

Jiang, Janet Hua; Puzzello, Daniela; Zhang, Cathy

Working Paper

How long is forever in the laboratory? Three implementations of an infinite-horizon monetary economy

Bank of Canada Staff Working Paper, No. 2021-16

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Jiang, Janet Hua; Puzzello, Daniela; Zhang, Cathy (2021) : How long is forever in the laboratory? Three implementations of an infinite-horizon monetary economy, Bank of Canada Staff Working Paper, No. 2021-16, Bank of Canada, Ottawa,
<https://doi.org/10.34989/swp-2021-16>

This Version is available at:

<https://hdl.handle.net/10419/241239>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

How Long is Forever in the Laboratory? Three Implementations of an Infinite-Horizon Monetary Economy

by Janet Hua Jiang,¹ Daniela Puzzello,² Cathy Zhang³

¹Currency Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

²Indiana University

³Purdue University

JJiang@bankofcanada.ca, dpuzzell@indiana.edu, cmzhang@purdue.edu



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

ISSN 1701-9397

©2021 Bank of Canada

Acknowledgements

We thank John Duffy, Guillaume Rocheteau, Yaroslav Rosokha, Neil Wallace, Randall Wright, and Sevgi Yuksel for their discussions and comments. Xinxin Lyu provided excellent research assistance. The views expressed in this paper are those of the authors and no responsibility should be attributed to the Bank of Canada.

Abstract

We compare three implementation schemes of an infinite-horizon monetary economy with discounting. Under the standard random termination scheme and its block variation, the economy lasts for an indefinite number of periods and the discounting factor is captured by the probability that the economy continues to the next period. These schemes rely on the belief that the experimenter can credibly implement a game that lasts an arbitrarily long time. We also propose a new method that does not rely on such a belief. Under this scheme, subjects participate in an experiment for a fixed number of periods where the discount factor is captured by a weighting factor that shrinks the payoffs over time. Dynamic incentives are preserved by paying subjects their continuation value, which is based on past market prices. The results show that dynamic incentives are preserved, and behavior is similar in all three implementations. Researchers may decide among these approaches, depending on the research question of interest and more practical concerns, such as the ease of implementation and the need to collect data for multiple supergames when the discount factor is high.

Topics: Central bank research, Economic models, Inflation and prices

JEL codes: C92, D83, E40

1 Introduction

Many economic models, such as those of repeated games, asset pricing and macroeconomics, adopt an infinite horizon with discounting. An important question for testing these models in the laboratory is how to implement an infinite horizon. The standard implementation of infinite-horizon environments uses random termination (RT), which interprets the discount factor as the probability that the game continues to the next period (see Roth and Murnighan 1978). There are clear advantages associated with this method. One is that it is straightforward to implement in the lab. Another, as observed by Mailath and Samuelson (2006), is that it is realistic: “viewing the horizon as uncertain in this way allows the model to capture some seemingly realistic features. One readily imagines knowing that a relationship will not last forever, while at the same time never being certain of when it will end.”

The RT method also has potential issues. For example, Mailath and Samuelson (2006) warn that “[u]nder this interpretation the distribution over possible lengths of the game has unbounded support. One does not have to believe that the game will last forever but must believe that it could last an arbitrarily long time.” Further, Selten et al. (1997) assert that “[i]nfinite supergames cannot be played in the laboratory. Attempts to approximate the strategic situation of an infinite game by the device of a supposedly fixed stopping probability are unsatisfactory since a play cannot be continued beyond the maximum time available.”¹ In other words, a potential issue is that the RT interpretation relies on the belief that if the game were to continue for an arbitrarily long time, then the experimenter could credibly implement such a game, which is clearly not possible.

The credibility issues associated with RT are especially pertinent for experimental monetary economics. First, this could jeopardize the existence of monetary equilibria where fiat money is valued (a concern voiced by Davis et al. 2020). In view of the limited recruitment time for a session, subjects should expect that fiat money will become useless when the session concludes. By backward induction, they would not have accepted money in prior periods and the monetary equilibrium would unravel. Second, high discount factors or continuation probabilities are often desirable for experimental monetary economics. For example, having many rounds within a monetary economy makes it easier for subjects to observe price trends and also allows researchers to better understand the impact of various monetary policies on inflation, output and welfare (e.g., Jiang et al. 2020).² Third, monetary models are often

¹Dal Bó and Fréchette (2018) discuss additional implementation issues associated with random termination methods.

²For another example, Jiang and Zhang (2018) study the acceptance of foreign currencies based on the

more complicated than repeated games, so it takes a longer time to go through the experimental instructions and complete one period of play. In addition, it may also be desirable to run multiple sequences so subjects can learn. As a result, it is more likely that the limited recruitment time may become binding in macro environments.

In this paper, we implement an infinite-horizon monetary economy through the standard random termination and its variation, the block random termination (BRT) developed by Fréchet and Yuksel (2017). Relative to most experimental studies on repeated games, we adopt a higher discount factor in order to obtain (on average) longer sequences. We also develop a new method that circumvents the credibility problem that is associated with random termination. Under this new scheme, subjects play a definite number of periods, say T periods with discounting, where their payoff shrinks by a certain amount each period. Specifically, subjects receive a fraction, equal to the discount factor $\beta \in (0, 1)$, of the payoff each period. The dynamic incentives in the infinite-horizon model are preserved by paying subjects a continuation value at the end of period T , which amounts to the present discounted value of the game from period T , thereafter, based on prior market outcomes (market prices). We call this method definite + discounting (DD) since it only relies on a definite number of periods and an interpretation of discounting.

The number of experimental studies based on infinite-horizon monetary economies has grown over the last two decades (see e.g., Duffy and Ochs 1999, 2002; Duffy and Puzzello 2014ab; Camera and Casari 2014; Jiang and Zhang 2018; Ding and Puzzello 2020; Rietz 2019; and Kamiya et al. 2019, among many others) and continues to grow. Omnipresent in the mind of these researchers is which implementation scheme to adopt. It is useful to evaluate and compare different methods of implementation schemes so as to guide the design of future experiments.

The underlying model for our experiment is a version of the infinite-horizon monetary environment of Lagos and Wright (2005) and Rocheteau and Wright (2005), which has a constant money supply. These models are micro-founded and hence amenable to experimental methods. Their key outcomes, including prices, output and welfare, can be directly observed or measured from the lab. Many applications have been based on these workhorse models, such as open market operations, currency competition, and forward guidance (see Williamson and Wright 2011; and Lagos, Rocheteau and Wright 2017 for surveys). Building an experimental infrastructure based on these models will pave the way for studying many interesting monetary issues in the laboratory.

infinite-horizon model by Matsuyama et al. (1993) and use a very long fixed horizon for the experiment. There, subjects are initially endowed with home tokens, and a long trading sequence or a high discount factor is desirable to give a chance for foreign tokens to circulate across groups.

In this paper, we explore whether credibility concerns associated with the two random termination schemes are of empirical relevance, and we compare which of the three methods is closest to implementing an underlying infinite-horizon model that will help inform and guide the design of future experiments. Our experimental results suggest that all three implementation schemes preserve the dynamic incentives that underlie our theoretical model. Most experimental economies function within a reasonable neighbourhood of their theoretical predictions. Among the three schemes, economic behavior is similar along some dimensions but differs along others. In particular, price dynamics and inflation rates are not significantly different across the three treatments and are broadly consistent with the theoretical prediction of zero inflation. Output is closer to the theoretical prediction in the RT and BRT treatments and not significantly different between the two. However, output is significantly lower in the DD treatment relative to the theoretical prediction and the other two treatments. Welfare is, in general, lower than the theoretical prediction in all three treatments. Across treatments, welfare is significantly lower in the BRT and DD treatments relative to the RT method.

To evaluate which implementation scheme is preferred overall, we assess the three methods along several dimensions. The findings from our experiments provide empirical support that our new method and the methods based on random termination are effective at avoiding issues associated with backward induction and are able to preserve the dynamic incentives in the underlying infinite-horizon model. Within the context of our experiment, the standard random termination method generates results that are most consistent with the theoretical prediction. The RT method is also relatively easy to implement, compared with the other two schemes. The other treatments require the experimenter to spend a non-negligible amount of time educating subjects on additional concepts; for example, the “block” in the BRT and the “continuation value” in the DD.

