a provider network in advance. However, one of the central problems for all customers in the real world is that their utility of a provider network depends on the numbers of members of this network. This leads to a recursive relationship between the buyer and seller expectations that is a key aspect of TSM, but which is simply assumed away in the model. The assumption is also unnecessary as Rochet and Tirole *assume* the system to converge towards a unique and stable equilibrium. In this case agents should arrive at the same result by just using past utilities as a proxy for the expected outcome of the current period - a far more realistic scenario which has been explored in the agent-based simulation in section 4.

3.3. Transformation of the system into a per-transaction mode

Rochet and Tirole transform all costs and utilities into a per-transaction mode. This means that all per-customer terms (the entry fees A^i , utility from network membership, B^i , and true costs, C^i) get divided by the expected number of transactions for that customer. This facilitates the derivation of an equilibrium and removes all remaining differences between entry fees and per-transaction fees which in turn generates problems for computational algorithms in which the provider uses the same optimization as a basis for her strategy: For the provider, there is no ex-ante difference between raising entry fees or transaction costs. However, ex-post, customers do not have an incentive to pay high per-transaction costs after gaining access to the network paying a still low (or negative) entry fee.

3.4. Negative per-transaction benefits for the customers cannot be considered

The customer's utility function (see equation 1 in the appendix) is assumed to be increasing in the number of customers from the other side, starting from a negative intercept (otherwise, no equilibrium solution would exist, see equation 5 in the appendix). This means, customers would not have a dominant incentive to join the network if there are no customers of the respective other type. In reality, network externalities dictate the provider to first build her user base (by setting negative fees, paying the first users to join the network) before network externalities emerge which can then be harvested by setting positive fees. The software industry provides many examples for such strategies. Further, it is also a potential business strategy to subsidize one side of the market (e.g. the buyers) entirely from the fees paid by the other side. Both cases are plausible and many examples in the real world exist, but they are obviously beyond the scope of the model.

3.5. Assumption of an unique equilibrium

For equilibrium treatments of TSM to be useful, the assumption that the system under investigation is near a unique and stable equilibrium must be plausible. For TSM this is generally not the case. Unfortunately, equilibrium approaches cannot deal with this problem adequately. Rochet and Tirole (Rochet and Tirole, 2006, p. 653) for instance

refer to the existence of such an equilibrium under “regularity conditions” – but it remains unclear what these conditions are.

This is particularly relevant since an algebraic derivation of these conditions is tremendously complex, even for the most simple possible case. A detailed analytical treatment of the problem is provided in appendix B. Here we focus on the key message: In many plausible situations we can show that there is no unique equilibrium and equilibrium models in the style of Rochet and Tirole are unable to say anything about how TSM work in such cases. For example, if some people use their credit cards all the time, while others use them only for specific purposes, such as making hotel reservations or booking flights, we will have a system of demand functions with four equilibria only two of which would be stable. See Figure 11 in Appendix B.

The Rochet-Tirole model generally neglects such cases. It’s central result – that TSM conform to the Lerner equation prescribing an equal ratio between prices and price-elasticities of demand for both seller and buyer prices – is difficult to test as all the relevant terms in the model mutually depend on each other and only price and demand can be isolated empirically. Given their methodological focus on equilibrium models, Rochet and Tirole are unable to provide any general conclusions beyond saying that higher numbers “captive” customers of type i lead to higher prices for type i : a larger number of already committed sellers prompts the provider to set lower prices for buyers and increase the price for sellers (Rochet and Tirole, 2006, p. 659).

This criticism applies to a wide range of equilibrium models and illustrates the fundamental problems of an equilibrium approach to a situation in which mutual dependency and continuous change are essential properties of the system under investigation.

4. An agent-based model of TSM

We now describe our ABM. Because the focus is on studying the price setting behavior of the provider, the customer agents are held as simple as possible. An overview over the parameters of the model can be found in table 1.

Symbol	Name	Description	Default value
A_{fixed}^i	Fixed entry fee	Relevant only in the optimization case (endogenous otherwise)	$(0, -200)$
$\widetilde{N^B}$	Number of buyers	-	10000
$\widetilde{N^S}$	Number of sellers	-	2000
$\widetilde{N^P}$	Number of providers	-	1, 10
C_{fixed}	Provider fixed cost	-	10000
C^i	Per customer fixed costs	Costs incurred to the provider for every customer.	50
c^i	Per transaction costs	Costs incurred to the provider for every transaction.	25
t_{max}	Nb. of time steps	-	500
-	Nb. of transactions per iteration	-	30000
-	Subscription threshold	Customers do not try to subscribe to further providers if their revenue is below the threshold	400
\bar{p}	Mean reservation price	$p_B^r \propto \mathcal{U}(\bar{p}, 1.5\bar{p})$ and $p_B^r \propto \mathcal{U}(0.5\bar{p}, \bar{p})$	1000
c_R	Roaming costs	Costs incurred to customers when using the network of a provider they have no subscription with.	100
F	Entry fee range	Range of admissible entry fees.	$(-3000, 5000)$
F	Transaction fee range	Range of admissible transaction fees.	$(-1000, 1010)$
-	Ema-Faktor	Exponential moving average factor governing the adaptation speed of the reinforcement learning mechanism	0.01

Table 1: The relevant fixed parameters of the ABM. Note that the main drivers for the results of the model are the particular provider decision making algorithm chosen.

4.1. The customers

4.1.1. Construction of the agents

The customer agents are instantiated either as buyers or sellers. Their type remains constant over time. While we also studied the impact of varying group sizes on the model output (see section 5 and figure 2), we use, where not indicated otherwise, a standard setting with 2000 sellers and 10000 buyers. Agents start in t_0 with a wealth of zero and are heterogeneous concerning their per-transaction benefits, b^i .⁵ This benefit changes every round and is drawn from a uniform distribution.

There are two reasons for these assumptions: Firstly, this largely fulfills the assumption of Rochet and Tirole (2003, 2006) with one stable equilibrium. Secondly, the main interest of the present research is not the influence of particular distributions of preferences, but the pricing behavior of the providers.

At t_0 all customers are subscribed to one of the providers chosen randomly for each customer.

4.1.2. Decision making of the agents

For a constant number of transactions in every round, two customers (one buyer and one seller) are chosen randomly. Each of those transactions is conducted only if a positive transaction benefit (shared equally afterwards⁶) remains for the two customers after subtraction of per-transaction prices they would incur in conducting the transaction.

In every round, agents may decide to join one of the provider's networks. In the model by Rochet and Tirole, the decision is made according to the following formula (which has been adjusted for the discrete case):

$$b_k^i + \frac{B^i - C^i}{N_{t-1}^j} \geq p_{t-1}^i = a_k^i + \frac{A^i - C^i}{N_{t-1}^j}.$$
$$b_k^i + \frac{B^i}{N_{t-1}^j} \geq p_{t-1}^i = a_k^i + \frac{A^i}{N_{t-1}^j}.$$

The condition is most conveniently verified by assessing whether or not the agent

incurred losses by participating in the network (of this particular provider) in the last period. We therefore translate these function into the following decision algorithm for agent k (assuming k is a seller and that the number of buyers in the network has not yet been updated; we also assume that if the agent is not subscribed to any provider, one is randomly assigned with a moderate probability of 10%⁷):

Algorithm 1 Decision making algorithm for customer k .

```

if  $providerlist_k \neq \{\}$  then
  for all  $P \in providerlist_k$  do
    if revenue from transactions using provider  $P$  is  $< 0$  then
      remove  $P$  from  $providerlist_k$ 
    end if
  end for
end if
if  $providerlist_k = \{\}$  and with probability 0.1 or  $revenue_k > 400$  then
  Select random provider  $P_0$ 
  Add  $P_0$  to  $providerlist$ 
end if

```

A particular provider is selected randomly by the agent with the probability to select provider i being a monotonly increasing function of the share of customers already subscribed to provider i .

4.2. The provider agent

We consider three different price setting strategies for the provider (see table 2).

