Adoption of a New Payment Method: Experimental Evidence*

Jasmina Arifovic† John Duffy‡ Janet Hua Jiang§

Revised Draft: February, 2023

Abstract

We develop a framework for studying the introduction of a new payment method in a controlled laboratory environment, where consumers (buyers) and merchants (sellers) can learn to coordinate their adoption decisions over time. The underlying game exhibits network adoption effects as emphasized by the theoretical literature. We elicit players' beliefs about the adoption decisions of the other side of the market so that we can directly test for network effects. We investigate how the additional fixed cost of adopting the new payment method, relative to its savings on per transaction costs, affects merchant's decisions to adopt the new payment method and how that in turn affects buyer's adoption decisions. We find that a low fixed cost favors quick adoption of the new payment method by all participants, while for a sufficiently high fixed cost, merchants gradually learn to reject the new payment method. We also find strong evidence of network effects and that the fixed costs are important for the strong response of seller acceptance decisions to buyer adoption decisions. An evolutionary learning model provides a good characterization of the dynamic adjustment paths found in our experimental data.

JEL classification numbers: E42, D14, D83, C72, C92.

Keywords: Payment methods, network effects, experimental economics.

*We thank the Editor and two anonymous referees for insightful comments and suggestions. We also thank Gabriele Camera, Kim P. Huynh, Gerald Stuber, Leo Van Hove, as well as participants at the Canadian Economic Association Meetings, the International Workshop on Theoretical and Experimental Macroeconomics, the Workshop on Theory and Experiments in Monetary Economics, the Economic Science Association World Meetings, the Econometric Society World Congress, the Conference on Computing in Economics and Finance, the Conference on Heterogeneous Agents and Agent-Based Modeling: the Intersection of Policy and Research, and seminars at Chapman University, Purdue University, and the University of California-Santa Barbara. Matthew Pentecost and William Ho provided excellent research assistance. We gratefully acknowledge funding support from the Bank of Canada. The views expressed in this paper are those of the authors and should not be attributed to the Bank of Canada.

†Department of Economics, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada. Email: arifovic@sfu.ca.

‡Corresponding author. Department of Economics, University of California, Irvine, 3151 Social Science Plaza, Irvine, CA 92697-5100. Email: duffy@uci.edu. Tel: 1-949-824-8341.

§Currency Department, Bank of Canada, 234 Wellington Street, Ottawa, ON, K1A 0G9, Canada. Email: jjiang@bankofcanada.ca.

1 Introduction

The payments industry has undergone significant changes in the past few decades. Alongside traditional means of payments such as cash and checks, various new means of payments have entered the payment landscape. Examples include debit and credit cards; general purpose pre-paid cards such as Visa and MasterCard gift cards; public transportation cards that have expanded into a method of making retail payments such as the Octopus card (in Hong Kong); online money transfer schemes such as Paypal and Swish (in Sweden); mobile payments such as Venmo (in the U.S.) and WeChat Pay (in China); and virtual crypto-currencies such as Bitcoin and Ethereum. In addition, new technologies are also being incorporated into existing payment methods to generate new ways to pay; examples include contactless debit and credit cards, Android Pay and Apple Pay. Recently, many central banks around the world have begun discussing or testing the issuance of another new payment method, central bank digital currency, or CBDC (see Boar and Wehrli 2021).

In this paper we propose a framework for studying the adoption of a new payment method in a controlled laboratory environment. In our framework, the "new" payment method is new in the sense that it has yet to be well-established as a widespread or default payment method, and not necessarily that it has only recently been introduced.[1] Buyers must decide on whether and how to allocate income to that new payment method and sellers must decide whether to accept it for payment so that both the use and acceptance of the new payment method is a more speculative act as compared with the "old", default payment method, for example cash. While the choice between cash and (say) electronic payments is one way of thinking about the choice between old and new payment methods, a more future-oriented example would be the choice between electronic payments (old) and CBDC (new), which will also involve an acceptance decision by sellers and a portfolio choice by buyers.

We view our experimental approach as an important and useful complement to theoretical modeling and empirical studies of the adoption of new payment methods. The theoretical literature on retail payments emphasizes network adoption effects, which arise because for the consumer (merchant), the benefits of adopting (accepting) a new payment method increase only if more merchants accept (more consumers use) that payment method. However, the existence of network externalities leads to multiple equilibria, which poses a problem for making theoretical predictions. Since equilibrium selection is ultimately an empirical question, our experimental study sheds light on which equilibrium is more likely to obtain under certain conditions.[2] In addition, while the theory often focuses on equilibrium analysis, our experimental approach sheds light on the *process* through which an equilibrium is reached in addition to the payoff considerations necessary for a new payment method to take off and become fully adopted.[3]

[1] We might also refer to the new payment method as an "alternative" or "non-default" payment method since in our setup the new payment method is always available as an alternative to the old default payment method.

[2] For theories on retail payments, see Rochet and Tirole (2002), Wright (2003), McAndrews and Wang (2012), Chiu and Wong (2014); Lotz and Zhang (2016); Li et al. (2020); and Lotz and Vasselin (2019). See Hall and Kahn (2003) for a more general discussion of the process of new technology adoption.

[3] Li et al. (2020) develop a dynamic model of payments system evolution, where the dynamics are driven by exogenous changes in consumer incomes and adoption costs and the payment provider (network)'s endogenous R&D investment. In

There is a large empirical literature that uses survey data to explore individual choices among different means of payments.[4] Limited by the data available, these studies focus mainly on the decisions made by just *one side* of the payment system, either the consumer side or the merchant side. Consumer and merchant surveys are rarely conducted concurrently and, because of high costs, are usually run at a low frequency, e.g., every few years. Merchant surveys are scarcer because it is more difficult to recruit merchants to participate in surveys. Therefore, it is a challenging task to study the feedback effects between the two sides of the payment system. The feedback from merchants to consumers has been made relatively easier thanks to improvements to recent consumer surveys that include questions on payment instruments accepted at the point of sale. Nonetheless, as pointed out in Huynh et al. (2014), special care must be taken to handle the endogeneity of card acceptance: consumers' choice of vendors may depend on the cash they have on hand. The empirical evidence on feedback in the other direction, from consumers to merchants, is mostly indirect and relies on a combination of consumer and merchant surveys conducted in different years (see Bounie et al., 2017a, for a study using French surveys; and Huynh et al., 2019, for a study using Canadian surveys).[5]

Complementing theoretical and empirical studies, we build an experimental framework to study how payment adoption decisions on both sides of the market co-evolve dynamically over time. The underlying payment adoption game captures the network adoption effects as emphasized by the theoretical literature, and, as a consequence, our network adoption game has multiple equilibria. In addition to exploring payment adoption choices, we elicit players' *beliefs* about market conditions so that we can directly investigate the feedback effects between the two sides. For example, we can examine how buyers' payment adoption decisions depend on their beliefs about sellers' decisions, and how buyers' beliefs respond to sellers' acceptance choices in the past. We also study feedback in the other direction: how sellers' acceptance decisions depend on their beliefs about buyers' choices as well as acceptance decisions by other sellers. The latter channel is especially important because it is difficult to test with field data.

The varied features of new payment methods call for a broad research agenda. A new payment method can take many forms and there are a myriad of potential factors that can affect its adoption. It can be challenging to clearly identify the factors that contribute to the success or failure of a particular case. In this regard, the advantage of a laboratory study is that there is better control over the environment

our experiment, subjects interact in a stationary physical environment, but their beliefs and choices evolve as they learn to coordinate with one another over time.

[4]For example, Arango et al. (2016) conduct an analysis using the Bank of Canada 2009 Method of Payment Survey, Koulayev et al. (2016) use the Federal Reserve Bank of Boston 2008 U.S. Consumer Payment Choice Survey, and von Kalckreuth et al. (2014) use the Deutsche Bundesbank 2008 Payment Habits in Germany Survey. Recently, Bagnall et al. (2016) studied consumers' use of cash by harmonizing payment diary surveys from seven countries.

[5]Survey data can be subject to errors due to insufficient incentives for truthful or careful reporting, or misunderstandings about the questions posed; these problems are alleviated to some extent in an incentivized experimental study where subjects are quizzed on their understanding of the rules prior to play and are paid based on their performance in accordance with the model's payoff structure. In addition to studies using survey data, other research – including Klee (2008), Cohen and Rysman (2013), and Wang and Wolman (2016) – use scanner data to study payment choices at the point of sale. Compared with survey data, scanner data are less prone to errors and misreporting. However, scanner data are less useful for the study of payment adoption.

enabling isolation of the relevant factors. In this paper, we explore the effects of the cost structure of new, prepaid payment schemes, and its interaction with network adoption effects. A potential limitation of our analysis is that we only focus on *prepaid schemes* that are denominated in the national currency. Many new payment methods are consistent with this restriction including the Octopus card, Venmo and We-Chat pay. CBDC will also most likely adopt this same type of scheme (see Kumhof and Noone 2018). Still, for simplicity, our new payment method does not include the use of non-prepaid schemes such as credit cards or payment methods not denominated in the national currency that could be subject to exchange rate risk, e.g., cryptocurrencies.

The cost structure of such payment schemes is obviously a crucial factor that affects adoption, and we attempt to capture two reasonable cost features manifested by new payment methods in the field. The first is that new payment methods are more efficient for both consumers and merchants in terms of per transaction costs. Such a cost-saving motive lies behind the various attempts to introduce all new payment methods: after all, if a new payment method did not offer any cost savings, there would be no reason to expect it to be adopted (or even introduced). There is also evidence that new developments in electronic payment technology can greatly speed up transactions and save on various handling costs associated with traditional payment methods.[6]

Second, as a natural force that works against the adoption of a new payment method, sellers must incur a fixed adoption or startup cost (besides per transaction costs) in order to process the new additional payment method to cover costs for renting a terminal, managing the system, and training employees in the use of the new payment system. The evidence suggests that these fixed startup costs can be substantial and play a pivotal role in adoption of a new payment method.[7] For instance, a survey on the costs of point-of-sale payments in 2014 by the Bank of Canada (see Kosse et al. 2017) suggests that the fixed costs of debit card processing by merchants average about 20% of the total cost (which includes fixed costs, transaction-linked variable costs and value-linked variable costs). This fixed cost is even more important for small retailers, which tend to have a smaller number of transactions and could partly explain the relatively low debit card acceptance rate among small and medium-sized retailers (68%) relative to larger retailers (97%). In the case of Mondex and Visa Cash in the 1990s, Bátiz-Lazo and Moretta (2016) suggest that the cost of setting up a chip-based acceptance infrastructure to secure nationwide universality was the main hurdle to widespread acceptance and usage.[8]

[6]According to a study by Polasik et al. (2013), who analyze the speed of various payment methods from video material recorded in the biggest convenience store chain in Poland, a transaction using contactless cards in offline mode without slips costs on average 25.71 seconds, a significant reduction over a cash transaction, which costs 33.34 seconds. More recently, the Bank of Canada conducted a similar time duration study. Based on 5,891 recorded transactions, the study finds that contactless debit and credit card transactions take about 15 seconds, whereas these could extend to almost 26 seconds when using the chip and PIN or the swipe technology (see Kosse et al. 2017).

[7]Consumers also bear some fixed adoption costs, such as acquiring payment balances. However, the cost borne by consumers is relatively small compared with the cost borne by merchants. For simplicity and better control, we assume that buyers do not bear any fixed cost to adopting the new payment method in our experiment.

[8]Li, McAndrews and Wang (2019) study the payments system evolution with a dynamic model in a two-sided market setting and calibrate their model to U.S. payment card data. A substantial fixed merchant adoption cost (about \$500 in 1997) is needed for their model to match the card adoption pattern in the data.

Given this cost structure of the two payment methods, buyers and sellers in our environment play a two-stage game. In the first stage, both agent types make simultaneous payment decisions. Buyers allocate their budget (endowment income) between the existing and the new payment methods. The portfolio choice of buyers is a realistic design feature. For example, buyers need to allocate payment balances between cash and bank account balances when choosing between cash and debit card payments. In the future, they may need to allocate between bank and CBDC account balances or to choose between debit card and CBDC payments. On the other side of the payment market, sellers decide whether or not to pay a fixed cost to accept the new payment method (they always accept the existing payment method as a result of custom or due to legal restrictions).[9] The second stage consists of multiple rounds of meetings where buyers and sellers trade with each other. In each meeting, trade is successful if the buyer can use a payment method accepted by the seller.

Due to network effects, the game underlying the experimental setup admits two symmetric pure strategy Nash equilibria. In one of these, the new payment method is not adopted and all transactions continue to be carried out using the existing payment method. In the other symmetric equilibrium, the new payment method is adopted and completely replaces the existing payment method. Our treatment variable is the magnitude of the fixed cost of adoption (relative to the saving on per transaction cost), and we assess how it affects the adoption outcome and network effects.[10]

Our experiment reveals that, depending on the fixed cost of adoption of the new payment method and on the choices made by participants on both sides of the market, either symmetric equilibrium can be selected. More precisely, if sellers face a low fixed cost of adopting the new payment method, then the new payment method is quickly adopted by all participants, whereas for a sufficiently high fixed adoption cost, sellers gradually learn to refuse to accept the new payment method and transactions are increasingly conducted using the existing payment method. We also test directly the feedback effects between the two sides of the market. We find that buyers' adoption decisions depend on their beliefs about sellers' adoption decisions and these beliefs depend on historic outcomes. Similarly, we find a feedback effect running in the other direction, especially for more substantial values of the fixed adoption cost.

Finally, beyond observing dynamic behavior, we also model the dynamic process of equilibrium selection using an evolutionary learning model that has been used to explain experimental results in other contexts. We find that this model provides a good fit to the experimental data we collected for the original three experimental treatments of our design. Indeed, the impressive fit of this learning model to the experimental data motivated us to use that model to design and predict outcomes for a new (fourth) experimental treatment. We then carried out that additional experimental treatment and again report a good fit between the evolutionary model and the experimental data. Thus, a further innovation of this paper is to demonstrate that the evolutionary learning model provides a good characterization of buyers

[9] Note that there may be fixed costs associated with the existing payment method as well. However, the decision is about whether to incur the *additional* fixed cost to accept the new payment method besides the existing method.

[10] In the experiment, we fix the savings that the new payment method provides on per transaction costs relative to the existing payment method and we vary the magnitude of the fixed adoption cost for the new payment method. In all of our treatments, the theory admits multiple equilibria.

and sellers' payment choices so that simulations of that model can be used for *design* purposes. Based on that success, we feel confident enough to use our evolutionary model to predict what would happen over longer periods of time, beyond the time frame allowed by a human subject experiment.

The closest paper to this one is by Camera et al. (2016), who develop a model of payment choice between cash and e-money/cards and also conduct an experimental study. While we take inspiration from Camera et al. (2016), our project differs from theirs along several dimensions. In terms of experimental design, one major innovation is that we elicit players' beliefs about the likely behavior of participants on the other side of the market. This belief elicitation allows us to directly test the network adoption effects, filling a gap in the analysis using survey data. In terms of the explanatory factor of interest, our paper examines the effects of the fixed adoption costs borne by sellers relative to the potential savings earned on per transaction costs. Camera et al. (2016) study the effects of proportional seller fees and buyer rewards. In terms of experimental results, Camera et al. (2016) find there is little feedback effect between the two sides of the market, whereas we find the opposite. The fixed seller adoption cost in our setting could make sellers respond more strongly to buyers' payment allocation decisions. The continuous buyer portfolio choice could make buyers respond in a more flexible way to seller adoption. Other design features, like more readily available market-level information and less complex tasks, could also be factors affecting adoption (see the more detailed discussion following our experimental results in Section 5). Finally, in terms of explaining the experimental data, we take a further step and demonstrate that an evolutionary learning model explains the dynamics in our experimental data well and is useful for experimental design purposes.

More broadly, our study is related to a large body of experimental studies of coordination games that involve multiple equilibria following initial studies by Van Huyck et al. (1990, 1991); see Devetag and Ortman (2007) for a survey.[11] However, payment choices have their own peculiarities that require a more tailored setting. First, most coordination games studied in the laboratory involve players all facing the same decision involving the same choice set (i.e. players are all on the same side). In the payment adoption game, consumers and merchants comprise two different sides of the market and face different costs of adopting the different payment options. Further the consumers make a continuous portfolio choice, whereas the merchant makes a zero-one acceptance choice. There is coordination both between the two sides and among players on the same side.[12] In addition, the framework developed here is more suitable for studying policy experiments specific to payments, e.g., the roles of subsidies and surcharges.

The rest of the paper is organized as follows. Section 2 describes the underlying game and equilibria. Section 3 describes our experimental design. The aggregate experimental results are presented in section 4. Individual buyer and seller behavior is examined in section 5. In section 6 we present an evolutionary

[11] In recent years, this literature has expanded into many areas such as bank runs (See Dufwenberg 2015 for a review), currency crises (Heinemann, Nagel and Ockenfels 2004), poverty traps (Lei and Noussair 2002), emergence of money as a medium of exchange (Duffy and Ochs 1999, 2002).

[12] Hossain, Minor and Morgan (2011) carry out an experimental study of platform competition that involves two sides of the markets. They find that when platforms are primarily vertically differentiated, markets inevitably tip to the more efficient platform even when platform coexistence is theoretically possible. They find strong evidence of equilibrium coexistence when platforms are primarily horizontally differentiated, so there is no single efficient platform.

learning model that can closely track our experimental findings and we use this model to design and predict behavior in a further experimental treatment. Finally, section 7 concludes with a summary and some directions for future research.

2 The Payment Adoption Game

In this section, we describe a simple model of the adoption of a new payment method that underlies our experimental study. Each market (or trading cycle) consists of two stages. In the first stage, buyers make a portfolio decision, splitting their budget between the new and the existing (old) payment methods. Simultaneously, sellers decide whether or not to accept the new payment method; the old payment method must always be accepted. In the second stage, buyers meet pairwise with sellers and engage in transactions involving either payment form, depending on the buyer's portfolio and the seller's acceptance of payment 2. For simplicity and ease of experimental implementation, we assume homogeneous buyers and sellers, an exogenously given spending budget for the buyer and fixed terms of trade.

We model the new payment method as being more efficient for both buyers and sellers in terms of per transaction costs, but sellers must pay a fixed cost to process it. As in other models of payment competition, the model we adopt for our experimental study has multiple equilibria as it involves simultaneous moves by both buyers and sellers. By contrast, in a sequential move game, the multiplicity of equilibrium problem might be attenuated as the first moving side, buyers or sellers, would have greater power to select the payment regime. However, the coordination problem among agents on the same side of the market as to whether or not they would accept a new payment method would remain.

2.1 Physical Environment

There is a large number of buyers (consumers) and sellers (merchants) in the market, each of unit measure. Each seller $i \in [0, 1]$ is endowed with a technology that allows them to costlessly produce units of good i. The seller derives zero utility from consuming his/her own good and instead tries to sell his/her good to buyers. The per unit price of the good is fixed at 1, which is also the seller's utility gain from each transaction.[13] In each period, each buyer $j \in [0, 1]$ visits all sellers once in a random order, and would like to purchase and consume one and only one unit of each good produced by each seller. The buyer is endowed with just enough income to make the desired purchases. The buyer's utility from consuming each good is u.

There are two payment instruments: payment 1, which is the existing payment method, and payment 2, which is a new payment method. Each transaction using the existing payment method 1 incurs a cost to buyers denoted by τ_1^b, and a cost to sellers denoted by τ_1^s,. Similarly, the per transaction costs for use of the new payment method 2 are denoted by τ_2^b and τ_2^s for buyers and sellers, respectively. Sellers must

[13]The use of exogenous prices greatly simplifies our analysis, but a more realistic setting would allow such prices to be endogenously determined enabling sellers to pass through payment costs to buyers.

pay a fixed cost, $F > 0$, that enables them to accept the new payment method, for example, to rent a terminal to process transactions using the new payment method.

Buyers and sellers play a two-stage game. In the first stage, sellers decide whether or not to accept the new payment method at the one-time fixed cost of F. Payment method 1, being the traditional (and legally recognized) payment method, is universally accepted by all sellers. Simultaneous with the sellers' decision, buyers make a portfolio choice as to how to divide their income endowment between the two payment methods. After sellers have made their acceptance decisions and buyers have made their portfolio decisions, the second stage starts. In the second stage, buyers go shopping visiting all stores (sellers) in a random order. When a buyer enters store i, s/he observes the seller's payment acceptance choice. The buyer then chooses which payment method to use to purchase one unit of good i, subject to having both methods available; if the buyer has only one payment method available, then the buyer can only offer that payment method to the seller. A trade is successful only if the means of payment the buyer has available and offers to the seller is accepted by the seller. Otherwise, there is no trade. At the end of the game, sellers spend their money balances on a general good. One unit of the general good costs one dollar and entails one unit of utility. For simplicity, we assume buyers do not wish to consume the general good and any unspent money does not yield them any extra utility; we also assume that the sellers do not value their own unsold goods. The model's implications do not hinge on these assumptions.

As a result, the payoff for the buyer is: number of payment 1 transactions $\times$ $(u - \tau_1^b)$ + number of payment 2 transactions $\times$ $(u - \tau_2^b)$; and the payoff for the seller is: number of payment 1 transactions $\times$ $(1 - \tau_1^s)$ + number of payment 2 transactions $\times$ $(1 - \tau_2^s)$ - F if accepting payment 2.

In what follows, we make the following four assumptions about costs: (1) $\tau_2^b < \tau_1^b$ and $\tau_2^s < \tau_1^s$; (2) $u - \tau_1^b > 0$; (3) $F \leq \tau_1^s - \tau_2^s + \tau_1^b - \tau_2^b$; and (4) $F \leq 1 - \tau_2^s$. Assumption (1) states that the new payment method saves on per transaction costs for both buyers and sellers. Under Assumption (2), buyers prefer trading with the existing payment method to not trading. Assumption (3) implies that the net benefit of investing in the ability to process the new payment method is positive for society if all transactions are carried out in the new payment method. Assumption (4) ensures the existence of an equilibrium where only the new payment method is used.

2.2 Equilibrium

We solve the model by backward induction, first examining the payment usage choice in the second stage and then the portfolio and payment acceptance choices in the first stage. We will focus on symmetric equilibria, where all buyers make the same portfolio choice decision and all sellers make the same acceptance decision. Let $0 \leq m^b \leq 1$ be the balance allocated to the new payment method by the buyer, and let $0 \leq m^s \leq 1$ denote the fraction of sellers who accept the new payment method. If $0 < m^s < 1$, then sellers play a mixed strategy, accepting the new payment method with probability m^s.

The buyer payment usage choice in the second stage is straightforward. If the seller accepts payment 2, then the buyer uses payment 2 if s/he has a positive payment 2 balance (as payment 2 involves a lower per transaction cost), and uses payment 1 if s/he has only payment 1 left. If the seller does not accept payment 2, then the buyer uses payment 1 if s/he has a positive payment 1 balance, and does not trade if s/he runs out of payment 1 (even if s/he has a positive payment 2 balance). In other words, in each meeting, the buyer first tries to use payment 2 if possible, then payment 1, and does not trade if neither payment is viable.

To solve the payment adoption problem in the first stage, we derive the best response of buyer's portfolio decision to seller's acceptance decision, and vice versa. The equilibrium is then given by the intersection of the two best response curves. We summarize the best response functions in Propositions 1 and 2, illustrate them in Figure 1 and discuss the intuition. The detailed derivation and proofs are in online Appendix A.

Proposition 1 *The buyers' optimal portfolio decision is to mimic the sellers' acceptance decision :*

$$m^b(m^s) = m^s.$$

If $m^b > m^s$, then each buyer makes m^s purchases using the new payment method and $1 - m^b$ purchases using the existing payment method. Buyers are not able to transact with a fraction $m^b - m^s$ of sellers because of payment mismatches (buyers want to use payment method 2 but sellers accept only payment method 1). The optimal choice of each buyer is to reduce m^b to m^s so as to minimize the probability of no-trade outcome. If $m^b < m^s$, then each buyer makes m^b transactions using the new payment method and $1 - m^b$ transactions using payment method 1 (among which $m^s - m^b$ are with sellers who also accept the new payment method). The optimal choice of each buyer is to increase their payment 2 balances to m^s so as to take advantage of the lower per transaction cost associated with payment 2 without increasing the risk of no trading.

Proposition 2 *The seller's best response $m^s(m^b)$ depends on the parameterization of F versus $\tau_1^s - \tau_2^s$ and is described as follows.*

1. *If $F < \tau_1^s - \tau_2^s$, then*

$$m^s(m^b) = \begin{cases} \frac{m^b(\tau_1^s - \tau_2^s)}{F} & \text{if } m^b \leq \frac{F}{\tau_1^s - \tau_2^s}, \\ 1 & \text{if } m^b \geq \frac{F}{\tau_1^s - \tau_2^s}. \end{cases}$$

2. *If $F > \tau_1^s - \tau_2^s$, then*

$$m^s(m^b) = \begin{cases} 0 & \text{if } m^b \leq 1 - \frac{(1-\tau_2^s) - F}{1 - \tau_1^s}, \\ 1 - \frac{(1-m^b)(1-\tau_1^s)}{(1-\tau_2^s) - F} & \text{if } m^b \geq 1 - \frac{(1-\tau_2^s) - F}{1-\tau_1^s}. \end{cases}$$

3. *If $F = \tau_1^s - \tau_2^s$, then*

$$m^s(m^b) = m^b.$$

For low fixed cost, or $F < \tau_1^s - \tau_2^s$, the seller's best response curve lies above the 45 degree line ($m^b < m^s$). Along the 45 degree line, there is no payment mismatch in the sense that all sellers who accept payment 2 trade with payment 2 in all meetings, and all those who do not accept payment 2 trade with payment 1 in all meetings. Given $F < \tau_1^s - \tau_2^s$, accepting payment 2 gives a higher payoff than rejecting it if $m^b > m^s$; as a result, more than m^b sellers would want to accept payment 2. Second, the seller's best response curve has a kink at $m^b \le F/(\tau_1^s - \tau_2^s)$. If buyers allocate $m^b \le F/(\tau_1^s - \tau_2^s)$ to payment 2, the number of transactions using payment method 2 is not large enough to recover the fixed acceptance cost for all sellers. As a result, sellers play a mixed strategy: $m^s = m^b(\tau_1^s - \tau_2^s)/F$ fraction of sellers accept both payment methods, the rest accept only payment method 1, and all sellers earn the same expected payoff. If $m^b \ge F/(\tau_1^s - \tau_2^s)$, then it is a dominant strategy for sellers to accept the new payment method: each seller makes more than $F/(\tau_1^s - \tau_2^s)$ sales in payment 2 to warrant the fixed investment to accept payment 2.

For high fixed cost, or $F > \tau_1^s - \tau_2^s$, the seller's best response curve lies below the 45 degree line. Along the 45 degree line, every seller accepting payment 2 transact with payment 2 and every seller not accepting payment 2 transact with payment 1 in all meetings. Given $F > \tau_1^s - \tau_2^s$, accepting payment 2 leads to a lower payoff than rejecting it. Again, the seller's best response curve has a kink. It is a dominant strategy for sellers not to accept the new payment method if buyers allocate too little money to payment 2, or $m^b \le \hat{m}_b \equiv 1 - [(1 - \tau_2^s) - F]/(1 - \tau_1^s)$. If $m^b \ge \hat{m}_b$, then sellers play a mixed strategy, choosing to accept with probability $m^s(m^b) = 1 - (1 - m^b)(1 - \tau_1^s)/[(1 - \tau_2^s) - F]$, at which the expected payoff from accepting and rejecting payment 2 is equal.

Finally, note that when $F = \tau_1^s - \tau_2^s$, then the seller's best response lies on the 45 degree line, along which sellers accepting payment 2 and not accepting payment 2 earn the same expected payoff.

Combining the analysis above, we can characterize the symmetric equilibrium of the economy using Figure 1. There are at least two symmetric pure strategy equilibria. In one of these equilibria, $m^b = m^s = 1$: all sellers accept the new payment method, and all buyers allocate all of their endowment to the new payment method – call this the all-payment-2 equilibrium (this equilibrium always exists provided that $F \le 1 - \tau_2^s$). There is a second symmetric pure strategy equilibrium where $m^b = m^s = 0$ and the new payment method is not accepted by any seller or held by any buyer – call this the all-payment-1 equilibrium. In both equilibria, there is no payment mismatch, and the number of transactions is maximized at 1. In the case where $F = \tau_1^s - \tau_2^s$, there exists a continuum of possible equilibria in which $m^s \in (0, 1)$ and $m^b = m^s$. Corresponding to the symmetric equilibria where each seller chooses to accept the new payment with probability m^s, there are also *asymmetric* pure-strategy equilibria, where m^s fraction of sellers accept the new payment method whereas the remaining fraction do not, and where buyers best respond to this asymmetry in their payment allocation decisions. The asymmetric equilibria render the same payoffs as the corresponding symmetric equilibria with mixed strategies.

Note that buyers are always better off in the all-payment-2 equilibrium relative to the all-payment-1 equilibrium. The seller's relative payoff in the two equilibria, however, depends on the fixed cost, F and on the savings on per transaction costs from the use of payment 2. If $F = \tau_1^s - \tau_2^s$, then the seller's payoff

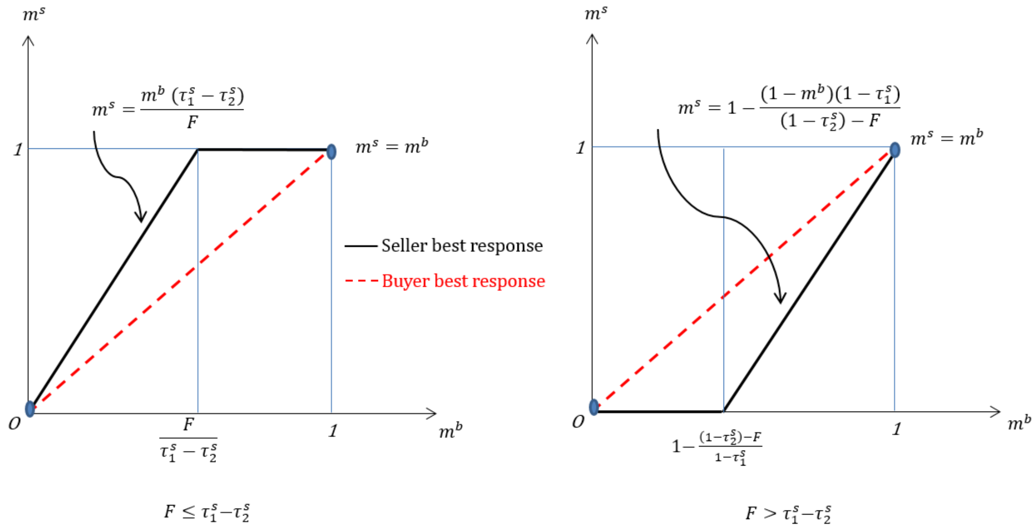

Figure 1: Best response functions of buyers and sellers

Notes. (1) Intersections of response curves reveal the equilibria. (2) Note that when $F = \tau_1^s - \tau_2^s$, the two best response functions coincide and there exists a continuum of possible equilibria in which $m^s \in (0, 1)$ and $m^b = m^s$.

is the same in both equilibria; if $F < \tau_1^s - \tau_2^s$, then the seller's payoff is higher in the all-payment-2 equilibrium than in the all-payment-1 equilibrium; finally, if $F > \tau_1^s - \tau_2^s$, then the seller's payoff is lower in the all-payment-2 equilibrium than in the all-payment-1 equilibrium.

While both the pure strategy, all-payment-1 or all-payment-2 equilibria always coexist in our set-up, we note that under simple, naïve best response dynamics for nearly all initial conditions, the all-payment-2 equilibrium would be the attractor in the case where $F \leq \tau_1^s - \tau_2^s$ (left panel of Figure 1) while the all-payment-1 equilibrium would be the attractor in the case where $F \geq \tau_1^s - \tau_2^s$ (right panel of Figure 1). Of course, it remains to be seen whether subjects behave according to such dynamics.