However, given the similarity in outcomes across the three schemes, one may find the other two alternatives desirable, depending on the question of interest. For instance, the BRT ensures that each sequence has a minimum number of periods (equal to the length of the block) and this could be useful for studying the emergence of an international currency as it takes time for a foreign currency to circulate. On the other hand, if the experimenter needs to collect data on multiple supergames for environments with high discount factors, then the DD method may be desirable as one can control the length of the fixed horizon. For example, one may want to run multiple supergames using a within-subjects design to study the adoption of a new technology, currency, or means of payment.³ Another advantage of

³For instance, the first supergame could capture the status quo. The new alternative could be introduced alongside the existing option in the second supergame. Then, the two options present in the second supergame

the DD scheme is that the fixed length facilitates comparisons among different supergames, sessions, and treatments.

The rest of this paper proceeds as follows. In Section 2, we discuss the related literature. Section 3 describes the theoretical environment and equilibrium. Section 4 focuses on the experimental design and procedures. We show and discuss the experimental results in Section 5. Section 6 concludes and mentions directions for future research.

2 Related Literature

The closest study to this paper is Fréchette and Yuksel (2017), who compare four implementations of an infinitely repeated prisoner’s dilemma (PD) experiment: (1) a standard random termination (Roth and Murnighan 1978); (2) a deterministic discounted play followed by a random termination (Sabater-Grande and Georgantzis 2002); (3) a block random termination; and (4) a deterministic discounted play followed by a coordination game (Anderson and Wengstrom 2012; Cooper and Kuhn 2014). The first method is the standard method used in experimental economics. The second and third methods also rely on random termination and have been used by researchers to implement infinite-horizon settings (see e.g., Duffy et al. 2019; Jiang et al. 2020; Sabater-Grande and Georgantzis 2002; and Wilson and Wu 2017). The fourth method features a fixed number of deterministic discounted rounds, followed by a coordination game induced by considering the grim trigger strategy and the strategy of always defecting, both of which are two particular strategies that are commonly adopted in the infinitely repeated game.

In this paper, we compare different implementation schemes in the context of infinite-horizon monetary models. We examine the standard RT and BRT as well as a new approach that can sustain dynamic incentives even if subjects play the monetary game for a fixed number of periods. There is a key difference between our new approach and the continuation-game approach used for repeated games experiments. Sustaining dynamic incentives in the continuation game requires that subjects build a particular link between their choice of strategy in the game and the history of play in the prior rounds of the game. An example of such a strategy is that if nobody has defected in the rounds prior to the coordination game then the subjects will use the grim trigger strategy, otherwise they will always use the defect strategy. The experimental evidence from Cooper and Kuhn (2014) and Fréchette and Yuksel (2017) suggests that in the game of the prisoner’s dilemma (PD), subjects perceive

could remain available in future supergames so as to explore the evolution associated with the adoption of different alternatives.

the continuation of the game as disjointed from the previous play and their behavior in the coordination game tends not to depend on the history of the play in the deterministic discounted part of the game.⁴ In our DD approach, subjects do not need to play a separate continuation game or build such a link; the continuation value is directly calculated from past market outcomes. In other words, in our approach, the continuation value automatically connects past outcomes by extrapolating the past to the future.⁵

In terms of the findings, Fréchette and Yuksel (2017) find that inter-temporal incentives and cooperation are best preserved under standard random terminations in repeated PD games. Our experimental results are consistent with theirs: we find that methods that rely on random termination effectively preserve dynamic incentives in the monetary economy; we also find that the behavior in the standard random termination scheme is close to the theoretical predictions. However, given that the standard random termination overall produces results that are comparable with the other methods, the most appropriate method depends on the research question of interest.

Marimon and Sunder (1993) and Lim et al. (1994) develop two alternative approaches to maintaining dynamic incentives in the context of overlapping generations (OG) monetary experiments. The first approach (which is used in both papers) adds to the economy forecasters whose only task is to forecast the price for the current period. The forecasters are rewarded for their forecasting accuracy. To end the game, upon gathering price predictions at the start of a period that has not previously been announced, the experimenter states that the game has reached an end and uses the mean predicted price to convert the monetary asset into real assets (that generate utility). The second approach (used in Lim et al. 1994) converts money into real assets at the average trading price in the last period of the game.⁶ Our DD method is similar to the second approach in that the terminal value of money is determined by past market outcomes. However, there are important differences between our approach to experimental settings and designs and theirs. The agents in our framework are infinitely lived as opposing to living for just two periods. The discounting factor embedded in our new method (which is absent in the OG experiments) makes our new method more comparable to the random termination treatments. Finally, in Marimon and Sunder (1993) and Lim et al. (1994), the subjects do not know when the game will end. In our design, the

⁴Davis et al. (2020) evaluate the role of money in a treatment with a finite-horizon game, followed by a hawk-dove game (unrelated to the money game) to circumvent the backward induction issue. They find a similar result that subjects fail to connect the two games. That is, subjects do not appear to condition their play in the hawk-dove game on what had previously happened in the money game.

⁵The continuation value is endogenously determined in our DD implementation. Thus, if subjects had not given value to money in the first T periods, then the continuation value would be zero. Similarly, if money was given value endogenously, then it would be reflected in the continuation value.

⁶Lim et al. (1994) find that the experimental results are robust to the two termination methods.

subjects are informed about the exact number of periods they will play the game.

Davis et al. (2020) criticize experiments that rely on the implementation of infinite-horizon settings on the same grounds as Selten et al. (1997). As a result, they explore the implementation of finite-horizon monetary models that are immune to critiques against experiments that are based on these models. In one model, the agents trade sequentially *without* knowing their positions in the line, which neutralizes backward induction so that monetary equilibria can be sustained. In another model, they consider a finitely repeated game where subjects face a hawk-dove game at the end. They then compare the value of the fiat money in these environments with an environment where agents know their positions and the monetary equilibrium does not survive backward induction. They find that the introduction of money leads to higher trade and efficiency in all economies, regardless of whether the theory predicts that money is welfare improving. Their experiment offers a useful framework within which to investigate why subjects value fiat money. However, at this stage, finite-horizon monetary models are not well-suited to study questions that are associated with inflation or the impact of monetary policies more generally.

Our paper is also related to Bruttel and Kamecke (2012), who propose an approach that is similar to ours so as to avoid the end-game effects inherent in indefinitely repeated PD games. They consider two mechanisms to compute continuation values: one that is based on the explicit construction of subjects' memory-1 strategies ("strategy") and one that is an implicit construction performed by a computer that is based on a Moore machine that has a memory of only one round ("moore"). They compare these implementation methods with random terminations. They find that cooperation rates tend to be higher under the condition of a memory-1 "strategy" and a "moore" than under the random termination method.

On a related note, Romero and Rosokha (2018) propose an approach that allows the experimenter to run indefinitely repeated games that use a high continuation probability. Their approach entails subjects directly constructing strategies (including those longer than memory-1) that are used to partially or fully automate action choices in the experiment. This automation speeds up action choices and allows more repetitions to be conducted within the time limit of an experimental session (which is equivalent to running a longer session). Their approach is useful in settings where subjects strategically interact with a counterparty and it is straightforward to construct plans of action. It would be useful for future research to (i) see whether this approach can be adapted to the competitive market settings adopted in our study, and (ii) use this approach in combination with our DD method to implement infinitely repeated PD games.⁷

⁷See section 5.5 for more details on how our method can be extended to infinitely repeated PD games.

3 Theoretical Framework

Our experimental economy is based on the infinite-horizon monetary models of Lagos and Wright (2005) and Rocheteau and Wright (2005). The model is micro-founded and thus well suited for laboratory implementations. Agents trade in two alternating markets and have quasilinear preferences in one market. These quasilinear preferences make the distribution of money balances degenerate and the model highly tractable. As a result, the model has served as a workhorse upon which many applications, such as open market operations, currency competition, forward guidance, have been developed. Building an experimental framework based on this model will pave the road for studying many interesting monetary issues in the laboratory in the future.

Specifically, in this paper we implement a simple version of the competitive markets monetary model of Rocheteau and Wright (2005); however, whereas the money supply in their model can be constant or can change, we implement a model with a constant money supply.⁸ The alternating centralized market structure is reminiscent of Bewley (1980) and Townsend (1980), where agents periodically alternate between the roles of buyers and sellers. The Bewley-Townsend model has the advantage that agents have symmetric preferences, while the quasilinear preferences in Rocheteau and Wright (2005) simplify welfare analysis since only consumption and output in the first market are relevant for welfare.

3.1 Environment

Time is discrete and continues forever. There are two types of agents, labelled type A and type B, each of size N . Each period consists of two markets, A and B, that open in sequence. In each market, there is a divisible good and a perishable good, called good A in market A and good B in market B. In market A (B), type A (B) agents want to consume but cannot produce, while type B (A) agents can produce but do not want to consume (in the following, we index the goods or markets with a subscript and the agent's type with a superscript). All agents discount between periods with a constant discount factor $\beta = (1 + \rho)^{-1} \in (0, 1)$, where $\rho > 0$ is the rate of the time preference. Instantaneous utilities for type A and B agents are given by

$$U^A = u(x_A) - x_B,$$

$$U^B = -x_A + v_0 + x_B.$$

⁸We also use this environment in a previous paper (Jiang et al. 2020) but with positive money growth. To focus on the termination scheme in this study, we chose to abstract away from positive inflation and study the case with a constant money supply.