In the first case the provider optimizes his profits using a differential evolution algorithm. This algorithm helps her compute the optimal entry fees and per interaction fees to be charged from the buyers and sellers. The differential evolution algorithm allows us to study a generative version of the Rochet-Tirole model and to compare it against the other cases as a benchmark. It is certainly the closest replication of a rational choice setting that we can get in a computational context. For such a *quasi rational choice* setting to be tractable, many unrealistic and implausible assumptions such that the provider knows the exact shape of the demand function he faces are required. This setting represents, however, a benchmark against which the other, more realistic versions of the model can be compared. It therefore allows us to test whether the inadequacy of the assumptions

necessary in the equilibrium context does matter for the outcome of the model. If it doesn't, the use of equilibrium models *a la* Rochet and Tirole could still be justified on the grounds that assumptions in economic models are always to a certain degree unrealistic, but that this 'unrealisticness' does not matter too much.

In order to go beyond the unrealistic assumptions of the original model and to get alternatives to this rational choice benchmark, we implement the model with two versions of a reinforcement learning algorithm for the provider. Providers that use the reinforcement learning algorithm to set entry fees and per interaction fees adapt their price setting strategies according to their success in the past. Such behavior does not involve many of the unrealistic assumptions necessary in the benchmark (such as complete knowledge of the demand functions) and is compatible with the basic notion of bounded rationality. We study the model with and without a satisficing element in the strategy algorithm. Because the computational simplicity of the reinforcement learning mechanism, we are also able to study the behavior of the model with more than one provider.

4.2.1. Instantiation of the provider agent

The provider gets instantiated with a list recording the subscriptions and the profits. Depending on which case of table 2 we consider, she is assigned a strategy according to which she sets the entryfees and per transaction fees. This can be either a differential evolution algorithm that should mimic an optimal and rational choice (to replicate the results of Rochet and Tirole) or a reinforcement learning algorithm, which represents a more realistic scenario of how the providers set their prices.

4.2.2. Price setting strategies for the provider

Provider decision using a differential evolution algorithm Every round the provider maximizes her profits using a differential evolution algorithm (Storn and Price, 1997). Such optimization algorithms require the provider to be informed about the distribution of the reservation prices and the types of the customer agents. Therefore, this decision

Provider decision making	Description	Cases considered			
		One provider	Ten provider	fixed entry fees	endoge- nous entry fees
Optimiza- tion algorithm	The provider knows the distribution of the customers reservation prices and uses a differential evolution algorithm to optimize his profit every round.	X		X	
Reinforce- ment learning	The provider changes the transaction and subscription fees for the customers depending on whether his revenues have risen or fallen compared with the previous round.	X	X		X
Reinforce- ment learning with satisficing	In this case the provider stops changing his strategy after having reached a satisfying level of profits.	X	X		X

Table 2: An overview of the cases considered in our simulation model.

making algorithm is not very realistic and used only to replicate the results of Rochet and Tirole in a generative and agent-based framework and to serve as an ideal benchmark case that can be compared to other decision making concepts. It illustrates, however, some of the fundamental difficulties of Rochet and Tirole's model that remain unaddressed in an analytical framework: While the algorithm is rather successful in finding optima on not too rugged target functions and works very well for per-transaction fees for buyers and sellers (a^s, a^B) as free variables (with fixed entry fees A^S, A^B), it fails in cases where the provider must set all four variables (A^S, A^B, a^s, a^B).⁸ This indicates that assuming a provider acts in an optimizing way is a very misleading assumption, rather than a useful abstraction of reality.

Provider decision algorithm using reinforcement learning A more realistic conception of the decision making process is to use a reinforcement learning algorithm⁹: in this case, the distribution of b^i and B^i remain unknown to the provider. Instead, she considers her success in the past and adjusts her prices accordingly.

There are three different ways the provider can adjust her prices: firstly, to continue using the current price, secondly, to increase the prices towards some maximum value and, thirdly, to decrease the prices towards a (negative) minimum value.¹⁰ All three possibilities are assigned a certain probability weight that gets adjusted according to the success of the provider compared to past revenues.

Practically, the provider decides according to the following procedure:

1. It firstly checks whether she can match her operating costs. If not, then the current strategy gets discouraged strongly and the strategy is set back to a sustainable (though not necessarily very successful) initial value.¹¹
2. It then takes the revenue of the previous round as a relative success measure.¹²
3. The probability weights of all strategies are adjusted, with the probability weight of the current strategy increased if the strategy was successful, otherwise decreased.
4. The provider chooses new prices according to the current probability weights.

Provider decision algorithm using reinforcement learning with satisficing The above reinforcement learning strategy may be extended to include satisficing in the sense of Simon (1955). In this case, before step 3 in the above algorithm, the provider checks whether a satisficing level of revenue and number of customers is reached. If this is the case, the strategy is not changed and the provider simply applies the same strategy again. If no satisficing level is reached, the provider acts identically as above.

4.3. Summary of the simulation

Here we summarize the sequence of events for the whole simulation. Note that a well documented version of the source code is also available.

1. Create the agents
2. Set global provider choice functions containing probability weights for the choice of providers by customers.
3. Create initial allocation of buyer and seller agents to providers (random uniform).
4. Iteration for the specified time steps. Each time step consists of six stages:
 - 1st Stage: Realization of the transactions.
 - i. Choose new reservation prices for all buyers and sellers.
 - ii. Conduct the specified number of transactions
 - A. Seller and buyer are randomly chosen
 - B. Compare the reservation prices p_B^r and p_S^r . Continue only if $p_B^r > p_S^r$.
 - C. Identify cheapest available channel of transmission between buyer and seller (depending on their subscriptions) and the corresponding transaction cost c^t .
 - D. If $\frac{p_B^r - p_S^r}{2} > c^t$, the transaction is conducted, buyers and sellers receive $\frac{p_B^r - p_S^r}{2}$ each, the providers are awarded the transaction costs.

2nd Stage: Customers pay the fees to the provider

- 3rd Stage: Providers pay fixed costs and per-customer costs.
- 4th Stage: Data gets collected and evaluated.
- 5th Stage: Unsubscription decisions by the customers and adjustment of the providers' network weights.
- 6th Stage: Subscription decisions of the customers.
- i. A customer subscribes to a provider if either of the following is true:
 - If she has more revenue left than a threshold value.
 - If the customer currently has no provider she has to subscribe to a provider with probability of 0.1. This is necessary to ensure that customers do not immediately join the same unfavorable provider.
 - ii. A random provider is chosen.
- 7th Stage: Reset weights of providers and customer's period variables (money, income, revenue,...)
- 8th Stage: Provider set entry fees and transactions fees for buyers and sellers.

5. Simulation results

5.1. Optimal provider decision using a differential evolution algorithm

Simulation results for rationally optimizing providers are shown in figures 1 (buyer and seller transaction prices) and 3 (welfare distribution).

As expected, the provider consistently chooses the same prices in every time period. The unbiased random distortion arise from imperfections in the approximative optimization algorithm. This leads to a stable overall result after a potential initial adjustment of the size of the networks (agents with low b^i leaving). The result is also invariant with respect to the number of providers as all rational optimizers will compute the same optimal prices: as the increasing returns brought about by the network externalities favor large competitors, all but one of these providers will fade into obscurity while one will take over

the entire population and continue identically to the case of just one provider. Since the rational optimizers are by design perfectly informed (see section 4.2), they outcompete heuristic decision-makers when confronted with them.

The rationally optimizing provider is included as a benchmark case and as a point of comparison to Rochet and Tirole’s model. While it would be as difficult to scrutinize whether the Lerner formula holds (Rochet and Tirole’s central result) as it is in reality, Rochet and Tirole’s interpretation of the same as more ”captive” customers of type i leading to higher prices for this type can be verified. The present ABM does not allow for ”captive” customers, but higher numbers of buyers relative to that of sellers do indeed lead to higher prices for buyers and lower ones for sellers. Figure 2 shows the prices set by rational optimization with differential evolution for different numbers of buyers compared to a constant 2000 sellers (with a constant expected potential number of transactions per customer). The result of the equilibrium model is nicely reproduced and has an intuitive interpretation: the smaller group has a more powerful position than the larger group, as they are important for the provider to build up a network. This, however, is but a tiny part of the rich dynamics that result from the more realistic heuristic decision mechanisms below.