3 Experimental Design

The experiment was designed to match the model as closely as possible, but without the continuum of buyers and sellers of unit mass. For each session of our experiment, we recruited 14 inexperienced subjects and randomly divided them up equally between the buyer and seller roles, so that each market had exactly seven buyers and seven sellers.[14] These roles were fixed for the duration of each session to enable subjects to gain experience with a particular role. The subjects then repeatedly played a market game that approximates the model presented in the previous section. See the experimental instructions in online Appendix B.

[14]Tacit collusion among groups of this size is difficult, see, e.g., the weak-link experiment of Van Huyck et al. (1990) involving 14-16 subjects. Social or other-regarding preferences are also less likely to play a role in such a large group of subjects, see, e.g., Roth et al. (1991).

Specifically, subjects participated in a total of 20 markets per session. Each market consisted of two stages, a payment choice and belief formation first stage and a second stage consisting of seven trading rounds.

In this first stage, each buyer was endowed with seven experimental money (EM) units (as there were seven sellers) and these buyers had to decide how to allocate their seven EM between the two payment methods. To avoid any biases due to framing effects, we used neutral language throughout, referring the existing payment method as "payment 1" and the new payment method as "payment 2".[15] Thus, in the first stage, buyers allocated their 7 EM between payment 1 and payment 2, with only integer allocation amounts allowed, e.g., 3 EM in the form of payment 1 and 4 EM in the form of payment 2. Each seller was endowed with seven units of goods (as there were seven buyers). Sellers were required to accept payment 1 but must decide in this first stage whether or not to accept payment 2 for the duration of that market. Sellers who decided to accept payment 2 had to pay a one-time fixed fee of T EM per market that enabled them to accept payment 2 in all trading rounds of that market. This fixed cost was deducted from the seller's sales profits at the end of the market.[16] As explained below, T is related to the fixed costs of adopting the new payment method (F) described in section 2 and serves as our main experimental treatment variable. Note that subjects could revisit their payment choices in the first stage of each mew market.

In addition to making payment choices in the first stage, subjects were also asked to forecast other participants' payment choices for that market. We elicited these forecasts because we wanted to better understand subjects' decision-making process and possible feedback effects between buyers and sellers. Specifically, buyers were asked to forecast how many of the seven sellers would choose to accept payment 2 in the forthcoming market. Sellers were asked to supply two forecasts: (1) the average amount of the seven EM units that all seven buyers would allocate to payment 2, and (2) how many of the other six sellers would choose to accept payment 2 in the forthcoming market. These forecasts (or beliefs) were incentivized; subjects earned 0.5 EM per correct forecast in addition to their earnings from buying/selling goods (also in EM). Each seller's forecast of the average amount of EM that buyers had allocated to payment 2 was counted as correct if it lay within ± 1 of the realized value. The other two forecasts were counted as correct only if they precisely equaled the realized value. We chose this incentive scheme in the interest of simplicity as belief elicitation was not the main focus of our experiment and we did not want to devote time to explaining more complicated elicitation procedures involving proper scoring rules. Note that no participant observed any other sellers' or buyers' payment choices or forecasts in this first stage; that is, all first stage choices and forecasts were private information and were made simultaneously.

Following completion of the first stage of each market, play immediately proceeded to the second, "trading" stage for the market, which consisted of a sequence of seven trading rounds. In these seven rounds, each subject anonymously met with each of the seven subjects who was in the opposite role

[15] Framing could affect the adoption decision, which we leave for future research.

[16] One interpretation is that sellers get a zero-interest loan to finance the fixed cost at the beginning of each new market and they must repay the loan at the end of the market.

to him/herself, sequentially and in a random order for each market.[17] In each meeting, the buyer and seller tried to trade one unit of good for one unit of payment (recall that the terms of trade in our model are fixed). Specifically, when each buyer met each seller (and not earlier), the buyer learned whether the seller accepted payment 2 or not, and then the buyer alone decided which payment method to use, conditional on the buyer's remaining balance for that market of either payment 1 or payment 2. Sellers were passive in these trading rounds, simply accepting payment 1 from the buyer or payment 2 if the seller had paid the one-time fee to accept payment 2 in that market, depending on the choice of the buyer. Thus, provided that a buyer had some amount of a payment type that the seller accepted, trade would be successful. For each successful transaction, both parties to the trade earned 1 EM less some transaction costs for the trading round, where the transaction costs depended on whether payment 1 or 2 was used, as detailed later. Notice that the only instances in which a transaction could not take place (was never successful) were those in which the buyer had only payment 2 and the seller did not accept payment 2. In those cases, no trade could take place and both parties earned 0 EM for the trading round. As discussed in the previous section, the choice in the trading stage is straightforward: buyers should try to use payment 2 if the seller accepts it, and try to use payment 1 if the seller does not accept payment 2. While designing the experiment, we considered the option of automating the trading stage. We opted for the current setup because we thought that going through the trading stage would help subjects to better grasp the consequences of their payment choices in the first stage. For instance, if a buyer only had payment 2 and the seller they met did not accept payment 2, the buyer would have to click on a "no trade" button to advance the program, which would make such no trade possibilities more salient.[18]

Following completion of the seven trading rounds of the second stage of a market, that market was over. Provided that the 20th market had not yet been completed, play then proceeded to a new two-stage market where buyers and sellers had to once again make payment choices and forecasts in the first stage and then engage in seven rounds of trading behavior in the second stage. Buyers were free to change their payment allocations and sellers were free to change their payment 2 acceptance decisions from market to market but only in the first stage of the market; the choices made in this first stage were then in effect for all seven rounds of the second trading stage of the market that followed.

In each market, subjects earn EM from forecasting in stage 1 and from the trading surplus from

[17]The set up that the buyer visits the seller in a random order in the trading stage does not affect the analysis of the single-market payment adoption game outlined in the previous section. However, subjects in our experiment interact in multiple markets, and this set up helps to preserve the anonymity of the seller across markets and prevents buyers from using strategies targeting a particular seller, for example, penalizing a particular seller who did not accept payment 2 in the past market. The buyers' portfolio choice considers only the aggregate acceptance decisions, and not each individual seller's decision.

[18]Subjects' choices suggest that they understood the stage-2 game very well. Among the 7,651 meetings where payment 2 can be used, 98.4% result in transactions using payment 2 (and the rest are conducted using payment 1). Among the 5,941 meetings where only payment 1 can be used, all transactions are fulfilled with payment 1. For future work, one design choice would be to let subjects choose in stage 2 for a few markets, and then automate their stage 2 choices to save time. Related to this, in our current design, both payment instruments are available from the first market, and the "newness" of payment 2 is captured by that sellers must decide whether to adopt it whereas they always accept payment 1. We could create an additional sense of "newness" by separating each session into two parts, letting subjects play with only payment 1 in the first part, and introducing payment 2 in the second part. This would allow us to capture the role played by habits. Automating stage 2 would save time and make this two-parts design more feasible in terms of recruitment time constraint.

transacting in the seven rounds in stage 2; only these *earned* EMs will be converted to cash payment at the end of the session. In particular, while buyers were endowed with 7 EM at the start of each market, this endowment must be allocated between payment 1 and payment 2 to facilitate transactions with sellers, and unused payment balances had no redemption value and could not be carried over to a subsequent market.[19] Similarly, although sellers were endowed with 7 units of goods at the start of each market, they could earn EMs in stage 2 only by transacting with buyers (minus the fixed cost if they accept payment 2), and unsold goods did not entitle them to additional EM earnings.

To facilitate decision making, we provided subjects with two pieces of information in stage 1, at the same time that buyers were asked to make their payment allocation choices and sellers were asked to make their payment 2 acceptance decisions and both types had to form forecasts as described above. The first piece of information provided to subjects consisted of payoff tables. The buyer's payoff table reported the buyer's market earnings if the buyer allocated between $0 \sim 7$ EM to payment 2 (and his/her remaining EM to payment 1) and if $0 \sim 7$ sellers accepted payment 2. The seller's payoff table reported the expected market earnings the seller could get from the two options (accept/reject payment 2) in cases where *all* buyers choose to allocate between $0 \sim 7$ EM to payment 2, and where $0 \sim 6$ of the other six sellers choose to accept payment 2. In addition to these payoff tables, sellers also had access to a "what if" calculator that computed their expected earnings in asymmetric cases where the seven buyers made different payment allocation choices.

The second piece of information that we provided subjects in stage 1 (beginning with market 2 and every market thereafter) was a history of outcomes in all past markets, including the subject's payment choice, the number of transactions using each of the two payment methods, the number of no-trade meetings, market earnings from trading, and the number of their correct forecasts. In addition, we reported an aggregate market-level statistic: the number of sellers who had chosen to accept payment 2 in the prior market. We provided the latter information so that sellers could learn about other sellers' payment 2 acceptance decisions in the just completed market; since all 7 buyers visit all 7 sellers and learn whether each seller accepts payment 2 or not, buyers had this economy-wide piece of information by the end of each market. By providing this same information also to sellers, we made sure that both sides of the market had symmetric information about seller choices. In addition, since we only have subjects in the laboratory for a limited time, providing this information to sellers helps to speed up learning.

For simplicity, we set the per transaction cost to be the same for all buyers and sellers, i.e., $\tau_1^b = \tau_1^s = \tau_1 = 0.5$, and $\tau_2^b = \tau_2^s = \tau_2 = 0.1$. Thus, consistent with assumption A1, it was always the case that $\tau_1 - \tau_2 = 0.4$ in all treatments of our experiment. We also set the utility gained from a sale or purchase, $u = 1$, so that consistent with assumption A2, $u - \tau_1^b > 0$. Our only treatment variable was the once-per-market fixed cost, T, that each seller had to pay to accept payment 2, which corresponds to the parameter F in the model with a continuum of agents via the transformation $T = 7F$. We initially

[19]Further, the EM payments earned in each market were not transferable to subsequent markets. Instead, earnings were recorded and paid out only at the end of the experiment following the completion of the 20th market. Thus, buyers started each new market with exactly 7 EM and had to make payment allocations anew in the first stage of each market in order to earn EM in that market.

Table 1: Payoffs Earned in the Two Equilibria

	All-payment-1 equilibrium		All-payment-2 equilibrium	
T	Buyer	Seller	Buyer	Seller
1.6	3.5	3.5	6.3	4.7
$2.8 = 7(\tau_1 - \tau_2)$	3.5	3.5	6.3	3.5
3.5	3.5	3.5	6.3	2.8

chose three different values for this main treatment variable: $T = 1.6$, 2.8, and 3.5, respectively.[20] Later, in section 6 we consider a fourth value for this treatment variable, $T = 4.5$.

In terms of theoretical predictions, under all of our different treatment conditions, there *always exist* two symmetric pure strategy Nash equilibria, one in which no seller accepts payment 2 and all buyers allocate all of their endowment to payment 1 and another such equilibrium in which all sellers accept payment 2 and all buyers allocate all of their endowment to payment 2. In the treatment with $T = 7(\tau_1 - \tau_2)$, there are also asymmetric equilibria, where $m^s \in \{1, 2, ..., 6\}$ sellers accept payment 2, whereas $7 - m^s$ do not, and all buyers allocate the same m^s units of their endowment to payment 2 and $7 - m^s$ to payment 1. These asymmetric equilibria characterize dual payment outcomes, where *both* payment methods are used to make transactions.[21]

Note that, given the lower transaction cost from using the new payment method, buyers are always better off in the all-payment-2 equilibrium relative to the all-payment-1 equilibrium. However, sellers' relative payoffs depend on the fixed cost, T. If $T < 7(\tau_1 - \tau_2)$, as in our $T = 1.6$ treatment (and represented graphically in the left panel of Figure 1), then the seller's payoff (like the buyer's payoff) is higher in the all-payment-2 equilibrium than in the all-payment-1 equilibrium. If $T > 7(\tau_1 - \tau_2)$, as in our $T = 3.5$ treatment (and represented graphically in the right panel of Figure 1), then the seller's payoff is lower in the all-payment-2 equilibrium than in the all-payment-1 equilibrium. Finally, if $T = 7(\tau_1 - \tau_2)$, as in our $T = 2.8$ treatment, then the seller's payoffs are the same in both equilibria. Table 1 reports the net payoffs (in EM) per market, that buyers and sellers earn in the two pure strategy equilibria of each treatment.

While our focus is on the two symmetric pure strategy equilibria, we note that there may also exist *asymmetric* equilibria where some fraction of sellers accept payment 2 whereas the remaining fraction do not, and buyers best respond by adjusting their portfolios of payment 1 and 2 so as to perfectly match this distribution of seller choices. These asymmetric equilibria are always present in the $T = 2.8$ treatment, where $T = 7(\tau_1 - \tau_2)$. In particular, any outcome where $m^s \in \{1, 2, ..., 6\}$ sellers accept

[20] The ratio $\tau_1/\tau_2 = 5$ and the various values for the fixed cost, T, were chosen to make the transaction and setup cost differences sufficiently salient to our subjects (in terms of their earnings) and are not meant to be empirically accurate. For reference, Welte and Molnar (2021) use industry data and data from the Bank of Canada 2015 Retailer Survey on the Cost of Payment Methods (Kosse et al. 2017) to calculate the costs associated with card payments for small and medium size retailers. On average, total acquirer charges represent 1.8% of the processes card sales. Of those charges, about 0.5% are fixed costs, and 1.3% are transaction-based fees.

[21] Note that because the population is finite in the experiment, the symmetric equilibria where sellers play mixed strategies accepting the new payment method with probability $m^s/7$ disappear, whereas the asymmetric equilibria survive.

payment 2, whereas $7 - m^s$ do not, and all buyers allocate the same m^s units of their endowment to payment 2 and $7 - m^s$ to payment 1 comprises an asymmetric equilibrium for this treatment.[22] These asymmetric equilibria characterize dual payment outcomes, where *both* payment methods are used to make transactions.

This multiplicity of equilibrium possibilities is one motivation behind our experimental study; equilibrium selection is clearly an empirical question that our experiment can help to address. We hypothesize that, as the transaction cost to sellers of accepting payment 2 increases from $T = 1.6$ to $T = 2.8$ and on up to $T = 3.5$, coordination on the all-payment-2 equilibrium will become less likely and coordination on the all-payment-1 equilibrium will become more likely; however, this remains an empirical question, as both symmetric equilibria always co-exist. In addition, we are interested in understanding the dynamic process of equilibrium selection in terms of the evolution of subjects' beliefs and choices.

The experiment was computerized and programmed using the z-Tree software (Fischbacher, 2007). At the beginning of each session, each subject was assigned a computer terminal, and written instructions were handed out explaining the payoffs and objectives for both buyers and sellers. See online Appendix B for example instructions used in the experiment. The experimenter read these instructions aloud in an effort to make the rules of the game public knowledge. Subjects could ask questions in private and were required to successfully complete a quiz to check their comprehension of the written instructions prior to the start of the first market. Communication among subjects was prohibited during the experiment.

In each session, subjects interact in 20 markets involving seven trading rounds each, or a total of 140 trading rounds per session. In each trading round, subjects could earn 1 EM ($u = 1$) less the transaction cost associated with their payment choice, and in the first stage, they could earn 0.5 EM per correct forecast. Following completion of the 20th market, subjects were paid their cumulative EM earnings from all rounds of all markets at the known and fixed rate of 1 EM = $0.15 and in addition they were paid a $7 show-up payment. Each session lasted for about two and half hours. The average earnings were between $15 and $25.

We have four sessions for each of our three original treatment conditions, $T = 1.6$, $T = 2.8$ and $T = 3.5$. As each session involved 14 subjects with no prior experience participating in our study, we have data from $4 \times 3 \times 14 = 168$ subjects. The experiment was conducted in two locations: Simon Fraser University (SFU), Burnaby, Canada, and at the University of California, Irvine (UCI), USA, using undergraduate student subjects. Specifically, two sessions of each of our three treatments (one-half of all sessions) were run at SFU and UCI, respectively. Our aim in conducting the experiment at two different locations was to assess whether our results would replicate with different subject pools and experimenters conducting the sessions. Despite our use of these two different subject pools, we did not find significant differences in either buyer or seller behavior across these two locations, as we show later in the paper,

[22]Because of our use of a finite population size of 14 subjects, there are also some asymmetric pure strategy equilibria of this same variety in the $T = 1.6$ and $T = 3.5$ treatments as well. However, these asymmetric equilibria are fewer in number than in the $T = 2.8$ treatment and they would disappear completely as the population size got larger and we approached the continuum of agents as in the theory, whereas the set of asymmetric equilibria in the $T = 2.8$ case would continue to grow and would eventually reach a continuum, as in the theoretical model.

enabling us to pool our data from both locations.

4 Aggregate Experimental Results

In this section, we present and discuss our experimental results at the aggregate level. The next section will address individual-level behavior.

Figures 2 to 4 show the time series on payment choices and transaction methods in each of the four sessions of our three treatments. In all of these figures, the horizontal axis indicates the number of the market, running from 1 to 20. Each figure (one per session) has two panels. In the first (left) panel the series labeled "BuyerPay2," shows the percentage of the buyers' endowment allocated to payment 2 in each market, averaged across the seven buyers of each session. In this same panel, the percentage of sellers accepting payment 2 in each market is indicated by the series labeled "SellerAccept." The second (right) panel of Figures 2 to 4 show three time series: (1) the frequency of meetings in each market that resulted in transactions using payment 1 labeled as "Pay1," (2) the frequency of meetings in each market that resulted in transactions using payment 2 labeled as "Pay2," and (3) circles indicating the frequency of no-trade meetings in each market labeled as "NoTrade."

Tables 2 and 3 display various statistics for each of the 12 sessions of the experiment. For each statistic, we show the treatment-level average in bold face. Table 2 provides statistics on five variables that summarize payment choice and usage and transactions. The first part of Table 2 reports on the percentage of the buyers' endowment that was allocated to payment 2 averaged across the seven buyers (bPay2%) and the percentage of the seven sellers accepting payment 2 (sAccept%). In particular, we report the session mean across the 20 markets, and the value in the first and the last markets of these two statistics. The second part of Table 2 reports on the percentage of meetings that resulted in trade using payment 2 (pay2Meetings%), trade using payment 1 (pay1Meetings%) or no trade (noTradeMeetings%). Table 3 provides the same set of statistics on payoff for buyers and sellers, measured as a percentage of the payoffs that could be earned in the all-payment-2 equilibrium (Part 1 of Table 3) and the all-payment-1 equilibrium (Part 2 of Table 3).

We also examine whether there are treatment differences using pairwise Wilcoxon rank-sum tests on data concerning payment choice and usage. These results are presented in Table 4 in online Appendix C. We focus on five variables: (1) "bPay2%," the percentage of endowment allocated to payment 2 averaged across the seven buyers, (2) "sAccept%," the percentage of the seven sellers accepting payment 2, (3) "pay2Meetings%," the percentage of meetings that resulted in trades using payment 2, and (4) "No-tradeMeetings%," the percentage of meetings that result in no trade.For each of these four variables, we compare the session average and its value in the first market in treatment 1 versus treatment 2, e.g., T=1.6 versus T=2.8. Each session is treated as an independent observation, so we have four observations on each variable for each treatment. For each pair-wise test, Table 4 reports the rank sums of the two treatments, the z-value and the p-value.

Using the data reported in Figures 2 to 4, Tables 2, 3, 4 and C.2, we summarize the results from the aggregate-level analysis of our data as a number of findings.

Finding 1 *Across the three treatments, as T is increased from 1.6 to 2.8 to 3.5 , there are significant decreases in the buyer's choice of payment 2, the seller's acceptance of payment 2, and successful transactions involving payment 2.*

Support for Finding 1 comes from Table 4. The rank-sum test results indicate that the buyer's choice of payment 2 (bPay2%), the seller's acceptance of payment 2 (sAccept%) and successful payment 2 transactions (pay2Meetings%) are significantly higher in the $T = 1.6$ treatment as compared with either the $T = 2.8$ or $T = 3.5$ treatments (p-value< 0.05). Further, these variables are larger in the $T = 2.8$ treatment as compared with the $T = 3.5$ treatment. Importantly, Table 4 also indicates that, initially, using the statistics in the first market of each session ("First market"), there are no significant differences in these three variables across our three treatments, indicating that all sessions started out with roughly similar initial conditions and the choices of buyers and sellers then evolved over time to yield the differences summarized in Finding 1. The next three findings summarize aggregate behavior in each of the three experimental treatments.

Finding 2 *When $T = 1.6$, the experimental economies converge to (or nearly converge to) the all-payment-2 equilibrium.*

Support for Finding 2 comes from Table 2 and Figure 2, which report on the evolution of behavior in the four sessions of the $T = 1.6$ treatment. In all four sessions, subjects move over time in the direction of the all-payment-2 equilibrium, which in this case represents a strict Pareto improvement for both sides. Furthermore, sellers do not suffer much loss from paying the fixed cost to accept payment 2: they can recover the fee if each buyer allocates just 4 EM (57%) or more to payment 2. As a result, sellers maintain high levels of acceptance, and this steadily high acceptance rate encourages buyers to quickly catch up, which, in turn, reinforces the incentives for acceptance of payment 2. As a result, toward the end of all four $T = 1.6$ treatment sessions, subjects have achieved convergence or near-convergence to the payment 2 equilibrium.[23]

Finding 3 *When $T = 2.8$, there is a mixture of outcomes consistent with the greater multiplicity of equilibria in this treatment.*

Support for Finding 3 comes from Table 2 and C.2 and Figure 3, which report on the evolution of behavior in the four sessions of the $T = 2.8$ treatment. In contrast with the $T = 1.6$ treatment, the data reported in Table 2 and Figure 3 reveal a mixture of outcomes across the four sessions of the $T = 2.8$

[23]Finding 2 is also supported by the convergence test. As shown in Table C.2, the estimated long-run expected values of bPay2%, sAccept%, and pay2Meetings%, are not significantly different from 100 for sessions 1, 2 and 4. For session 3, the estimated long-run values of bPay2% and pay2Meetings%, are also not significantly different from 100, but the seller's choice of payment 2, sAccept%, is statistically different from 100, though the magnitude of the difference is small ($\leq 5\%$).

treatment. In particular, we observe that the experimental economy either lingers in the middle ground between the all-payment-1 and all-payment-2 equilibria (sessions 1, 3 and 4) or appears to be very slowly converging toward the all-payment-2 equilibrium (session 2).

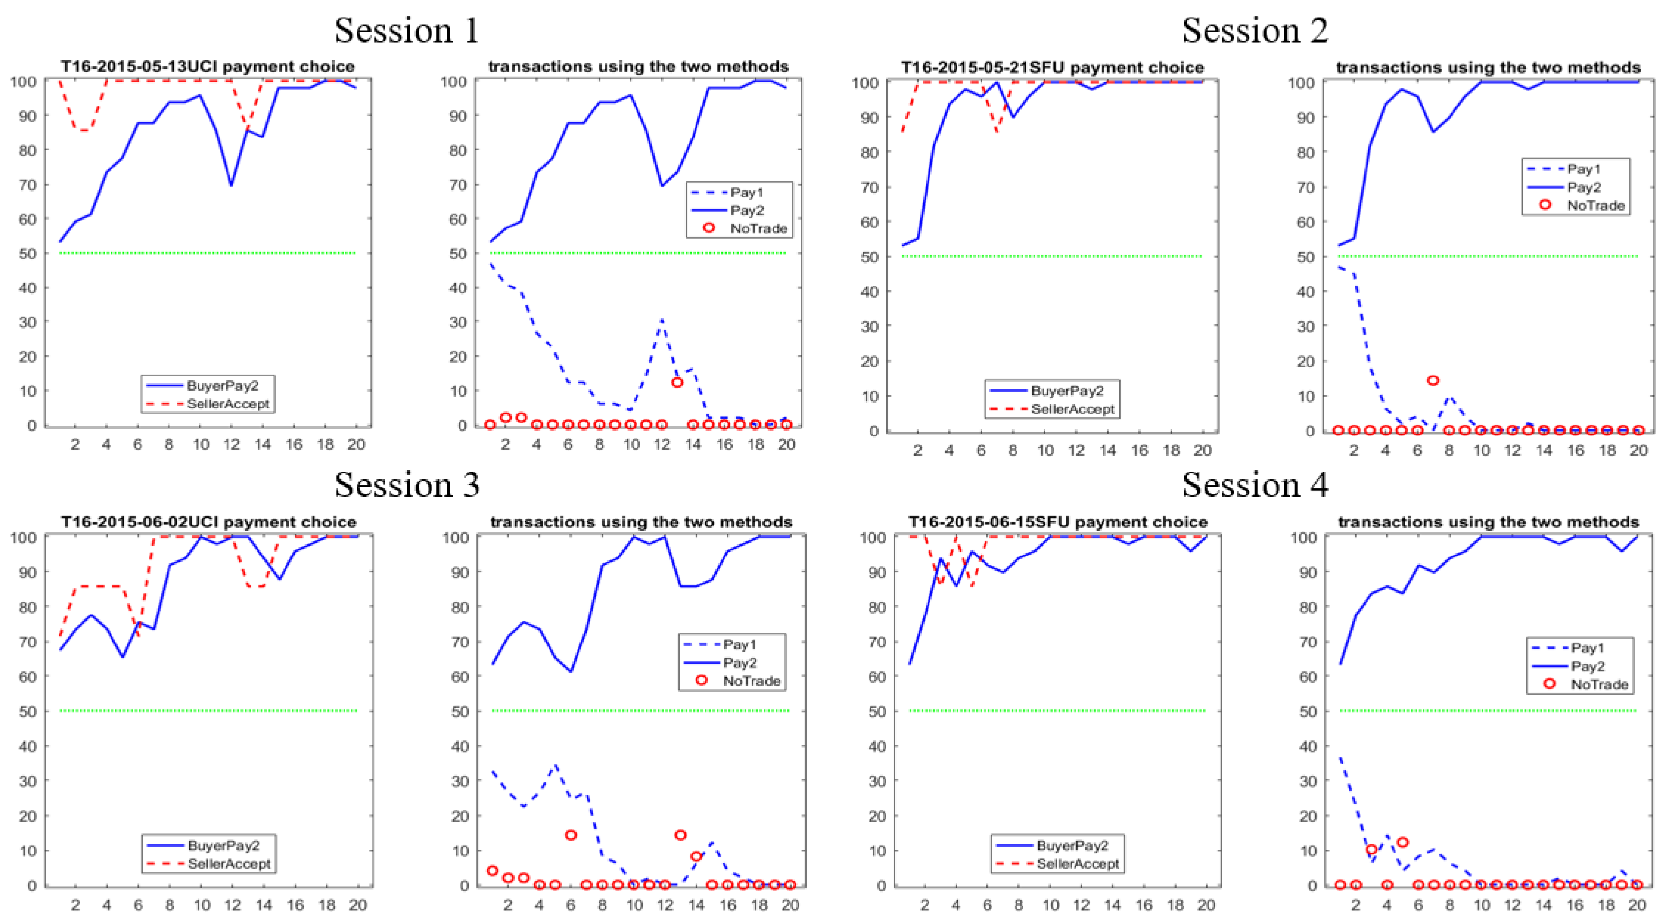

Figure 2: Payment Choice and Usage T=1.6

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first shows payment choices: the blue solid line represents the percentage of money allocated to payment 2, averaged across the seven buyers; the red dashed line represents the percentage of the seven sellers accepting payment 2. The second figure describes payment usage. The solid blue line is the percentage of meetings using payment 2; the dashed blue line is the percentage of meetings using payment 1; red circles are the percentage of meetings where no trade takes place.

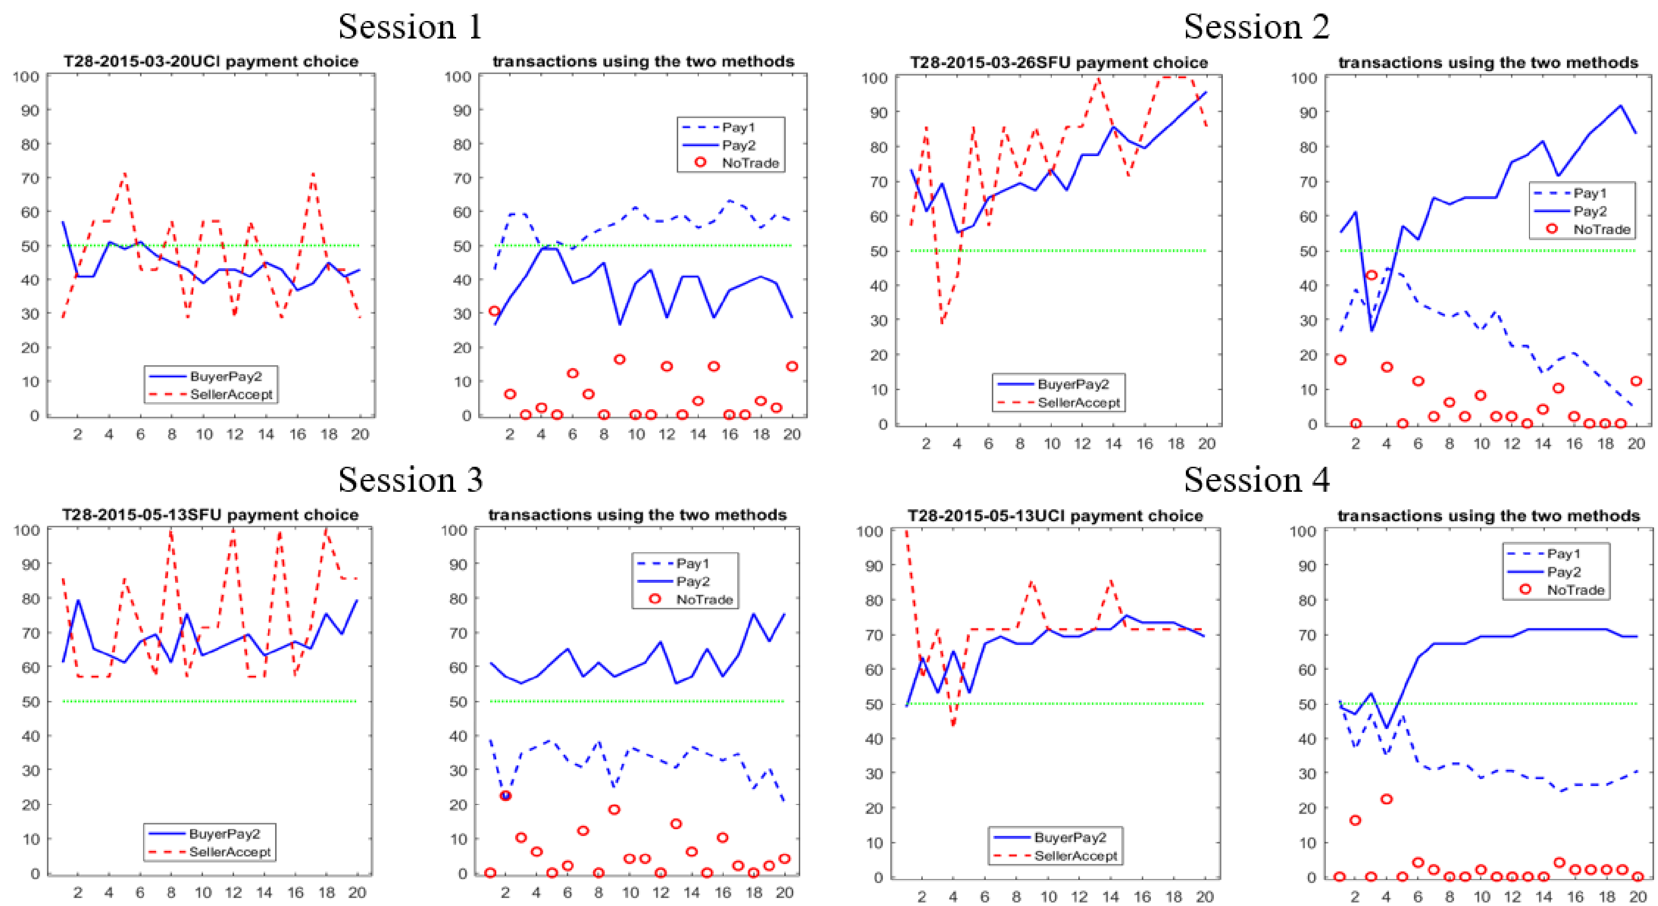

Figure 3: Payment Choice and Usage T=2.8

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first shows payment choices: the blue solid line represents the percentage of money allocated to payment 2, averaged across the seven buyers; the red dashed line represents the percentage of the seven sellers accepting payment 2. The second figure describes payment usage. The solid blue line is the percentage of meetings using payment 2; the dashed blue line is the percentage of meetings using payment 1; red circles are the percentage of meetings where no trade takes place.