Type A agents get utility $u(x_A)$ from consuming x_A units of good A, where $u'(0) > 0$, $u''(0) < 0$ and $u'(0) = \infty$, and incur disutility x_B from producing x_B units of good B. For type B agents, the disutility from producing good A is x_A and the utility from consuming good B is $v_0 + x_B$.⁹ The first-best level of output in market A is x_A^* such that $u'(x_A^*) = 1$.

Lack of commitment, no formal enforcement, and private trading histories restrict the emergence and sustainability of credit arrangements and the lack of a double coincidence of wants rules out barter.¹⁰ There is a single intrinsically useless asset, called money, that could serve as a medium of exchange. Money is divisible and storable in any amount, m_t , and the money supply is fixed at M .

3.2 Monetary Equilibrium

We focus on stationary equilibria where real variables are constant over time. Because the money supply is constant, the price level is also constant over time. In the analysis below, we omit the time subscript and use the accent “ \wedge ” to denote the variables in the next period. As is standard, we start backwards by characterizing the agents’ decision problems in market B and combining these with the choices made in market A to characterize the equilibrium. Here we summarize the value functions and describe the equilibrium allocations and prices. For additional details on solving for a stationary monetary equilibrium, see Appendix A.

In market B, agents can trade good B and money in a competitive market where the price of good B is p_B . The value function of a type i agent who enters market B with money holdings m simplifies to

$$\max_{\hat{m}} W^i(m) = \max_{\hat{m}} \left\{ \frac{-\hat{m} + m}{p_B} + v^i + \beta V^i(\hat{m}) \right\},$$

where $v^A = 0$ and $v^B = v_0$. The optimal choice for \hat{m} , the money balance taken to the following market B, solves

$$\beta \frac{\partial V^i(\hat{m})}{\partial \hat{m}} - \frac{1}{p_B} \leq 0, \text{ with equality if } \hat{m} > 0.$$

⁹We introduce the term v_0 here and in our earlier study (Jiang et al. 2020) to roughly equalize payoffs between type A and B subjects in the lab. Introducing this term does not affect equilibrium predictions. Further, we keep the functional forms linear for type B agents for the sake of higher tractability. See also Rocheteau and Nosal (2017).

¹⁰Notice that only aggregate outcomes— i.e., prices—are observable. However, since the population is finite in the lab, informal enforcement schemes are theoretically possible (see Aliprantis et al. 2007; and Araujo et al. 2012). However, in a similar framework as ours, Duffy and Puzzello (2014a) find that laboratory outcomes in economies of 6 and 14 subjects are closer to predictions of monetary equilibrium and do not find support for the emergence of informal enforcement schemes.

As usual, the value function $W^i(m)$ is linear in m and the choice of money holdings in the next period is independent of one's current money holdings.

In market A, agents can trade good A using money in a competitive market at market price p_A . Type B agents, who are producers in market A, incur a linear production cost to produce x_A units of good A. Their decision problem is

$$V^B(m) = \max_{x_A} \left\{ -x_A + \frac{(m + x_A p_A)}{p_B} + W^B(0) \right\},$$

which uses the envelope result $\frac{\partial W(m)}{\partial m} = \frac{1}{p_B}$. The first-order condition of type B's problem implies

$$p_A = p_B. \quad (1)$$

Type A agents, who are consumers in market A, can buy and consume x_A units of good A. Their value function in market A is

$$\begin{aligned} V^A(m) &= \max_{x_A} \left\{ u(x_A) + \frac{(m - p_A x_A)}{p_B} + W^A(0) \right\} \\ \text{subject to } p_A x_A &\leq m. \end{aligned}$$

If the cash constraint does not bind, then $u'(x_A) = p_A/p_B$ which, when combined with type B's decision, gives $x_A = x_A^*$. If the cash constraint binds, then $x_A = m/p_A$.

We can now combine the agents' decision problems from markets A and B to solve for the stationary monetary equilibrium. For type B agents, the net marginal value of carrying money to the next market A is equal to $\frac{1}{p_B}[-1 + \beta] < 0$. Money carried by type B agents to market A will be idle and can be used to purchase good B in the next market B. Since $\beta < 1$, holding idle balances is costly and hence type B agents will spend all of their money in market B and enter market A with zero balances.

For type A agents, the net marginal value of carrying money to the next market A is equal to $\frac{1}{p_B}[-1 + \beta u'(x_A)]$. Since $u'(0) = \infty$, type A agents take positive money balances to market A. In equilibrium, the net marginal benefit of carrying money is zero and output in market A (per consumer or producer) is

$$u'(x_A) = 1/\beta. \quad (2)$$

Since $\beta < 1$, type A agents will carry just enough money to spend in market A (or they spend all their money on good A and enter market B with zero balances).

Now we summarize the equilibrium price and quantity. In market A, each type A consumes x_A and each type B produces x_A , where x_A solves equation (2). The price level, which is the same for both markets according to equation (1), is given by the money market clearing condition

$$p_A = p_B = \frac{M}{Nx_A}. \quad (3)$$

In market B, the amount of consumption by each type B agent (which is the same as each type A agent's production) is

$$x_B = \frac{M}{Np_B} = \frac{M}{Np_A} = x_A. \quad (4)$$

4 Experimental Design

In this section, we describe the experimental design we use to implement the monetary model outlined in the previous section.

4.1 Treatments and Discounting Schemes

Our experiment considers three treatments that are characterized by the way we implement an infinite horizon with exponential discounting. The first treatment is the standard indefinite-horizon implementation by Roth and Murnighan (1978); we label it as a “random termination” or **RT** for short. In each period of the RT treatment, the economy continues to a new period with probability β and ends with complementary probability $1 - \beta$. This implementation relies on the interpretation of the discount factor as the probability of continuation.

The second treatment is a variation of the block random termination treatment proposed by Fréchette and Yuksel (2017). We label it as a “block random termination” or **BRT** for short. In this treatment, subjects make decisions in a block of T periods and learn whether the sequence has ended within the block only after the whole block has been played. If the sequence has not ended within the block, from period $T + 1$ on, then at the end of each period the subjects receive live information about whether the sequence will continue to a new period. Note that in Fréchette and Yuksel (2017) subjects always play in blocks of T periods and, if time allows, after a sequence ends they will start a new block of T periods. We modified the original procedure in our experiment because it guarantees at least T periods of data for each sequence and, at the same time, allows us to fit more sequences into a session:

after the first block, the sequence can stop anytime instead of at the end of another T -period block (see also Duffy et al. (2019)).

Finally, the third treatment features a definite horizon with discounting and integrated continuation value. This method does not rely on the experimenter’s credibility of implementing an arbitrarily long session. We label this treatment “definite + discounting,” or **DD** for short, since subjects make decisions for a definite number of periods, followed by the present discounted value of future payoffs. In this setting, subjects are informed that they will make decisions for T periods with discounting, where their payoffs shrink in each period. Specifically, in periods 1 through T , subjects receive a fraction, equal to the discount factor β , of each period’s payoff.¹¹ In period $T + 1$, the payoffs are assigned as follows: The market A price is computed as the average of market A prices in periods 1 through T . Similarly for market B prices.¹² Further, in each market of period $T + 1$, buyers automatically bid all of their tokens. Given the prices and the token bids, we compute consumption as the ratio of the buyer’s bid divided by the average price. Each producer produces an amount that is equal to the average consumption. We can then assign payoffs in period $T + 1$. At the end of period $T + 1$, we compute the continuation value or present discounted value from period $T + 2$, thereafter, based on the average consumption and output in the first $T + 1$ periods. In Appendix B, we describe in more detail the calculation of the continuation value and we check and verify that both the steady state monetary and nonmonetary equilibria of the infinite-horizon model from Section 3 remain in equilibria for the economy underlying the DD treatment.¹³

4.2 Parameterization and the Market Game

Our treatment variable is the infinite-horizon implementation method: RT, BRT or DD. We conduct four sessions for each treatment and adopt a between-subjects design where each session of the experiment consists of a new group of subjects that are making decisions under only one of the three implementation environments.

To determine the sequence lengths for each session of the two treatments with RT and BRT, we predrew four different sets of random numbers, one for each session. We chose to predraw the random numbers instead of generating random numbers in real time, during the

¹¹For example, the weighting factor applied to period 1’s payoff is 1, to period 2’s payoff is β , to period 3’s payoff is β^2 , etc.