Interestingly, the rational optimization mechanism only works well with fixed entry fees, i.e. with only the transaction fees being subject to optimization. Simultaneous optimization of all four terms fails, non-optimal prices are set, and the network collapses. This is because most parameter settings then lead to empty networks and thus a (constant) non-existing profit with nothing to optimize. The field of successful values is small and accidentally hit by the initial population only in rare cases. This is yet another powerful argument against the use of models that simply *assume* that such unique and stable equilibria exist and can be reached by the provider.

In summary, our model is able to transfer the basic model by Rochet and Tirole into an agent-based framework and to study the importance of several particular assumptions. The difficulties we encounter cast serious doubts on whether the equilibrium-based approach is useful in the context of TSM.




Figure 1: Transaction costs charged by the rational optimization provider with fixed subscription costs for both types of customers (upper panel 0, lower panel 200), averages over 100 runs.




Figure 2: Transaction prices after 40 iterations depending on the the number of buyers compared to a constant number of 2000 sellers, averages over 100 runs.

5.2. Reinforcement learning and satisficing

Using the reinforcement learning algorithm for the decision making of the providers allows us both to drop unrealistic assumptions - e.g. the perfect information about demand functions and the ability to perform the optimization - and to let the provider set the entryfees for the customers endogenously. However, the results are more volatile and path dependent than for the differential evolution algorithm case. The system does not converge towards an equilibrium. Instead it follows certain quasi-periodical patterns; once the different runs are out of sync, the corridor of results widens, although the simulation still indicates a clear corridor in which the results can be found.

Different from perfect optimization, the behavior varies across customers. A selection mechanism works on the providers' strategies in the form of customers joining networks with more favorable price setting and generating both network externalities and revenue there. The selection mechanism does, however, only function if there is competition.

Consequently, simulation runs with a only single provider generally result in malicious pricing behavior, destruction of the network and complete failure (see the number of




Figure 3: Development of the total revenue and its distribution resulting in settings with rational optimization providers with fixed subscription fee (upper panel 0, lower panel -200), averages over 100 runs.

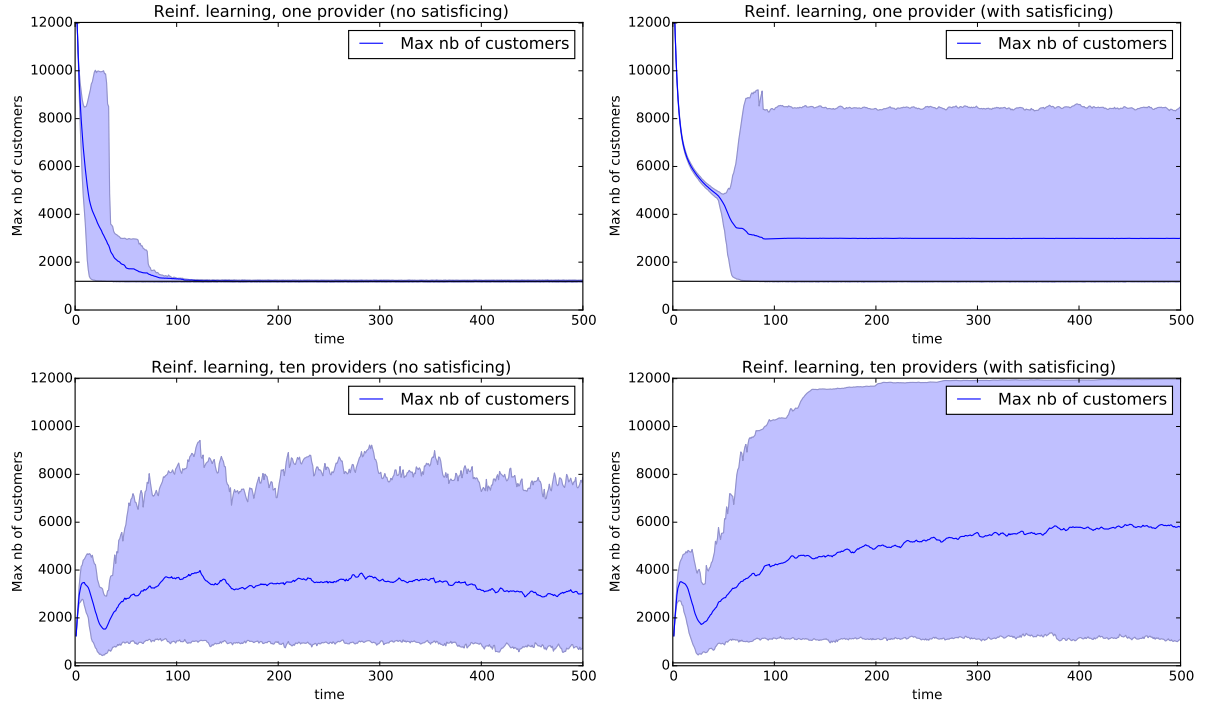


Figure 4: Development of the number of customers in the (largest) network of reinforcement learning providers, mean and 90% quantiles from 200 runs. The lower black lines indicate the lower bound of customers in the network from random trial and error joining of the providers' networks.




Figure 5: Development of the number of customers in the (largest) network of 10 competing providers following reinforcement learning with satisficing, mean and 90% quantiles from 200 runs.

customers as shown in figure 4). Without satisficing (upper left panel), only the 10% customers that randomly join the provider’s network before leaving it again in the following period remain. As this is an artifact of the simulation design, this strategy can be considered to be absolutely unsuccessful. With satisficing (upper right panel), the provider may succeed, but the network still breaks down in more than 60% of the runs. In case of multiple (in our case ten) competing providers, the network is successfully established in many runs both with and without satisficing. With satisficing (lower right panel), the average number of customers shows a clear upward trend, something that does not occur for the setting without satisficing (lower left panel) because reinforcement learning can lead to overshoots (driving prices too high or too low) if satisficing does not freeze the strategy at some satisfactory setting. It is important that the heuristic decision making provider is more successful if the range of possible successful price settings is larger, e.g. if *ceteris paribus* the fixed costs the provider faces are lower. This is shown for the case with ten provider employing reinforcement learning with satisficing in figure 5. If the viable range of successful strategies is lower, the probability to hit one of them by chance is lower and the expected time to the successful establishment of a network by any one of the providers increases.

As a single provider is usually not able to build up a functioning network, we are now focusing on the cases of ten providers and study the path dependent pricing decisions in more detail. Analogous to figures 1 and 3 for the rational optimization case, figures 6 and 8 show the development of buyer and seller prices as well as revenue and its distribution on buyers, sellers and providers for all reinforcement learning providers with (upper panel) and without (lower panel) satisficing. As seen in figure 6, both types of providers systematically charge a higher subscription fee for the larger group, the buyers. This corresponds to the pattern seen both in figure 2 for rational optimization and in Rochet and Tirole’s anecdotal evidence and inferred by them to be a result of the Lerner equation. However, the present result is more nuanced: The opposite pattern can be seen for transaction fees, here the customers in the smaller group, the sellers, are charged more. It is crucially important for the provider to have enough members of the smaller

group in the network such that sufficiently many interactions can take place. With this accomplished, a higher per-transaction fee can be charged, and it here, the sellers, the members of the smaller group, are charged slightly more, as they have a larger expected number of transactions than the buyers as a result of the different group sizes and the uniform interaction probability. Note that the same pattern can currently be observed for for credit card systems in the US with providers subsidizing the buyer side with generous benefits while charging the merchant side. Of course in this case, the merchants simply pass on the fees to the buyers in the form of extra charges for credit cards.