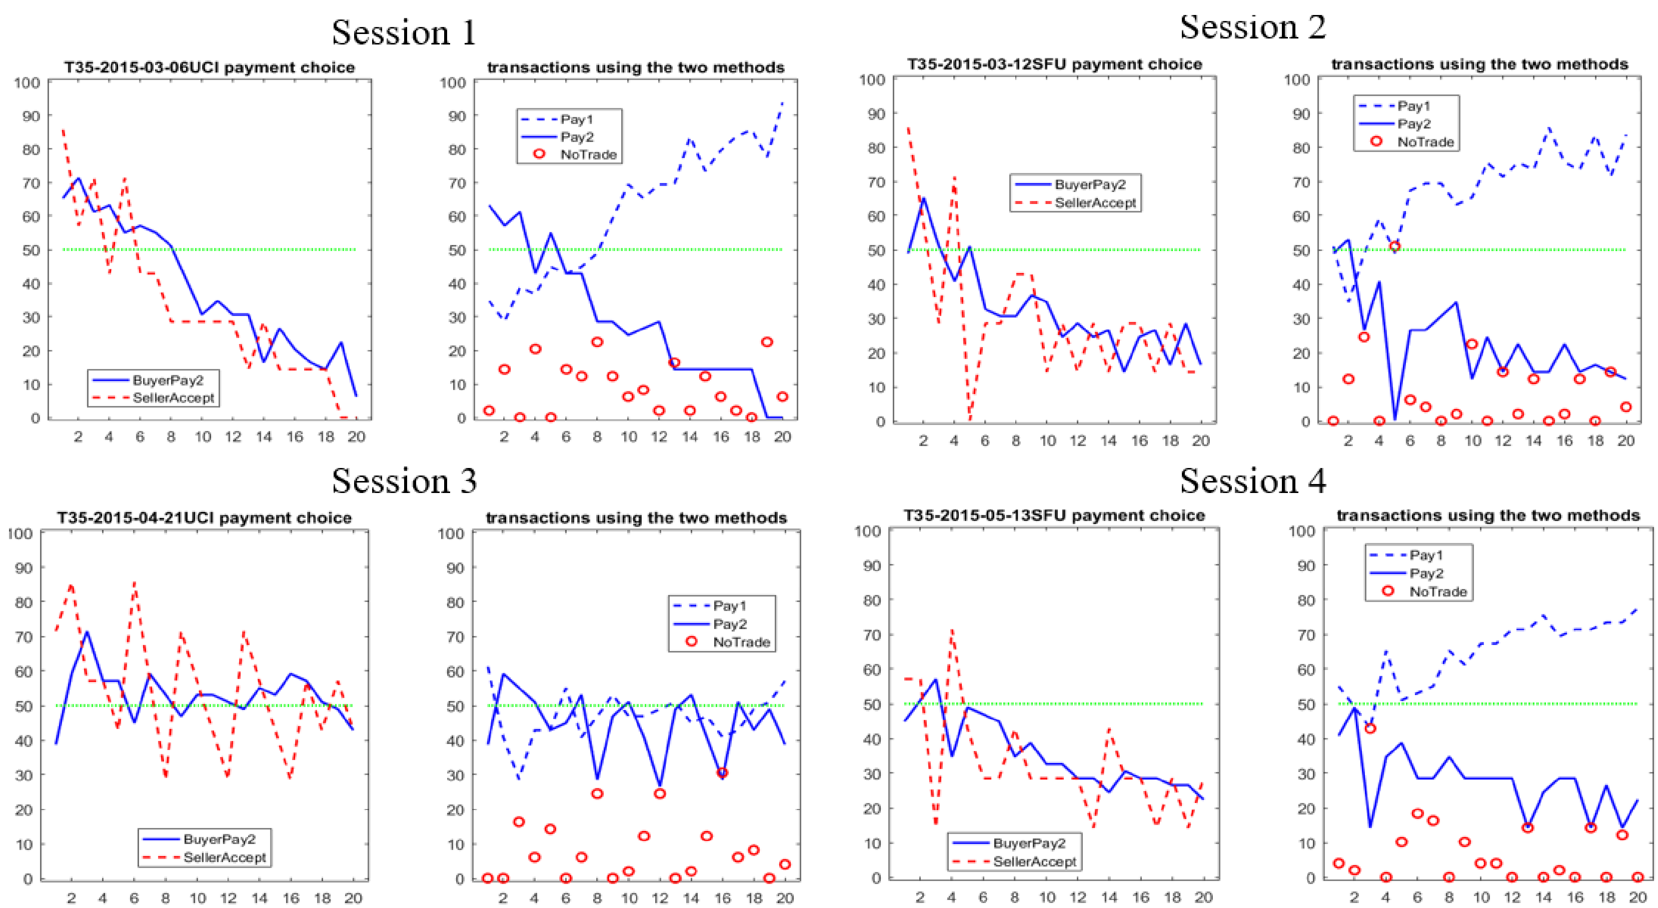

Figure 4: Payment Choice and Usage T=3.5

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first shows payment choices: the blue solid line represents the percentage of money allocated to payment 2, averaged across the seven buyers; the red dashed line represents the percentage of the seven sellers accepting payment 2. The second figure describes payment usage. The solid blue line is the percentage of meetings using payment 2; the dashed blue line is the percentage of meetings using payment 1; red circles are the percentage of meetings where no trade takes place.

Table 2: Payment Choice and Usage

		T=1.6					T=2.8					T=3.5				
		1	2	3	4	**all**	1	2	3	4	**all**	1	2	3	4	**all**
bPay2%	session mean	85	93	88	94	**90**	44	74	68	67	**63**	38	33	53	36	**40**
	first market	53	53	67	63	**59**	57	73	61	49	**60**	65	49	39	45	**49**
	last market	98	100	100	100	**99**	43	96	80	69	**72**	6	16	43	22	**22**
sAccept%	session mean	98	99	93	99	**97**	46	79	74	72	**68**	33	31	54	33	**38**
	first market	100	86	71	100	**89**	29	57	86	100	**68**	86	86	71	57	**75**
	last market	100	100	100	100	**100**	29	86	86	71	**68**	0	14	43	29	**21**
pay2 Meetings%	session mean	84	92	86	93	**89**	38	67	62	64	**58**	29	23	45	28	**31**
	first market	53	53	63	63	**58**	27	55	61	49	**48**	63	49	39	41	**48**
	last market	98	100	100	100	**99**	29	84	76	69	**64**	0	12	39	22	**18**
pay1 Meetings%	session mean	15	7	12	6	**10**	56	26	32	33	**37**	62	67	47	64	**60**
	first market	47	47	33	37	**41**	43	27	39	51	**40**	35	51	61	55	**51**
	last market	2	0	0	0	**1**	57	4	20	31	**28**	94	84	57	78	**78**
noTrade Meetings%	session mean	1	1	2	1	**1**	6	7	6	3	**6**	9	9	8	8	**9**
	first market	0	0	4	0	**1**	31	18	0	0	**12**	2	0	0	4	**2**
	last market	0	0	0	0	**0**	14	12	4	0	**8**	6	4	4	0	**4**

Table 3: Payoff

		T=1.6					T=2.8					T=3.5				
Part 1: payment-2 equilibrium as benchmark																
		1	2	3	4	**all**	1	2	3	4	**all**	1	2	3	4	**all**
buyer	session mean	93	96	93	96	**94**	69	82	80	82	**78**	64	61	71	64	**65**
	first market	79	79	81	84	**81**	50	70	83	77	**70**	83	77	73	71	**76**
	last market	99	100	100	100	**100**	60	86	87	86	**80**	52	59	71	66	**62**
seller	session mean	91	95	92	95	**93**	87	84	84	91	**86**	102	99	91	102	**98**
	first market	72	77	85	78	**78**	68	80	80	59	**72**	79	67	74	89	**77**
	last market	99	100	100	100	**100**	86	86	88	98	**89**	117	114	105	112	**112**
Part 2: payment-1 equilibrium as benchmark																
buyer	session mean	167	173	167	173	**170**	124	147	144	148	**141**	114	110	127	115	**116**
	first market	142	142	147	151	**146**	91	126	149	139	**126**	149	139	131	129	**137**
	last market	178	180	180	180	**180**	109	155	156	156	**144**	94	106	127	118	**111**
seller	session mean	122	128	124	128	**126**	87	84	84	91	**86**	82	79	73	82	**79**
	first market	97	103	114	105	**105**	68	80	80	59	**72**	63	53	60	71	**62**
	last market	133	134	134	134	**134**	86	86	88	98	**89**	94	91	84	89	**90**

Table 4: Rank-sum Test – Treatment Effect

		Rank-sum treatment 1	Rank-sum treatment 2	z-value	p-value	
T=1.6 versus T=2.8						
Session average	bPay2%	26	10	2.309	0.021	**
	sAccept%	26	10	2.323	0.020	**
	pay2Meetings2%	26	10	2.309	0.021	**
	No-tradeMeetings%	10	26	-2.309	0.021	**
First market	bPay2%	18	18	0	1	
	sAccept%	21.5	14.5	1.042	0.298	
	pay2Meetings2%	22	14	1.169	0.243	
	No-tradeMeetings%	15	21	-0.992	0.321	
T=2.8 versus T=3.5						
Session average	bPay2%	25	11	2.021	0.043	**
	sAccept%	25	11	2.033	0.042	**
	pay2Meetings2%	25	11	2.021	0.043	**
	No-tradeMeetings%	10	26	-2.309	0.021	**
First market	bPay2%	22.5	13.5	1.307	0.191	
	sAccept%	18.5	17.5	-0.149	0.882	
	pay2Meetings2%	18.5	17.5	0.145	0.885	
	No-tradeMeetings%	20	16	0.615	0.539	
T=1.6 versus T=3.5						
Session average	bPay2%	26	10	2.309	0.021	**
	sAccept%	26	10	2.337	0.019	**
	pay2Meetings2%	26	10	2.309	0.021	**
	No-tradeMeetings%	10	26	-2.309	0.021	**
First market	bPay2%	23	13	1.452	0.147	
	sAccept%	22.5	13.5	1.348	0.178	
	pay2Meetings2%	23	13	1.488	0.137	
	No-tradeMeetings%	16.5	19.5	-0.500	0.617	

Notes. (1) Combined sample size for each test is 8. (2) *p-value≤0.1; p-value≤0.05; p-value≤0.01.

Recall that when $T = 2.8$, sellers are indifferent between the two equilibria as their payoffs are the same in either equilibrium. Further, any outcome where all buyers allocate the same proportion of their endowment to payment 2 as the fraction of sellers accepting payment 2 is always an asymmetric equilibrium in this setting. Generally we observe that the fraction of sellers accepting payment 2 hovers above (with some volatility) the fraction of endowment that buyers allocate to payment 2.

If, over time, buyers increasingly insist on the new payment method 2, then sellers are likely to accommodate the buyers' choices by accepting it; this seems to be the case in session 2. In the final market of session 2 of the $T = 2.8$ treatment, we see from Table 2 that buyers' allocation of their endowment to payment 2 averages 96%, and six out of seven sellers (86%) accept payment 2. The number of payment 2 transactions in this session increases from 55% in the first market to 84% in the last market. It seems reasonable to conjecture that session 2 would converge more closely to the all-payment-

2 equilibrium if the session had continued beyond 20 markets.[24] Note that, compared with the $T = 1.6$ treatment sessions, the process of convergence toward the all-payment-2 equilibrium in session 2 of the $T = 2.8$ treatment is considerably slower, more erratic and incomplete.

By contrast, in the other three sessions of the $T = 2.8$ treatment there is little to no evidence of convergence toward either the all-payment-1 or all-payment-2 equilibrium. Table 2 reveals that for session 3 there is a small increase in the number of transactions using payment 2 from 61% in the first market to 76% in the final market, but the economy remains far away from the all-payment-2 equilibrium, or any other symmetric equilibrium. In session 1, the number of sellers accepting payment 2 fluctuates between 2/7 (28.5%) and 5/7 (71.4%). Consequently, buyers are not willing to take the lead by acquiring high payment 2 balances, fearing that they will not be able to trade in case some sellers reject payment 2. As a result, the buyer's average payment 2 allocation and the number of sellers accepting payment 2 both average between 40-50% throughout the entire session, but there is never coordination on the same rate. Finally, session 4 shows an upward trend in payment 2 usage in the first seven markets, but this trend abruptly halts thereafter, with the average number of payment 2 transactions barely changing from market eight onward (the value fluctuates between 67% and 71%). This outcome represents near (but imperfect) convergence to an interior asymmetric equilibrium where approximately five out of seven sellers (71%) are accepting payment 2 and buyers are allocating approximately 5 units of their 7 EM endowment to payment 2. This is the closest instance we have to a dual payments equilibrium outcome in our data.

Finding 4 *When $T = 3.5$, the experimental economy slowly converges to the all-payment-1 equilibrium, or lingers in the middle ground between the two pure strategy equilibria.*

Support for Finding 4 comes from Table 2 and Figure 4, which report on the evolution of behavior in the four sessions of the $T = 3.5$ treatment. When $T = 3.5$, sellers do better in the payment 1 equilibrium as compared with the payment 2 equilibrium; by contrast buyers always prefer the payment 2 equilibrium. Nevertheless, in each of the four sessions, more than 50% of sellers start out in the first market accepting payment 2, perhaps fearing that they will lose business in the case where some buyers show up with only payment 2 remaining. With experience, sellers learn to resist accepting payment 2 and to engage in a "tug-of-war" with buyers; the average acceptance rate over all four sessions declines from 75% in the first market to just 21% in the last market.

Sellers appear to be winning this contest in sessions 1, 2 and 4, pulling the economy back in the direction of the status quo, all-payment-1 equilibrium. For example, in session 1 of the $T = 3.5$ treatment, Table 2 reveals that buyer's payment 2 allocation falls from 65% in the first market to an average of just 6% in the last market. On the seller's side, six out of seven sellers in this session (86%) accept payment

[24]The formal convergence analysis (Table C.1) suggests that for session 2 of the $T = 2.8$ treatment, the estimated long-term value of the buyers' allocation of their endowment to payment 2 (bPay2%) is not significantly different from 100% in this session. However, the other two variables, the sellers' acceptance rate and percentage of payment 2 meetings, have not yet converged to 100% in a statistical sense. Nevertheless, it seems reasonable to conjecture that session 2 would pass the convergence test if the session had continued beyond 20 markets.

2 in the first market; by the last market of this session, no seller is accepting payment 2. The number of payment 1 transactions increases from 35% in the first market to 94% by the final, 20th market; over the same interval, payment 2 transactions fall from 63% to 0%. In session 3, the tug of war continues throughout the session and neither side is able to gain the upper hand; in that session, the average buyer's payment 2 allocation and the number of sellers accepting payment 2 consistently fluctuates around 50%, but there is no coordination on any asymmetric equilibria in this setting.[25]

Finding 5 *Mis-coordination in payment choices increases in T.*

From Tables 2 and 4, mis-coordination in payment choices becomes more severe with higher T, as manifested in the increasing frequency of no-trade meetings, which averaged 0.8% for $T = 1.6$, 5.6% for $T = 2.8$ and 8.6% for $T = 3.5$.

Finally, in terms of payoffs (see Table 3), buyers benefit relative to the all-payment-1 equilibrium, whereas sellers only benefit in the $T = 1.6$ treatment and sellers suffer in the other two treatments relative to the all-payment-1 equilibrium benchmark. The latter finding is summarized as follows.

Finding 6 *Sellers may choose to accept the new payment method (payment 2) even if doing so reduces their payoffs relative to the status quo where only payment 1 is used.*

Merchants often complain about high costs associated with accepting electronic payments but feel obliged to accept those costly payments for fear of upsetting or losing their customers (the so-called "must-take" phenomenon). A recent study by Bounie, François and Van Hove (2017b) finds that in the case of France in 2008, the must-take phenomenon applies to 5.8-19.8% of card-accepting merchants.[26] We observe a similar pattern in our experiment. For instance, when $T = 3.5$, although sellers understand that they will lose relative to the status quo if the economy moves to the all-payment-2 equilibrium, most sellers still begin the session accepting the new payment method (the average frequency of sellers accepting payment 2 in the first market is 68% in treatment $T = 2.8$ and 75% in treatment $T = 3.5$). Furthermore, in all sessions with $T = 2.8$ and 3.5, throughout all 20 markets, the seller's average payoff is always below 3.5, the payoff that they would earn if the economy were in the all-payment-1 equilibrium

[25] Table C.2 reveals that session 1 of the $T = 3.5$ treatment passes the formal convergence test, with all three variables bPay2%, sAccept% and pay2Meetings% not statistically different from 0. In sessions 2 and 4 of the $T = 3.5$ treatment, the trend works in favor of the seller, but the speed of convergence is slow; at the end of the session 2(4), 12% (22%) of transactions are still conducted in payment 2. Sessions 2 and 4 do not pass the convergence test; however, given the downward trend in payment 2 choice and usage, one can reasonably conjecture that convergence to the all-payment-1 equilibrium would be achieved with more market repetitions. Session 3 does not pass the convergence test either and it is not clear whether the session would converge even with more market repetitions.

[26] Some evidence in support of this claim comes from Evans (2011, p. vi), who observes that "...the US Congress passed legislation in 2010 that required the Federal Reserve Board to regulate debit card interchange fees; the Reserve Bank of Australia decided to regulate credit card interchange fees in 2002 after concluding that a market failure had resulted in merchants paying fees that were too high; and in 2007 the European Commission ruled that MasterCard's interchange fees violated the EU's antitrust laws."

(see Figures C.2 and C.3 in online Appendix C). Still, the “must-take” phenomenon is likely to be greater in practice than in our experiment which involves anonymous, no repeat-matching. More generally, sellers could develop reputations for accepting multiple payment options and the response of buyers to such reputations would likely increase pressure on reluctant sellers to adopt the new payment methods.

5 Individual Experimental Results

In this section we explore in further detail individual decision-making in our experiment. In particular, we examine the buyers’ portfolio decisions and the sellers’ payment 2 acceptance decisions taking into account treatment variables and buyers and sellers’ elicited beliefs. Since the impact of beliefs on behavior can result in endogeneity issues, we adopt an instrumental variables (IV) regression approach. In the first stage, we regress each individual’s belief concerning the other players’ payment choices for the current market on the outcomes that obtained in the previous market, together with four control variables: the market number, 1, 2, ..., 20 (“market”) to capture any time trend, a location dummy (“location”) equal to 1 if the data were collected at SFU, and two treatment dummies, T16 and T35, to pick up treatment level effects from the $T = 1.6$ and $T = 3.5$ treatments respectively (the baseline treatment is thus $T = 2.8$). In the second stage, we regress the payment choice on the estimated belief term and the control variables. The specific rationale for this two-stage regression is that individuals’ payment choices depend on their beliefs about other players’ choices, and their beliefs are in turn affected by historic outcomes.

Table 5 reports regression estimates for the buyer’s decision. The first-stage regression is a linear, random effects regression model of buyer’s beliefs. The dependent variable is “bBelief%”, the buyer’s own incentivized belief as to the number of sellers who would be accepting payment 2 in the current market. The main independent variable is “mktAcceptL%,” representing the percentage of sellers accepting payment 2 in the last market.[27] In addition, there are four control variables: market, location, and the treatment dummies, T16 and T35. The second-stage regression is a linear, random effects regression model of the buyer’s portfolio decision. The dependent variable is the percentage of the buyer’s endowment that he/she allocated to payment 2 (card choice%), and the explanatory variables are the estimated bBelief% term from the first stage and the four control variables.[28]

In the first-stage regression, the coefficient on MktAcceptL(%) is positive and statistically significant with a p-value close to zero, which suggests that buyers adjust upward their belief about seller acceptance of payment 2 in the current market after observing a higher past market seller acceptance of payment 2.

[27] Buyers learned the percentage of sellers accepting payment 2 in the prior market because they visited all seven sellers and were informed in every instance whether or not the seller accepted payment 2. In addition, we reported information on the percentage of sellers who accepted payment 2 in all prior markets on buyers’ stage 1 portfolio allocation screen. Thus buyers had ready access to the statistic mktAcceptL%.

[28] In the context of the payment choice game, the buyer’s payoff depends on his/her own portfolio choice and the sellers’ acceptance decisions, but not on other buyers’ choices. As a result, we elicited the buyer’s belief regarding the number of sellers accepting payment 2, but not the buyer’s belief about other buyers’ choices. For the same reason, the regressions for the buyer’s choice include the buyer’s belief on seller acceptance, but not beliefs about other buyers’ choices.

Compared with the baseline treatment $T = 2.8$, the buyer's belief about sellers' acceptance of payment 2 is 12% higher in treatment $T = 1.6$, and 7% lower in treatment $T = 3.5$; both coefficients are statistically significant with p-values close to zero. There is no significant time trend in buyers' belief (the variable market is insignificant). The location dummy variable is also insignificant, suggesting that there was no significant difference in buyers' beliefs between SFU and UCI, and rationalizing our pooling of the data from the two locations.

In the second-stage regression, only the coefficient on bBelief(%) is significant, with an estimate of 0.994, which is close to 1. This finding is consistent with the theory: the buyer's best response is to allocate $x\%$ of their endowment to payment 2 if they expect $x\%$ of sellers to accept payment 2. We summarize these results as follows.

Table 5: Buyer Payment 2 Choice

	(1) Stage1:bBelief(%)	(2) Stage 2:bPay2(%)
MktAcceptL(%)	**0.550*****	
	(0.037)	
bBelief(%)		**0.994*****
		(0.046)
market	**0.145**	**0.071**
	(0.120)	(0.080)
location	**0.929**	**1.409**
(SFU=1;UCI=0)	(1.356)	(1.226)
T16	**12.254*****	**0.189**
	(2.225)	(2.243)
T35	**-6.892*****	**-1.207**
	(1.755)	(1.752)

Notes. (1) *p-value≤0.1; p-value≤0.05; p-value≤0.01. (2) Each regression has $1,596 = 12 \times 7 \times 19$ observations. There are 12 sessions, each with 7 buyers and 20 markets. For each individual, we have 19 observations as the first-stage regression uses a lagged variable as an independent variable. (3) The regressions have clustered errors at the individual subject level.

Finding 7 *Buyers' portfolio choices depend on their beliefs about sellers' acceptance of payment 2, and these beliefs in turn depend on historic market outcomes and the value of T.*

Table 6 reports regression estimates from the seller's acceptance decisions. Again, we used an IV regression, for the same reasons given earlier for buyer choices. The first stage estimates the sellers' beliefs concerning other sellers' acceptance decisions and about buyers' portfolio choices using linear, random-effects models. The dependent variables are "sBeliefB(%)" and "sBeliefS(%)," which are the seller's incentivized beliefs about the buyers' average allocation of their endowment to payment 2 and the percentage of the other six sellers (excluding themselves) who would accept payment 2 in the current market. For the independent variables, besides the four control variables used in the regression on buyer payment choices, there are four variables that capture the historical data: "sOtherAcceptL(%)," which is the percentage of sellers who accepted payment 2 in the previous market (this historical information was revealed to sellers at the start of stage 1 of each new market when they had to make a new payment adoption decision); "sAcceptL," which is the seller's own acceptance decision in the previous market; "sPay2DealL(%)," which is the percentage of transactions the seller succeeded in conducting using payment 2 in the previous market; and "sNoDealL(%)," which is the percentage of no trade outcomes the seller encountered in the previous market. The second-stage regression is a probit, random-effects model of the seller's acceptance decision. The dependent variable, "sAccept," is a binary variable that is equal to 1 if the seller accepts payment 2, and 0 otherwise. The explanatory variables are the estimated belief terms from the first-stage regression, sBeliefB(%) and sBeliefS(%), and the four control variables. For the second-stage regression, Table 7 shows the marginal effects of the explanatory variables on the seller's probability (in %) of accepting payment 2.

From the first column in Table 6, we see that sellers adjust their beliefs about buyers' payment 2 balances upward if the sellers completed more payment-2 transactions (sPay2DealL(%)) or encountered more no-trade meetings (sNoDealL(%) in the previous market. This result is intuitive. If the seller accepted payment 2 in the previous market, then more payment-2 meetings imply that buyers carry higher payment-2 balances. If the seller did not accept payment 2 in the previous section, then more no-trade meetings also imply that buyers have higher payment-2 balances. In either case, the seller adjusts sBeliefB(%) upward. The negative coefficient before sAcceptL means that sellers who accepted payment 2 in the previous market tends to start at a lower base probability.[29] The positive coefficient before sOtherAcceptL(%) suggests that sellers expect buyers to increase their payment 2 balances after observing a higher market acceptance rate for payment 2 in the previous market. The coefficient on the location dummy is small and not statistically significant which rationalizes our pooling of the data from the two locations. There is no significant time trend in sBeliefB(%) as indicated by the insignificance of the market number variable. The coefficients on the two treatment dummies have the expected signs (positive for T16 and negative for T35), but are not statistically significant at the 10% significance level.

Sellers' beliefs about other sellers' payment 2 acceptance choices follow a similar pattern to their beliefs about buyers' allocation of endowments to payment 2. As the second column of Table 6 reveals, sellers expect the fraction of other sellers accepting payment 2 to increase if they completed more payment 2 transactions or encountered more no-trade meetings in the previous market. They also expect the higher acceptance rate in the previous market to persist into the current market (the coefficient on sOtherAcceptL% is 0.457 and the p-value is close to 0). The treatment dummies have the same signs as in the regression for sellers' beliefs about buyers' portfolio choices, but are statistically significant (the

[29] We have sAcceptL in the stage 1 regression specification so that sellers who accepted payment 2 and those who did not could have different intercepts. Those who accepted respond to the number of payment 2 transactions (and they have no no-trade meetings), and those who did not accept respond to the number of no-trade meetings (and they have no payment-2 meetings). Given that the sellers respond to different past variables, it is reasonable to not restrict that they have the same intercept. Note that although sAcceptL is negative, the predicted sBeliefB(%) tends to be higher if sAcceptL=1, because it depends on other variables in the regression specification too, such as the value and slope of sCardDeal(%) and sNoDeal(%).

p-value is 0.030 for T16 and 0.054 for T35). Sellers' beliefs about other sellers do not exhibit significant time trend or location effects.

Finally, as column 3 of Table 6 reveals, the seller's decision to accept payment 2 depends crucially on their beliefs about buyers' portfolio choice. In the second-stage regression, the coefficient on sBeliefB(%) is 1.008 and has a p-value close to 0. In words, the probability of a seller accepting payment 2 increases by 1% if they expect buyers to increase their allocation to payment 2 by 1%. The effect of sellers' beliefs about other sellers' payment 2 acceptance decisions is small in magnitude and not statistically significant. Thus it is principally sellers' beliefs about *buyer* behavior that matter most for their acceptance decisions. The treatment effects have the expected signs; relative to the $T = 2.8$ treatment, the sellers' acceptance rate is more than 15% higher in treatment $T = 1.6$, and the difference is significant at the 10% level. There is again an absence of any significant time trend or location effect. We summarize our findings for seller behavior as follows:

Finding 8 *Sellers' acceptance of payment 2 depends on T and on sellers' beliefs about buyer allocations to payment 2; these beliefs, in turn, depend on historic market outcomes.*

Summarizing the IV regression results, we have found that both buyers and sellers adjust their behavior in a manner that is consistent with our theory. Buyers' portfolio choices depend crucially on their beliefs about sellers' acceptance decisions, and, vice versa, sellers' acceptance decisions depend crucially on their beliefs about buyers' portfolio choices, suggesting that the payment choice in our experiment exhibit a strong *network effect*. The payment system evolves over time as subjects adjust their beliefs, and therefore their choices, in response to historic data. We have also confirmed the strong treatment effects (either directly on choices or indirectly through beliefs) that we report on earlier in finding 1. In addition, we have not found evidence for any strong location effects in the decisions made by buyers or sellers, despite conducting our experiment using student subjects at two different universities, SFU and UCI.[30]

Finally, we discuss what may have driven the strong network effects running in both directions in our paper, which are absent in the previous work by Camera et al. (2016) Regarding the network effects from sellers to buyers, one contributing factor may be the difference in buyers' adoption decisions. In our setting, buyers make a continuous portfolio allocation of funds between the two payment methods. In Camera et al. (2016) buyers must allocate *all* of their money to one of the two payment methods; this more restrictive money allocation choice may make buyers less sensitive to sellers' adoption decisions.[31]

[30]Complementary to the regression analysis, in online Appendix F, we provide histograms of buyers' and sellers' beliefs at the start of three markets (1, 10 and 20) for each treatment to show how these beliefs evolved over time. One observation is that the initial, market 1 beliefs tend not to differ significantly across the four treatments. However, these beliefs evolve over time and become significantly different from one another in the later markets across the different treatments. Given the strong dependence of payment decisions on beliefs, the dynamic pattern of beliefs also translates into payment adoption. As pointed out in the previous section, the payment adoption and use variables also started out with roughly similar initial conditions and also diverged over time.

[31]We also run IV regressions for buyers separately for each treatment (see table C.1 in the online Appendix C), and we find that, in all three treatments, buyers' portfolio choices depend significantly on their beliefs about sellers' acceptance decisions.

Table 6: Seller Acceptance of Payment 2

	(1) Stage 1:sBeliefB(%)	(2) Stage 1:sBeliefS(%)	(3) Stage 2:sAccept
sBeliefB(%)			**1.008*****
			(0.190)
sBeliefS(%)			**-0.093**
			(0.186)
sAcceptL	**-14.483*****	**-0.482**	
	(4.962)	(4.262)	
sPay2DealL(%)	**0.379*****	**0.225*****	
	(0.047)	(0.047)	
sNoDealL(%)	**0.370*****	**0.402*****	
	(0.074)	(0.054)	
sOtherAcceptL(%)	**0.300*****	**0.457*****	
	(0.044)	(0.053)	
market	**-0.012**	**0.041**	**0.149**
	(0.132)	(0.146)	(0.161)
location	**2.203**	**-0.913**	**0.080**
(SFU=1;UCI=0)	(2.607)	(1.571)	(3.833)
T16	**4.700**	**6.366****	**15.295***
	(2.989)	(2.927)	(8.984)
T35	**-2.759**	**-4.244***	**-4.394**
	(2.807)	(2.201)	(3.771)

Notes. (1) *p-value≤0.1; p-value≤0.05; p-value≤0.01. (2) Each regression has 1596=12x7x19 observations. There are 12 sessions, each with 7 sellers and 20 markets. For each individual, we have 19 observations as the first-stage regression uses a lagged variable as an independent variable. (3) The regressions have clustered errors at the individual subject level. (4) For the stage-2 regression, coefficient represent the marginal effect on the probability of sellers accepting payment 2.

Regarding the network effects from buyers to sellers, the most likely explanation is the existence of a fixed adoption cost in our setting, which makes sellers respond more strongly to buyers' payment allocation decisions. In particular, sellers would reject the new payment method if buyers hold only the existing method. In Camera et al. (2016), the fixed adoption cost is zero. Although the exclusive use of the existing payment method comprises a Nash equilibrium, it is not robust to trembles: sellers are indifferent between accepting and rejecting the new payment method in this equilibrium and have the incentive to accept the new payment method to minimize payment mismatch if there is no cost of doing so. This could explain why seller adoption is consistently high in Camera et al. (2016) regardless of buyers' choices. To provide some evidence about this conjecture, we run the IV regressions for sellers for each treatment separately. The results are shown in Table 7. There we see that, as in Table 6, seller's acceptance of payment 2 depends on their beliefs about buyer's acceptance of payment 2, but only when the fixed cost T is sufficiently high (greater than 1.6). We further observe that, by contrast with Table 6, when we disaggregate seller behavior by treatment, sellers' beliefs about the acceptance decisions of other *sellers* (the coefficient for sBeliefS(%) in stage 2 regression) now becomes significant in explaining the sellers' adoption decisions when T is sufficiently high (again greater than 1.6). Thus, the belief interaction effects tend to be significant only when the fixed costs are sufficiently high. We summarize these results as follows.