¹²We use the prices of all prior periods’ prices to compute the average price and minimize the amount of noise and possible strategic manipulations subjects make.

¹³In the steady state monetary equilibrium, the continuation values for type A and type B agents at the end of period T are $V^A = \frac{\beta^T}{1-\beta}(u(x_A) - x_B)$ and $V^B = \frac{\beta^T}{1-\beta}(x_B - x_A)$.

experiments, to allow for a more accurate comparison across the two random termination treatments. Indeed, we use the same set of sequence lengths in the RT and BRT treatments to control for the effect of different sequence lengths on behavior (e.g., Fréchette and Yuksel 2017; Duffy and Puzzello 2014b). See Table 1 for a summary of the sequence lengths used for each session.

The functional forms and parameter values are set as in Jiang et al. (2020).¹⁴ The discount factor is set to $\beta = 0.9$; in the RT and BRT methods, this is simply the probability that the economy continues for an additional period. Additionally, in the BRT treatment we set the block length equal to 10 as this is also the expected duration associated with a 0.9 probability of continuation. In the DD implementation, we let the subjects play for $T = 10$ periods. We then compute the payoffs in periods 1 through 10, period 11 and the continuation value, as explained in Section 4.1. For example, in periods 1 through 10, the payoffs are weighted by the appropriate discounting term for each period: 1 in period 1, 0.9 in period 2, 0.9² in period 3, etc.

The period utility functions for type A and type B agents are respectively

$$U^A = \underbrace{A \frac{x_A^{1-\eta}}{1-\eta}}_{\text{market A}} - \underbrace{x_B}_{\text{market B}} \text{ and } U^B = \underbrace{-x_A}_{\text{market A}} + \underbrace{v_0 + x_B}_{\text{market B}};$$

where $A = 2.6563$, $\eta = 0.37851$, and $v_0 = 8$. With these parameters, the equilibrium value of market A (and market B) output is 10. The parameter v_0 is chosen to roughly equalize the equilibrium expected payoffs for types A and B subjects. The total number of subjects is $2N = 10$, with the exception of one session where $2N = 8$ because fewer subjects showed up for that session. We focus on output and welfare in market A because agents' utilities are linear in market B and market A is the main determinant of the welfare effects. We define the welfare ratio, denoted W , as a measure of efficiency relative to the first-best quantity of output in market A, $x_A^* = 13.2$:

$$W \equiv \frac{\sum_i [u(x_{A,i}) - x_{A,i}]}{N[u(x_A^*) - x_A^*]}. \quad (5)$$

Welfare is the sum of the trading surpluses related to the consumption of each individual type A subject. Given these parameters, the equilibrium predictions for output, prices, inflation, and welfare are, respectively, $x_A = x_B = 10$, $p_A = p_B = 0.5$, and $W = 0.98$.

¹⁴Jiang et al. (2020) compare different inflationary monetary policies in a monetary economy where the infinite horizon is implemented using a BRT.

To implement competitive pricing in markets A and B, subjects participate in a market game along the lines proposed by Shapley and Shubik (1977) (see e.g., Arifovic 1996; Bernasconi and Kirchkamp 2000; Ding and Puzzello 2020; Duffy et al. 2011; Duffy and Puzzello 2014ab; among others, for implementations in other experiments). In both markets, producers submit a quantity to produce (x_A or x_B) while consumers submit a bid of tokens for good A or good B (b_A or b_B). Subjects independently make these decisions in isolation and do not observe the current actions of other participants. The market price in each market is then computed as

$$p = \frac{\sum_i b_i}{\sum_i x_i} = \frac{\text{Total Tokens Bid}}{\text{Total Amount Produced}},$$

where b_i and x_i are the individual bids and production decisions of consumers and producers, respectively, for subject i . If the total amount of tokens bid or the total amount produced is zero, then no trade takes place. If the price is positive, then buyers consume an amount equal to their bid divided by the market price and their point total increases as specified by the utility function in each market while their token total decreases by the amount bid. Producers lose points from production, as specified by the production function, but their token total increases by the amount produced times the market price.

4.3 Experimental Procedures

The experiments in this study were conducted at Indiana University and Purdue University in 2018 and 2019 (see Table 1).¹⁵ Participants were undergraduate students at these universities.¹⁶ We conducted four sessions in each treatment with a total of 118 subjects. No subject participated in more than one session of the project, although some subjects may have participated previously in other economics experiments.

The total length of a session ranges from 100 to 120 minutes, though all subjects are recruited for two hours. Participants receive a \$5 show-up payment plus earnings from the experiment. Subjects are paid for all periods of all sequences.¹⁷ Average earnings across all

¹⁵The data for the BRT treatment are from the baseline treatment of Jiang et al. (2020). The main focus of Jiang et al. (2020) is on welfare implications of different implementations of inflationary monetary policies, while here we focus on the different implementations of the infinite horizon.

¹⁶The demographic composition of the subjects are very similar across both universities, except that Indiana University offers slightly more liberal arts majors than Purdue due to the presence of engineering majors at the latter. However, the results are not noticeably different across the two universities.

¹⁷Note that in the BRT, if the sequence ends within the block, then the subjects are only paid for the periods before the ending period even though they have committed for the entire block. To reinforce this to these subjects, we review this item several times in the instructions and include questions in the post-

subjects and treatments are \$42.10. Notice that sometimes there are more sequences in the RT treatment, compared with the same repetition in the BRT treatment, since we do not need to run a block of 10 periods at the start of each new sequence. In the DD treatment, we are always able to run four sequences of 10 periods each.

In the experiments, a period consists of market A followed by market B. Subjects are equally divided into fixed roles of type A and type B agents. The mapping of the production and consumption decisions to points is described in detail in the written instructions and presented to the subjects in table form in both the instructions and on their computer screens. Subjects can also see the previous periods' prices for both markets, which allows them to view prices over time. See Appendix D for a sample screenshot.

Table 1: Session Characteristics

Treatment	Session	Date	Subjects	Location	Sequences
Random Termination (RT)	RT1	9/5/2019	10	Purdue	9, 15
	RT2	9/12/2019	10	Purdue	6, 8, 2, 16
	RT3	9/5/2019	10	Indiana	13, 10, 5, 11
	RT4	9/5/2019	10	Indiana	5, 6, 4, 1, 10, 13
Block Random Termination (BRT)	BRT1	8/3/2018	8	Purdue	9, 15
	BRT2	8/24/2018	10	Indiana	6, 8, 2, 16
	BRT3	8/29/2018	10	Indiana	13, 10, 5
	BRT4	9/5/2018	10	Purdue	5, 6, 4
Definite Horizon with Discounting (DD)	DD1	10/1/2019	10	Purdue	10, 10, 10, 10
	DD2	10/3/2019	10	Purdue	10, 10, 10, 10
	DD3	10/10/2019	10	Indiana	10, 10, 10, 10
	DD4	10/14/2019	10	Indiana	10, 10, 10, 10

Each session consists of instructions, a quiz on the instructions, the experiment, and subject payment. Upon entering the laboratory, participants are randomly assigned a computer station and given a written copy of the instructions. Participants read the instructions and complete a comprehension quiz on them, after which the experimenter reviews each question individually and begins the experiment. We purposely spend a large portion of time on this phase of the experiment (typically 45 minutes to an hour) to ensure each subject's comprehension. All parts of the experiment are programmed using the software z-Tree (Fischbacher 2007). See the appendix for the experimental instructions and quiz.

instruction comprehension quiz that tests their understanding of the block .

5 Experimental Results

We now discuss our main results. In our analysis, we include periods where subjects actively play the game; that is, all rounds for RT, all rounds in blocks of 10 periods (i.e., even rounds where the sequence ends within the block and subjects therefore play for periods for which they are unpaid) and all rounds after the block for the BRT, and 10 periods for the DD (i.e., even if subjects are paid for each hypothetical period 11 and onward). We focus on three outcome variables: output, the price level, and the welfare ratio. In the first subsections, we compare the level of each variable to their theoretical steady state prediction and identify the treatment effect. We then analyze the evolution of these variables across time. Finally, in the last subsection, we summarize the main findings and compare the relative advantages of the three implementation methods.

5.1 Output

We first discuss the output in the experimental economies. Table 2 provides the summary statistics for the output per producer for each session that took place across all trading periods. Notice that the output per producer is the same as the consumption per consumer because we have the same number of producers and consumers. In the analysis below, we first compare the output to that predicted in the theoretical steady state and we summarize the results in Finding 1. We then discuss the treatment effects and summarize these in Finding 2.