Transforming all the different types of fees and prices into one single per-transaction prices as done in Rochet and Tirole's approach overlooks this pattern which furthermore is of a strongly path-dependent nature (see figure 7). That is, it is statistically observable in averages and quantiles of the distribution of results of many simulation runs, but there are runs for which the reversed pattern emerges and persists partly or entirely over the course of the simulation. The same path-dependence is of course found in other variables including the size of networks, and the development of revenues.

Figure 8 shows the development of welfare (revenue) and its distribution. On average, this is lower than in the rational optimization case (see above, figure 3) as the rational optimizer has an unrealistic amount of information about the demand structure of the market and is able to set optimal prices. However, consistently increasing absolute revenue can be observed for the satisficing strategy (lower panels). It also attains disproportionately higher levels for cases with lower fixed costs.¹³ The development of revenue is also path-dependent: In some cases, the providers fail to successfully establish a functioning network resulting in the loss of the network externalities that would both continue to sustain the network and generate the bulk of the revenue.

In all cases, the smaller group, that of the sellers, is much better off as the provider cannot afford losing their support and only hesitantly increases transaction prices for them which also affects the buyer's ability to conduct cost-effective transactions. In real-world economic systems, the sellers may proceed to pass this entire fee on to the buyers (see above).




Figure 6: Development of subscription and transaction fees in settings of 10 competing providers following reinforcement learning with and without satisficing, averages over 200 runs.

Considering these results from a policy perspective suggests that overall welfare could be increased if the public agent supports providers in building up a well-functioning network that generates positive externalities for all agents, and then provides incentives for a sustainable satisficing strategy. This preliminary conclusion, however, illustrates the intricacy of TSM and the need for their further investigation from an agent-based perspective: the original Rochet Tirole model simply assumed away the complexity that makes such policy questions both difficult and interesting.

6. Conclusions

The present paper contributes to the study of TSM in two respects: Firstly, on the methodological level, we showed that studying two-sided markets with conventional equilibrium models bears the danger to overlook the most fundamental features of these markets. Secondly, on the substantial level, we introduced an agent-based model that

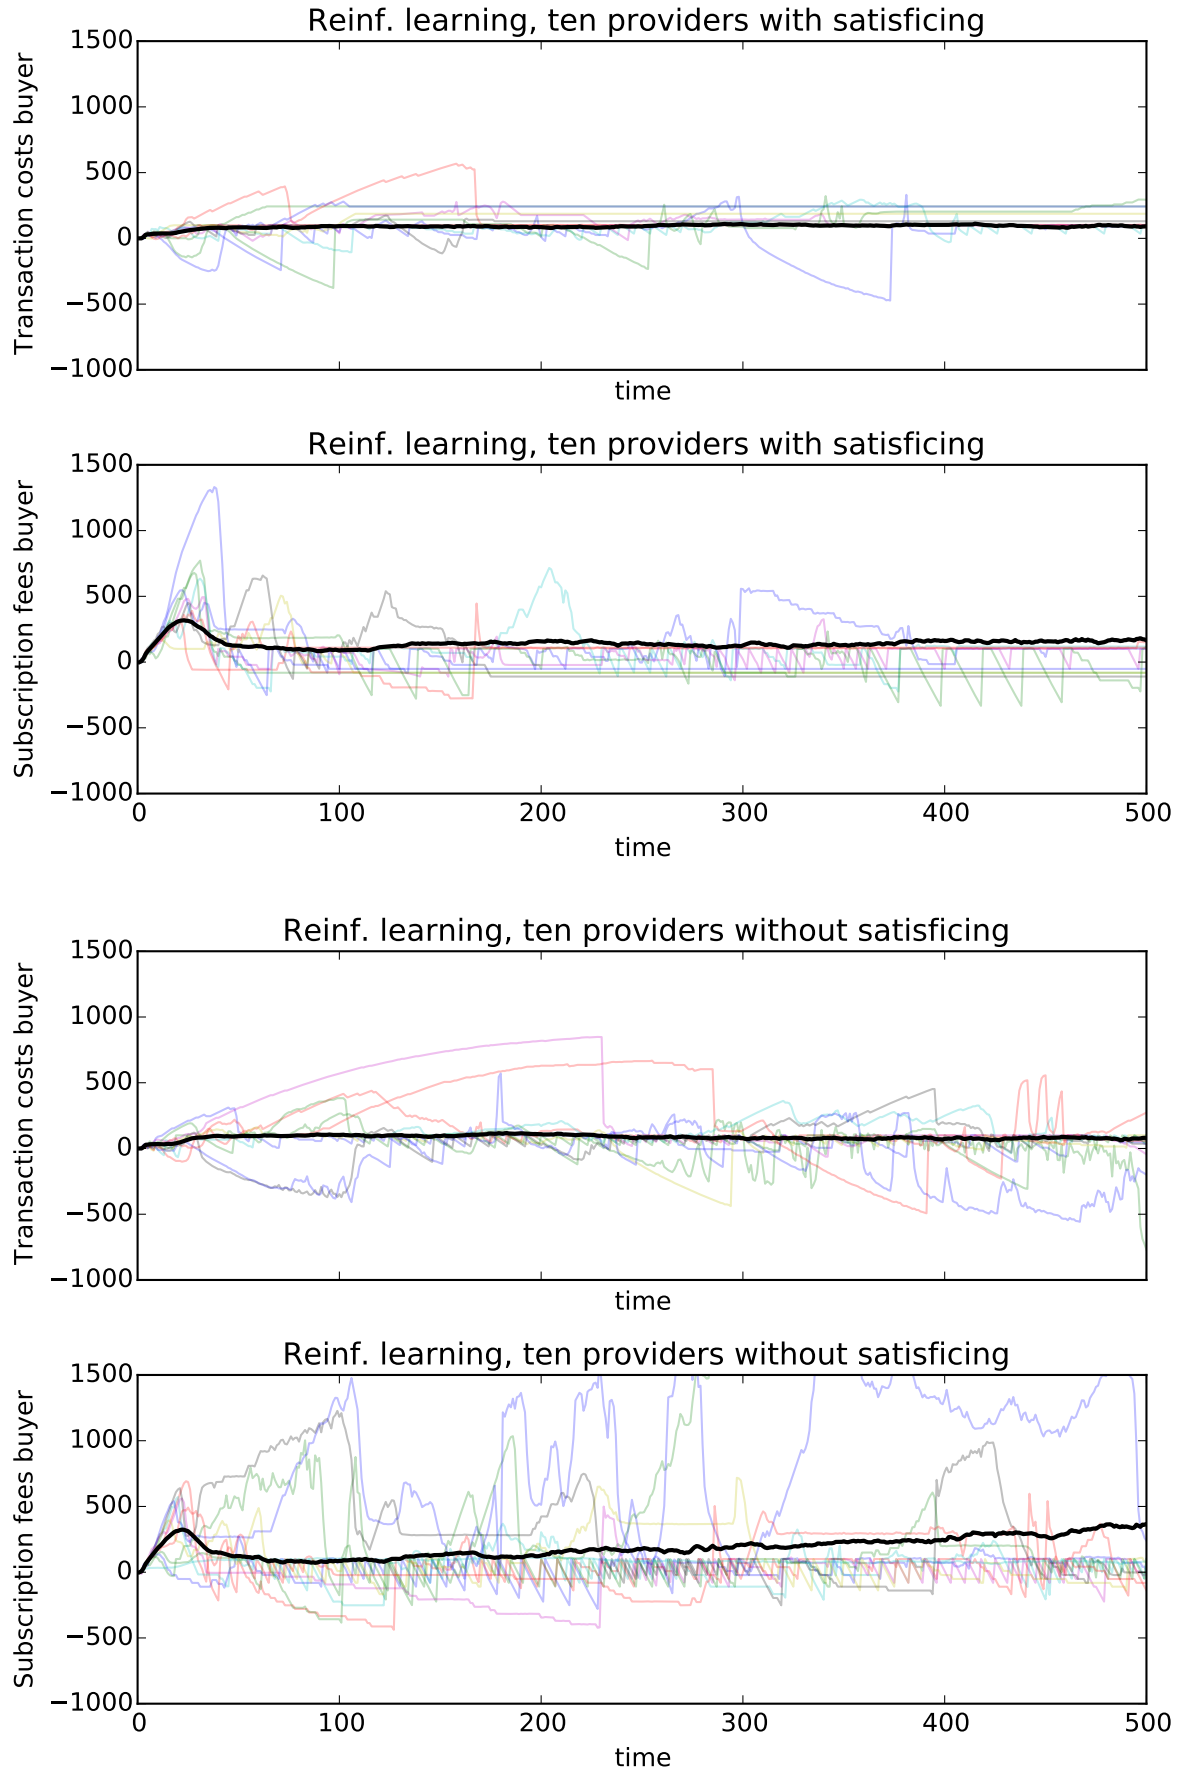


Figure 7: Development of buyer transaction prices in 30 example runs with 10 providers following reinforcement learning with and without satisficing.