Finding 9 *The fixed adoption cost is important for the network effect running from the buyer's portfolio choice to the seller's adoption choice. This effect is insignificant for the treatment with the lowest fixed cost (when $T = 1.6$), but becomes significant for more substantial fixed costs (when $T = 2.8$ or 3.5).*

Other design features may also have played a role in driving the difference in network effects between the two papers. In our experiment, buyers meet all sellers and vice versa in each period and subjects quickly learn about the aggregate market condition, which boosts interaction between the two sides. In Camera et al. (2016), subjects visit a different partner in every period and observe only their own trading histories. The lack of market-level information may slow down the response to the other side's adoption decision. Another possible factor is that in Camera et al. (2016), subjects make both payment and terms of trade decisions. Such multi-tasking may be too challenging for subjects and cause them to focus on the more salient fee/reward structure rather than the coordination problem. In our experiment, we fix the terms of trade so that subjects can concentrate on the choice of payment methods and the related coordination problem.

6 An Evolutionary Learning Model of Payment Choice

Whereas our theoretical model is static, the adoption of a payment method is inherently a dynamic process. Our experiment suggests that this dynamic process involves some learning over the repeated markets of our design. Toward a better understanding of this dynamic learning process, in this section we develop an evolutionary learning model to understand behavior in our payment adoption game and we compare simulation results using that model with our experimental data. After establishing that our evolutionary model provides a good fit to our experimental data, we use that model to predict outcomes for a fourth experimental treatment. That is, we use the evolutionary model for experimental design. We then carry out an additional experimental treatment and again find a good fit between the evolutionary model and this new experimental data.

6.1 The Individual Evolutionary Learning Model

The learning model we use is the individual evolutionary learning model (IEL), which has been successfully used to characterize and predict the behavior of human subjects in many different economic environments (see, for example, Arifovic and Ledyard, 2007, 2011, 2012, 2018; Arifovic, Boitnott and

Table 7: Seller Acceptance (1=Accept, 0=Reject), Probit with Random Effects by Treatment

		T=1.6	T=2.8	T=3.5
	sAcceptL(%)	**-15.252***	**-19.577*****	**4.043**
		(9.212)	(7.222)	(5.084)
	sPay2DealL(%)	**0.361** ***	**0.368** ***	**0.148** **
		(0.069)	(0.082)	(0.075)
Stage 1: sBeliefB(%)	sNoDealL(%)	**0.231** **	**0.224**	**0.404** ***
		(0.094)	(0.153)	(0.108)
	sOtherAcceptL(%)	**0.185**	**0.301*****	**0.262** ***
		(0.126)	(0.060)	(0.097)
	market	**0.387** **	**-0.147**	**-0.238**
		(0.191)	(0.269)	(0.359)
	location	**10.837** ***	**3.374**	**-5.225**
	(SFU=1;UCI=0)	(3.947)	(3.008)	(5.541)
	sAcceptL(%)	**12.152***	**-0.179**	**1.935**
		(6.627)	(5.956)	(7.893)
	sPay2DealL(%)	**0.173*****	**0.124****	**0.245*****
		(0.052)	(0.055)	(0.093)
Stage 1: sBeliefS(%)	sNoDealL(%)	**0.335*****	**0.195****	**0.523*****
		(0.086)	(0.079)	(0.099)
	sOtherAcceptL(%)	**0.304** ***	**0.297** ***	**0.501** ***
		(0.081)	(0.052)	(0.106)
	market	**0.337** ***	**0.195**	**0.025**
		(0.094)	(0.166)	(0.493)
	location	**4.878** *	**8.832** ***	**-9.433****
	(SFU=1;UCI=0)	(2.648)	(2.433)	(4.034)
	sBeliefB(%)	**0.253**	**1.345** ***	**1.369** ***
		(0.257)	(0.180)	(0.064)
Stage 2: sAccept	sBeliefS(%)	**0.067**	**-0.958** ***	**-0.683** ***
		(0.153)	(0.234)	(0.236)
	market	**0.109**	**0.955** **	**0.190**
		(0.157)	(0.400)	(0.455)
	location	**-1.597**	**11.637** *	**-0.350**
	(SFU=1;UCI=0)	(2.753)	(6.553)	(7.653)

Notes. (1) *p-value≤0.1; **p-value≤0.05; *** p-value≤0.01. (2) Each regression has 532=4x7x19 observations. There are 4 sessions, each with 7 buyers and 20 markets. For each individual, we have 19 observations as the first-stage regression uses a lagged variable as an independent variable. (3) The regressions have clustered errors at the individual subject level.

Duffy, 2019). IEL belongs to the evolutionary algorithm class of machine learning models and provides a flexible, dynamic approach to modelling strategic interactions over time.[32]

The environment in which we simulate the IEL model corresponds precisely to our experimental environment, i.e., it is inhabited by seven buyers and seven sellers, and it lasts for 20 markets (periods) with each period involving a first payment choice stage followed by a second stage of seven trading rounds. The sequence of events in each period and round exactly follows the experimental design. The artificial agents have the same amount of information that the human subjects had at every decision node. Each of the seven artificial buyers and sellers in our evolutionary model has as a set of J rules; each rule consists of a single number. For buyer $i \in \{1, 2, ..., 7\}$, a rule $m^i_{b,j}(t) \in \{0, 1, ..., 7\}$, where $j \in \{1, 2, ..., J\}$, represents the number of EM units the buyer places in payment 2 in market (period) t. For seller $i \in \{1, 2, ..., 7\}$, a rule $m^i_{s,j}(t) \in [0, 1]$, where $j \in \{1, 2, ..., J\}$, represents the probability that the seller accepts payment 2. When a particular seller rule is selected as an actual rule, a random number between 0 and 1 is drawn from a uniform distribution over $[0, 1]$. If that number is less than or equal to $m^i_{s,j}(t)$, the seller decides to accept payment 2. Otherwise, the seller decides not to accept payment 2 for that market.

The initial set of J rules (strategies) for all buyers and sellers is randomly chosen.[33] Then, for the initial market, one strategy is randomly chosen for each buyer and each seller from their initial set of J strategies and those are the strategies played in the first market. Thereafter, the updating of the buyers' and sellers' sets of rules and the strategy chosen by each player takes place at the end of each seven-round market and consists of four steps (further details are provided in online Appendix D):

1. Experimentation. Experimentation introduces new, alternative rules that otherwise might never have a chance to be tried. It ensures that a certain amount of diversity is maintained. This operation involves changing each element of the current set of rules to a new rule with some probability. The new rule is drawn from a normal distribution, with mean equal to the current rule.[34]

2. Foregone payoff calculation. As the rules that our IEL sellers and buyers use are simple, the foregone payoff calculation is what gives the algorithm its power. It evaluates how each strategy would have performed had it been used in the just completed period. For the buyer, the relevant variable for calculating the foregone payoff is the number of sellers who accepted payment 2, which subjects learn at the end of each market. For the seller, there are two relevant variables. The first

[32]There are certainly other learning approaches that could be used, such as reinforcement learning or experience weighted attraction. For a comparison between these alternative approaches and the IEL model we use, see, e.g. Arifovic and Ledyard (2004, 2011).

[33]To be precise, for the buyer, the initial set of rules is drawn from a uniform distribution from 0 to 7 scaled by the average bPay2% in Table 2 over 50. For the seller, it is drawn from a uniform distribution from 0 to 1 scaled by the average sAccept% over 50. For example, for $T = 2.8$, the scale is 60/50 for buyers, and 89/50 for sellers. The purpose of the scaling is to match the average starting point in the experimental sessions.

[34]If the draw from the normal distribution is larger (smaller) than the upper (lower) limit of the choice set, limit of the choice set.

is the number of sellers excluding seller i that accepted payment 2; this variable can be accurately inferred from the market-wide seller acceptance statistic. The second is the average allocation to payment 2 by all 7 buyers. In our experiment, this information is not public knowledge, and the seller must try to deduce it based on the available information. The deduction is seller-specific and depends on each seller's experience from the previous period.

3. Replication. Replication reinforces rules that would have been good choices in previous periods. Specifically, rules with higher foregone payoffs are more likely to replace those with lower foregone payoffs. We implement replication using the following tournament process: we randomly pick two members of the current set of J rules with replacement, compare their forgone payoffs, and place the rule with the higher payoff in the new set of rules; we repeat this process J times so that the new set of rules has J members.

4. Selection. Following experimentation, foregone payoff calculation and replication, we randomly choose one rule from the new set of J rules based on foregone payoffs; the probability with which a strategy is selected is proportional to its relative foregone payoff among the player's set of J rules.

6.2 Simulation Results and Comparison with the Experimental Data

We adapted the IEL model in the manner described above to capture the dynamics of our experimental data. There are three free parameters for the IEL model. We set the size of the rules, $J = 150$, the experimentation rate, $\mu = 0.35$, and the standard deviation of experimentation, $\sigma = 3.5$. The parameters are the same for buyers and sellers. These parameter choices are based on prior work by Arifovic and associates and were not chosen to fit our experimental data. The standard deviation of the experimentation rate was chosen to be the mid-point of the buyer's strategy space.

Figure 5 plots the path of the payment choice (buyer's allocation to payment 2, seller's acceptance of payment 2) averaged across the four experimental sessions (first row), as well as 50 simulated IEL sessions (second row) for each of the three treatments (T=1.6, 2.8 and 3.5). The IEL simulated paths exhibit patterns that are strikingly similar to those of the experimental sessions. In particular, we observe that for $T = 1.6$ treatment, there is quick convergence to the all-payment-2 equilibrium in both the simulated and experimental data averages. For the $T = 2.8$ treatment, there is a mixture of outcomes: the simulated economies can either linger in the middle ground between the all-payment-1 or all-payment-2 equilibria, or slowly converge to the all-payment-2 equilibrium, just as is the case in the experimental data. Finally, for the $T = 3.5$ treatment, the simulated sessions either slowly converge to the all-payment-1 equilibrium or linger in the middle ground between the two pure strategy equilibria. The average paths again match those of the experimental data.

Encouraged by the good fit between the experimental and IEL model simulations for the first three treatments, we used the IEL algorithm to run simulations with higher values for T. We were interested in understanding how large T needed to be in order for there to be coordination on the all-payment-1 equilibrium within the time frame of our experimental sessions (20 markets). That is, we used the IEL for experimental design. We found that the simulated economies almost always converged to the all-payment-1 equilibrium within 20 markets for $T = 4.0$ and higher. To further validate the ability of the IEL model to characterize our experimental data, we ran a new experimental treatment with $T = 4.5$ to confirm this prediction. We again ran 4 sessions of this new $T = 4.5$ treatment, each with 14 subjects, following the same procedures as for the other three treatments.

As the fourth column of Figure 5 reveals, the IEL model simulations for the T=4.5 treatment (average of 4 runs in row 2 and of 50 runs in row 3) are a close match to the experimental data (average of 4 sessions in row 1). As predicted by the IEL model, all four experimental sessions are very close to the all-payment-1 equilibrium by the end of the experiment. For a more disaggregated view of the experimental data, see Figure E.1 in online Appendix E, which (similar to Figures 2 to 4) graphs the time path of first stage payment choices and the second stage payment usage for the four new experimental sessions of the $T = 4.5$ treatment.

Since the IEL algorithm provides a dynamic model of behavior that approximates well the path taken by subjects in our experiment we use the IEL model to carry out another interesting analysis. As shown in the previous section, the experimental sessions with T=2.8 and T=3.5 often do not converge to either symmetric equilibrium within the time frame of 20 markets that we allowed in our experiment. That is, the experimental sessions are constrained in terms of the number of markets that we can run, by the need to avoid subject boredom, and to allow for meaningful subject payments. One useful exercise is to use the IEL model to gain some insights about whether the economy would eventually converge to one of the two symmetric pure strategy equilibria and, if yes, which equilibrium will be chosen and how long will that transition process take?[35] We first establish some criteria for declaring convergence of the IEL algorithm to an equilibrium outcome, and we also construct an index that measures the stability of that convergence outcome over next 100 periods (markets); the details of these definitions and the results of the simulations are provided in online Appendix D. Then, we ran 50 simulations of the IEL algorithm for each of the four values for T, and computed these measures. We find that for the $T = 1.6$ and $T = 2.8$ treatments, the IEL simulations eventually converge to the all-payment-2 equilibrium. The mean time to convergence to the all-payment-2 equilibrium (using our convergence criterion) for the $T = 1.6$ treatment is 8.5 time periods (markets) with a small standard deviation (1.9) and very high index of stability of the equilibrium outcome equal to 98%. On the other hand, for the $T = 2.8$ treatment, the mean times to convergence to the all payment 2 equilibrium is considerably longer, 76.9 markets with a larger standard deviation (st. dev. 47.8). However, once convergence occurs, the index of stability for the $T = 2.8$ treatment is 88%. For the $T = 3.5$ and $T = 4.5$ treatments, the IEL simulations eventually converge to the all payment 1 equilibrium. In the $T = 3.5$ treatment, this takes on average, 49.2 markets (std. dev. 28.3) with an

[35]To achieve convergence, we specify that the experimentation rate and standard deviation decay over time. See Appendix D for details.

index of stability of 90%. For the $T = 4.5$ treatment, the mean number of periods (markets) required for convergence to the all payment 1 equilibrium is low, just 7.4 periods with a standard deviation of only 1.6, and the value of the index of stability is equal to 98%. In Appendix D we also report on some further simulation exercises that explore the "tipping point" value of T for which adoption of the new payment method occurs or does not occur as well as the effects of reducing the seller's payment costs which increases their gains from trade.

6.3 Analysis of Treatment with T=4.5

In this subsection we briefly analyze the experimental data from the $T = 4.5$ treatment (more details are provided in Figure E.1 and Tables E.1 to E.6 in online Appendix E). The four experimental sessions of the new, $T = 4.5$ treatment converge very close to the all-payment-1 equilibrium. Averaged across the four sessions, Table E.1 reveals that by last market, buyers allocate just 7% of their endowment to payment 2 meaning that they are allocating an average of 93% to payment 1; on average just 14% of sellers (1 out of 7) choose to accept payment 2; and 93% of transactions are conducted with payment 1.[36] Compared with the other three treatments, payment mismatch is more serious in this new $T = 4.5$ treatment: averaged across the 20 markets and the four sessions, 12% of meetings result in no trade, versus 9% for $T = 3.5$, 6% for $T = 2.8$ and 1% for $T = 1.6$. These treatment differences are also significant according to rank-sum tests (see Table E.3) and our convergence test (see Table E.4).

The analysis of buyers' allocation of endowment to payment 2 as reported in Table E.5, which uses the same IV approach as in Table 5, suggests that buyers respond strongly to their beliefs about the sellers' acceptance decisions: the coefficient on the term bBelief(%) is 0.763, meaning that a 1% increase in this belief about seller acceptance of payment 2 induces a 0.763% increase in allocation of endowment to payment 2. Buyers adjust their beliefs about sellers' acceptance decisions after observing the market acceptance rate in the previous market. Over time, buyers reduce their allocations to payment 2: the coefficient on the market number variable is -0.623 (from market 1 to 20, the reduction is about 12%). Subjects at SFU tend to believe that more sellers will accept payment 2 (than subjects at UCI), but the magnitude of the difference is small at about 5%. The regression for sellers' acceptance decisions in Table E.6, following the same IV approach as Tables 6 and 7, suggests that sellers' beliefs again respond to historic market outcomes in ways that are similar to the first three treatments. Sellers' acceptance decisions respond to their beliefs. However, the coefficients on the (estimated) belief terms are not statistically significant for treatment $T = 4.5$, largely because many sellers quickly decide not to accept payment 2 due to the high fixed cost. This is similar to treatment $T = 1.6$, where many sellers quickly decide to accept payment 2 due to the low fixed cost.

[36] The buyer's payment choice passes the convergence test (to the all-payment-1 equilibrium) in all four sessions. The seller's payment choice passes this convergence test only in session 2; the other sessions do not pass the convergence test mainly because one or two sellers switched between accepting and rejecting payment 2. The percentage of payment-2 meetings pass the convergence test in three of the four sessions.

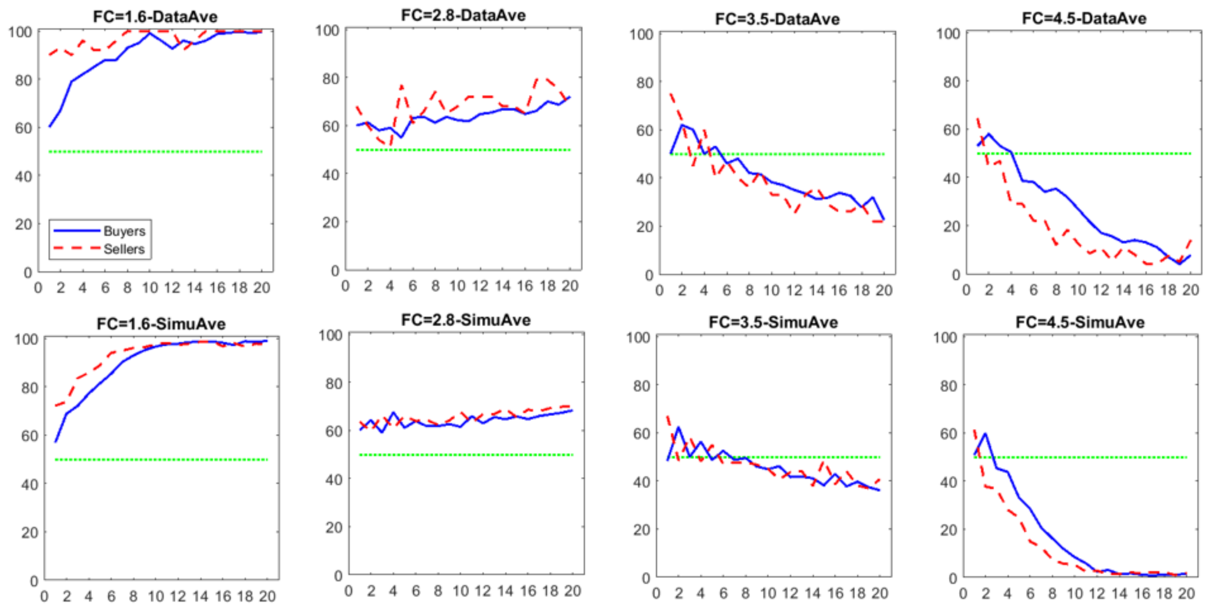

Figure 5: Payment Choice Averaged across Simulated Sessions

Notes. (1) The four figures in the first row represent the average across the four experimental sessions for each of the four treatments; and the four figures in the second row represent the average across 50 simulated sessions. (2) Horizontal axis: market. (3) The blue solid line represents the percentage of money allocated to payment 2, averaged across the seven buyers; the red dashed line represents the percentage of the seven sellers accepting payment 2.

7 Conclusion and Directions for Future Work

We have developed an experimental framework to understand factors that may contribute to the adoption of a new additional payment method when there already exists a payment method that all sellers accept. To complement theoretical and empirical studies on payment choice and adoption, we design our experiment so that we can observe how the two sides of the payment market, consumers and merchants, coordinate their adoption decisions over time. We elicit players' beliefs about the other side adoption decision to test directly the network effects, which are emphasized by the theoretical literature but pose a challenge for empirical works to directly test. Among the many possible factors that could affect the adoption of a new payment method, we focus on one important factor, the cost structure of the payment method, and in particular, the fixed cost that the seller must incur (relative to saving in per transaction cost) associated with the adoption of a new additional payment method. We use the experimental framework to study how this factor affects the adoption process and test directly the network adoption effects.

We find that the new payment method will take off if the fixed cost is low so that both sides benefit by switching to the new payment method. If the fixed cost is high such that the seller endures a loss in the equilibrium where the new payment method is used relative to the equilibrium where it is not accepted, some sellers nevertheless respond by accepting the new payment method initially, fearing to lose business, but they mostly eventually learn over time to resist the new payment method and pull the economy back to the old payment method. If neither side displays much willpower to move behavior toward one equilibrium or the other, then the economy may linger in the middle ground between the two equilibria. Importantly, we find evidence from our two-sided experiment involving both buyers and sellers of strong network effects. Buyers' portfolio choices depend on their beliefs about sellers' acceptance of the new payment method whereas sellers' acceptance of the new payment method depends on their beliefs about buyer's portfolio choices. Such direct evidence for network effects would be difficult to document or obtain outside of the laboratory. The fixed merchant adoption cost is important for generating the strong feedback effect from consumers to merchants.

We also find that an evolutionary learning model, derived from the machine learning literature, captures the dynamic process of payment choices well. Simulations of this evolutionary model provide a good fit to our experimental data, and we show how the learning model is useful for the design of new experimental treatments, by conducting an additional treatment to validate the model's prediction. Further, we show that our empirically validated evolutionary learning algorithm can be useful in assessing longer-run outcomes, beyond the time frame of a human subject experiment.

In this study we focus on the effect of the cost structure of new payment method. The framework we have developed in this paper can be used to address a number of interesting questions. For instance, (1) what are the effects of subsidies on the adoption of a new payment method? Our current results suggest that a sufficiently large subsidy to sellers in the $T = 3.5$ treatment would work to promote the adoption of the new payment method, e.g. a subsidy of 0.7 that reduced the fixed costs of adoption from 3.5 to 2.8. (2) In the present model we fix the terms of trade, but it would be of interest to allow sellers to offer discounts or impose surcharges and ask how that affects the adoption of the new payment method. (3) The baseline model also assumes a single new payment method. It would be of interest to further explore how the economy evolves if there is more than one new payment method. Imagine the simplest case featuring two new payment methods with the same setup costs. Competition between the two new payment methods may make it difficult for agents to coordinate on a single new payment method. As a result, it is possible that neither payment method would take off, or it may take a longer time for the economy to converge. (4) It would be interesting to isolate the desirable features of an existing payment method that a new payment method would need to have in order to make it more competitive. Anonymity? Robustness to network breakdowns? What is the effect of an incidence of identity theft? Suppose the economy has arrived at the equilibrium with the new payment method. Can the economy revert to the equilibrium with the old payment method, and if so, under what circumstances? (5) Suppose consumers receive their income in the form of either the old or the new payment method and assume there is a small cost of portfolio adjustment at the beginning of each market. Do these changes matter for the equilibrium that is selected? (6) The baseline model assumes that sellers always accept the old payment method. What would happen if we relax this assumption and allow sellers to decide whether to accept the old payment method? (7) The baseline model assumes that consumers cannot convert between the old and the new payment methods after the initial portfolio choice. As a result, if the consumer has only the new payment method but the seller does not accept it, there will be no trade. In such instances, we could modify the model to allow for conversion opportunities subject to a cost. Note that this new specification does not change the equilibrium outcomes. In equilibrium, buyers' portfolio choices will be fully consistent with the sellers' acceptance decisions and there is no need to convert. However, the conversion cost(s) may affect the speed of convergence to equilibria. We believe that all of these extensions of our model are worth examining theoretically and/or experimentally, but we must leave such an analysis to future research.

References

Arango, Carlos, Kim P. Huynh, and Leonard Sabetti (2016) "Consumer payment choice: Merchant card acceptance versus pricing incentives," Journal of Banking and Finance 55, 130-141.

Arifovic, Jasmina, Joshua F. Boitnott, and John Duffy (2019) "Learning correlated equilibria: an evolutionary approach," Journal of Economic Behavior and Organization 157, 171-190.

Arifovic, Jasmina and John Ledyard (2004) "Scaling Up learning models in public good games," Journal of Public Economic Theory, 6 (2), 2004, pp. 203–238.

Arifovic, Jasmina, and John Ledyard (2007) "Call Market Book Information and Efficiency," Journal of Economic Dynamics and Control 31, 1971-2000.

Arifovic, Jasmina, and John Ledyard (2011) "A behavioral model for mechanism design: Individual evolutionary learning," Journal of Economic Behavior and Organization 78, 374-395.

Arifovic, Jasmina, and John Ledyard (2012) "Individual evolutionary learning, other regarding preferences, and the voluntary contribution mechanism," Journal of Public Economics, 96, 808-823.

Arifovic, Jasmina, and John Ledyard (2018) "Learning to alternate," Experimental Economics 21(3), 692-721.

Bagnall, John, David Bounie, Kim P. Huynh, Anneke Kosse, Tobias Schmidt, Scott Schuh, and Helmut Stix (2016), "Consumer cash usage: a cross-country comparison with payment diary survey data," International Journal of Central Banking.

Bátiz-Lazo, Bernardo and Tony Moretta (2016) "Mondex and VisaCash: A First (Failed) Attempt at an Electronic Purse," in: Batiz-Lazo B., Efthymiou L. (Eds.) The Book of Payments. Palgrave Macmillan, pp. 177-185.

Boar, Cordruta, Henry Holden and Amber Wadsworth (2020) "Impending Arrival - A Sequel to the Survey on Central Bank Digital Currency."Bank for International Settlements Working Paper, NO. 107.

Bounie, David, Abel François, and Leo Van Hove (2017a) “Consumer payment preferences, network externalities, and merchant card acceptance: An empirical investigation,” Review of Industrial Organization 51(3), 257-290

Bounie, David, Abel François, and Leo Van Hove (2017b) “Merchant acceptance of payment cards: ‘must take or wanna take’?” Review of Network Economics, 117-146.

Camera, Gabriele, Marco Casari, and Stefania Bortolotti (2016), “An experiment on retail payment systems,” Journal of Money, Credit and Banking 48, 363-392.

Chiu, Jonathan, and Tsz-Nga Wong (2014) “E-money: efficiency, stability and optimal policy,” Bank of Canada Staff Working Paper 2014-16.

Devetag, Giovanna and Ortmann, Andreas (2007) “When and why? A critical survey on coordination failure in the laboratory,” Experimental Economics 10, 331–344.

Duffy, John and Jack Ochs (1999) “Emergence of money as a medium of exchange: an experimental study,” American Economic Review 89, 847-77.

Duffy, John and Jack Ochs (2002) “Intrinsically worthless objects as media of exchange: experimental evidence,” International Economic Review 43, 637-73.

Dufwenberg, Martin (2015). “Banking on experiments?” Journal of Economic Studies 42(6), 943-971.

Cohen, Michael, and Marc Rysman (2013) “Payment choice with consumer panel data,” Working Paper 13-6, Federal Reserve Bank of Boston.

Evans, David S. (2011) “Interchange fees: the economics and regulation of what merchants pay for cards,” Boston: Competition Policy International.

Fischbacher, Urs (2007) “z-Tree: Zurich toolbox for ready-made economic experiments,” Experimental Economics 10(2), 171-178.

Hall, Bronwyn H. and Beethika Khan (2003), “Adoption of new technology,” In: Jones, Derek C. (ed.), New Economy Handbook, Amsterdam: Elsevier Science.

Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels (2004) “The theory of global games on test: experimental analysis of coordination games with public and private information,” Econometrica 72, 1583-1599.

Hossain, Tanjim, Dylan Minor, and John Morgan (2011) “Competing matchmakers: an experimental analysis,” Management Science 57(11), 1913-1925.

Huynh, Kim P., Gradon Nicholls, and Oleksandr Shcherbakov (2019)“Explaining the interplay between merchant acceptance and consumer adoption in two-sided markets for payment methods,” Bank of Canada Staff Working Paper 2019-32.

Huynh, Kim P., Philipp Schmidt-Dengler, and Helmut Stix (2014) “The role of card acceptance in the transaction demand for money,” Bank of Canada Staff Working Paper 2014-44.

Klee, Elizabeth (2008), “How people pay: evidence from grocery store data,” Journal of Monetary Economics 55, 526–541.

Kosse, Anneke, Heng Chen, Marie-Hélène Felt, Valéry Dongmo Jiongo, Kerry Nield, and Angelika Welte (2017) “The costs of point-of-sale payments in Canada,” Bank of Canada Staff Discussion Paper 2017-4.

Koulayev, Sergei, Marc Rysman, Scott Schuh, and Joanna Stavins (2016) “Explaining adoption and use of payment instruments by U.S. consumers,” RAND Journal of Economics 47(2), 293-323.

Kumhof, Michael, and Clare Noone (2018) “Central bank digital currencies – design principles and balance sheet implications,” Bank of England Staff Working Paper No. 725.

Lei, Vivian and Charles N. Noussair (2002) “An experimental test of an optimal growth model,” American Economic Review 92, 549-70.

Li, Bin, James McAndrews, and Zhu Wang (2020), “Two-sided market, R&D, and payments system evolution,” Journal of Monetary Economics 115, 180-199.

Lotz, Sebastien and Françoise Vasselin (2019) “A new monetarist model of fiat and e-Money,” Economic Inquiry 57(1), 498-514.

Lotz, Sebastien, and Cathy Zhang (2016) “Money and credit as means of payment: a new monetarist approach,” Journal of Economic Theory 164, 68-100.

McAndrews, James, and Zhu Wang (2012), “The economics of two-sided payment card markets: pricing, adoption and usage,” Working Paper 12-06, Federal Reserve Bank of Richmond.

Polasik, Michal, Jakub Górka, Gracjan Wilczewski, Janusz Kunkowski, Karolina Przenajkowska, and Natalia Tetkowska (2013) “Time efficiency of point-of-sale payment methods: empirical results for cash, cards and mobile payments,” In: Cordeiro J., Maciaszek L., Filipe J. (eds) Enterprise Information Systems. ICEIS 2012. Lecture Notes in Business Information Processing, vol 141. Springer, Berlin, Heidelberg.

Rochet, Jean-Charles, and Jean Tirole (2002) “Cooperation among competitors: some economics of payment card associations,” RAND Journal of Economics 33(4), 549-570.

Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir (1991) "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," American Economic Review, 81, 1068-95.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil, (1990) "Tacit coordination games, strategic uncertainty, and coordination failure," American Economic Review 80, 234-48.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil, (1991) "Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games," Quarterly Journal of Economics 106, 885-910.

Von Kalckreuth, Ulf, Tobias Schmidt, and Helmut Stix (2014) "Choosing and using payment instruments: evidence from German microdata," Empirical Economics 46(3), 1019-1055.

Wang, Zhu, and Alexander Wolman (2016) "Payment choice and currency use: insights from two billion retail transactions," Journal of Monetary Economics 84, 94-115.

Welte, Angelika and Jozsef Monar (2021) "The market for acquiring card payments from small and medium sized Canadian merchants," Business Economics 56, 87-97.

Wright, Julian (2003), "Optimal card payment systems," European Economic Review 47, 587-612.

Online Appendix to:
"Adoption of a New Payment Method: Experimental Evidence"

by: Jasmina Arifovic (Simon Fraser University)

John Duffy (University of California, Irvine)

Janet Hua Jiang (Bank of Canada)

A Best Response Functions in the Payment Adoption Stage

A.1 Buyer's Portfolio Decision

In this subsection, we prove Proposition 1 and analyze the buyer's portfolio decision, m^b, conditional on the seller's adoption decision, m^s. We will carry out the analysis in two cases: (1) $m^b \geq m^s$, and (2) $m^b \leq m^s$.

If $m^b \geq m^s$, then each buyer makes m^s purchases using the new payment method and $1 - m^b$ purchases using the existing payment method. Buyers are not able to transact with a fraction $m^b - m^s$ of sellers because of payment mismatches (buyers want to use payment method 2 but sellers only accept payment method 1). The buyer's expected payoff in this case is:

$$\pi^b = \underbrace{m^s(u - \tau_2^b)}_{\text{transactions using payment method 2}} + \underbrace{(1 - m^b)(u - \tau_1^b)}_{\text{transactions using payment method 1}}.$$

Note that

$$d\pi^b/dm^b = -(u - \tau_1^b) < 0.$$

It follows that for this case, the optimal choice of each buyer is to reduce m^b to m^s so as to minimize the probability of a payment mismatch or no-trade outcome.

If $m^b \leq m^s$, then each buyer makes m^b transactions using the new payment method and $1 - m^b$ transactions using payment method 1 (among which $m^s - m^b$ are with sellers who also accept the new payment method). The buyer's expected payoff is now given by:

$$\pi^b = \underbrace{m^b(u - \tau_2^b)}_{\text{transactions using payment method 2}} + \underbrace{(1 - m^b)(u - \tau_1^b)}_{\text{transactions using payment method 1}}.$$

In this case we have that

$$d\pi^b/dm^b = -\tau_2^b + \tau_1^b > 0.$$

Thus, if $m^b \leq m^s$, then buyers should increase their payment 2 balances to m^s so as to minimize transaction costs.