Finding 1. *Market output deviates slightly from the steady state prediction in the RT and BRT treatments and is significantly lower than the steady state prediction in the DD treatment.*

As shown in Table 2, the output per producer is lower than its steady state prediction of 10 for most of the experimental sessions. The highest market A output is 12.711 in session BRT1, and the highest market B output is 12.949 in session RT4. The lowest outputs for markets A and B are both observed in session DD3, at 4.511 and 5.233, respectively. Table 3 estimates the deviation of the period's average output from its steady state prediction of 10. In the RT treatment, the observed output is close to the steady state prediction in both markets (the deviations are -0.508 in market A and -0.335 in market B). In the BRT treatment, the output in market A is only slightly below the steady state prediction (the deviation is -0.200 and is statistically insignificant at the 10% level); the deviation is more significant at -1.312 in market B. In the treatment for the DD, the output per producer

is significantly lower than that in the steady state prediction in both markets, in both an economic and statistical sense. The deviation is -2.314 for market A and -2.871 for market B, both of which are statistically significant at the 1% level.

Finding 2. *Comparing the RT and BRT treatments, market A's output per producer is not significantly different and market B's output per producer is lower in the BRT. Relative to the RT and BRT treatments, in both markets the output per producer is significantly lower in the DD.*

To identify the differences in market output across treatments, we regress the period market output per producer on two treatment dummies, with robust standard errors, for each market. The constant term from the regression can be interpreted as the estimated market output per producer in the RT treatment, while the coefficients on the two dummy treatment variables correspond to the marginal effect of the two other implementation schemes. The regression results are provided in Table 4.

Table 4 shows that market A's output is not significantly different between the RT and BRT treatments. The difference in market B's output is more significant, with output in the BRT 0.978 lower than that in the RT (statistically significant at the 5% level). Relative to the other two treatments, output is lower in the DD treatment. The output per producer in market A is 1.807 lower in the DD relative to the RT and 2.114 lower relative to the BRT during the period. In market B, the output per producer is 2.536 lower in the DD relative to the RT and 1.558 lower relative to the BRT; both differences are statistically significant at the 1% level.¹⁸

5.2 Price

Table 5 provides summary statistics for the prices in the two markets for each session. Table 6 provides a comparison of the price level with the theoretical steady state prediction, which we summarize in Finding 3. Table 7 presents the treatment effects, which we summarize in Finding 4.

Finding 3. *The price levels in both market A and market B are lower than the theoretical steady state prediction in treatments RT and BRT. Price levels are close to the steady state prediction in treatment DD.*

The theoretical steady state price level is $p_A = p_B = 0.5$. Table 5 shows that in the

¹⁸Note that the output is notably lower in the DD3 session relative to the other sessions. It appears that, in this session, the subjects coordinate on a lower level of output.

Table 2: Market Output per Producer-Summary Statistics

Session	Obs.	Market A		Market B	
		Mean	std.	Mean	std.
RT1	24	7.766	2.625	6.721	0.518
RT2	32	9.938	2.362	11.110	0.526
RT3	39	8.597	1.783	7.007	0.306
RT4	39	11.085	2.306	12.949	0.513
RTall	134	9.492	0.220	9.665	0.328
BRT1	25	12.711	2.529	13.036	0.328
BRT2	46	7.484	3.232	5.600	0.397
BRT3	33	11.594	3.033	9.745	0.405
BRT4	30	8.950	2.142	8.025	0.476
BRTall	134	9.800	0.303	8.688	0.301
DD1	40	9.426	2.100	9.097	0.428
DD2	40	8.783	1.912	6.192	0.271
DD3	40	4.511	1.167	5.233	0.400
DD4	40	8.023	2.804	7.996	0.516
DDall	160	7.686	0.222	7.129	0.237

Table 3: Market Output per Producer Relative to Steady State Prediction

	Obs.	Market A	Market B
RT	134	-0.508** (0.220)	-0.335 (0.328)
BRT	134	-0.200 (0.303)	-1.312*** (0.301)
DD	160	-2.314*** (0.222)	-2.871*** (0.237)

Notes. (1) Standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01.

RT and BRT treatments, the average price level (across all periods) is lower than 0.5 in all sessions except for session BRT2. In the RT treatment, the average price is 0.355 in both markets. In the BRT treatment, the average price is 0.359 in market A and 0.425 in market B. In both treatments, the deviation from the steady state prediction is statistically significant at the 1% level. For treatment DD, the average price is 0.469 in market A and 0.506 in market B; that is, it is closer to the theoretical predictions than in the other two treatments.

The lower price level in the RT and BRT relative to the theory can be attributed to the fact that buyers do not spend all of their money, which also implies that some money is

Table 4: Regression of Market Output per Producer on Treatment Dummies

Variables	Market A Output per Producer	Market B Output per Producer
BRT	0.307 (0.374)	−0.978** (0.446)
DD	−1.807*** (0.312)	−2.536*** (0.405)
Constant	9.492*** (0.220)	9.665*** (0.328)
Observations	428	428
R-squared	0.094	0.089

Notes. (1) Robust standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01.

held by sellers and thus is not spent.¹⁹ The price level in the DD is, on average, close to the theoretical prediction, largely because of the low output relative to the other two treatments, especially in the DD3 session.

In terms of across-treatment comparisons (see Table 7), the two treatments with random termination, RT and BRT, are close to each other, while the DD treatment has significantly higher prices than the other two. Between the RT and BRT, market A’s price is not significantly different in either magnitude or statistical significance. In market B, although the difference in the price between the RT and BRT is statistically significant at the 5% level, the magnitude of this difference is small (higher by 0.070 in the BRT). The difference between the DD and the other two treatments is more substantial. For example, relative to the RT, the price in market A is 0.114 lower and the price in market B is 0.151 lower in the DD.

5.3 Welfare

In this subsection, we focus on welfare in the experimental sessions. Table 8 provides summary statistics for each session. Similar to the analysis on output, we first evaluate whether the theoretical point prediction describes the experimental sessions well (Table 9 and Finding 5), and then we turn our attention to the differences between the three treatments (Table 11 and Finding 6).

Finding 5. *Welfare is significantly lower than the steady state prediction in all three treatments.*

¹⁹Tables C.1 and C.2 in the Appendix compare token spending ratios and the fraction of money held by buyers across treatments. The spending ratios range from 71% to 80% and the fractions of money held by buyers range from 77% to 86%. There are no significant differences in these two variables between treatments.

Table 5: Price-Summary Statistics

Session	Obs.	Market A		Market B	
		Mean	std.	Mean	std.
RT1	24	0.399	0.045	0.455	0.047
RT2	32	0.377	0.016	0.335	0.018
RT3	39	0.278	0.009	0.328	0.014
RT4	39	0.387	0.016	0.336	0.017
RTall	134	0.355	0.011	0.355	0.012
BRT1	25	0.278	0.012	0.266	0.012
BRT2	46	0.535	0.031	0.689	0.059
BRT3	33	0.244	0.018	0.271	0.017
BRT4	30	0.283	0.020	0.323	0.031
BRTall	134	0.359	0.017	0.425	0.027
DD1	40	0.403	0.018	0.472	0.055
DD2	40	0.314	0.013	0.454	0.029
DD3	40	0.705	0.032	0.639	0.038
DD4	40	0.455	0.035	0.458	0.037
DDall	160	0.469	0.017	0.506	0.021

Table 6: Price Relative to Steady State Prediction

	Obs.	Market A	Market B
RT	134	-0.145*** (0.011)	-0.145*** (0.328)
BRT	134	-0.141*** (0.017)	-0.075*** (0.027)
DD	160	-0.031* (0.017)	0.006 (0.021)

Notes. (1) Standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01

According to Table 8, for the average period, the welfare ratios in all 12 experimental economies are lower than the theoretical prediction of 0.98. The welfare ratio is the highest in session RT4 at 0.902 and the lowest in session DD3 at 0.678. More formally, Table 9 provides estimates for the deviations of the period's welfare ratios from the steady state prediction (0.98) for each treatment. Welfare is significantly lower in all treatments relative to its theoretical prediction. Recall that the deviation of the period's average output is not significantly different from the point prediction in the BRT treatment. The lower welfare relative to the steady state prediction is partly due to some consumption dispersion among consumers and across time.

To measure the extent of the consumption dispersion for the three treatments, we calcu-

Table 7: Regression of Market Price on Treatment Dummies

Variables	Market A Price	Market B Price
BRT	0.004 (0.020)	0.070** (0.030)
DD	0.114*** (0.021)	0.151*** (0.024)
Constant	0.355*** (0.011)	0.355*** (0.012)
Observations	428	428
R-squared	0.079	0.058

Notes. (1) Robust standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01.

late the coefficient of the variation in individual consumers' consumption in market A, which we report in Table 10. The results suggest that the dispersion is greatest in treatment DD, followed by the BRT and finally the RT.