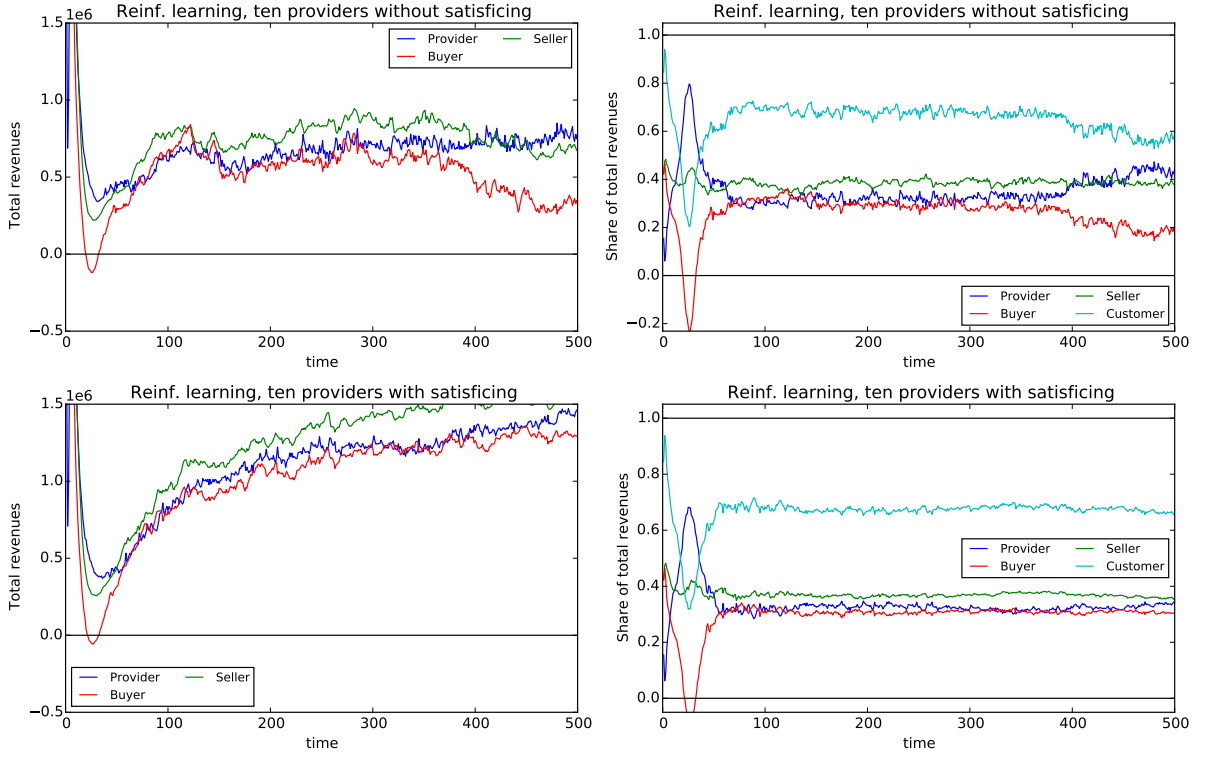


Figure 8: Development of provider, seller, and buyer revenue with 10 competing reinforcement learning providers with (lower panels) and without (upper panels) satisficing, averages over 200 runs. The negative share of buyers in the cases without satisficing indicates the unsustainability of this strategy.

illustrates both the usefulness and necessity of a computational approach to the study of TSM.

With regard to our methodological contribution: The most promising results obtained in conventional equilibrium models are equilibrium conditions - such as the Lerner equation in Rochet and Tirole's approach - which can, given abundant confidence in the result, be interpreted to account for specific pricing patterns - such as higher prices the customer group with more "captive" users.

While we have shown that it is possible to reproduce this in a much more realistic agent-based setting with little difficulty, the universality of this result is highly doubtful for two reasons:

Firstly, the entire equilibrium construct rests on an array of incredibly bold assumptions as discussed in section 3. Our computational approach showed that this unrealisticness matters because relaxing the strong assumptions changes the dynamics predicted by the model significantly.

Secondly, agent-based simulations in section 5 also showed that once the assumption of rational optimization with perfect information is relaxed, the dynamics of the system become highly path-dependent. Universal prediction *a la* Rochet and Tirole is not appropriate for such systems.

With regard to our second contribution, our ABM was able to generate the following insights: First, heuristic decision algorithms, as considered here in the form of reinforcement learning with and without Simonian satisficing, are able to distinguish effectively between user groups and to set subscription and transaction fees accordingly. The dominant pattern was found to be high subscription fees for the side with larger numbers of agents (buyers in our case) in order to subsidize the smaller group while transaction fees were higher for the smaller group that on average conducts more transactions.

Second, in order for the heuristic algorithms to work efficiently, a certain selection pressure is essential. In our case, this was provided by competition between several providers. In real-world economic systems, alternative technological paths are also feasible, e.g. if the dominant video game platform providers charge excessively high prices, sooner

or later, another technology will emerge in which the price-setting behavior happens to be more moderate.

Third, we found relentless and myopic optimization, even with (or especially with) heuristic decision algorithms to be susceptible to unsustainable decisions, specifically overshooting prices. A simple modification, introducing satisficing into the decision mechanism resolves this. Note that the decision mechanisms considered here are still excessively simple, if not trivial. Nevertheless, their application to TSM represents a new contribution to the literature and illustrates the point that decision makers using extremely simple heuristics can achieve reasonable results. In real-world systems it is often observed that providers spend long years and unbelievable amounts of money on building up an installed base of users to generate network externalities and attract more users (this being known as the "burn rate" during the new economy bubble) before starting to make profits. More successful examples include Google and Facebook, less successful (and hence today also less well-known) ones include WorldCom, GeoCities, or Tiscali.

Fourth, a larger field of viable solutions which result in the successful establishment of a network have a profound impact on the speed with which heuristic strategies are able to find them. This translates directly into the probability of the emergence of a network which would then be self-sustainable as a result of the network externality. Factors constraining the size of this field are in our case operating costs of the providers; in real-world systems, this would be complemented by technological aspects, such as the regulatory environment, and cultural factors.

Finally, as compared to an unrealistic rational optimizer with perfect information, the welfare-efficiency in terms of revenue generated for providers and both types of customers is rather modest in heuristic optimization cases (and the rational optimizer may not even constitute the welfare optimum). While more sophisticated heuristic strategies may perform better, it is not likely that they will get close to the benchmark. As a complication, any number of interfering effects may occur in real systems: technological lock-ins with huge sunk costs, unsuccessful providers with a high initial "burn rate" which then translates into losses for investors or for the public, or the complete failure to establish a working network.

The complicated two-sided network externality and the disparity between the provider's incentives and public interest and will in any case prevent any efficiency-optimizing effect that may be present in other economic systems. That is not to say that TSM systems are *per se* inefficient, merely that regulation can have and did have a positive effect on this.