From the analysis above it follows that buyers' optimal portfolio decision is to mimic the sellers' acceptance decision:

$$m^b(m^s) = m^s.$$

A.2 Seller's Payment 2 Acceptance Decision

In this subsection we prove Proposition 2 and derive the seller's acceptance decision conditional on the buyer's portfolio decision, m^b. We will carry out our analysis under two parameter settings: (1) $F \leq \tau_1^s - \tau_2^s$, and (2) $F > \tau_1^s - \tau_2^s$. For each parameter setting, similar to the discussion of the buyer's choice, we analyze the seller's decision in two cases: $m^b \geq m^s$ and $m^b \leq m^s$.

Parameter Setting (1): $F \leq \tau_1^s - \tau_2^s$. If $m^b \geq m^s$, then each seller who accepts the new payment method engages in a unit measure of payment-method-2 transactions (remember that buyers use the new payment method whenever the seller accepts it), and has a payoff of

$$\pi_2^s = 1 - \tau_2^s - F.$$

Sellers who only accept payment method 1 engage in an average of $(1-m^b)/(1-m^s) \leq 1$ transactions using payment method 1 (the total payment method 1 balance in the economy is $1-m^b$ and this is divided among the $1-m^s$ sellers who only accept payment 1). Sellers who accept only payment 1 thus have a payoff of

$$\pi^s = \frac{1-m^b}{1-m^s}(1-\tau_1^s).$$

In this case,

$$\begin{aligned}
(\pi_2^s - \pi_1^s)|_{m^b \geq m^s} &= 1 - \tau_2^s - F - \frac{1-m^b}{1-m^s}(1-\tau_1^s) \\
&= (\tau_1^s - \tau_2^s - F) + \frac{(1-\tau_1^s)(m^b - m^s)}{1-m^s}.
\end{aligned}$$

As long as $m^b > m^s$, we have $\pi_2^s > \pi_1^s$, i.e., each seller who accepts the new payment method is able to trade for the new payment method in all meetings, which makes it profitable to pay the fixed cost, F, to accept the new payment method. As a result, π^s will increase. In equilibrium, it must be the case that $m^b \leq m^s$.

If $m^b \leq m^s$, the payment method 2 balance in the economy can support m^b payment 2 transactions, which are divided among m^s sellers who accept payment 2. Each seller who accepts payment 2 can trade in all meetings, among which m^b/m^s will be payment 2 transactions, and the remaining $1 - m^b/m^s$ will be payment 1 transactions. The expected payoff of a seller who accepts the new payment method is therefore:

$$\begin{aligned}
\pi_2^s &= \underbrace{\frac{m^b}{m^s}(1-\tau_2^s)}_{\text{transactions using payment method 2}} + \underbrace{\left(1 - \frac{m^b}{m^s}\right)(1-\tau_1^s)}_{\text{transactions using payment method 1}} - F \\
&= (1-\tau_1^s) + \frac{m^b}{m^s}(\tau_1^s - \tau_2^s) - F.
\end{aligned}$$

Sellers who accept only payment 1 engage in payment 1 transactions in all meetings and have a payoff of

$$\pi_1^s = 1 - \tau_1^s.$$

In this case,

$$(\pi_2^s - \pi_1^s)|_{m^b \le m^s} = \frac{m^b}{m^s}(\tau_1^s - \tau_2^s) - F.$$

If $m^b \ge F/(\tau_1^s - \tau_2^s)$, then it is a dominant strategy for sellers to accept the new payment method: each seller makes more than $F/(\tau_1^s - \tau_2^s)$ sales in payment 2 to warrant the fixed investment to accept payment 2. If $m^b \le F/(\tau_1^s - \tau_2^s)$, the number of transactions using payment method 2 is not large enough to recover the fixed acceptance cost for all sellers. As a result, sellers play a mixed strategy: $m^s = m^b(\tau_1^s - \tau_2^s)/F$ fraction of sellers accept both payment methods, and the rest accept only payment method 1. All sellers earn the same expected payoff ($\pi_1^s = \pi_2^s$).

To summarize, if $F \le \tau_1^s - \tau_2^s$, then given the buyer's strategy m^b, the seller's strategy is such that

$$m^s(m^b) = \begin{cases} \frac{m^b(\tau_1^s - \tau_2^s)}{F} & \text{if } m^b \le \frac{F}{\tau_1^s - \tau_2^s}, \\ 1 & \text{if } m^b \ge \frac{F}{\tau_1^s - \tau_2^s}. \end{cases}$$

Note that if $F = \tau_1^s - \tau_2^s$, then $m^s(m^b) = m^b$.

Parameter Setting (2): $F > \tau_1^s - \tau_2^s$. Suppose that $m^b < m^s$. Then, sellers who do not accept the new payment method earn a higher payoff (i.e., $\pi_2^s - \pi_1^s < 0$). As a result, m^s will decrease. In equilibrium, it must be the case that $m^b \ge m^s$.

If $m^b \ge m^s$, it is a dominant strategy for sellers not to accept the new payment method if $m^b \le \hat{m}_b \equiv 1 - [(1 - \tau_2^s) - F]/(1 - \tau_1^s)$. If $m^b \ge \hat{m}_b$, then sellers play a mixed strategy, choosing to accept with probability $m^s(m^b) = 1 - (1 - m^b)(1 - \tau_1^s)/[(1 - \tau_2^s) - F]$, which solves $(\pi_2^s - \pi_1^s)|_{m^b \ge m^s} = 0$.

To summarize, under the parameter setting $F > \tau_1^s - \tau_2^s$, given the buyer's strategy m^b, the seller's strategy is such that

$$m^s(m^b) = \begin{cases} 0 & \text{if } m^b \le 1 - \frac{(1-\tau_2^s)-F}{1-\tau_1^s}, \\ 1 - \frac{(1-m^b)(1-\tau_1^s)}{(1-\tau_2^s)-F} & \text{if } m^b \ge 1 - \frac{(1-\tau_2^s)-F}{1-\tau_1^s}. \end{cases}$$

B Experimental Instructions, T-2.8 Treatment (other instructions similar)

Welcome to this experiment in economic-decision making. Please read these instructions carefully as they explain how you earn money from the decisions that you make. You are guaranteed $7 for showing up and completing the study. Additional earnings depend on your decisions and on the decisions of other participants as explained below. You will be earning experimental money (EM). At the end of the experiment, you will be paid in dollars at the exchange rate of 1 EM = $0.15.

There are 14 participants in today's experiment: 7 will be randomly assigned the role of buyers and 7 the role of sellers. You will learn your role at the start of the experiment, and remain in the *same* role for the duration of the experiment. Buyers and sellers will interact in 20 "markets" to trade goods for payment. There are two payment methods, payment 1 and payment 2.

Each market consists of two stages. The first is the payment choice stage. Each buyer is endowed with 7 EM and decides how to allocate it between the two payment methods. Each seller is endowed with 7 units of goods. Sellers have to accept payment 1, but can decide whether or not to accept payment 2. Sellers who decide to accept payment 2 have to pay a one-time fee of 2.8 EM. No participant observes any seller's choice at this stage.

The second stage is the trading stage, which consists of a sequence of 7 rounds. In these 7 rounds, you meet with each of the 7 participants who are in the opposite role to yourself sequentially and in a random order. In each meeting you try to trade one unit of good for one unit of payment. The buyer decides which payment to use and the trade is successful if and only if the seller accepts the payment offered by the buyer. For each successful sale or purchase, you earn 1 EM less some transaction costs. The transaction cost to both sides is 0.5 EM if payment 1 is used, and 0.1 EM if payment 2 is used. If the buyer offers payment 1 (which is always accepted by sellers), then trade is successful and both the buyer and the seller earn a *net* payoff of 1-0.5=0.5 EM. If the buyer offers payment 2 and the seller has decided to accept payment 2 in the first stage, then trade is again successful and both earn a *net* payoff of 1-0.1=0.9 EM. If the buyer has only payment 2 and the seller has decided not to accept it, then no trade can take place and both earn 0 EM. At the end of the market, unspent EMs or unsold goods have no redemption value and do not entitle you to extra earnings.

Task summary

<table>
<tr><td rowspan="2">Market 1</td><td>Stage 1: Payment choice
Buyers allocate 7 EM between the two payments
Sellers decide whether to accept payment 2 at a one-time fee of 2.8 EM</td></tr>
<tr><td>Stage 2: Trading (7 rounds)
Each buyer meets each of the 7 sellers in a random order
Trade with payment 1 ➔ net payoff of 0.5 EM
Trade with payament 2 ➔ net payoff of 0.9 EM
No trade ➔ net payoff of 0 EM</td></tr>
<tr><td rowspan="2">Market 2</td><td>Stage 1: Payment choice</td></tr>
<tr><td>Stage 2: Trading (7 rounds)</td></tr>
<tr><td>…</td><td>…</td></tr>
<tr><td rowspan="2">Market 20</td><td>Stage 1: Payment choice</td></tr>
<tr><td>Stage 2: Trading (7 rounds)</td></tr>
</table>

More Inforamtion for Sellers

As a seller, your earnings in a market (in EM) is calculated as

Option I	Accept payment 2	Number of payment 1 transactions x 0.5 + **Number of payment 2 transactions x 0.9 – 2.8**
Option II	Not accept payment 2	Number of payment 1 transactions x 0.5

The benefit to sellers of accepting payment 2 is to increase the likelihood that you sell goods to buyers (remember no trade can take place if the buyer has only payment 2 and you do not accept it), and to reduce transaction costs and therefore increase net earnings by 0.4 EM each time a buyer pays in payment 2. The cost to sellers of accepting payment 2 is that you have to pay a one-time fee of 2.8 EM at the beginning of the market even if no buyers offer to pay you with payment 2 in that market.

Which option leads to higher earnings depends on all other 13 subjects' decisions. Table 1 on page 7 lists the average market earnings for the seller from the two options (accept / reject payment 2) in cases where all buyers choose to allocate between 0~7 EM to payment 2, and where 0~6 of the other 6 sellers choose to accept payment 2. As you can see, either option can give higher earnings depending on other participants' decisions. During the experiment, please keep Table 1 at hand for reference. In addition, you can use a "what if" calculator on the computer screen to compute the average earnings in situations where buyers make different payment allocations.

Your earnings from accepting payment 2 tend to increase if more buyers allocate more money to payment 2, and if fewer sellers accept payment 2. The opposite is true if you reject payment 2.

More Information for Buyers

As a buyer, your earnings in a market are calculated as

Number of payment 1 transactions x 0.5 + Number of payment 2 transactions x 0.9

As a buyer, the benefit of allocating more money to payment 2 is that you save 0.4 EM each time you use payment 2 instead of payment 1. The cost is the risk that you may not be able to trade if the seller does not accept payment 2 and you run out of payment 1 (which is always accepted). Your market earnings depend on your own payment allocation and the 7 sellers' decisions on acceptance of payment 2. Table 2 on page 7 lists the buyer's market earnings if the buyer allocates 0~7 EM to payment 2 (and the rest to payment 1) and if 0~7 sellers accept payment 2. You should allocate more money to payment 2 if you expect more sellers to accept it. Table 2 will also be on your computer screen when you make payment decisions.

Forecast

At the start of each market before making payment decisions, you are asked to forecast other participants' choices for that market. Buyers forecast how many of the 7 sellers will choose to accept payment 2. Sellers forecast (1) the average amount of EM that all 7 buyers will allocate to payment 2, and (2) how many of the other 6 sellers will accept payment 2. You earn 0.5 EM per correct forecast in addition to your earnings from buying/selling goods.

Earnings

At the end of the experiment, you will be paid your earnings in cash and in private. Your earnings in dollars will be: Total earning (trading + forecasting) in EM x 0.15 + 7 (show-up fee).

Computer Interface

You will interact anonymously with other participants using the computer workstations. You will see three types of screens (Figures 1-6 show sample screens).

Payment choice screen, Figures 1-2. This is where you make payment choices depending on whether you are a buyer (Figure 1) or a seller (Figure 2). Each screen has 4 parts. The upper portion summarizes information about previous markets. To the left of the blank column are your own activities, including your payment choice, the number of transactions using each of the two payment methods, the number of no-trade meetings, market earning from trading, and the number of correct forecasts that you made. To the right of the blank column, there is an aggregate market-level statistic, the number of sellers who accepted payment 2.

The middle section provides information about your average potential earnings from trading in each market. The buyer screen (Figure 1) shows Table 2. The seller screen (Figure 2) has a "what if" calculator. A seller can type in the number of buyers choosing to allocate 0~7 EM to payment 2 and the number of other 6 sellers accepting payment 2 (the default value is 0 in all fields; the first 8 fields must add up to 7; enter an integer 0~6 in the last field), press the "Calculate" button to create a record showing the average market earnings from accepting payment 2 and not accepting it, as well as the average buyers' allocation to payment 2 in that scenario. For example, if you would like to check your potential average earnings in the situation where **5** buyers allocate 2 EM to payment 2, **2** buyers allocate 3 EM, and **3** of the other six sellers accept payment 2, type in "**5**" in the field "# buyers with pay2=2", "**2**" in the field "# buyers with pay2=3", and "**3**" in the field "# other sellers accept pay2." You can create as many records as you wish at the start of each market.

In the lower-left section, you forecast what other participants will do in the new market. Enter an integer within the indicated range for each forecast. The seller's forecast of buyer's average payment 2 allcoation is counted as correct if it lies within ±1 of the realized value.

In the lower-right section, you choose how to split your 7 EM between the two payment methods if you are a buyer (Figure 1), and whether to accept payment 2 at a one-time fee of 2.8 EM if you are a seller (Figure 2).

Trading screen, Figures 3-4**.** In each of the 7 trading rounds, buyers decide whether to buy a unit of the seller's good using either payment 1 or payment 2. This decision depends on the buyer's remaining balances of payment 1 and payment 2, and whether or not the seller has agreed to accept payment 2; this information is shown on the buyer's computer screen (see the lower left box in Figure 3). Sellers do not choose at this stage, and can click on the "OK" button to review information on the waiting screen (see Figure 4). From round 2 on, the upper section of the screen reviews your activities in the previous round and in the current market up until then.

Waiting screen, Figures 5-6. At any point in the experiment if you finish your decision sooner than other participants, you will see a waiting screen with information on previous markets and your potential market earnings similar to what you observe on the payment choice screen.

Finally, sellers who invest in the one-time fixed cost to accept payment 2 may have a negative "market earnings" in one or a few rounds. As a result of this, you may see a message screen explaining the situation. After you have been alerted to this situation, you can click on the "continue" button on the screen to proceed.

Figure 1: buyer's payment allocation screen

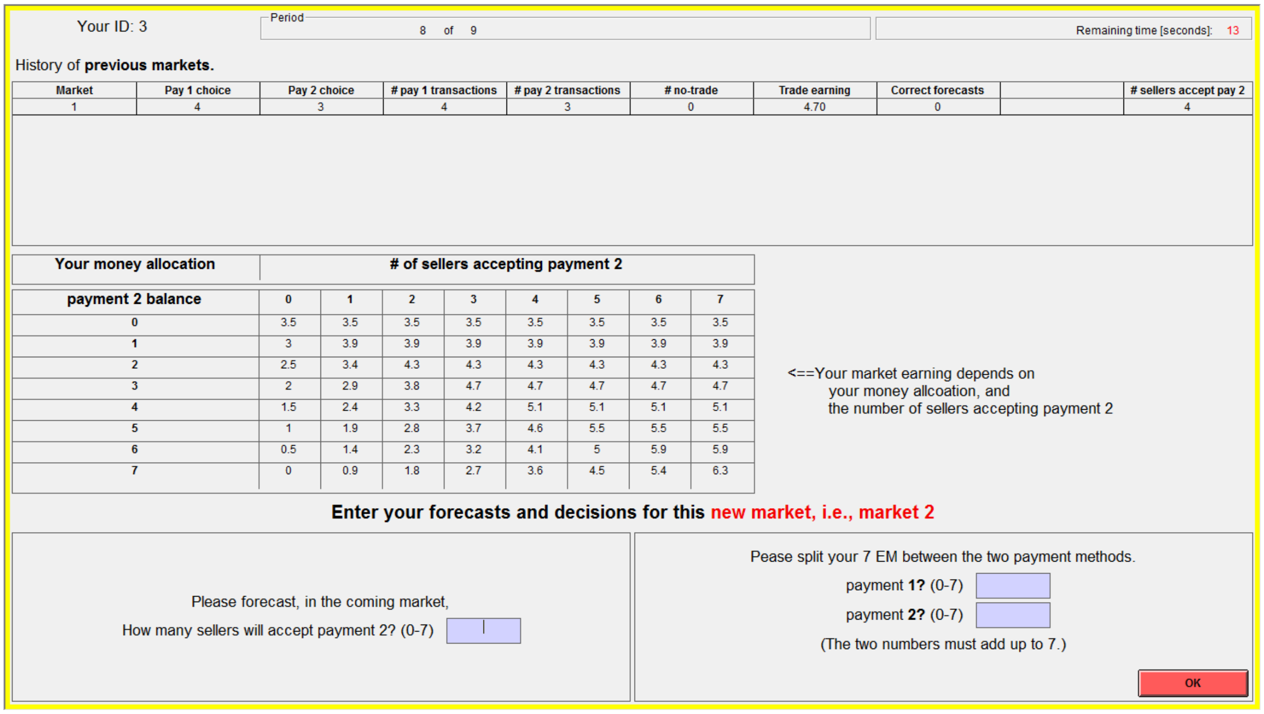

Your ID: 3

Period 8 of 9

Remaining time [seconds]: 13

History of **previous markets.**

Market	Pay 1 choice	Pay 2 choice	# pay 1 transactions	# pay 2 transactions	# no-trade	Trade earning	Correct forecasts		# sellers accept pay 2
1	4	3	4	3	0	4.70	0		4

Your money allocation	# of sellers accepting payment 2							
payment 2 balance	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
0	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5
1	3	3.9	3.9	3.9	3.9	3.9	3.9	3.9
2	2.5	3.4	4.3	4.3	4.3	4.3	4.3	4.3
3	2	2.9	3.8	4.7	4.7	4.7	4.7	4.7
4	1.5	2.4	3.3	4.2	5.1	5.1	5.1	5.1
5	1	1.9	2.8	3.7	4.6	5.5	5.5	5.5
6	0.5	1.4	2.3	3.2	4.1	5	5.9	5.9
7	0	0.9	1.8	2.7	3.6	4.5	5.4	6.3

<==Your market earning depends on
your money allcoation, and
the number of sellers accepting payment 2

Enter your forecasts and decisions for this new market, i.e., market 2

Please forecast, in the coming market,
How many sellers will accept payment 2? (0-7)

Pease split your 7 EM between the two payment methods.
payment **1?** (0-7)
payment **2?** (0-7)
(The two numbers must add up to 7.)

OK

Figure 2: seller's payment 2 acceptance screen

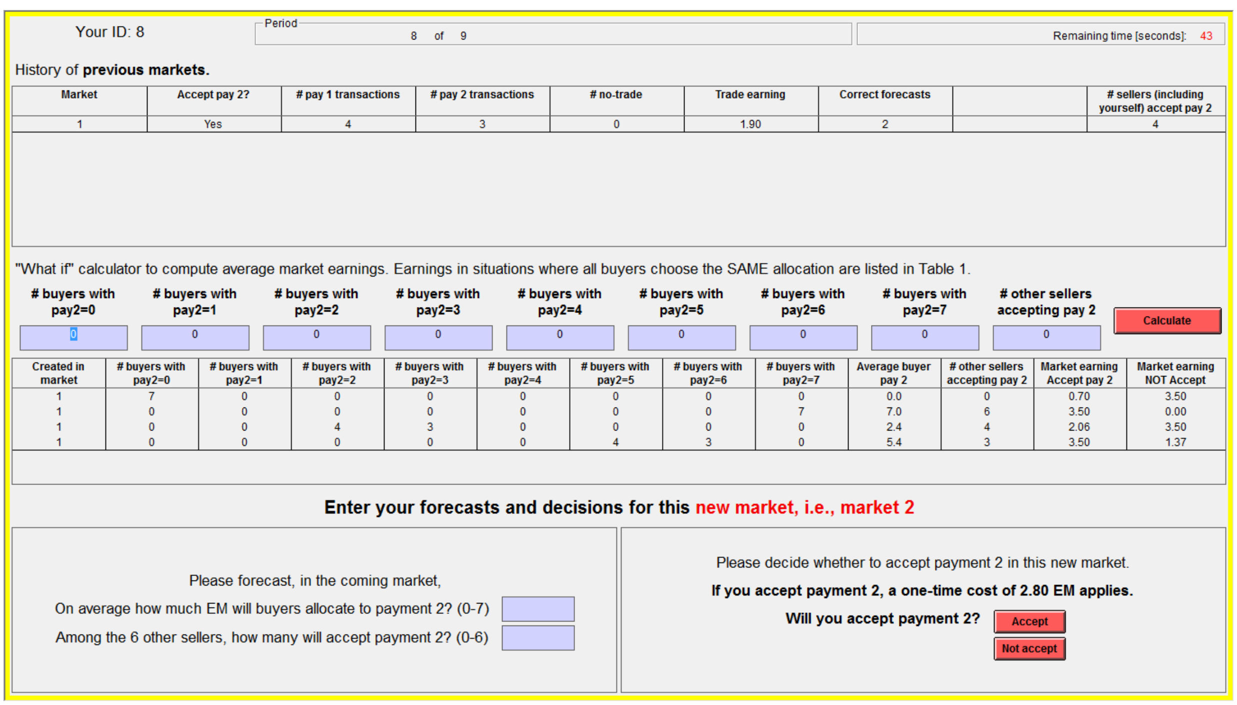

Your ID: 8

Period 8 of 9

Remaining time [seconds]: 43

History of **previous markets.**

Market	Accept pay 2?	# pay 1 transactions	# pay 2 transactions	# no-trade	Trade earning	Correct forecasts		# sellers (including yourself) accept pay 2
1	Yes	4	3	0	1.90	2		4

"What if" calculator to compute average market earnings. Earnings in situations where all buyers choose the SAME allocation are listed in Table 1.

# buyers with pay2=0	# buyers with pay2=1	# buyers with pay2=2	# buyers with pay2=3	# buyers with pay2=4	# buyers with pay2=5	# buyers with pay2=6	# buyers with pay2=7	# other sellers accepting pay 2	
0	0	0	0	0	0	0	0	0	Calculate

Created in market	# buyers with pay2=0	# buyers with pay2=1	# buyers with pay2=2	# buyers with pay2=3	# buyers with pay2=4	# buyers with pay2=5	# buyers with pay2=6	# buyers with pay2=7	Average buyer pay 2	# other sellers accepting pay 2	Market earning Accept pay 2	Market earning NOT Accept
1	7	0	0	0	0	0	0	0	0.0	0	0.70	3.50
1	0	0	0	0	0	0	0	7	7.0	6	3.50	0.00
1	0	0	4	3	0	0	0	0	2.4	4	2.06	3.50
1	0	0	0	0	0	4	3	0	5.4	3	3.50	1.37

Enter your forecasts and decisions for this new market, i.e., market 2

Please forecast, in the coming market,
On average how much EM will buyers allocate to payment 2? (0-7)
Among the 6 other sellers, how many will accept payment 2? (0-6)

Please decide whether to accept payment 2 in this new market.
If you accept payment 2, a one-time cost of 2.80 EM applies.
Will you accept payment 2?

Accept
Not accept

Figure 3: buyer's trading screen

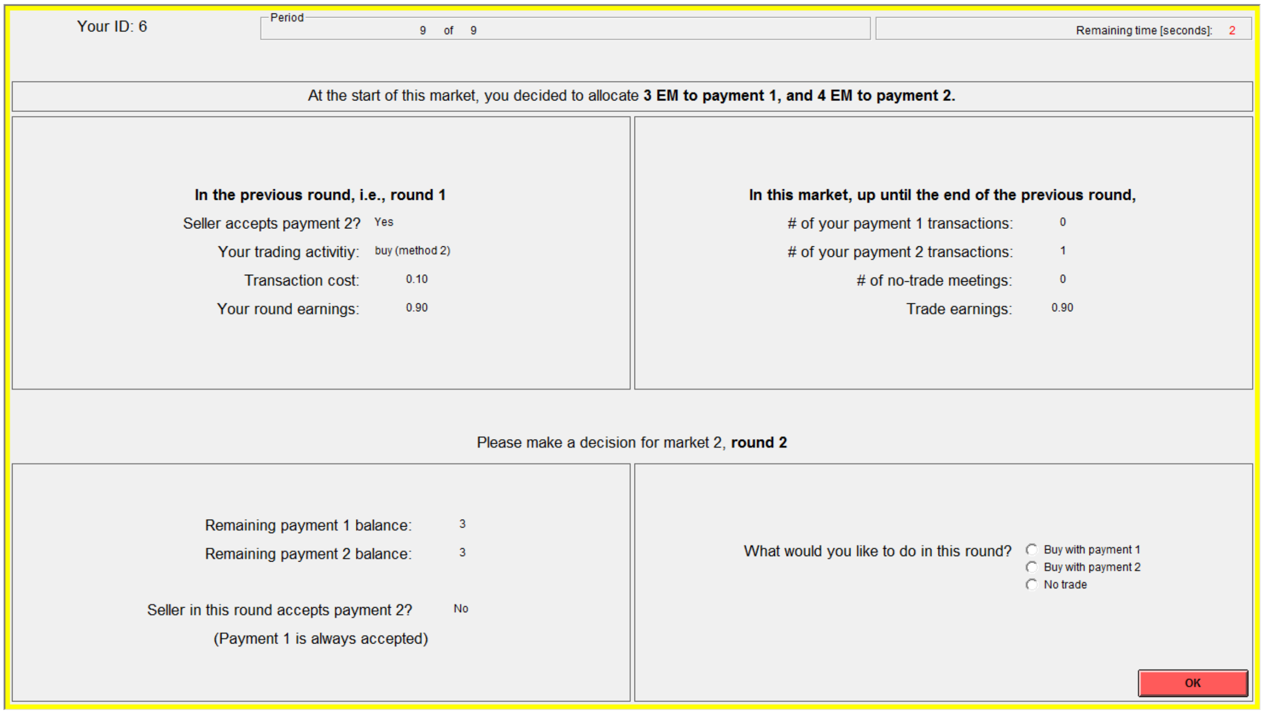

Figure 4: seller's trading screen

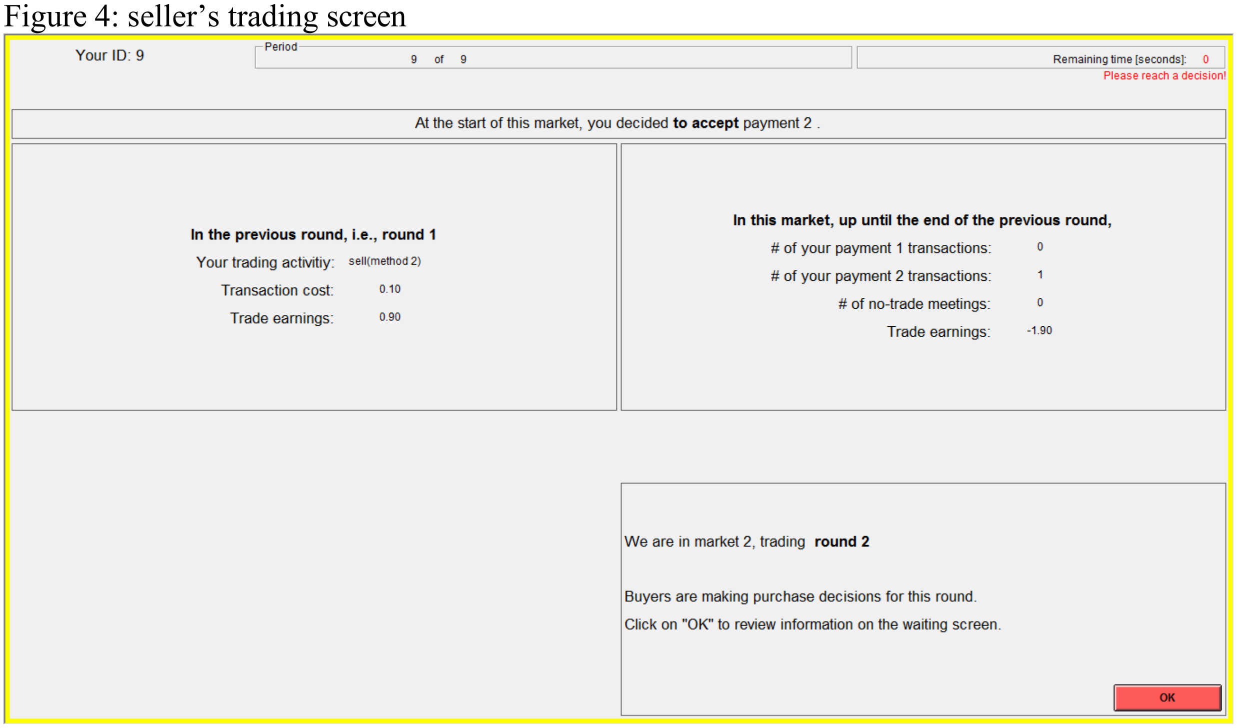

Figure 5: buyer's waiting/information screen

Your ID: 2

Period

8 of 9

History of **previous markets.**

Market	Pay 1 choice	Pay 2 choice	# pay 1 transactions	# pay 2 transactions	# no-trade	Trade earning	Correct forecasts		# sellers accept pay 2
1	5	2	5	2	0	4.30	0		4

Your money allocation	# of sellers accepting payment 2							
payment 2 balance	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
0	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5
1	3	3.9	3.9	3.9	3.9	3.9	3.9	3.9
2	2.5	3.4	4.3	4.3	4.3	4.3	4.3	4.3
3	2	2.9	3.8	4.7	4.7	4.7	4.7	4.7
4	1.5	2.4	3.3	4.2	5.1	5.1	5.1	5.1
5	1	1.9	2.8	3.7	4.6	5.5	5.5	5.5
6	0.5	1.4	2.3	3.2	4.1	5	5.9	5.9
7	0	0.9	1.8	2.7	3.6	4.5	5.4	6.3

<==Your market earning depends on
your money allcoation, and
the number of sellers accepting payment 2

This is the waiting screen ...

We are in **market 2** the payment choice stage ...

You just decided to allocate **2 EM to payment 1, and 5 EM to payment 2.**

Some participants are still making their decisions.
While waiting, you can review the information on the screen.

Figure 6: seller's waiting/information screen

Your ID: 8

Period

8 of 9

History of **previous markets.**

Market	Accept pay 2?	# pay 1 transactions	# pay 2 transactions	# no-trade	Trade earning	Correct forecasts		# sellers (including yourself) accept pay 2
1	Yes	4	3	0	1.90	2		4

The "What if" records you have created.

Created in market	# buyers with pay2=0	# buyers with pay2=1	# buyers with pay2=2	# buyers with pay2=3	# buyers with pay2=4	# buyers with pay2=5	# buyers with pay2=6	# buyers with pay2=7	Average buyer pay 2	# other sellers accepting pay 2	Market earning Accept pay 2	Market earning NOT Accept
1	7	0	0	0	0	0	0	0	0.0	0	0.70	3.50
1	0	0	0	0	0	0	0	7	7.0	6	3.50	0.00
1	0	0	4	3	0	0	0	0	2.4	4	2.06	3.50
1	0	0	0	0	0	4	3	0	5.4	3	3.50	1.37

This is the waiting screen ...

We are in **market 2** the trading stage ...

You just decided **NOT to accept** payment 2 in this market

Some participants are still making their decisions.
While waiting, you can review the information on the screen.

Table 1: Seller's average market earnings

- This table considers the case where ***all buyers choose the same payment allocation;*** use the "what-if" calculator for cases where buyers make different allocations.
- The earnings for ***accept****ing* payment 2 are in the ***upper-left*** corner,
- The earnings for ***not accept****ing* payment 2 are in the ***lower-right*** corner.