Finding 6. *Welfare is significantly lower in the BRT and the DD relative to the RT.*

To identify the treatment differences, we regress the welfare ratio on the treatment dummies and report the results in Table 11. These results suggest that welfare is significantly lower in the BRT and DD treatments relative to the RT. The difference is 0.036 between the BRT and RT, 0.063 between the DD and the RT, and 0.027 between the BRT and the DD. The difference in the first two pairs (last pair) is significant at the 1% (5%) significance level.

Related to the remarks on consumption dispersion in Finding 5, the difference in the consumption dispersion also contributes to the differences in the welfare ratios across the different treatments. For example, the period's average output in the BRT is not significantly different from that in the RT. However, the difference in the welfare ratio is significant at the 1% level.

5.4 Dynamics

In previous subsections, we examined the level of output, the price levels and welfare in the experimental economies relative to their theoretical steady state predictions and discussed the differences between the three treatments. In this subsection, we investigate how the experimental economy evolves over time within and across sequences in a session.

First, we focus on the time trends in the output and prices and compare the patterns across the three implementation schemes. For this purpose, we regress the market output

Table 8: Welfare Ratio: Summary Statistics

Session	Obs.	Mean	std.
RT1	24	0.804	0.096
RT2	32	0.867	0.078
RT3	39	0.856	0.053
RT4	39	0.902	0.064
RTall	134	0.863	0.007
BRT1	25	0.898	0.040
BRT2	46	0.731	0.104
BRT3	33	0.887	0.065
BRT4	30	0.850	0.072
BRTall	134	0.827	0.009
DD1	40	0.849	0.063
DD2	40	0.837	0.068
DD3	40	0.678	0.083
DD4	40	0.836	0.081
DDall	160	0.800	0.008

Table 9: Welfare Ratio Relative to Steady State Prediction

	Obs.	Welfare Ratio -0.98
RT	134	-0.116^{***} (0.007)
BRT	134	-0.152^{***} (0.009)
DD	160	-0.179^{***} (0.008)

Notes. (1) Standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01.

per producer and the price for the period and its interaction terms with the two treatment dummies, the BRT and the DD. The coefficient for the period captures the time trend in the RT treatment, and the coefficients of the interaction terms capture the difference in the time trends between the BRT and the DD relative to the RT. For the price, we use the log term because the coefficient for the period can be interpreted as the inflation rate. In the regressions, we also include session-sequence dummies so that each sequence can have a different level (the trend is assumed to be the same within each treatment). The results for the output are shown in Table 12 and the results for the price level are shown in Table 13. We summarize the results in Findings 7 and 8.

Finding 7. *Average outputs for market A and market B exhibit mild negative trends in all three treatments. In market A, the difference in the time trends in output across the three*

Table 10: Individual Consumption in Market A: Summary Statistics by Treatment

Treatment	Obs.	Mean	std	CoV.
RT	670	9.492	6.197	0.652
BRT	645	9.687	6.763	0.698
DD	900	7.686	5.790	0.753

Table 11: Regression of Welfare Ratio on Treatment Dummies

Variables	Welfare Ratio
BRT	−0.036*** (0.011)
DD	−0.063*** (0.011)
Constant	0.863*** (0.007)
Observations	428
R-squared	0.068

Notes. (1) Robust standard errors in parentheses. (2) * p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01.

treatments is insignificant. In market B, the output exhibits a significantly steeper trend in the DD relative to the BRT.

Table 12 shows that all three treatments experience a mild downward trend in output per producer in markets A and B (see Figure E.1 in Appendix E for more details on the patterns for the output per producer over time for each individual session). Market A’s output decreases by 0.214 per period in the RT, slightly less (by 0.016) in the BRT and slightly more (by 0.025) in the DD. Market B’s output decreases by 0.141 per period in the RT treatment and slightly less (by 0.038) in the BRT. In the DD, market B’s output exhibits a noticeably steeper downward trend: the slope over the period is 0.133 lower than for the RT (although the difference is not statistically significant at the 10% level). The difference in the trend in market B’s output between the BRT and the DD is 0.171 and is statistically significant at the 5% level.

Finding 8. *Prices in markets A and B exhibit mild positive trends or inflation in all three treatments. The differences in market A’s inflation across the three treatments are insignificant. Market B’s inflation rate is significantly higher in the DD relative to the RT.*

Table 13 suggests that all three treatments experience mild upward trends in prices, or positive inflation in markets A and B (the coefficient on the time period captures the growth rate of the price level and, hence, is an estimate of the inflation rate). In the RT treatment, the inflation rate is 2.3% in market A and 2.7% in market B. In the BRT, the inflation rate is

2.6% in market A and 4.5% in market B. In the DD, the inflation rate is 3.5% in market A and 6.1% in market B.²⁰ In addition, there are no significant differences in the estimated inflation rates across the three treatments. The market A inflation is not significantly different across the three treatments. In market B, the DD exhibits higher inflation than the RT by 3.4% and the difference between the DD and the BRT is insignificant at the 10% level.

Table 12: Time Trend in Market Output

Variables	Market A Output per Producer	Market B Output per Producer
Period	−0.214*** (0.056)	−0.141** (0.060)
Period*BRT	0.016 (0.087)	0.038 (0.079)
Period*DD	−0.025 (0.079)	−0.133 (0.092)
Observations	428	428
R-squared	0.619	0.667

Notes. (1) Robust standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01. (3) The regression also includes a constant and session-sequence dummies, and their coefficients are omitted.

Table 13: Time Trend in Prices with Session-sequence dummies

Variables	log (Market A Price)	log (Market B Price)
Period	0.023*** (0.008)	0.027*** (0.007)
Period*BRT	0.003 (0.012)	0.018 (0.012)
Period*DD	0.011 (0.012)	0.034** (0.013)
hline Observations	428	428
R-squared	0.633	0.144

Notes. (1) Robust standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01. (3) The regression also includes a constant and session-sequence dummies and their coefficients are omitted.

From the above analysis, the output and the price level each exhibit a mild trend within a single sequence in each of the three treatments. We interpret this as evidence that our experimental economies are broadly stationary and dynamic incentives are well-preserved in all three implementation schemes. In what follows, we conduct additional analyses on dynamic incentives to check whether subjects behave differently in the last sequence relative

²⁰In their market A, Duffy and Puzzello (2020) also observe a mild inflation rate in the constant money supply treatment. In their setting, prices in market A are formed via generalized Nash bargaining.

to earlier sequences and in period 10 relative to other periods. For each treatment and for each variable (price, output and spending ratio), we run three regressions, one for all of the sequences, one for the next-to-last sequence and one for the last sequence. The across-sequences comparison helps us identify the constraint that is imposed by the recruitment time, which potentially affects the last sequence of the RT and BRT treatments. It also captures the learning effect across sequences. In each regression, besides the time period, we introduce a dummy variable, *period10*, which is equal to 1 for period 10, to identify whether the subjects behave differently during this period. For the DD treatment, period 10 is the last trading period before continuation payoffs are computed. For the BRT, the realizations of the random draws that determine the actual duration of the sequences are revealed in period 10. For the RT, period 10 could play a special role in that it is the expected duration of a typical sequence. The results are presented in Tables D.1 to D.3 in Appendix D.

In the RT treatment, inflation tends to be significant only in the last sequence, which can be attributed to a decreasing output and an increasing spending ratio. A possible explanation is that as the end of the recruitment time gets closer, subjects expect the sequence to end soon and respond by reducing their output and increasing their spending.

Relative to the RT, the BRT shares the similarity that the actual experimental economy lasts for an indefinite number of periods and the recruitment time constraint may be perceived as more likely to bind for the last sequence. Related to this is the feature the BRT shares with the RT, which is a more notable downward trend in output A in the last sequence compared with earlier sequences.²¹

On the other hand, the block design in the BRT may introduce behavioral changes in period 10 when the random realizations in the first 10 periods are revealed and the sequence may end at the end of period 10. From the coefficient on the dummy variable *Period10*, we see that the subjects in the BRT tend to increase their spending and decrease their production in period 10, while the same tendency is not observed in the RT.