While our proposed ABM is still exceedingly simple, the first results were already promising and may serve as a starting point for further research. Several aspects naturally warrant further investigation: the effect of non-homogeneous interaction patterns or a non-trivial network structure between customers, more sophisticated strategies on the part of the providers, and possible welfare-enhancing policy measures.

The methodological conclusion of this article is that all these important aspects require an investigation from an agent-based rather than equilibrium perspective. This will also facilitate the empirical assessment of the theories, since ABM enable us to investigate a substantially wider variety of hypotheses both on the assumption and the conclusion side.

Acknowledgements

Appendix

A. Rochet and Tirole's Equilibrium Model of Two-Sided Markets

Here we provide a detailed revision of the canonical TSM model as proposed by Rochet and Tirole (2003, 2006).

A.1. Setting

There are 3 types of agents, buyers, sellers, and one single provider. As all buyers, $\widetilde{N^B}$, will potentially interact with all sellers, $\widetilde{N^S}$, the provider must be able to facilitate all $\widetilde{N^B}\widetilde{N^S}$ transactions.

Buyers and sellers are denoted by superscript $i \in \{B, S\}$. They will receive a total benefit from being part of the provider's network of B^i and a per-transaction benefit of b^i . The provider charges them a membership fee of A^i and a transaction fee of a^i while incurring a per-customer cost of C^i and a per-transaction cost c . Call the number of buyers and sellers in the network N^B and N^S respectively resulting in $N^B N^S$ potential transactions of which the number of actual transactions is a constant share. This is a relatively strong assumption justified by Rochet and Tirole (2006, p. 652) with being common practice.

This entails that from the agent's point of view all agents of the other type (buyers from the point of view of sellers and vice versa) are homogeneous; the only heterogeneity allowed in this model is that agents have different b^i and B^i , which are assumed to be continuously distributed (Rochet and Tirole, 2003, p. 999). Unfortunately, the particular type of the distributions has not been further specified by the authors, a fact that causes difficulties for the computational replication of their findings (see below). Further, the per-transaction utilities are assumed to be constant over all transactions an agent participates in independently of the transaction partner and any other factor. Rochet and Tirole make a point of highlighting that this features non-rivalry (Rochet and Tirole, 2006, p. 653).

A.2. Buyer/Seller Decision Problem

The buyer's and seller's utilities U^i consist of the net network utility, $b^i - a^i$, multiplied by the network size in terms of members of the other type's (potential transaction partners), N^j , and the net standalone utilities, $B^i - A^i$ (Rochet and Tirole, 2006, p. 653)

$$U^i = (b^i - a^i)N^j + B^i - A^i. \quad (1)$$

Agents are assumed to be able to calculate their utility in advance. They will only become part of the network if their utility is larger than zero, hence the network size can be computed from the distribution of utilities, i.e. share of the agents whose utilities are larger than zero $N^i = Pr(U \geq 0)\widetilde{N}^i$ (Rochet and Tirole denote $Pr()$ as the absolute frequency instead of the CDF, hence $N^i = Pr(U \geq 0)$) (network participation condition).

Rochet and Tirole then proceed to eliminate the network-level by and transform the buyer's and seller's decision function into a per-transaction mode by dividing network access fees and standalone benefits by the number of potential transaction partners (Rochet and Tirole, 2006, p. 653) (or actual transaction partners, (Rochet and Tirole, 2006, p. 654, footnote 24)). In order to be able to more easily transform also the providers optimization problem (equation 3) into a per-transaction form, they first subtract the per-member cost for the provider of maintaining the network, C^i , thereby obtaining something like the (per-transaction) value added ($B^i - C^i$ or $\frac{B^i - C^i}{N^j}$) or (per-transaction) profit ($A^i - C^i$ or $\frac{A^i - C^i}{N^j}$). It does, however, not appear to change either the decision problems or the equilibria in any substantial way. This yields a quasi-per-interaction-price (Rochet and Tirole, 2006, p. 653)

$$p^i = a^i + \frac{A^i - C^i}{N^j}.$$

and a network participation condition

$$\begin{aligned}
N^i &= \widetilde{N}^i Pr(U \geq 0) \\
&= \widetilde{N}^i Pr((b^i - a^i)N^j + B^i - A^i \geq 0) \\
&= \widetilde{N}^i Pr\left(b^i - a^i + \frac{B^i - A^i}{N^j} \geq 0\right) \\
&= \widetilde{N}^i Pr\left(b^i - a^i + \frac{B^i - A^i}{N^j} + a^i + \frac{A^i - C^i}{N^j} \geq a^i + \frac{B^i - C^i}{N^j}\right) \\
&= \widetilde{N}^i Pr\left(b^i - a^i + \frac{B^i - A^i}{N^j} + a^i + \frac{A^i - C^i}{N^j} \geq p^i\right) \\
&= \widetilde{N}^i Pr\left(b^i + \frac{B^i - C^i}{N^j} \geq p^i\right)
\end{aligned}$$

The network sizes N^i thus take the form of functions of p^i and N^j (as all other parameters, $\widetilde{N}^i, b^i, B^i, C^i$, are constants) which may be interpreted as demand functions

$$N^i = D^i(p^i, N^j) \left(= \widetilde{N}^i Pr\left(b^i + \frac{B^i - C^i}{N^j} \geq p^i\right) \right)$$

such that a system

$$N^i = D^i(p^i, N^j) \quad i, j = B, S, j \neq i \quad (2)$$

is obtained in which the N^j can be internalized, thus $N^i = n^i(p^i, p^j)$ with $i \in \{B, S\}$.

The functions n^i are characterized by their partial derivatives (with respect to the two p^i) which are obtained (Rochet and Tirole, 2006, p. 653) from total differentiation of

$$N^i = D^i(p^i, N^j) = D^i(p^i, D^j(p^j, N^i))$$

$$\begin{aligned}
dN^i &= \frac{\partial D^i}{\partial p^i} dp^i + \frac{\partial D^i}{\partial p^i} \left(\frac{\partial D^j}{\partial p^j} dp^j + \frac{\partial D^j}{\partial N^i} dN^i \right) \\
dN^i &= \frac{\partial D^i}{\partial p^i} dp^i + \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j} dp^j + \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i} dN^i \\
1 = \frac{dN^i}{dN^i} &= \frac{\partial D^i}{\partial p^i} \frac{dp^i}{dN^i} + \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j} \frac{dp^j}{dN^i} + \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i} \\
1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i} &= \frac{\partial D^i}{\partial p^i} \frac{dp^i}{dN^i} + \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j} \frac{dp^j}{dN^i}
\end{aligned}$$

From this, $\frac{dN^i}{dp^i}$ can be obtained by letting $dp^j = 0$ (this eliminates $\frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j} \frac{dp^j}{dN^i}$), resolving

for $\frac{dp^i}{dN^i}$, and taking the inverse

$$\begin{aligned}
1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i} &= \frac{\partial D^i}{\partial p^i} \frac{dp^i}{dN^i} \\
\frac{\partial D^i}{\partial p^i} \frac{dp^i}{dN^i} &= 1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i} \\
\frac{dp^i}{dN^i} &= \frac{1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i}}{\frac{\partial D^i}{\partial p^i}} \\
\frac{dN^i}{dp^i} &= \frac{\frac{\partial D^i}{\partial p^i}}{1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i}}.
\end{aligned}$$

Analogously, $\frac{dN^i}{dp^j}$ can be obtained by letting $dp^i = 0$, resolving for $\frac{dp^j}{dN^i}$, and taking the inverse

$$\begin{aligned}
1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i} &= \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j} \frac{dp^j}{dN^i} \\
\frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j} \frac{dp^j}{dN^i} &= 1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i} \\
\frac{dp^j}{dN^i} &= \frac{1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i}}{\frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j}} \\
\frac{dN^i}{dp^j} &= \frac{\frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial p^j}}{1 - \frac{\partial D^i}{\partial p^i} \frac{\partial D^j}{\partial N^i}}.
\end{aligned}$$

The infinitesimal expressions $\frac{dN^i}{dp^i}$ and $\frac{dN^i}{dp^j}$ can then be inferred to take the same form.