All buyer's allocation to **payment 2**	# of **other 6 sellers** accepting payment 2							
	0	1	2	3	4	5	6	
0	0.7 3.5	0.7 3.5	0.7 3.5	0.7 3.5	0.7 3.5	0.7 3.5	0.7 3.5	←if accept pay 2 ←if not accept pay 2
1	3.5 3.0	2.1 3.5	1.6 3.5	1.4 3.5	1.3 3.5	1.2 3.5	1.1 3.5	
2	3.5 2.5	3.5 2.9	2.6 3.5	2.1 3.5	1.8 3.5	1.6 3.5	1.5 3.5	
3	3.5 2.0	3.5 2.3	3.5 2.8	2.8 3.5	2.4 3.5	2.1 3.5	1.9 3.5	
4	3.5 1.5	3.5 1.8	3.5 2.1	3.5 2.6	2.9 3.5	2.6 3.5	2.3 3.5	
5	3.5 1.0	3.5 1.2	3.5 1.4	3.5 1.8	3.5 2.3	3.0 3.5	2.7 3.5	
6	3.5 0.5	3.5 0.6	3.5 0.7	3.5 0.9	3.5 1.2	3.5 1.8	3.1 3.5	
7	3.5 0	3.5 0	3.5 0	3.5 0	3.5 0	3.5 0	3.5 0	

Table 2: Buyer's market earning

Your allocation to **payment 2**	# of sellers accepting payment 2							
	0	1	2	3	4	5	6	7
0	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5
1	3.0	3.9	3.9	3.9	3.9	3.9	3.9	3.9
2	2.5	3.4	4.3	4.3	4.3	4.3	4.3	4.3
3	2.0	2.9	3.8	4.7	4.7	4.7	4.7	4.7
4	1.5	2.4	3.3	4.2	5.1	5.1	5.1	5.1
5	1.0	1.9	2.8	3.7	4.6	5.5	5.5	5.5
6	0.5	1.4	2.3	3.2	4.1	5.0	5.9	5.9
7	0.0	0.9	1.8	2.7	3.6	4.5	5.4	6.3

C Additional Analysis, Figures and Tables

In this Appendix, we provide additional analysis, figures and tables. Figures C.1 to C.3 plot the payoffs against market for each of the three treatments. Table C.1 shows the regression analysis with individual buyer data separately for each of the three treatments.

We also formally test whether a session converges to either of two symmetric strategy equilibria by estimating the process followed by three variables, the percentage of the buyers' endowment allocated toward payment 2 averaged across the seven buyers (bPay2%), the percentage of sellers accepting payment 2 (sAccept%), and the percentage of meetings that resulted in trade using payment 2 (Pay2Meetings%), over time. In particular, we run the following regression for each session and for each of these three variables:

$$y_{j,s} = \lambda_j y_{j,s-1} + \mu_j + \epsilon_{j,s}, \tag{C.1}$$

where $y_{j,s}$ is the value of the variable being tested in market s for session j. From (C.1), we say that the variable converges to its payment-1-only equilibrium value if the estimate of the long-run expected value for y_j, $\frac{\mu_j}{1-\lambda_j}$, is not significantly different from 0. Similarly, we say that the variable converges to its payment-2-only equilibrium value if $\frac{\mu_j}{1-\lambda_j}$ is not significantly different from 100. Table C.2 reports the estimates and standard errors for $1-\lambda_j$, $100-\frac{\mu_j}{1-\lambda_j}$ and $\frac{\mu_j}{1-\lambda_j}$; the p-values indicate whether the estimated variable is significantly different from 0. Thus, if $100-\frac{\mu_j}{1-\lambda_j}$ (alternatively $\frac{\mu_j}{1-\lambda_j}$) is significantly different from zero, then we can reject the hypothesis of convergence of that variable to the all-payment-2 (all-payment-1) equilibrium.

Table C.1: Buyer Payment 2 Choice (%) with Random Effects by Treatment

		T=1.6	T=2.8	T=3.5
	MktAcceptL(%)	**0.824*****	**0.352** ***	**0.392** ***
		(0.109)	(0.049)	(0.047)
Stage 1: bBelief(%)	market	**0.779** ***	**0.312** **	**-0.983** ***
		(0.203)	(0.153)	(0.187)
	location	1.401	**7.576** ***	-5.877
	(SFU=1;UCI=0)	(2.339)	(2.824)	(2.778)
	bBelief(%)	**0.828** ***	**0.916** ***	**1.013** ***
		(0.146)	(0.110)	(0.084)
Stage 2: bPay2(%)	market	**0.354** **	**0.104**	**0.037**
		(0.155)	(0.132)	(0.170)
	location	**3.944** *	**3.315**	**-1.091**
	(SFU=1;UCI=0)	(2.333)	(2.695)	(1.877)

Notes. (1) *p-value≤0.1; **p-value≤0.05; *** p-value≤0.01. (2) Each regression has 532=4x7x19 observations. There are 4 sessions, each with 7 buyers and 20 markets. For each individual, we have 19 observations as the first-stage regression uses a lagged variable as an independent variable. (3) The regressions have clustered errors at the individual subject level.

Table C.2: Test of Convergence to Symmetric Equilibrium

Treatment	Session		bPay2%			sAccept%			Pay2Meetings%		
			$1-\lambda$	$100-\frac{\mu}{1-\lambda}$	$\frac{\mu}{1-\lambda}$	$1-\lambda$	$100-\frac{\mu}{1-\lambda}$	$\frac{\mu}{1-\lambda}$	$1-\lambda$	$100-\frac{\mu}{1-\lambda}$	$\frac{\mu}{1-\lambda}$
T=1.6	1	Coef.	0.255***	6.404	93.596***	0.792***	2.256	97.744***	0.213**	5.464	94.536***
		Std.Err.	0.092	6.152	6.152	0.243	1.596	1.596	0.088	7.698	7.698
	2	Coef.	0.346***	0.167	99.833***	1.059***	0.794	99.206***	0.338***	0.749	99.252***
		Std.Err.	0.094	4.268	4.268	0.288	1.209	1.209	0.094	4.401	4.401
	3	Coef.	0.180*	2.782	97.218***	0.566***	4.863**	95.137***	0.182*	4.114	95.886***
		Std.Err.	0.107	9.271	9.271	0.130	2.314	2.314	0.098	9.198	9.198
	4	Coef.	0.504***	2.394	97.606***	1.118***	1.504	98.496***	0.336**	1.658	98.342***
		Std.Err.	0.143	2.828	2.828	0.288	1.131	1.131	0.136	4.561	4.561
T=2.8	1	Coef.	0.793***	56.802***	43.198***	1.118***	52.632***	47.368***	0.949***	61.648***	38.352***
		Std.Err.	0.259	1.638	1.638	0.266	3.127	3.127	0.260	1.787	1.787
	2	Coef.	0.085	12.881	87.119***	0.718***	19.711***	80.289***	0.224**	26.806***	73.194***
		Std.Err.	0.128	25.604	25.604	0.181	4.894	4.894	0.107	8.209	8.209
	3	Coef.	1.266***	32.104***	67.896***	1.244***	26.316***	73.684***	0.611*	37.652***	62.348***
		Std.Err.	0.255	1.020	1.020	0.207	2.809	2.809	0.332	2.856	2.856
	4	Coef.	0.586***	31.034***	68.966***	1.178***	29.096***	70.904***	0.172	29.749***	70.251***
		Std.Err.	0.174	2.246	2.246	0.322	2.987	2.987	0.180	1.799	11.799
T=3.5	1	Coef.	0.051	121.447	-21.447	0.310*	79.948***	20.052	0.111	98.947***	1.053
		Std.Err.	0.097	123.719	123.719	0.164	14.086	14.086	0.122	38.359	38.359
	2	Coef.	0.301**	72.188***	27.812***	0.918***	72.516***	27.484***	0.762***	78.476***	21.524***
		Std.Err.	0.142	6.520	6.520	0.179	4.073	4.073	0.173	2.983	2.983
	3	Coef.	0.891***	46.161***	53.839***	1.030***	46.573***	53.427***	1.097***	55.102***	44.898***
		Std.Err.	0.268	2.010	2.010	0.217	3.574	3.574	0.248	2.034	2.034
	4	Coef.	0.271	68.059***	31.941***	1.012***	68.403***	31.597***	0.920***	72.909***	27.091***
		Std.Err.	0.189	7.278	7.278	0.238	3.640	3.640	0.245	2.441	2.441

Notes. (1) *p-value≤0.1; p-value≤0.05; p-value≤0.01. (2) Number of observations: 19. (3) The estimate of the long-run equilibrium value of bPay2% is negative from the unconstrained estimation, but it is not significantly different from zero. A constrained estimation that restricts $\mu \geq 0$ will give $\hat{\mu} = 0$.

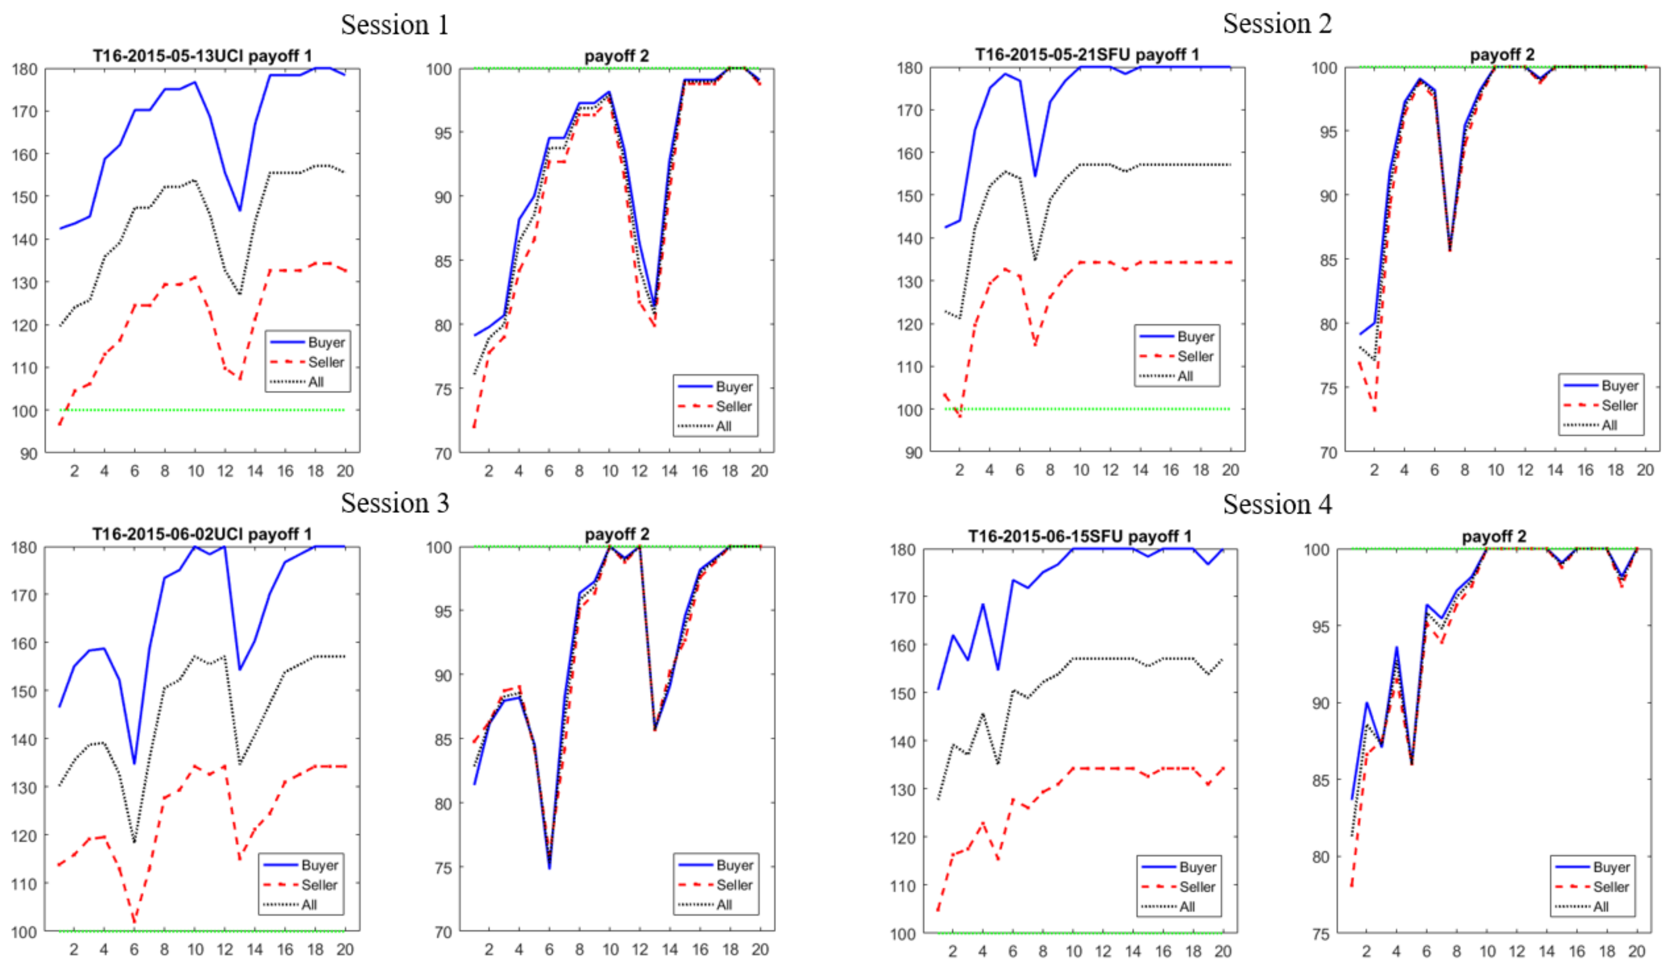

Figure C.1: Payoff T=1.6

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first figure shows subjects' payoffs normalized by the payoffs in the payment-1 equilibrium as the benchmark (100%). The blue solid line represents sellers, the red dashed line represents buyers and the black line marked with circles represents buyers and sellers together. The second figure shows payoffs normalized by payoffs in the payment-2 equilibrium as the benchmark.

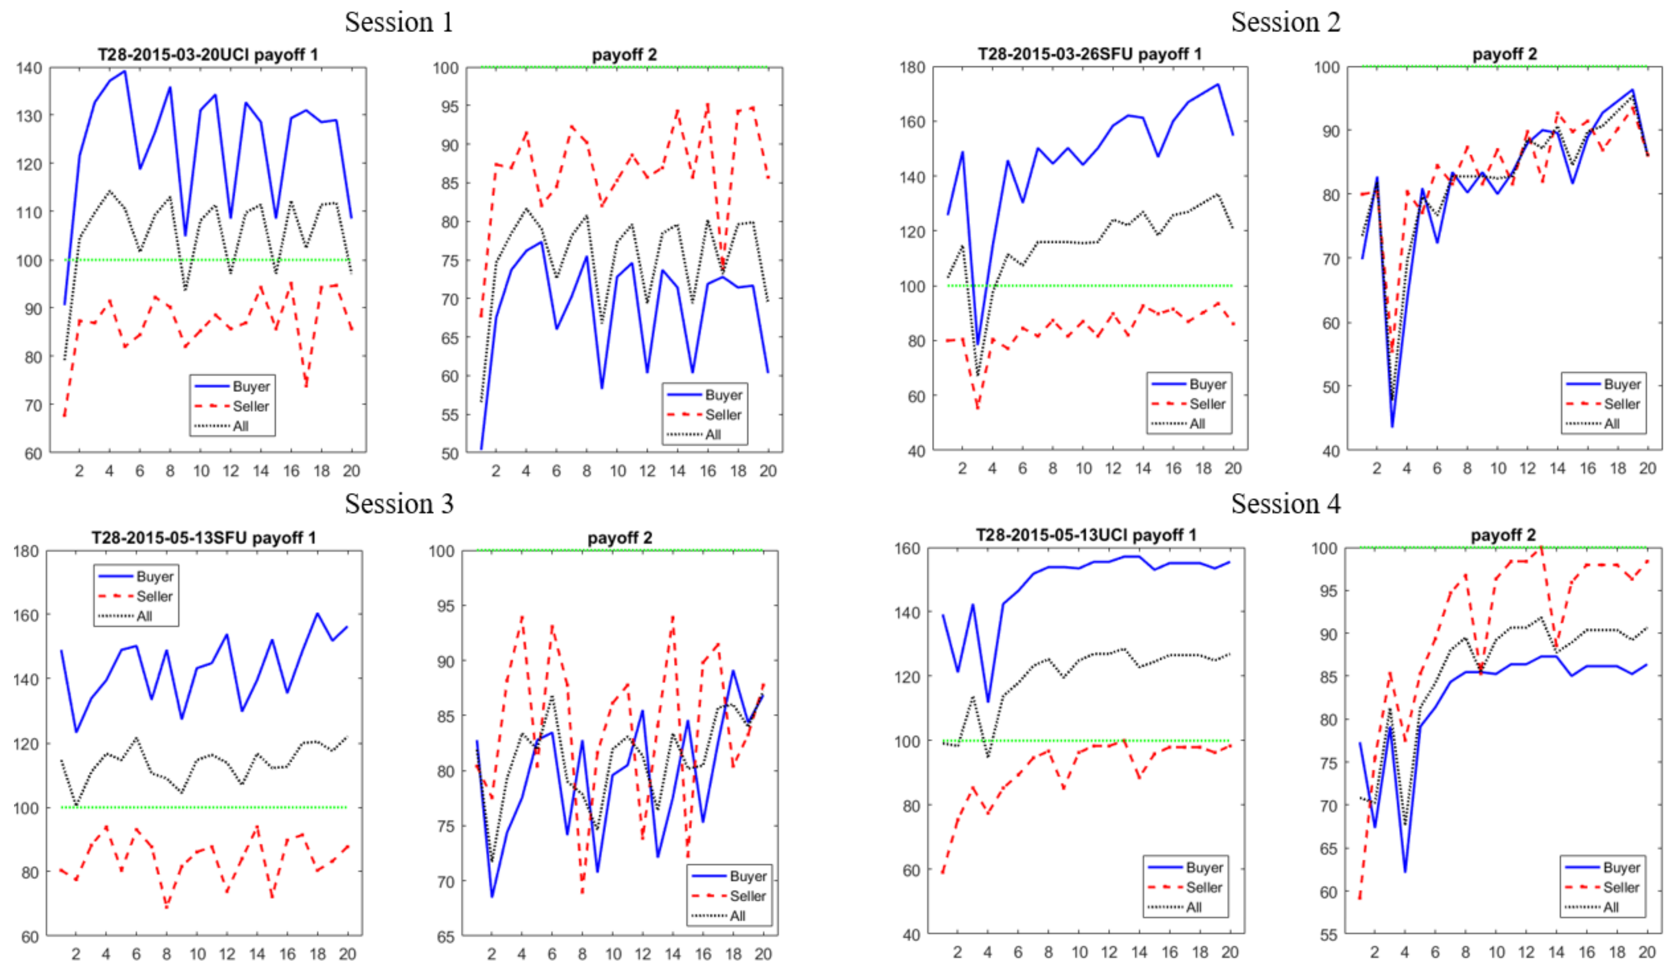

Figure C.2: Payoff T=2.8

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first figure shows subjects' payoffs normalized by the payoffs in the payment-1 equilibrium as the benchmark (100%). The blue solid line represents sellers, the red dashed line represents buyers and the black line marked with circles represents buyers and sellers together. The second figure shows payoffs normalized by payoffs in the payment-2 equilibrium as the benchmark.

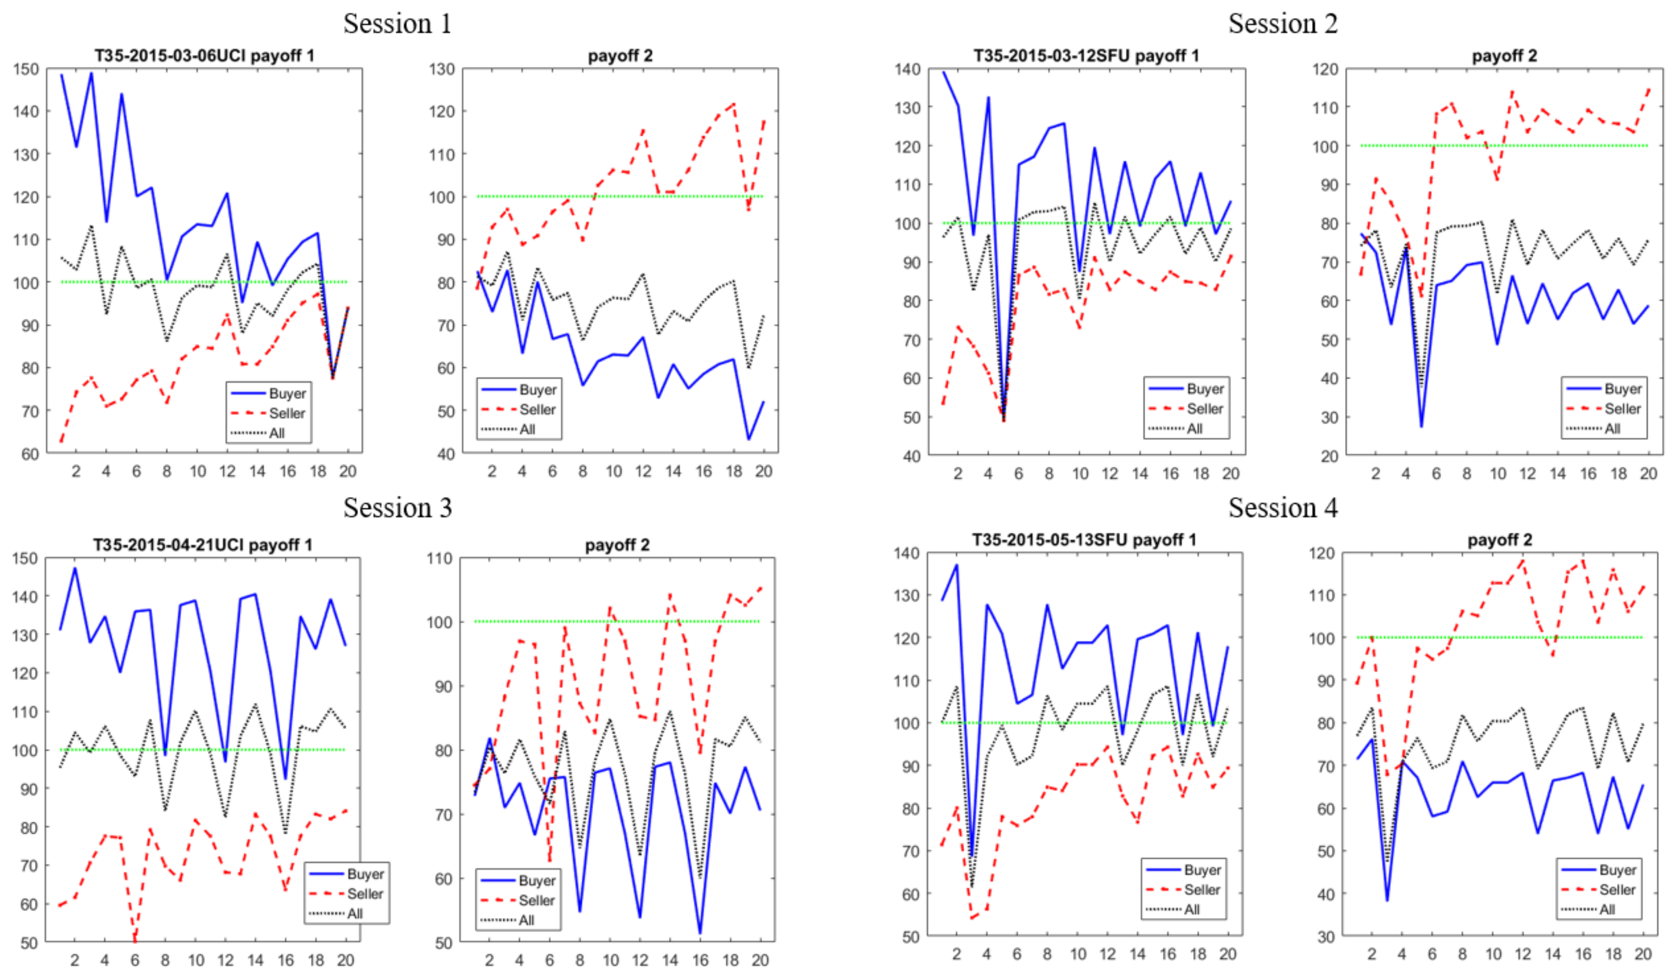

Figure C.3: Payofff T=3.5

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first figure shows subjects' payoffs normalized by the payoffs in the payment-1 equilibrium as the benchmark (100%). The blue solid line represents sellers, the red dashed line represents buyers and the black line marked with circles represents buyers and sellers together. The second figure shows payoffs normalized by payoffs in the payment-2 equilibrium as the benchmark.

D Individual Evolutionary Learning Model

In this Appendix, we provide more detailed information about the IEL and we report on additional simulation results using that model.

Regarding IEL, we first discuss how to calculate the foregone payoff. Second, we describe how sellers update their expectations about the average amount that buyers allocated to payment 2, which they use to calculate the foregone payoff. Third, we describe the process of experimentation. Third, we describe our IEL convergence criteria, and present the results of average times to convergence and stability of equilibrium outcomes for the four values of fixed cost that were used in our experimental economies. Finally, we use the IEL to explore the "tipping point" for T regarding adoption of the new payment method, and the effects of reducing the seller's payment costs which increases their gains from trade.

Forgone Payoff Calculation

First we describe how to calculate the foregone payoff calculation depends on whether an agent is a buyer or a seller.

At the end of each period, buyers know the number of sellers who actually accepted payment 2 in that period, $s_a(t) \in \{0, 1, ..., 7\}$. For buyers, the foregone payoff of each rule $m^i_{b,j}(t)$ in buyer i's set at the end of market t is computed in the following way.

$$\Pi^i_{b,j}(t) = (7 - m^i_{b,j}(t))(u - \tau_1) + \min\left[s_a(t), m^i_{b,j}(t)\right](u - \tau_2).$$

Note that in our simulation, we assume that the buyer adopts the following (payoff-maximizing) strategy (given her initial payment allocation): if the buyer meets a seller who accepts payment 2, then the buyer uses payment 2 if he still has payment 2 left and she uses payment 1 otherwise; if the seller does not accept payment 2, then the buyer uses payment 1 if he still has some payment 1 left and she does not trade otherwise. Note that if $m^i_{b,j}(t)$ is larger than $s_a(t)$, then $s_a(t)$ is used in the calculation and there is missing trade in some rounds, where the buyer has only payment 2 left and the seller does not accept it.

The computation of sellers' foregone payoffs is more complex. For each seller, we use two variables to compute the foregone payoffs of all the rules in a seller's set. The first variable is $s_a^{-i}(t) \in \{0, 1, ..., 6\}$, the number of sellers excluding seller i that accepted payment 2. Recall that this information was provided to sellers at the end of each market of the experiment. The second variable is $s^f\bar{m}^i_b(t) \in \{0, 1, ..., 7\}$, the forecast of seller i of the average allocation to payment 2 by all 7 buyers. Note that in our experiment buyers' allocations to payment 2 are not public knowledge, but we did elicit sellers' forecasts of buyers' average allocation to payment 2. Thus, our artificial agents, like the human subjects, must form an expectation of this value. The updating of this expectation is seller-specific and depends on each seller's experience from the previous period; the details of this updating are given in the next section "Seller Belief Formation." After we have $s_a^{-i}(t)$ and $s^f\bar{m}^i_b(t)$, we use them to evaluate foregone payoffs for all of the rules in a seller's rule set in three steps.

First, we calculate the expected number of transactions that would have been completed using payment 2 provided that the seller had accepted payment 2:

$$n^{i,a}(t) = 7 * \min\left\{s^f\bar{m}^i_b(t)/(s_a^{-i}(t) + 1), 1\right\}, \tag{D.1}$$

and the expected profit from doing so (note that $7 - n^{i,a}(t)$ transactions use payment 1):

$$\pi_s^{i,a} = n^{i,a}(t)(u - \tau_2) + (7 - n^{i,a}(t))(u - \tau_1) - T. \tag{D.2}$$

Second, we calculate the expected number of transactions involving payment 1 that would have taken place if the seller did not accept payment 2:

$$n^{i,n}(t) = 7 * \min\left\{(7 - s^f \bar{m}_b^i(t))/(7 - s_a^{-i}(t), 1\right\}, \tag{D.3}$$

and the expected payoff from doing so (note that transactions can only be carried out with payment 1):

$$\pi_s^{i,n}(t) = n^{i,n}(t)(1 - \tau_1). \tag{D.4}$$

Finally, for each rule j that is in seller i's rule set, we calculate the expected foregone payoff as the weighted average of $\pi_s^{i,a}(t)$ and $\pi_s^{i,n}(t)$:

$$\Pi_{s,j}^i(t) = m_{s,j}^i(t)\pi_s^{i,a}(t) + (1 - m_{s,j}^i(t))\pi_s^{i,n}(t). \tag{D.5}$$

Seller Belief Formation

Below we describe the process by which sellers update their expectations about the average amount that buyers allocated to payment 2, $s^f \bar{m}_b^i(t)$, which is used to calculate the foregone payoff for sellers. This is carried out in four steps.

Step 1. Infer the boundaries on the initial payment 2 allocation of each trading partner (one for each of the seven rounds of transaction) in the past market.

Note that in the experiment, a seller does not know a buyer's initial payment allocation and must infer it with a limited information set, which includes (i) whether the seller herself accepted payment 2 or not, (ii) how many of the other six sellers chose to accept payment 2, (iii) in which round (out of seven) she meets the buyer, and (iv) whether the transaction uses payment 1, payment 2 or fails to take place. In many situations, the seller will not be able to exactly pinpoint the buyer's initial choice, but she can always infer either the lower bound (call it a *Min inference*) or the upper bound (call it a *Max inference*) of the buyer's payment 2 allocation.

Let r refer to the current round, M to the number of sellers ($M = 7$), and $s_a^{-i}(t)$ the number of other sellers who accepted payment 2 in the past market. Table D.1 summarizes what the seller can infer about her round r trading partner's initial payment 2 allocation (assuming that the buyer behaves optimally). There are four cases to consider, with the first two applying to a seller who did not accept payment 2, and the other two cases applying to a seller who accepted payment 2.

In case A, the seller chose to not accept payment 2 and no transaction occurred for the round in question. The seller knows that the buyer in that round had no payment 1 available – this implies a maximum that the buyer allocated to payment 1 at the beginning of the market, or a minimum allocated to payment 2. The value of this minimum is given by $8 - r + (r - 1) * s_a^{-i}(t)/(M - 1)$. We illustrate this case with two examples, one assuming that all other sellers accepted payment 2, or $s_a^{-i}(t) = 6$, and the other assuming that none of the other sellers accepted payment 2, or $s_a^{-i}(t) = 0$ (and we do the same with the other three cases).

Table D.1: Seller's guess about buyer's initial payment 2 allocation

Case	Transaction	P2 Accepted?	Buyer initial P2 allocation
A	None	No	Min: $m^b \geq 8 - r + (r-1) * s_a^{-i}(t)/(M-1)$
B	Use P1	No	Max: $m^b \leq 7 - r + (r-1) * s_a^{-i}(t)/(M-1)$
C	Use P1	Yes	Max: $m^b \leq (r-1) * s_a^{-i}(t)/(M-1)$
D	Use P2	Yes	Min: $m^b \geq 1 + (r-1) * s_a^{-i}(t)/(M-1)$

Notes: P1 (P2) is short for payment 1 (payment 2).

Example A1. Suppose $s_a^{-i}(t) = 6$. The seller knows that no matter which round it is, the buyer had not needed to spend payment 1 until the buyer met the seller herself. If the buyer did not have payment 1, it is because he chose to allocate all his 7 units to payment 2. In this example, the seller can infer exactly that the buyer chose $m^b = 7$.

Example A2. Suppose $s_a^{-i}(t) = 0$. The seller's guess depends on the round of the transaction. If $r = 1$, then the seller knows exactly that the buyer's payment 2 allocation was $8 - 1 = 7$ (and the initial balance of payment 1 is 0). If $r = 2$, then the buyer's initial payment 1 balance could be either 0 (in which case he did not trade in the first round), or 1 (in which case he paid with payment 1 in the first round); equivalently, his initial payment 2 allocation was either 6 or 7 (so the minimum is $8 - 2 = 6$). For each round, going further in time, the seller has less and less precise information about the buyer's initial choice because she does not know how many times the buyer has used payment 1 in the previous rounds.