Another difference is that subjects play for at least 10 periods in each sequence in the BRT and this stability in the length of play may facilitate learning within and across sequences. The increasing spending ratio within the next-to-last sequence could suggest within-sequence learning. The stabilization of the spending ratio and inflation in market A and the output

²¹We took steps to minimize the end-game effects in the RT and BRT treatments. For example, we recruited subjects for a longer time than the duration of play, only started a sequence if there was ample time to finish it, and did not announce the number of sequences ahead of time. However, even though the subjects do not know which sequence is the last one, it is possible that their perceptions are affected by the implementation method used.

in market B in the last sequence in the BRT may suggest across-sequence learning.²²

In the DD, there seems to be a slight end-game effect: in period 10 relative to previous periods, the output tends to drop in market B and the spending ratio tends to increase in market A, which translates into positive inflation. The end-game effect becomes weaker in the last sequence relative to the earlier sequences, suggesting that some learning is happening. As for the time trend, compared with the other two treatments, the DD exhibits a more persistent decreasing trend in output, which could be attributed to salient discounting in this treatment (as opposed to implicit discounting through the continuation probability in treatments RT and BRT). To investigate this further, in the appendix, we analyze the impact of period 10 on inflation and output (see Tables D1-D3) and find that while output drops slightly in period 10 in the BRT and DD, there are still positive output levels in period 10 in these two treatments; for instance, in period 10, market A's average output per producer is 8.3 (versus 8.8 in period 9 and 10.1 in periods 1-9) in the BRT, and 6.7 (versus 7.3 in period 9 and 7.8 in periods 1-9) in the DD.²³

Our analysis here suggests that dynamic incentives are preserved in all three implementations. Output and price exhibit only mild trends within a sequence, suggesting that the experimental economy is broadly stationary. Even in the last sequence of the RT and BRT, where the recruitment time constraint could be more binding, the negative trend in output and the positive trend in inflation are still mild. This contrasts with the finite-economy setup in McCabe (1989) and Davis et al. (2020) where, when the theory predicts that money should not be valued, the value of money sharply decreases as the subjects gain experience in observing the end of a horizon. In terms of treatment differences between the three implementation schemes, the recruitment time constraint seems to affect the last sequence relative to earlier sequences in the RT and BRT. The BRT and DD introduce some end-block and end-game effects in period 10. On the other hand, there is evidence that the more stable environment in these two treatments promotes learning both within and across sequences.

²²The increasing spending ratio in the last sequence of the RT more likely intensifies spending as the recruitment time is closer to its end. In the RT, the spending ratio in the earlier sequences does not show an upward trend (in fact, it shows an insignificant negative trend), which suggests insignificant learning within sequences in the earlier sequences. By contrast, the spending ratio in the BRT increases within a sequence in the earlier sequences, suggesting within-sequence learning, while it starts to stabilize in the last sequence, suggesting across-sequence learning.

²³For comparison, in the RT, market A's average output in period 10 is 8.1 versus 8.3 in period 9 and 9.8 in periods 1-9.

5.5 Discussion

To summarize the main results, we find that monetary trade is sustained and the dynamic incentives underlying the theoretical model are broadly preserved in all three implementation schemes. The experimental outcomes fall within a reasonable neighbourhood of the theoretical prediction. Output is slightly lower than the theoretical prediction in the RT and BRT treatments (and are not significantly different between the two) but are significantly lower in the DD. Welfare is lower than the theoretical prediction in all treatments, with the RT being the closest to the theory, followed by the BRT and then the DD.

From these findings, can we conclude which method is superior? For the implementation of infinite-horizon models, an important criterion is the preservation of the dynamic incentives underlying the theoretical model. Other more practical criteria include the ease of implementation in the lab and the method's ability to allow for relatively long sessions or supergames.

Within the context of our experimental monetary economy with a constant money supply, all three methods are able to preserve the dynamic incentives that sustain monetary trade. The standard random termination method generates behavior that is closest to the theoretical predictions. The relative superiority, however, is only marginal relative to the BRT and not overwhelming relative to the DD. Given the broad similarity of the results among the three methods, one may choose any of these approaches, depending on the research questions of interest. In the following, we discuss the relative advantages of each method.

As for more practical concerns, the RT is relatively straightforward to implement compared with the other two schemes. For instance, one needs to explain additional concepts in the BRT (i.e., the block) relative to the RT (i.e., only the sequence termination scheme). Indeed in our case, we were able to run more sequences in the RT treatment relative to the BRT. For the DD, one also needs to explain the concept of the continuation value and assess whether the subjects understand the automated calculations. Another advantage of the RT is that the stake of the decisions is constant over time, while it decreases with the other two; in the DD, the payoffs are discounted over time, while in the BRT, the probability that a later period is counted for the payment is smaller within the block. If it is mentally challenging to make the right decisions, then the subjects may have weaker incentives to do their best in later periods. This concern, however, may be mitigated by considering that the task tends to become easier as the subjects learn and gain experience as the sequence proceeds. A related point is that although in theory the probability of continuation is equivalent to the discount factor, the explicit discounting in the DD may look more salient than

the implicit discounting in the RT and BRT. The relatively lower output in the DD could also be attributed to this consideration.

The BRT design ensures that the experimenter observes behavior in a minimum number of periods in each supergame. This could be a useful feature for certain questions. For example, the literature on price bubbles in experimental asset markets has shown that it takes time for a bubble to occur. The standard random termination often generates short sequences that are not informative for studies of this issue. Another example where the BRT could be useful is an inflationary monetary economy, where observing the path of past price changes would help subjects infer the inflation rate. The stability of the environment in the BRT relative to the RT helps subjects learn within and across sequences.

There are also benefits associated with the DD method. First, it is not subject to the credibility criticism placed on approaches that are based on random terminations since, by design, dynamic incentives are ensured by the value of continuation. From the discussion in the previous subsection, there is some evidence that the recruitment time constraint affects subjects' behavior in the last sequences of the RT and the BRT. Although there is some evidence that there are end-game effects in the DD with the continuation value in place, the effect weakens over time as subjects learn across sequences (and the stability of the DD treatment tends to promote learning). Second, the fixed rounds of decision making in the DD facilitates comparisons between different sequences, sessions and treatments. Third, if the experimenter needs to collect data on multiple sequences for settings with high discount factors, then the DD method may be desirable as one can control the length of the fixed horizon. Some research questions need to use a within-subjects design that runs sequences that feature different treatments in the same sessions. For example, to study the adoption and usage of a new currency or a means of payment, one may need to run at least three sequences: the first with only the existing currency, the second with a new currency alongside the existing currency, and then a third repeating the second to investigate whether the first-mover advantage will weaken or disappear over time.

Finally, in this paper we implement the DD method for a monetary economy with competitive markets and can conveniently use past price levels for the calculation of the continuation value. We believe this method can also be adapted to the implementation of infinitely repeated PD games by incorporating strategy elicitation methods in the design. The experimenter, independently of the discount factor used in the experiment, can choose the number of periods, say T , during which subjects construct strategies. The period payoffs are discounted appropriately in each period up to period T . The question then is how to compute the continuation value of the game after T ; there are different options researchers

can use, depending on the strategy elicitation method one chooses to adopt; for example, Bruttel and Kamecke (2012); Dal Bó and Fréchette (2019); Romero and Rosokha (2018); or Romero and Rosokha (2019). The common feature of these elicitation methods is that subjects are allowed to construct strategies they wish to play. However, there are also important differences. For example, Bruttel and Kamecke (2012); Dal Bó and Fréchette (2019); and Romero and Rosokha (2019) allow subjects to construct memory-1 strategies, while Romero and Rosokha (2018) allow subjects to construct longer memory strategies. Regardless of the chosen method, strategies constructed in period T can then be used to determine the continuation value of a game and, thus, assign the discounted sum of payoffs that is implied by the strategy chosen. We leave for future research a more exhaustive study of how different elicitation methods affect the computation of continuation values and, in turn, possibly affect cooperation rates in repeated PD games.

6 Conclusions

The implementation of infinite-horizon models often relies on random termination which interprets the discount factor as the probability that a game continues to the next period. This approach is subject to the criticism that it is challenging for the experimenter to credibly implement a game that lasts an arbitrarily long time, given the limited length of an experimental session.

In this paper, we propose a new implementation scheme to address such credibility issues. In our proposed scheme, subjects play a game for a fixed number of periods and the discount factor is interpreted as a weighting factor that shrinks payoffs over time. Dynamic incentives are preserved by paying subjects a continuation value that is based on prior outcomes. We apply our proposed scheme in an infinite-horizon monetary economy and compare outcomes with two commonly used implementations that are based on random termination.

We find that dynamic incentives are preserved and behavior is similar in all three implementations. Given the broad similarity of the results among the three methods, one may choose any of the approaches, depending on the research question of interest and the more practical criteria, such as the ease of implementation in the lab and the ability of the method to allow for relatively more or longer supergames. For instance, the random termination treatment is relatively easy to explain to subjects and implement in the lab, while the block random treatment ensures that each supergame has a minimum number of periods, which makes it particularly useful for research questions where it is important to have a sufficiently large number of periods per sequence or supergame. The benefit of the

DD is that one can control the length of the fixed horizon, which makes it easier to conduct within-subjects treatments in the same session when the underlying game involves a high discount factor.