A.3. Provider Decision Problem

The platform's profit (Rochet and Tirole, 2006, p. 654) is given by:

$$\begin{aligned}
\pi &= (A^B - C^B)N^B + (A^S - C^S)N^S + (a^B + a^S - c)N^B N^S \\
&= \frac{A^B - C^B}{N^S} N^B N^S + \frac{A^S - C^S}{N^B} N^B N^S + (a^B + a^S - c)N^B N^S \\
&= \left(\frac{A^B - C^B}{N^S} + a^B + \frac{A^S - C^S}{N^B} + a^S - c \right) N^B N^S \\
&= (p^B + p^S - c)N^B N^S \\
&= (p^B + p^S - c)n^B(p^B, p^S)n^S(p^B, p^S) \\
&= (p^B + p^S - c)D^B(p^B, N^S)D^S(p^S, N^B) \\
&= p^B D^B D^S + p^S D^B D^S - c D^B D^S
\end{aligned} \tag{3}$$

Here, $N^B N^S$ is the number of potential transactions and thus Rochet and Tirole's proxy variable for the actual volume of transactions, V .

Neglecting the $\frac{\partial D^i}{\partial N^j} \frac{\partial N^j}{\partial p^i}$ terms (in their earlier paper the demands were defined as $D^i(p^i)$,

i.e. not depending on N^j (Rochet and Tirole, 2003, p. 996)), maximization of equation 3 yields the first order conditions:

$$\begin{aligned}\frac{\partial \pi}{\partial p^B} &= D^B(p^B)D^S(p^S) + p^B \frac{\partial D^B(p^B)}{\partial p^B} D^S(p^S) + p^S \frac{\partial D^B(p^B)}{\partial p^B} D^S(p^S) + c \frac{\partial D^B(p^B)}{\partial p^B} D^S(p^S) = 0 \\ &= D^B(p^B)D^S(p^S) + (p^B + p^S - c) \frac{\partial D^B(p^B)}{\partial p^B} D^S(p^S) = 0\end{aligned}$$

$$\begin{aligned}\frac{\partial \pi}{\partial p^S} &= D^B(p^B)D^S(p^S) + p^B \frac{\partial D^S(p^S)}{\partial p^S} D^B(p^B) + p^S \frac{\partial D^S(p^S)}{\partial p^S} D^B(p^B) + c \frac{\partial D^S(p^S)}{\partial p^S} D^B(p^B) = 0 \\ &= D^B(p^B)D^S(p^S) + (p^B + p^S - c) \frac{\partial D^S(p^S)}{\partial p^S} D^B(p^B) = 0.\end{aligned}$$

Therefore

$$\begin{aligned}(p^B + p^S - c) \frac{\partial D^B(p^B)}{\partial p^B} D^S(p^S) &= (p^B + p^S - c) \frac{\partial D^S(p^S)}{\partial p^S} D^B(p^B) \\ \frac{\partial D^B(p^B)}{\partial p^B} D^S(p^S) &= \frac{\partial D^S(p^S)}{\partial p^S} D^B(p^B) \\ \frac{\frac{\partial D^B(p^B)}{\partial p^B}}{D^B(p^B)} &= \frac{\frac{\partial D^S(p^S)}{\partial p^S}}{D^S(p^S)}\end{aligned}\tag{4}$$

Taking the buyer/seller price elasticities of demand,

$$\begin{aligned}\eta^B &= -\frac{\frac{dD^B}{D^B}}{\frac{dp^B}{p^B}} = -p^B \frac{dD^B/dp^B}{D^B} \\ \eta^S &= -\frac{\frac{dD^S}{D^S}}{\frac{dp^S}{p^S}} = -p^S \frac{dD^S/dp^S}{D^S}\end{aligned}$$

and inferring that the infinitesimals take the same form,

$$\eta^i = -p^i \frac{\frac{\partial D^i}{\partial p^i}}{D^i} \quad i = B, S$$

the expression

$$\frac{\eta^i}{p^i} = \frac{\frac{\partial D^i}{\partial p^i}}{D^i}$$

can be substituted into condition 4, obtaining the equilibrium condition (Rochet and

Tirole, 2003, p. 996)

$$\frac{\eta^B}{p^B} = \frac{\eta^S}{p^S}$$

which, together with Lerner's formula (with demand elasticity $\eta = p \frac{\partial V / \partial p}{V}$)

$$\frac{p - c}{p} = -\frac{1}{p \frac{\partial V / \partial p}{V}} = \frac{1}{\eta}$$

$$\frac{p}{p - c} = p \frac{\partial V / \partial p}{V} = \eta$$

provides

$$\frac{1}{p - c} = \frac{\eta}{p} = \frac{\eta^B}{p^B} = \frac{\eta^S}{p^S}.$$

Note that this may be written also in numerous other forms, see e.g. Rochet and Tirole (2006, p. 654). Also note that while this takes the form of a general law, it is identical to equation 4 with merely the definition of the elasticities applied. It would also be exceedingly difficult to verify empirically or computationally as the derivatives with respect to p^i are usually not known.

A.4. The Unique Nash-Equilibrium Between Buyers and Sellers

The buyer and seller demand system (2) is four dimensional with the dimensions being N^B , N^S , p^B , and p^S . The dimensions N^i are endogenous and the p^i exogenous. The network sizes N^i depend mutually on each other and would continue to adjust until a Nash equilibrium is found. In the Nash equilibrium no user¹⁴ has an incentive to unilaterally deviate. Depending on the shape of the curves $N^i(p^i, N^j)$, there may be multiple or no equilibria (for given sets of exogenous p^i). However, Rochet and Tirole exclude this quite real possibility (see figure 9) by stating that "under regularity conditions, the system has a unique solution characterizing memberships N^B and N^S as functions of (p^B, p^S) " (Rochet and Tirole, 2006, p. 653). Unfortunately, it does not become clear what these regularity conditions are, and what kind of distribution for the b^i is required to assure the equilibrium. From a dynamical systems perspective it is also not clear whether a potentially existing equilibrium would be stable, or whether it could be reached from reasonable initial




Figure 9: Determination of critical network sizes of network n^j required for user i to join (denoted n^{j*}) depending on the utility parameters of user i , B^i , b^i , A^i , and a^i .

conditions. We are able to explore these questions with our ABM introduced in section 4.

From the utility function (1) it then follows that for users of type i a critical size of the network j , n^{j*} , can be found for which they are indifferent to join or not to join the network with

$$n^{j*} = \frac{-(B^i - A^i)}{b^i - a^i} \quad (5)$$

where $-(B^i - A^i)$ is the axis intercept of the utility as function of n^j and $(b^i - a^i)$ is its slope (see figure 9). Generally, it would be reasonable to assume the slope to be positive,¹⁵ which means that for all agents, either the critical size n^{j*} is a minimum critical size (right part of the function in figure 9) or it does not exist and the agent will always join the network (left part of the function in figure 9). Since Rochet and Tirole assume that B^i is often negative (Rochet and Tirole, 2006, p. 652) (and imply a non-negative entry fee A^i), we take the first case (there exists a minimum critical value n^{j*}) for all users as a starting point in this example. (Otherwise, the demand functions $D^i(p^j, N^j)$ would simply have a

positive intercept, $D^i(p^j, 0) > 0$.)

B. Analytical treatment of the problem of a unique equilibrium in the Rochet/Tirole model

As mentioned in section 3.5, to obtain a system with a unique and stable equilibrium as assumed by Rochet and Tirole, the network size (demand) functions of the buyers and sellers need to have a concave shape with a high initial slope such that there exists exactly one Nash equilibrium. Here we will provide an exhaustive analytical treatment of the "regularity conditions" that are necessary to ensure the unique and stable equilibrium. These assumptions are *not* discussed and made transparent in the original papers by Rochet and Tirole and are highly dubious in many situations where TSM play a role. This fact substantiates our claim that a sole focus on equilibrium models for the analysis of TSM is misplaced.