In case B, the seller did not accept P2 and a transaction took place with P1. Since the buyer had P1 to use, there was a minimum that this buyer allocated to P1, or a maximum allocated to P2. The value of this maximum is given by $7 - r + (r-1) * s_a^{-i}(t)/(M-1)$.

Example B1. Suppose $s_a^{-i}(t) = 6$. The seller knows that the buyer initially allocated at least 1 unit of payment 1, but this is all she can infer: the buyer could have more payment 1 in hand to spend in later rounds and/or might have spent some payment 1 in previous rounds. Equivalently, the seller can only infer that the buyer had initially at most 6 units in payment 2.

Example B2. Suppose $s_a^{-i}(t) = 0$. The seller can infer that the buyer initially allocated at least r units of payment 1 if he still has payment 1 to spend in round r. Equivalently, the buyer allocated at most $7 - r$ units in payment 2 initially. As the round number increases, the seller acquires more precise information about the buyer's initial choice. If the buyer still had payment 1 in round 7, then the seller knows exactly that the buyer chose to allocate all his money to payment 1 and nothing to payment 2.

In case C, the seller accepted payment 2 but the buyer used payment 1. This implies that the buyer had no payment 2 left. This sets a maximum on how much the buyer allocated to payment 2. The value of the maximum is given by: $(r-1) * s_a^{-i}(t)/(M-1)$.

Example C1. Suppose $s_a^{-i}(t) = 6$. The seller knows that the buyer ran out of payment 2 in previous rounds, but she is not sure how many times the buyer had used payment 2 in the previous rounds: the number can vary from 0 to $r - 1$. As the round goes on, the seller's inference becomes less accurate.

Example C2. Suppose $s_a^{-i}(t) = 0$. The seller can infer exactly that the buyer allocated nothing to payment 2 irrespective of which round it is because the buyer could not spend payment 2 in previous rounds.

In case D, the seller accepted payment 2, and the buyer paid using payment 2. This implies that there is a minimum amount that the buyer allocated to payment 2. The value of the limit is given by

$1 + (r-1) * s_a^{-i}(t)/(M-1)$.

Example D1. Suppose $s_a^{-i}(t) = 6$. Since the buyer still had payment 2 in round r, the seller can infer that the buyer had at least r units of payment 2 initially. As the round number increases, the seller has more accurate information about the buyer's portfolio choice. In round 7, the seller knows that the buyer chose to allocate all his money to payment 2.

Example D2. Suppose $s_a^{-i}(t) = 0$. The seller's inference about the buyer's initial portfolio choice is very imprecise: she only knows the buyer has at least 1 unit of payment 2 to spend in this round, but has no idea whether the buyer had used payment 2 in previous rounds and still had more payment 2 to spend in later rounds.

Step 2. Evaluate the accuracy of the seven inferences in step 1. A Min (Max) inference with a larger (smaller) bound implies a smaller set of possible values for the buyer's P2 (and P1) allocation and is more accurate. We use the *certainty value* (CE) to quantify the accuracy of these inferences. The CE is calculated as 8 minus the number of elements in the set of possible values of the buyer's payment 2 allocation; as a result, a smaller set is awarded a higher CE. For example, an inference with $m^b \geq 7$ implies a set with a single element $\{7\}$, so its CE is $8 - 1 = 7$. An inference with $m^b \leq 4$ implies the set $\{0, 1, 2, 3, 4\}$ with 5 elements, so its CE is $8 - 5 = 3$. Equivalently, the following formula can be used to calculate the CE:

$$CE = \begin{cases} x \text{ for inference } m^b \geq x, \\ 7 - x \text{ for inference } m^b \leq x. \end{cases}$$

Step 3. Use the lower and upper bounds for the seven inferences to estimate the lower and upper bounds for the "average" buyer. The "average" lower (upper) bound is calculated as the sum of the lower (upper) bounds of all Min (Max) inferences, weighted by their CEs calculated in step 2.

Step 4. The final step is to use the "average" lower and upper bounds to calculate the expectation about the average buyer's P2 allocation, $s^f \bar{m}_b^i(t)$, as the weighted sum of the two "average" bounds, with the weight of the lower (upper) average bound being given by the number of Min (Max) inferences.

Experimentation

In this section, we describe the parameterization of the experimentation rate and the standard deviation of experimentation. The rate of experimentation μ_t is

$$\mu_t = \frac{0.35}{1 + (t-1)/5000},$$

where t is the current period at which experimentation occurs. The initial rate of experimentation, μ_1, is set to 0.35. The rate of experimentation slowly decreases as t increases. The standard deviation of experimentation σ_t is

$$\sigma_t = \frac{3.5}{1 + (t-1)/5000}$$

The initial value of 3.5 is the midpoint of the buyer's strategy choice set of allocating 0 to 7 EM to payment 2. The standard deviation of experimentation also decays slowly over time.

Convergence Criteria

We define the following convergence criteria.

If, in a given period t, 85% of buyers' aggregate units are placed in payment 1, and 85% of sellers do not accept payment 2, then we classify that outcome as the payment 1 equilibrium outcome.

Similarly, if, in a given period t, 85% of buyers' aggregate units are placed in payment 2, and 85% of sellers accept payment 2, then the outcome in that period is classified as the payment 2 equilibrium outcome.

We declare the *time to convergence* as the first period in which either of the above criteria is satisfied in a simulation that criterion continues to be satisfied at least twice more in the following 10 periods.

The *mean time to convergence* reported on in Table D.2 measures the average time to convergence over 50 simulations for different values of T. We also report in that table the standard deviation, median, minimum and maximum of time to convergence from our simulation exercises.

After the time to convergence is recorded, we continue running the simulation for at least another 100 periods to examine the stability of the convergence result. Our *stability index*, reported in the final column of Table D.2 gives the percentage of periods when a convergence criterion is met out of the 100 periods following the recorded convergence period.

The first four rows of Table D.2 reports on convergence times and stability for the four values of T considered in our experiment. We see that convergence to the all payment 2 equilibrium obtains for $T = 1.6$ and 2.8 while the all payment 1 equilibrium is the convergent outcome under higher values, $T = 3.5$ and 4.5.

Additional Results

We use the IEL model to conduct two sets of additional analysis. First, we explore where the "tipping point" for T lies regarding adoption of the new payment method. As our simulations suggest that the new payment method is adopted by the end when $T = 2.8$, while it is discarded by the end when $T = 3.5$, the tipping point value of T must lie between these two values. Therefore, we ran simulations of the IEL model for values of T ranging from 2.9 to 3.4, holding all else constant. The results are shown in the middle section of Table D.2. We observe that in this region for the fixed adoption cost T, the level of stability is low as revealed by our stability index. However, for $T = 3.2$ a majority (66%) of simulations converge to the all payment 2 equilibrium while for $T = 3.3$ a majority (42%) of simulations converge to the all payment 1 equilibrium, so we may regard $T = 3.2$ as approximating the tipping point.

Second, we examine the effects of reducing the seller's payment costs which increases their gains from trade. Intuitively, this change should promote further acceptance of payment 2 as sellers try to secure more trade. The simulation results support this intuition. Section 3 of Table D.2 reports the convergence results as we reduce the seller's payment cost terms, τ_1, τ_2, and T, all by 50% (doing so keeps $T/(\tau_1 - \tau_2)$ constant relative to the reference treatment in our experiment). Here we also consider the additional case of $T = 5.5/2 = 2.75$.

As these simulation results reveal, the reduction in sellers' costs increases the region for which the all payment 2 equilibrium is achieved. For example, when $T = 3.5/2 = 1.75, \tau_1 = 0.5/2, \tau_2 = 0.1/2$, all 50 simulated economies converge to the equilibrium where all players use the new payment method in the end, while they all converged to the equilibrium with use of the old payment method by the end when the cost terms are taken from the corresponding treatment of our experiment. When $T = 4.5/2, \tau_1 = 0.5/2, \tau_2 = 0.1/2$, the lower seller costs result in 68% of runs converging to use of the new payment method, while they all converged to the equilibrium with the old payment method when $T = 4.5, \tau_1 = 0.5, \tau_2 = 0.1$ as in the corresponding treatment in our experiment. For large enough values for T, namely $T = 5.5/2, \tau_1 = 0.5/2, \tau_2 = 0.1/2$, we still observe convergence to the all payment 1 equilibrium.

Table D.2: Convergence Times and Stability Analysis

Seller cost terms			Equilibrium	Time to Convergence					Stability
T	τ_1^s	τ_2^s	Outcome	Mean	Median	Min	Max	Std	Index
Section 1: parameters as in the experiment									
1.6	0.5	0.1	payment 2	8.5	8	5	13	1.9	98%
2.8	0.5	0.1	payment 2	76.9	69	12	233	47.8	88%
3.5	0.5	0.1	payment 1	49.2	40	15	119	28.3	90%
4.5	0.5	0.1	payment 1	7.4	7.5	3	11	1.6	98%
Section 2: T values between 2.8 and 3.5									
2.9	0.5	0.1	payment 2	122.1	108	23	344	83.7	88%
3.0	0.5	0.1	payment 2 96%	140.4	139.5	15	375	89.0	79%
3.1	0.5	0.1	payment 2 86%	162.7	112.5	17	532	136.2	67%
3.2	0.5	0.1	payment 2 66%	191.4	127.5	13	755	174.7	55%
3.3	0.5	0.1	payment 2 42%	200.9	180	13	536	133.0	39%
3.4	0.5	0.1	payment 2 6%	121.5	107	14	280	69.6	54%
Section 3: Seller cost reduced by 50% relative to the experiment									
0.8	0.25	0.05	payment 2	9	8.5	5	17	2.4	99%
1.4	0.25	0.05	payment 2	34	33.5	10	85	17.3	90%
1.75	0.25	0.05	payment 2	57.3	54	9	151	30.6	87%
2.25	0.25	0.05	payment 2 68%	146.4	103	18	573	123.3	63%
2.75	0.25	0.05	payment 1	18.8	17	8	36	6.7	93%

Notes. (1) For each set of parameter values we run a set of 50 simulations. The number of markets is 400, except for $T = 2.9$ to 3.0, where it is 500, and for $T = 3.1, 3.2, 3.3, 2.25$, where it is 1000. The large number of markets guarantees that *all* 50 simulated economies meet the convergence criterion at least 100 markets before the end of the simulation so that we can calculate the stability index for each simulated economy. (2) Unlike in the simulations with 20 markets as shown in Figure 5, for these simulations, the initial rules are drawn from the uniform distribution from 0 to 7 for buyers, and the uniform distribution from 0 to 1 for sellers. We do not scale the initial rules to match the average starting values of bPay2$\%$ and sAccept$\%$ in Tables 2 and E.1, as the simulation results are very robust to the starting values. In addition, these simulations involve parameter values not used in our experiment and therefore we do not have empirical counterparts of average bPay2$\%$ and sAccept$\%$.

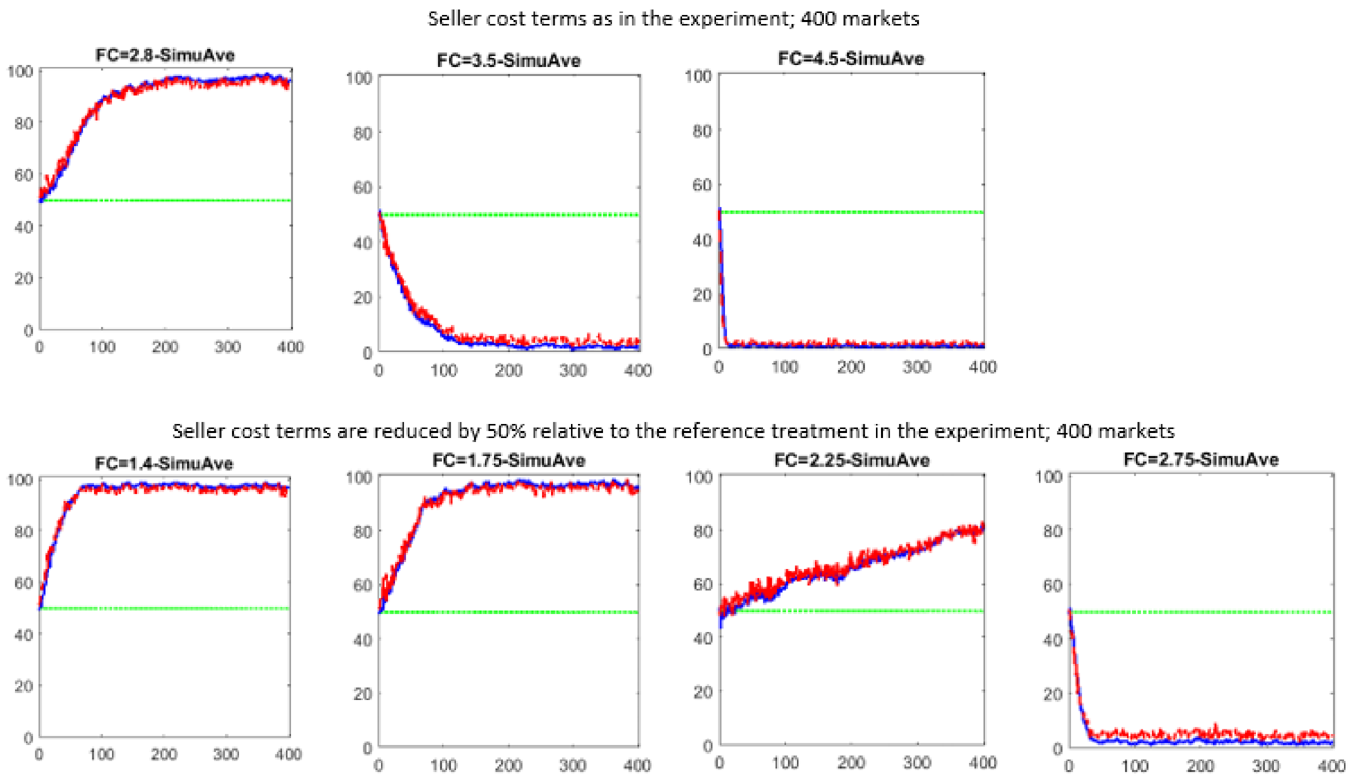

Figure D.1: Effects of Reducing Seller Payment Costs

Notes. The figure shows how the simulated IEL economies change in response to a decrease in seller payment costs. In each figure, the horizontal axis indicates the market number 1-400, the red line is the percentage of sellers accepting payment 2 and the blue line is the percentage of buyers' endowment allocated to payment 2. As a reference, the first row shows the average across 50 simulated economies for 400 markets when the seller's cost terms are exactly as in the experiment. The second row shows the simulation results where all the sellers' cost terms are reduced by 50% relative to the reference treatment in the experiment. We skip treatment with $T = 1.6$ as the figures before and after the reduction in cost terms are almost the same. For $T/(\tau_1 - \tau_2) = 5.5/0.4$, we only show the case with reduced cost terms as we do not have a corresponding experimental treatment.

Figure D.1 shows the simulated path of buyers' allocation to payment 2 and sellers' adoption of payment 2. The simulations shown in this figure are means from 50 runs over 400 markets (the long time horizon allows us to observe the convergence pattern). The top panel of the figure shows simulation results for the parameterizations used in our experiment, while the bottom panel shows the corresponding long-run outcomes from cutting the seller's costs in half. Consistent with the results reported in Section 3 of Table D.2, we find that when we cut all of the sellers' costs in half, the "tipping point" for which the long-run outcome sticks to the old (instead of switching to the new payment method) now lies at a higher value for $T/(\tau_1 - \tau_2)$, somewhere between $T = 4.5/0.4$ and $T = 5.5/0.4$.

E Graphs and Tables for Treatment T=4.5

Table E.1: Payment Choice and Usage $T = 4.5$

	Session	1	2	3	4	**all**
bPay2%	Session mean	21	22	29	34	**26**
	first market	47	47	67	47	**52**
	last market	8	4	8	8	**7**
sAccept%	Session mean	16	15	18	24	**18**
	first market	57	57	71	71	**64**
	last market	29	0	0	29	**14**
pay2Meetings%	Session mean	12	12	14	20	**14**
	first market	41	41	61	45	**47**
	last market	8	0	0	8	**4**
pay1Meetings%	Session mean	79	78	71	66	**74**
	first market	53	53	33	53	**48**
	last market	92	96	92	92	**93**
noTradeMeetings	Session mean	9	10	15	15	**12**
	first market	6	6	6	2	**5**
	last market	0	4	8	0	**3**

Table E.2: Payoff $T = 4.5$

Part 1: payment-2 equilibrium as benchmark						
	Session	1	2	3	4	**all**
buyer	session mean	55	55	53	56	**55**
	first market	70	70	79	74	**74**
	last market	59	53	51	59	**56**
seller	session mean	153	155	142	137	**147**
	first market	103	103	99	82	**97**
	last market	136	187	179	136	**159**
Part 2: payment-1 equilibrium as benchmark						
buyer	session mean	100	99	96	101	**99**
	first market	127	127	143	134	**132**
	last market	107	96	92	107	**100**
seller	session mean	79	80	73	71	**76**
	first market	53	53	51	42	**50**
	last market	70	96	92	70	**82**

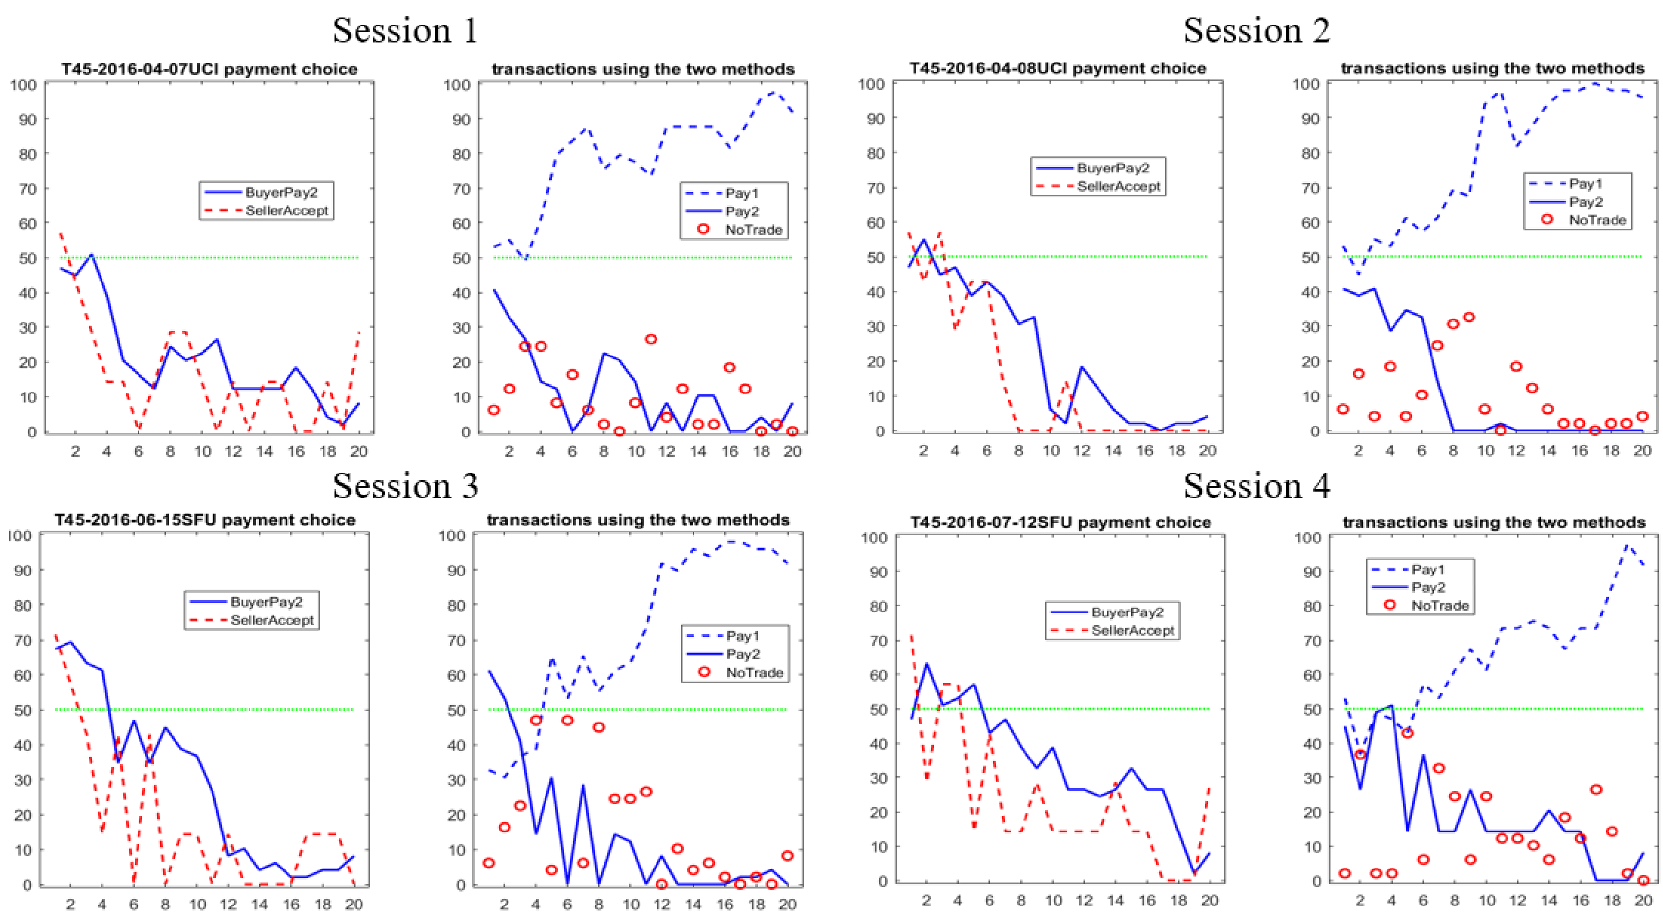

Figure E.1: Payment Choice and Usage T=4.5

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first shows payment choices: the blue solid line represents the percentage of money allocated to payment 2, averaged across the seven buyers; the red dashed line represents the percentage of the seven sellers accepting payment 2. The second figure describes payment usage. The solid blue line is the percentage of meetings using payment 2; the dashed blue line is the percentage of meetings using payment 1; red circles are the percentage of meetings where no trade takes place.

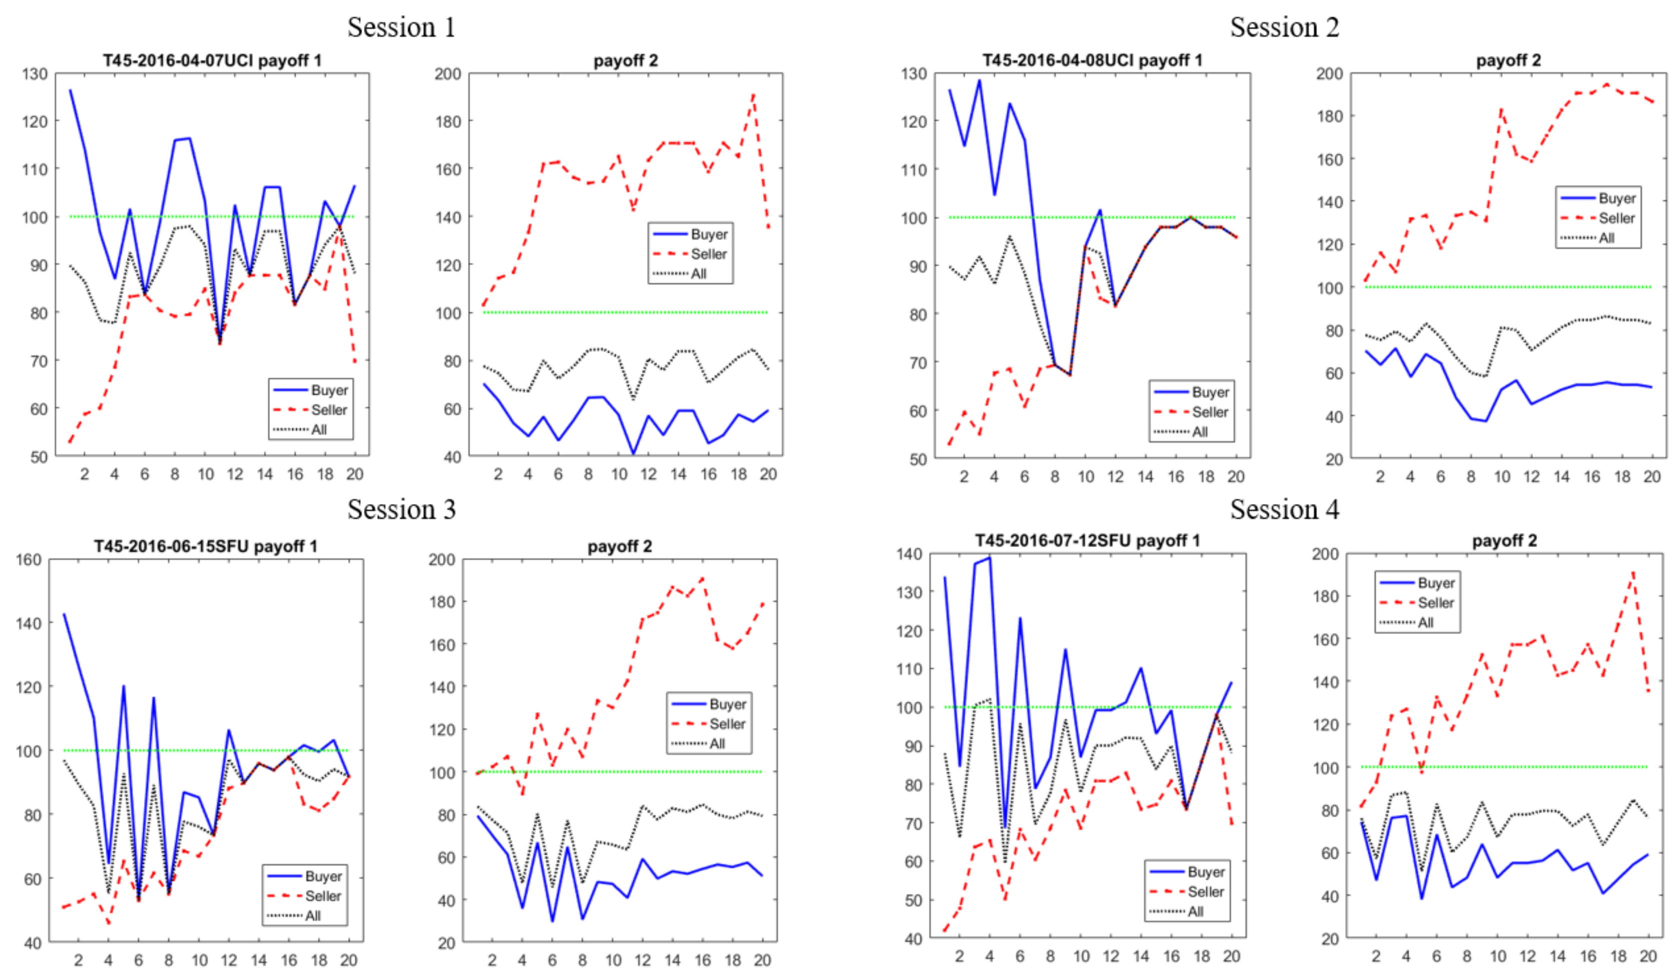

Figure E.2: Payoff T=4.5

Notes. (1) Horizontal axis: market. (2) There are two figures for each session. The first figure shows subjects' payoffs normalized by the payoffs in the payment-1 equilibrium as the benchmark (100%). The blue solid line represents sellers, the red dashed line represents buyers and the black line marked with circles represents buyers and sellers together. The second figure shows payoffs normalized by payoffs in the payment-2 equilibrium as the benchmark.

Table E.3: Rank-sum Test (T=4.5 versus other three treatments)

		rank-sum treatment 1	rank-sum treatment 2	z-value	p-value
T=1.6 versus T=4.5					
Session average	bPay2%	26	10	2.309	0.021**
	sAccept%	26	10	2.323	0.020**
	pay2Meetings2%	26	10	2.309	0.021**
	No-tradeMeetings%	10	26	-2.309	0.021**
First market	bPay2%	22.5	13.5	1.348	0.178
	sAccept%	25	11	2.097	0.036**
	pay2Meetings2%	24	12	1.764	0.078*
	No-tradeMeetings	11	25	-2.124	0.034**
T=2.8 versus T=4.5					
Session average	bPay2%	26	10	2.309	0.021**
	sAccept%	26	10	2.309	0.021**
	pay2Meetings2%	26	10	2.309	0.021**
	No-tradeMeetings%	10	26	-2.309	0.021**
First market	bPay2%	23	13	1.479	0.139
	sAccept%	19	17	0.298	0.766
	pay2Meetings2%	19.5	16.5	0.438	0.661
	No-tradeMeetings%	18	18	0.000	1.000
T=3.5 versus T=4.5					
Session average	bPay2%	25	11	2.021	0.043**
	sAccept%	26	10	2.323	0.020**
	pay2Meetings2%	26	10	2.309	0.021**
	No-tradeMeetings%	10	26	-2.309	0.021**
First market	bPay2%	16	20	-0.592	0.554
	sAccept%	16	20	1.222	0.222
	pay2Meetings2%	18	18	0.000	1.000
	No-tradeMeetings%	11.5	24.5	-1.947	0.052*

Notes. (1) Combined sample size for each test is 8. (2) *p-value≤0.1; p-value≤0.05; p-value≤0.01.

Table E.4: Test of Convergence to Symmetric Equilibrium $T = 4.5$

Session		bPay2%			sAccept%			Pay2Meetings%		
		$1-\lambda$	$100-\frac{\mu}{1-\lambda}$	$\frac{\mu}{1-\lambda}$	$1-\lambda$	$100-\frac{\mu}{1-\lambda}$	$\frac{\mu}{1-\lambda}$	$1-\lambda$	$100-\frac{\mu}{1-\lambda}$	$\frac{\mu}{1-\lambda}$
1	Coef.	0.220	87.686***	12.314	0.587***	86.772***	13.228***	0.423**	92.353***	7.647
	Std.Err.	0.139	10.495	10.495	0.211	5.569	5.569	0.194	5.711	5.711
2	Coef.	0.118	96.471***	3.529	0.253*	96.091***	3.909	0.134	103.78***	-3.778
	Std.Err.	0.099	22.867	2.867	0.152	14.602	14.602	0.137	23.643	23.643
3	Coef.	0.124	95.312***	4.688	0.651***	86.976	13.024**	0.464***	92.655***	7.345
	Std.Err.	0.081	22.428	22.428	0.152	5.130	5.130	0.121	5.246	5.246
4	Coef.	0.088	87.654***	12.346	0.689***	79.964***	20.036***	0.504***	83.644***	16.356***
	Std.Err.	0.128	40.628	40.628	0.163	0.741	0.741	0.152	4.681	4.681

Notes. (1) *p-value≤0.1; p-value≤0.05; p-value≤0.01. (2) Number of observations: 19.

Table E.5: Buyer Payment 2 Choice (%) $T = 4.5$

	(1) Stage1:bBelief(%)	(2) Stage 2:bPay2(%)
MktAcceptL(%)	**0.545*****	
	(0.053)	
bBelief(%)		**0.763*****
		(0.085)
market	**-1.536*****	**-0.623***
	(0.235)	(0.333)
location	**5.294*****	**3.920**
(SFU=1;UCI=0)	(2.204)	(2.428)

Notes. (1) *p-value≤0.1; p-value≤0.05; p-value≤0.01. (2) Each regression has 532=4x7x19 observations. There are 4 sessions, each with 7 buyers and 20 markets. For each individual, we have 19 observations as the first-stage regression uses a lagged variable as an independent variable. (3) The regressions have clustered errors at the individual subject level.

Table E.6: Seller Acceptance of Payment 2 $T = 4.5$

	(1) Stage 1:sBeliefB(%)	(2) Stage 1:sBeliefS(%)	(3) Stage 2:sAccept
sBeliefB(%)			**0.816**
			(0.665)
sBeliefS(%)			**-0.293**
			(0.923)
sAcceptL(%)	**0.089**	**4.712**	
	(10.836)	(6.207)	
sPay2DealL(%)	**0.258*****	**0.167***	
	(0.129)	(0.089)	
sNoDealL(%)	**0.378*****	**0.238*****	
	(0.094)	(0.077)	
sOtherAcceptL(%)	**0.242*****	**0.391*****	
	(0.088)	(0.095)	
market	**-1.270*****	**-1.333*****	**-0.103**
	(0.480)	(0.433)	(0.606)
location	**-0.959**	**-5.953***	**1.663**
(SFU=1;UCI=0)	(2.852)	(3.251)	(6.569)

Notes. (1) *p-value≤0.1; p-value≤0.05; p-value≤0.01. (2) Each regression has 532=4x7x19 observations. There are 4 sessions, each with 7 sellers and 20 markets. For each individual, we have 19 observations as the first-stage regression uses a lagged variable as an independent variable. (3) The regressions have clustered errors at the individual subject level. (4) For the stage-2 regression, coefficient represent the marginal effect on the probability of sellers accepting payment 2.

F Beliefs

In this Appendix, we provide some additional results regarding subjects' beliefs. Figures F.1 to F.3 show the histogram of bBelief, sBeliefB, sBeliefS at the beginning (market 1), in the middle (market 10) and at the end (market 20) to show how these beliefs evolved over time. One observation is that the initial, market 1 beliefs tend not to differ significantly across the four treatments. However, these beliefs evolve over time and become significantly different from one another in the later markets across the different treatments. Given the strong dependence of payment decisions on beliefs, the dynamic pattern of beliefs also translates into payment adoption. As pointed out in the main text, the payment adoption and use variables also started out with roughly similar initial conditions and also diverged over time.

Our experimental results suggest that sellers are more willing to adopt the new payment method than buyers when T is 1.6 and 2.8, and in the beginning of the 20 markets when T=3.5 and 4.5. Below we ask whether this can be attributed to differences between buyers' beliefs and sellers' beliefs. To investigate this question, we show in Table F.1 the average bBelief among all buyers and sBeliefB among all sellers across all 20 markets. We carry out a Kolmogorov-Smirnov (K-S) test, where the null hypothesis is that the two belief distributions follow the same cumulative density function (CDF). The last column of Table F.1 reports the p-values from this test. Table F.2 shows the results for the same exercise, but for the first market only (columns 2-4) and for the first two markets only (columns 5-7). Figures F.4-F.6 graph the CDFs of the beliefs used in these comparisons.

First, to check whether sellers are more willing to adopt the new payment method than buyers when T is not too big (i.e., when $T = 1.6$ and 2.8) and whether this can be attributed to differences in beliefs, we examine the first two rows of Table F.1. We see that that bBelief tends to be higher than sBeliefB when $T = 1.6$, while the distribution of the two belief terms is not significantly different from one other for $T = 2.8$. The experimental results that sellers are leading adoption in these two treatments therefore cannot be attributed to differences in buyer and seller beliefs.

Second, to see whether sellers are more willing to adopt the new payment method than buyers at beginning of the 20 markets when T is large, we look at table F.2. The table shows that bBelief and sBeliefB are not significantly different from one another for the first (or the first two) market(s). Again, it seems that the willingness of sellers to adopt the new payment is not driven by more optimistic beliefs, but is instead driven by the fear of losing transactions.

This result is not too surprising given that sellers face different choices than buyers. Sellers alone make a binary adoption decision; not accepting the new payment method runs the risk of losing transactions. Buyers, on the other hand, make a *portfolio* choice and they tend to split their endowment evenly between the two payment methods to test the waters in the early markets (from Table 2 and E.1, "payment choice and usage", the buyer's allocation to payment 2 in the first market is 59% when $T = 1.6$, 60% when $T = 2.8$, 49% when $T = 3.5$ and 52% when $T = 4.5$).

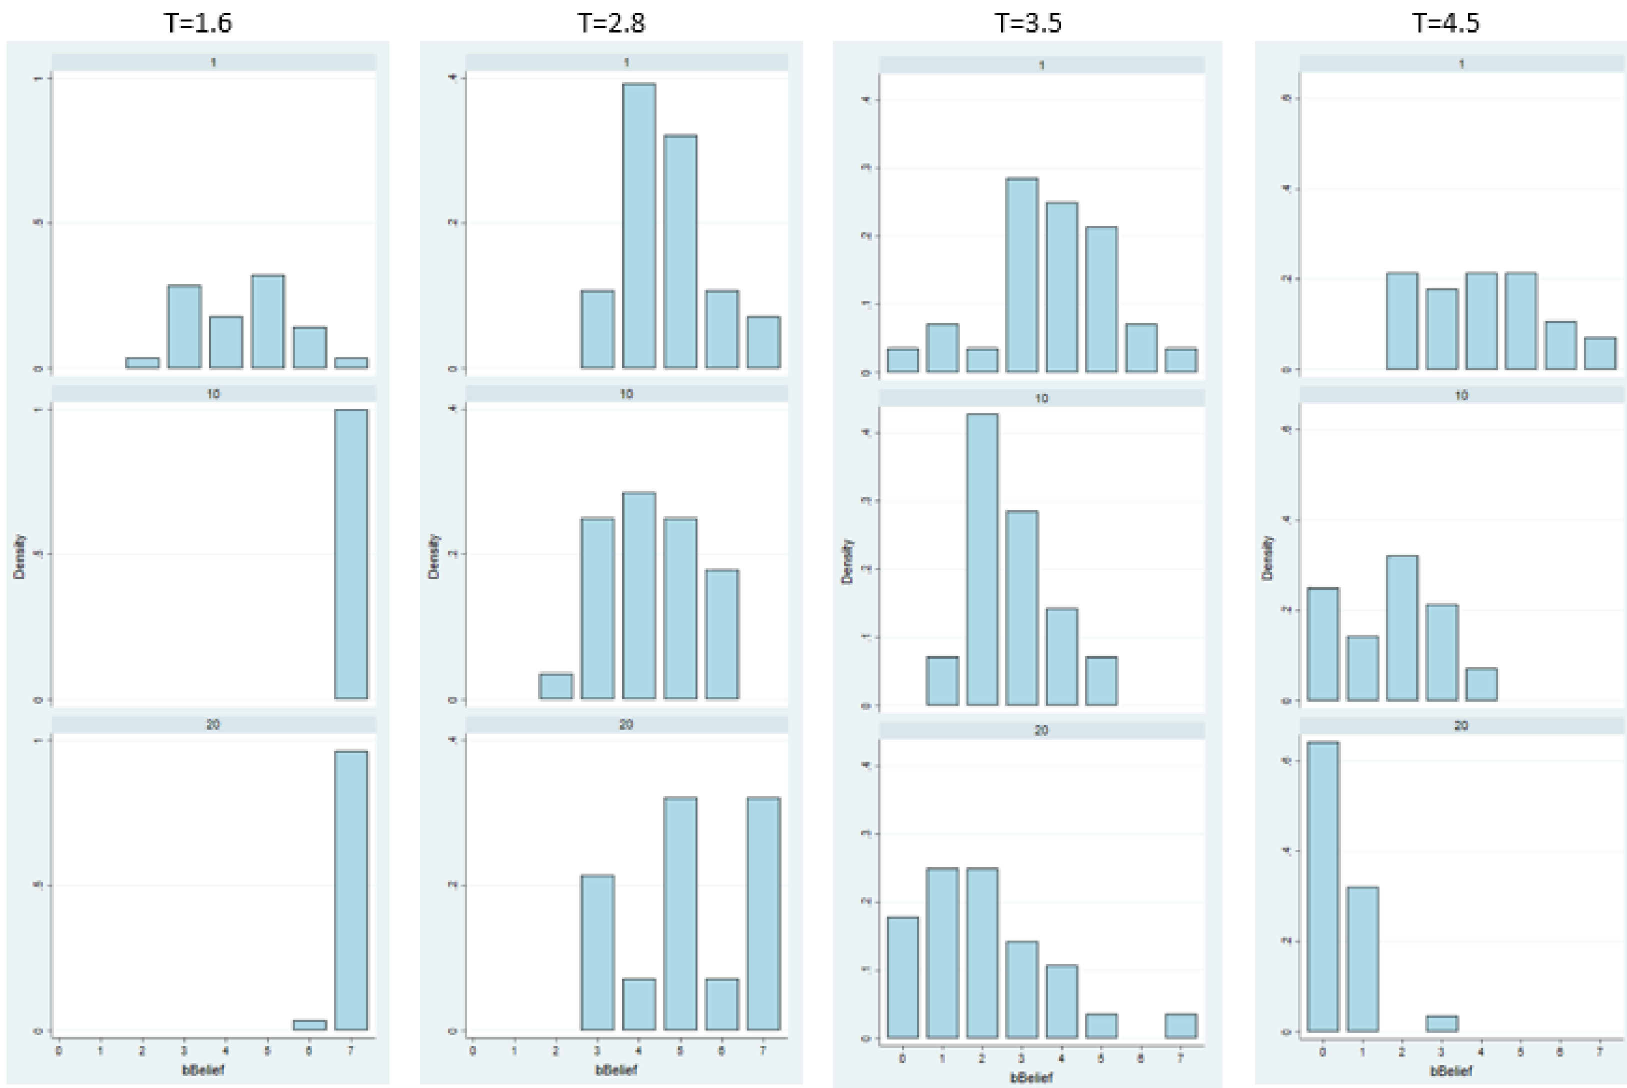

Figure F.1: Histograms of bBelief by treatment

Notes. (1) Each column in this figure represents one treatment. For each treatment, we show histograms of bBelief in market 1 (top row), market 10 (middle row) and market 20 (bottom row). (2) Each histogram shows the probability density distribution of beliefs across the 28 subjects who participated as buyers in the four sessions of the same treatment.

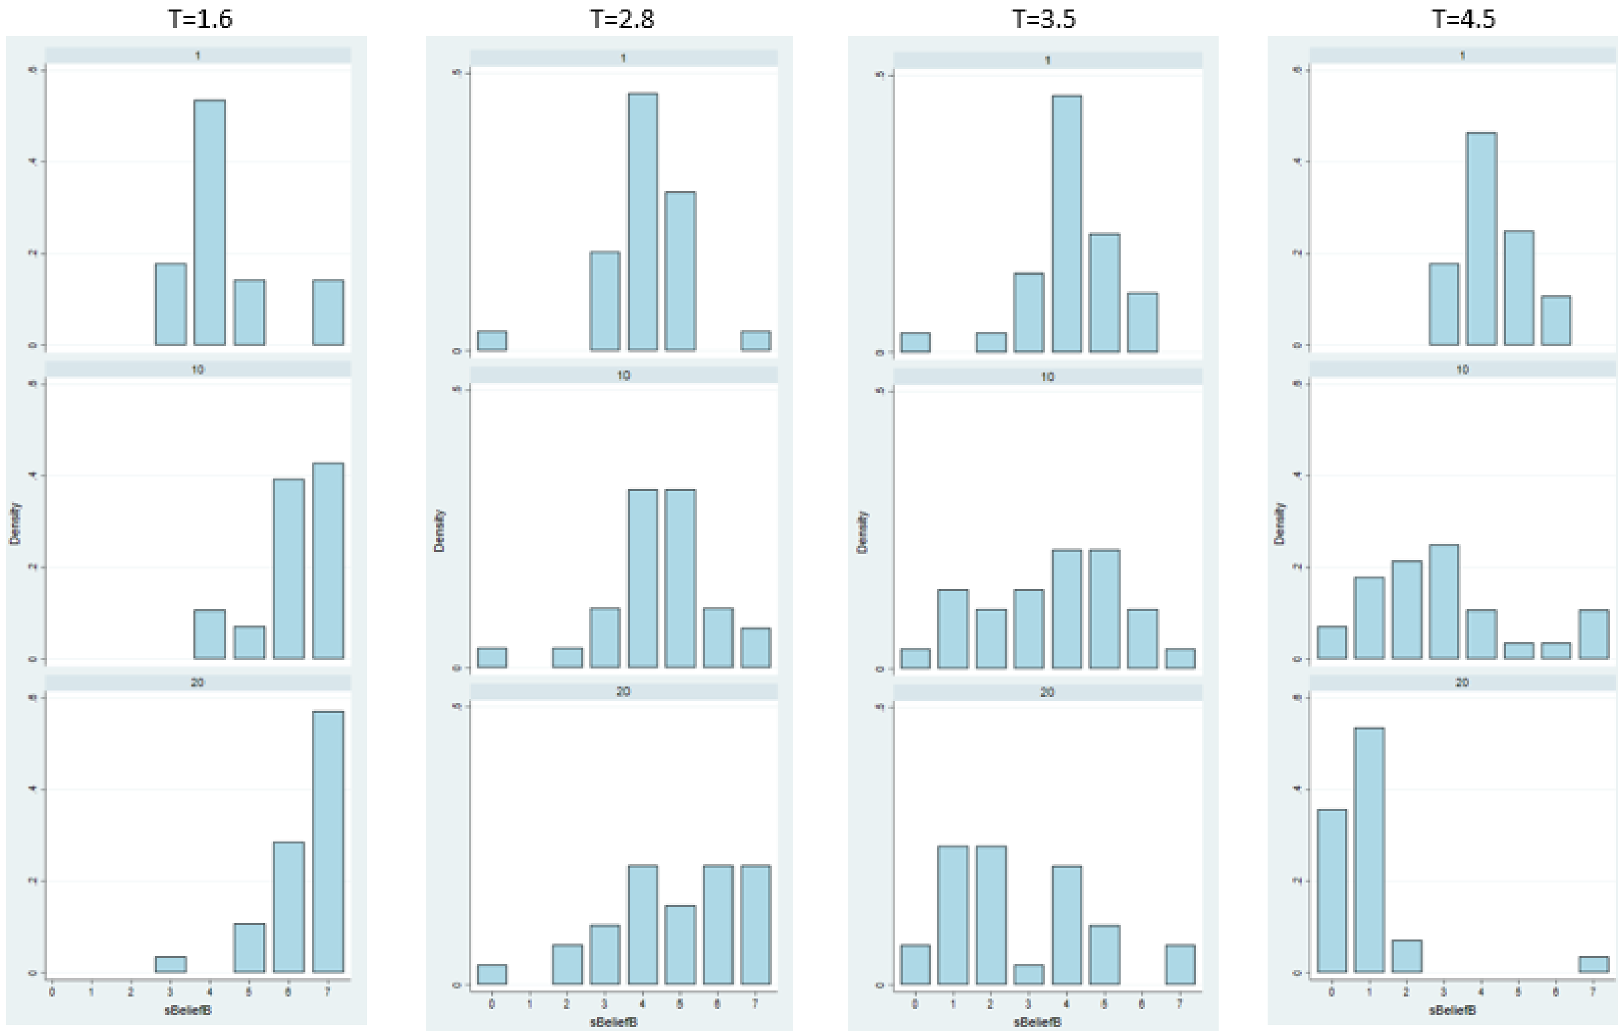

Figure F.2: Histogram of sBeliefB by treatment

Notes. (1) Each column in this figure represents one treatment. For each treatment, we show histograms of sBeliefB in market 1 (top row), market 10 (middle row) and market 20 (bottom row). (2) Each histogram shows the probability density distribution of beliefs across the 28 subjects who participated as sellers in the four sessions of the same treatment.

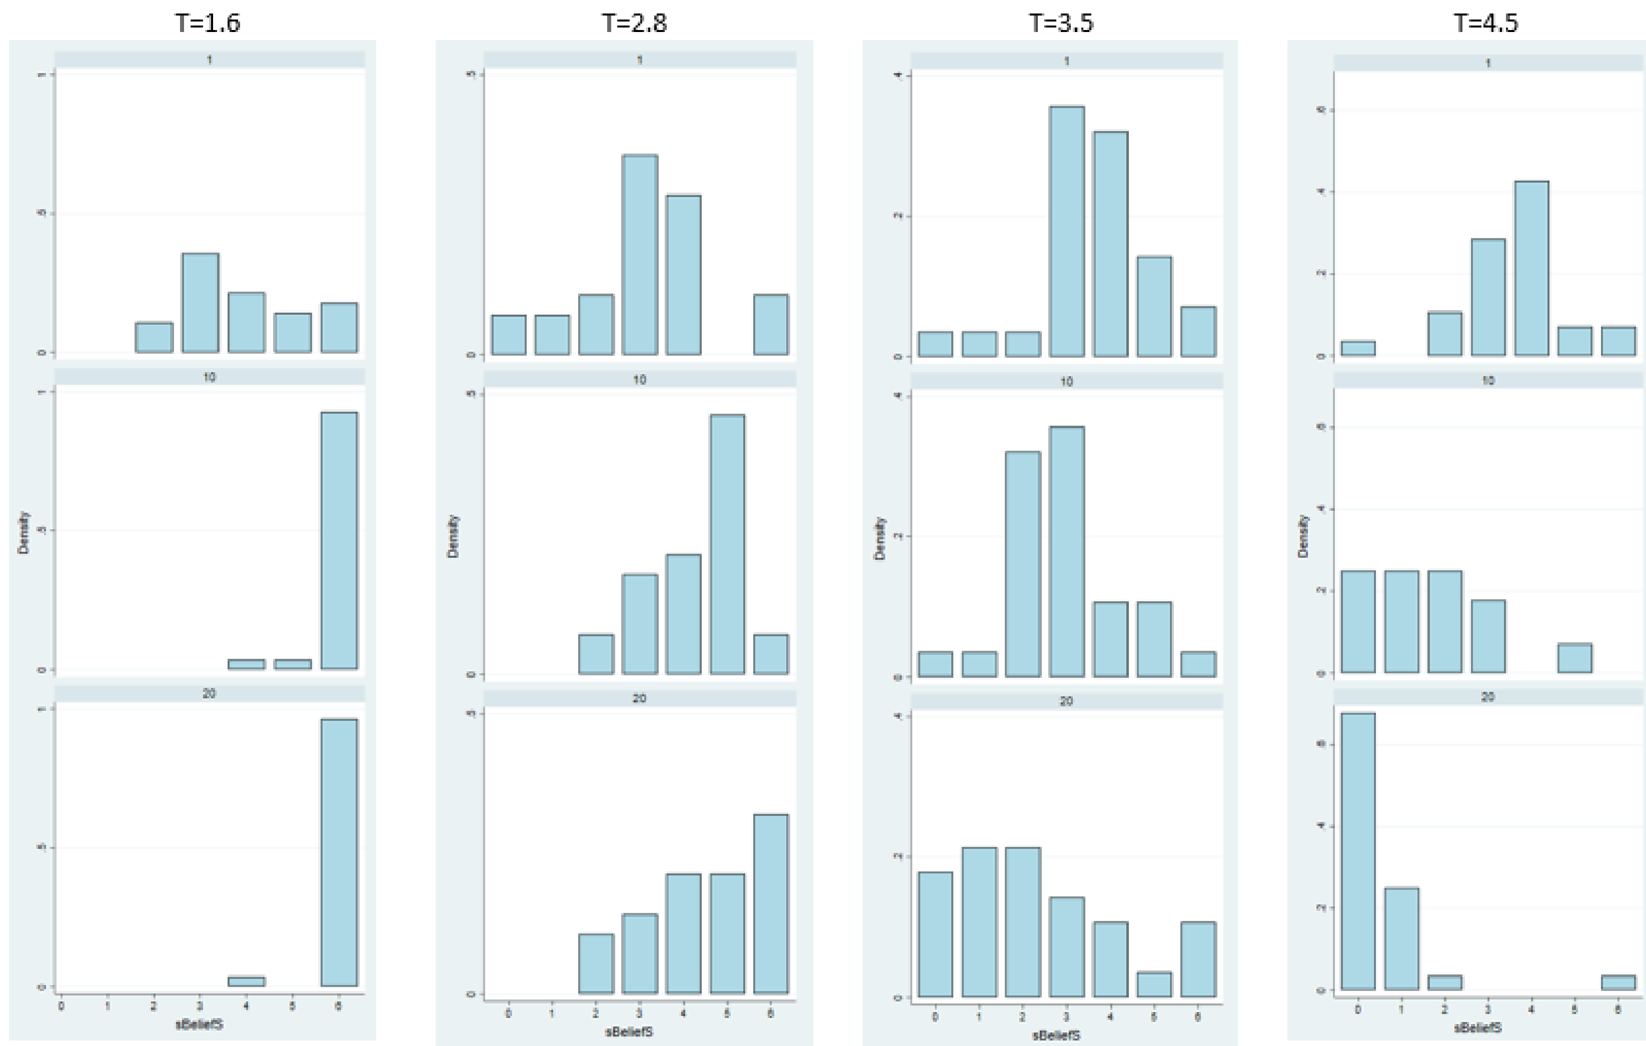

Figure F.3: Histogram of sBeliefS by treatment

Notes. (1) Each column in this figure represents one treatment. For each treatment, we show histograms of sBeliefS in market 1 (top row), market 10 (middle row) and market 20 (bottom row). (2) Each histogram shows the probability density distribution of beliefs across the 28 subjects who participated as sellers in the four sessions of the same treatment.

Table F.1: Differences in Buyer and Seller Beliefs (All Markets)

T	bBelief	sBeliefB	p-value of K-S test
1.6	6.44	5.86	0.000
2.8	4.58	4.48	0.320
3.5	3.01	3.59	0.000
4.5	1.78	2.51	0.000

Table F.2: Differences in Buyer and Seller Beliefs in Early Markets

	Market 1			Markets 1-2		
T	bBelief	sBeliefB	p-value of K-S test	bBelief	sBeliefB	p-value of K-S test
1.6	4.36	4.39	0.541	4.79	4.45	0.230
2.8	4.64	4.07	0.763	4.64	4.30	0.334
3.5	3.75	4.07	0.541	4.07	4.20	0.999
4.5	4.04	4.29	0.541	4.13	4.05	0.979

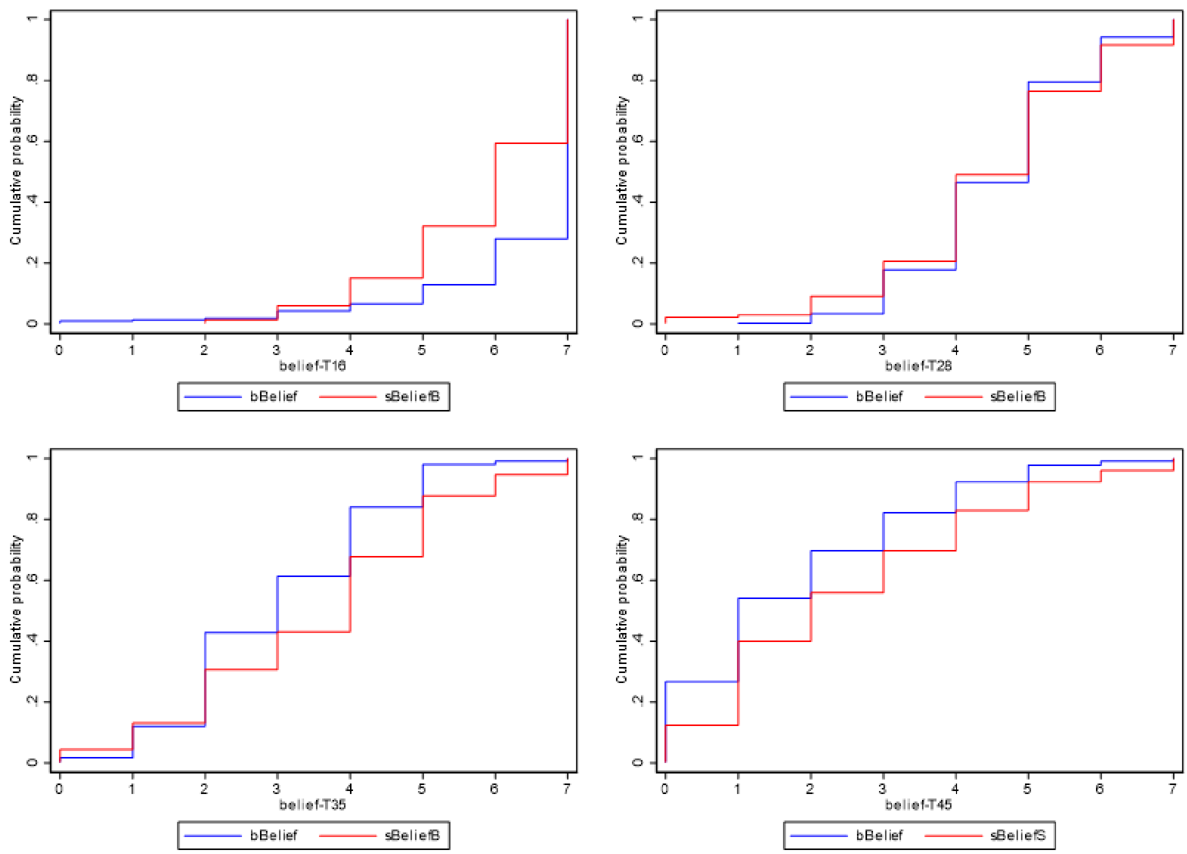

Figure F.4: Comparison of bBelief and sBeliefB (All Markets), by Treatment

Notes. (1) Each figure represents one treatment. The CDF for bBelief is in blue, and for sBeliefB is in red. Each line is the cumulative density distribution of beliefs in the 20 markets by the 28 buyers or sellers in the four sessions of the same treatment.

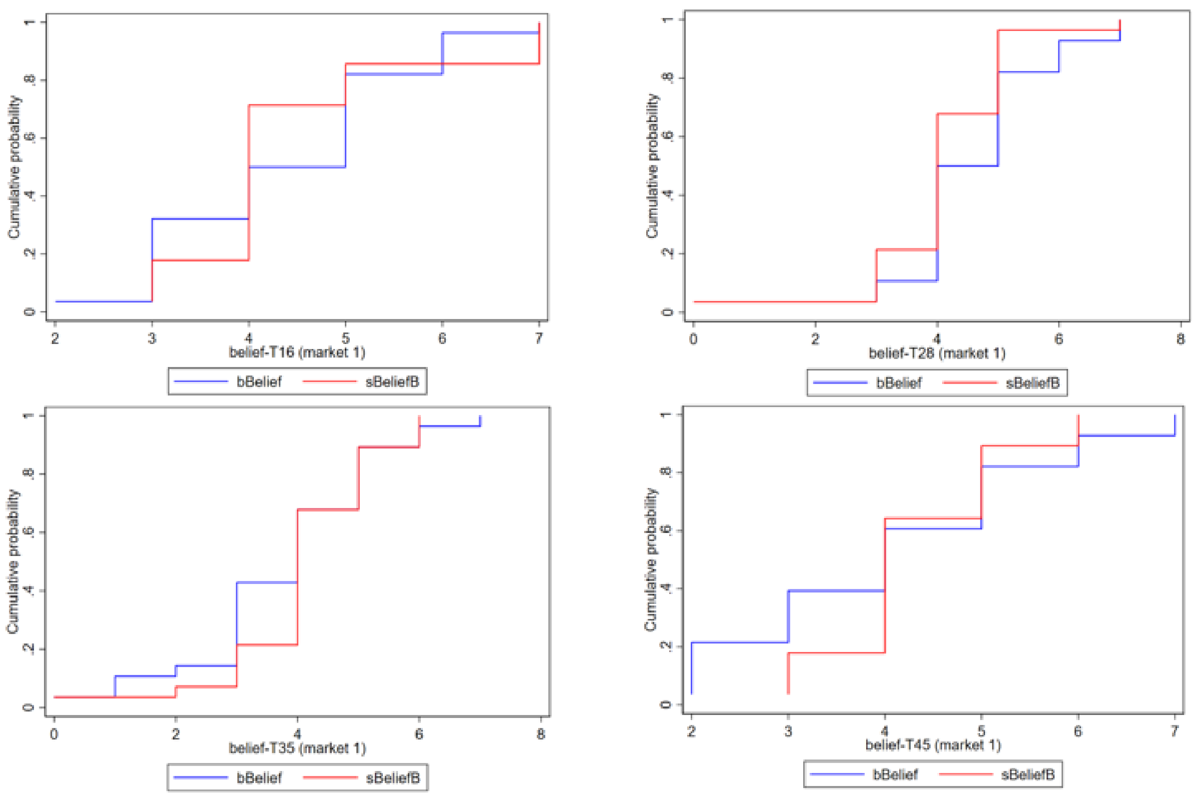

Figure F.5: Comparison of bBelief and sBeliefB (Market 1 only), by Treatment

Notes. (1) Each figure represents one treatment. The CDF for bBelief is in blue, and for sBeliefB is in red. Each line is the cumulative density distribution of beliefs in the first market by the 28 buyers or sellers in the four sessions of the same treatment.

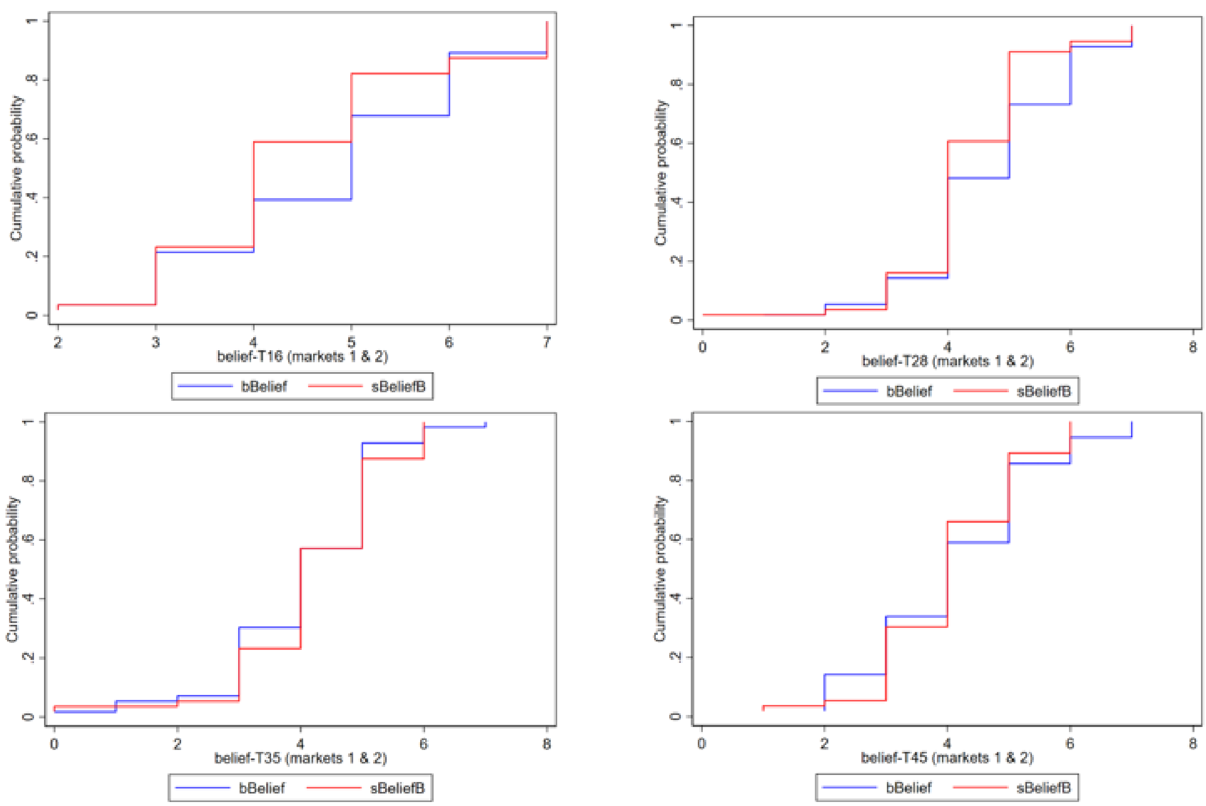

Figure F.6: Comparison of bBelief and sBeliefB (Markets 1 and 2 Only), by Treatment

Notes. (1) Each figure represents one treatment. The CDF for bBelief is in blue, and for sBeliefB is in red. Each line is the cumulative density distribution of beliefs in the first two markets by the 28 buyers or sellers in the four sessions of the same treatment.