The network size (demand) functions are essentially cumulative distribution functions multiplied by the a factor Z^i such that they give absolute values, not shares. If they are to be concave and monotonically increasing ($\frac{\partial f(x)}{\partial x} > 0$, $\frac{\partial^2 f(x)}{\partial^2 x} < 0$) the corresponding probability distribution function must be positive and monotonically falling ($f(x) > 0$, $\frac{\partial f(x)}{\partial x} < 0$). Since this is the distribution function over the minimum critical sizes of network N^j , n^{j*} , the required shape is conveniently produced by the function for n^{j*} (see equation (5) in appendix A) if a simple uniform distribution, $s \sim \mathcal{U}(s_{min}, s_{max})$ is supplied for the per transaction benefits b . We obtain (for $s << s_{max}$)

$$b^i(s) = s$$

$$n^{j*}(s) = \frac{-(B^i - A^i)}{b^i(s) - a^i} = \frac{-(B^i - A^i)}{s - a^i}$$

As noted above, values $s < a^i$ would be self-defeating since in that case, the customer would incur more and more losses with each transaction (and would therefore refrain from conducting any transactions). Let us therefore assume for this computation an $s_{min} = a^i$.³ The resulting function is given in figure 10b; it is a bijective map between b^i and the minimum number of customers of the other type for the present i customer to

participate, n_j^* . Since it is bijective, we can rearrange parts of the map such that we have the part which results in participation, $a^i \leq b^i(=s) \leq s_{max}$ first and sort it to become a monotonically increasing function (i.e., effectively turn it around since it is monotonically decreasing in figure 10b), yielding figure 10c, or, functionally:


$$n^{j*}(s_{max} - s) = \frac{-(B^i - A^i)}{(s_{max} - a^i) - b^i(s)} = \frac{-(B^i - A^i)}{(s_{max} - a^i) - s} \quad (6)$$

Figure 10c could now also be read the other way around (s as dependent on n^{j*}), which now gives the exact part of the distribution that will join the network if n^{j*} customers of the other type also participate. That gives the demand function $D^i(n^j)$.⁴ Functionally, this is the inverse of equation 6,


$$D^i(n^j) = s(n^{j*}) = s_{max} - a^i + \frac{B^i - A^i}{n^j},$$

as depicted in figure 10d. In this case it results in two equilibria, one of which is stable. Note that increasing a^i shifts the right (stable) equilibrium leftward, closer to the other one.


As discussed in the main text, there are many situations in which these assumptions would be violated. Figure 11 shows the example discussed in section 3.5 in which we have two groups of customers, some who use their credit cards all the time, and other who use them only for particular purposes. This plausible situation results in a system of demand functions with four equilibria two of which would be stable.




(a) Distribution of b^i



(b) Critical values n^{j*} as a function of b^i (with constant B^i, A^i, a^i)




(c) Critical values n^{j*} in increasing order




(d) Resulting system of demand functions $n^i = D(n^j), n^j = D(n^i)$


Figure 10: Demands and equilibria resulting from uniformly distributed b^i (and constant B^i)




(a) Distribution of b^i



(b) Critical values n^{j*} as a function of b^i (with constant B^i, A^i, a^i)



(c) Critical values n^{j*} in increasing order



(d) Resulting system of demand functions $n^i = D(n^j), n^j = D(n^i)$

Figure 11: Demands and equilibria resulting from a two-peaked distribution of b^i (and constant B^i)

References

- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3):668–691.
- Arthur, W. B., Ermoliev, Y. M., and Kaniovski, Y. M. (1987). Path dependent processes and the emergence of macro-structure. *European Journal of Operational Research*, 30:294–303.
- Bendor, J., Mookherjee, D., and Ray, D. (2001). Reinforcement learning in repeated interaction games. *The B.E. Journal of Theoretical Economics*, 1(1):Article 3.
- Choi, J. P. (2010). Tying in two-sided markets with multi-homing. *The Journal of Industrial Economics*, 58(3):607–626.
- David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2):332–337.
- Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137.
- Katz, M. L. and Shapiro, C. (1985). Network externalities, competition and compatibility. *American Economic Review*, 75(3):424–440.
- Koloseni, D., Lampinen, J., and Luukka, P. (2013). Differential evolution based nearest prototype classifier with optimized distance measures for the features in the data sets. *Expert Systems with Applications*, 40(10):4075 – 4082.
- Liu, J. and Lampinen, J. (2005). A fuzzy adaptive differential evolution algorithm. *Soft Computing*, 9(6):448–462.
- Meyer, T. G. (2012). *Path Dependence in Two-sided Markets: A Simulation Study on Technological Path Dependence with an Application to Platform Competition in the Smartphone Industry*. PhD thesis, Free Uni-

- versity of Berlin, Berlin, Germany. Available online http://www.diss.fu-berlin.de/diss/servlets/MCRFileNodeServlet/FUDISS_derivate_000000012507/Meyer_2012_Path_dependence_in_two_sided_markets.pdf.
- Peitz, M., Rady, S., and Trepper, P. (2011). Experimentation in two-sided markets. CEPR Discussion Paper No. DP8670. Available at SSRN: <http://ssrn.com/abstract=1965918>.
- Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029.
- Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3):645–667.
- Shy, O. (2011). A short survey of network economics. *Review of Industrial Organization*, 38(2):119–149.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118.
- Storn, R. and Price, K. (1997). Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359.

Notes

¹More extensive lists of examples are given in, e.g. Rochet and Tirole (2003) or Meyer (2012).

²For a brief historical overview, see section 5 in Shy (2011).

³ a^i is the transaction fee set by the provider, shifting a^i into the distribution would just remove a part of the distribution, so we would for the present computation be interested in just the remaining part, $\mathcal{U}(a^i, s_{max})$.

⁴Note that this would have to be normalized by the range of the uniform distributed variable s ; the shape of the function would, however, not change.

⁵Rochet and Tirole's model also allows for a membership return B_i . We have this attribute included in our model but set it to zero as it is done by Rochet and Tirole for most cases considered by them.

⁶This is due to the observation that providers typically do not interfere with the terms of the transaction between their customers. Consequently, this is modeled as buyer's and seller's reservation prices for the transaction. Note that this results in a slight change compared to the case with uniformly distributed b^i that was found to approximate the Rochet-Tirole case of the unique Nash-equilibrium in the demand functions. Allowing transfers between b^B and b^S results in b^i being distributed \mathcal{U}^2 .

⁷We do not set this probability to 100% because in the case of only one provider, this would enable the provider to exploit the agents since they would be forced to join her network again as soon as they left it in order to avoid the associated losses.

⁸The parameters of the differential evolution are chosen in accordance with the literature (Liu and Lampinen, 2005; Koloseni et al., 2013).

⁹See Bendor et al. (2001) for a discussion of this formal approach and its historical development.

¹⁰Price corrections use a slow geometric moving average over past and current target values. Corrections for the different kinds of fees are assessed separately (i.e. they may move into different directions at the same time).

¹¹This is to avoid vastly negative revenues caused by negative prices or entry fees. This is necessary as the success of a strategy in step 2 is evaluated in relative terms compared to the last round; negative but slightly increasing revenue would therefore be deemed to indicate a successful strategy.

¹²A (multiplicative) combination of the revenue and number of customers in the last round can also be used; it will lead to very similar results.

¹³Note that for the case without satisficing, the buyer revenue is sometimes negative but this again is a result of (by chance occurring) dysfunctional networks with exceptionally high prices that are present in a minority of runs but have an impact on the average.

¹⁴And since the users react mechanically, the number of users could game theoretically be seen as strategies of a single buyer-agent and a single seller-agent reacting to each other's strategies.

¹⁵Otherwise, the agent's utility would fall with the number of transactions which, in turn, would raise the question, why she would then undertake these transactions and why she would join the network in the first place, though she might do that because of a generous transfer in the form of a negative access price A^i .