Our study contributes to the experimental literature by comparing different infinite-horizon implementation methods in the context of a monetary economy. Our DD approach can also be modified to enrich the random termination approach that is commonly adopted in the laboratory. For example, in the event the economy is still going on and the recruitment time has come to an end, the experimenter could invoke an interpretation of β as the discount factor, compute the continuation value and pay subjects accordingly. This adaptation blends the random termination and discounting interpretations. Our new method can also be adapted to implement infinitely repeated prisoner dilemma games. We plan to pursue this implementation in future research. Another avenue for future research is to consider an implementation along the lines proposed by Lim, Prescott, and Sunder (1986) and Marimon and Sunder (1993), where additional subjects are recruited to elicit forecasts for future prices.

References

Aliprantis, C., G. Camera and D. Puzzello (2007). “Contagion Equilibria in a Monetary Model.” *Econometrica*, 75 (1), 277-282.

Andersson, O. and E. Wengstrom (2012). “Credible communication and cooperation: Experimental Evidence from Multi-Stage Games.” *Journal of Economic Behavior and Organization*, 81(1), 207–219.

Araujo, L., B. Camargo, R. Minetti and D. Puzzello (2012) “Essentiality of Money with Centralized Trade.” *Journal of Monetary Economics*, 59, 612-621.

Arifovic, J. (1996). “The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies.” *Journal of Political Economy*, 104, 510–541.

Bernasconi, M. and O. Kirchkamp (2000). “Why do monetary policies matter? An experimental study of saving and inflation in an overlapping generations model.” *Journal of Monetary Economics*, 46, 315–343.

Bewley, T. (1980). “The Optimum Quantity of Money.” In *Models of Monetary Economies*, ed. by J. H. Kareken and Neil Wallace, 169–210.

Bruttel, L., Kamecke, U. (2012). “Infinity in the Lab. How Do People Play Repeated Games?.” *Theory and Decision*, 72, 205–219.

Camera, G. and M. Casari (2014). “The Coordination Value of Monetary Exchange: Experimental Evidence.” *American Economic Journal: Microeconomics*, 6, 629–660.

Cooper, D. and K. Kuhn (2014). “Communication, Renegotiation, and the Scope for Collusion.” *American Economic Journal: Microeconomics*, 6(2), 247–278.

Dal Bó, P. and G. Fréchette (2018). “On the Determinants of Cooperation in Infinitely Repeated Games: A survey.” *Journal of Economic Literature*, 56 (1), 60–114.

Dal Bó, P. and G. Fréchette (2019). “Strategy Choice in the Infinitely Repeated Prisoner’s Dilemma.” *American Economic Review*, 109, 3929–52.

Davis, D., O. Korenok, P. Norman, B. Sultanum and R. Wright (2020). “Playing with Money.” *Journal of Economic Behavior & Organization*. Available at <https://doi.org/10.1016/j.jebo.2020.06.031>.

Ding, S. and D. Puzzello (2020). “Legal Restrictions and International Currencies: An Experimental Approach.” *Journal of International Economics*, 126, 103342.

Duffy, J., J. Jiang and H. Xie (2019). “Experimental Asset Markets with An Indefinite Horizon.” Working Paper. Available at SSRN: <https://ssrn.com/abstract=3420184>.

Duffy, J., A. Matros, and T. Temzelides (2011). “Competitive Behavior in Market Games: Evidence and Theory.” *Journal of Economic Theory*, 146, 1437–1463.

Duffy, J. and J. Ochs (1999). “Emergence of Money as a Medium of Exchange: An Experimental Study.” *American Economic Review*, 89(4): 847–877.

Duffy, J. and J. Ochs (2002). “Intrinsically Worthless Objects as Media of Exchange: Experimental Evidence.” *International Economic Review*, 43, 637–673.

Duffy, J. and D. Puzzello (2014a). “Gift Exchange versus Monetary Exchange: Experimental Evidence.” *American Economic Review*, 104(6), 1735–1776.

Duffy, J. and D. Puzzello (2014b). “Experimental Evidence on the Essentiality and Neutrality of Money in a Search Model.” *Experiments in Macroeconomics, Research in Experimental Economics*, 17, Emerald.

Duffy, J. and D. Puzzello (2020). “The Friedman Rule: Experimental Evidence,” *Mimeo*.

Fischbacher, U. (2007). “Z-Tree: Toolbox for Ready-Made Economic Experiments.” *Experimental Economics*, 10, 171–178.

Fréchette, G. and S. Yuksel (2017). “Infinitely repeated games in the laboratory: four perspectives on discounting and random termination.” *Experimental Economics*, 20, 279–308.

Jiang, J., D. Puzzello, and C. Zhang (2020). “Inflation and Welfare in the Laboratory.”

Mimeo.

Jiang, J. and C. Zhang (2018). “Competing currencies in the laboratory.” *Journal of Economic Behavior and Organization*, 154, 253–280.

Kamiya, K., H. Kobayashi, T. Shichijo, and T. Shimizu (2019). “Efficiency of Monetary Exchange with Divisible fiat Money: An experimental Approach,” *Mimeo.*

Lagos, R., G. Rocheteau, and R. Wright (2017). “Liquidity: A New Monetarist Perspective.” *Journal of Economic Literature*, 55(2), 371–440.

Lagos, R. and R. Wright (2005). “A Unified Framework for Monetary Theory and Policy Analysis.” *Journal of Political Economy*, 113(3), 463–486.

Lim, S., E. Prescott, and S. Sunder (1994). “Stationary Solution to the Overlapping Generations Model of Fiat Money: Experimental Evidence,” *Empirical Economics*, 19, 255–277.

Mailath, G. and L. Samuelson (2006). *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

Marimon, R. and S. Sunder (1993). “Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence.” *Econometrica*, 61(5), 1073–1107.

Matsuyama, K., N. Kiyotaki, and A. Matsui (1993). “Toward a Theory of International Currency.” *Review of Economic Studies*, 60, 283–307.

McCabe, K. (1989). “Fiat Money as a Store of Value in an Experimental Market.” *Journal of Economic Behavior and Organization*, 12, 215–231.

Rietz, J. (2019). “Secondary Currency Acceptance: Experimental Evidence with a Dual Currency Search Model.” *Journal of Economic Behavior and Organization*, 166, 403–431.

Rocheteau, Guillaume and Ed Nosal (2017). *Money, Payments and Liquidity*, MIT Press.

Rocheteau, G. and R. Wright (2005). “Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium.” *Econometrica* 73, 175–202.

Romero, J. and Y. Rosokha (2018). “Constructing Strategies in the Indefinitely Repeated Prisoner’s Dilemma Game.” *European Economic Review*, 104, 185–219.

Romero, J. and Y. Rosokha (2019). “Mixed Strategies in the Indefinitely Repeated Prisoner’s Dilemma.” *Mimeo.*

Roth, A. and K. Murnighan (1978). “Equilibrium Behavior and Repeated Play of the Prisoner’s Dilemma.” *Journal of Mathematical Psychology*, 17(2), 189–198.

Sabater-Grande, G. and N. Georgantzis (2002). “Accounting for Risk Aversion in Re-

peated prisoners' Dilemma Games: An experimental test." *Journal of Economic Behavior & Organization*, 48(1), 37–50.

Selten, R., M. Mitzkewitz, and G. Uhlich (1997). "Duopoly Strategies Programmed by Experienced Players." *Econometrica*, 65 (3), 517-555.

Shapley, L., and M. Shubik (1977). "Trade Using One Commodity as a Means of Payment." *Journal of Political Economy*, 85, 937-968.

Townsend, Robert (1980). "Models of Money with Spatially Separated Agents." In *Models of Monetary Economies*, ed. by J. H. Kareken and Neil Wallace, 265–304.

Williamson, S. and R. Wright (2011). "New Monetarist Economics: Models." *Handbook of Monetary Economics*, vol. 3A, B. Friedman and M. Woodford, eds., Elsevier, 25-96.

Wilson, A. and H. Wu (2017). "At-will relationships: How an option to walk away affects cooperation and efficiency." *Games and Economic Behavior*, 102, 487-507.

A Steady State Monetary Equilibrium

We focus on stationary equilibria where real variables are constant over time. Because the money supply is constant, the price level is also constant over time. In the analysis below, we omit the time subscript and use “ \wedge ” to denote the variables used in the next period. As is standard, we start backwards by characterizing agents’ decision problems in market B and combine that with agents’ choices in market A to characterize the equilibrium.

Market B Decision Problems

In market B, agents can trade good B and spend money in a competitive market where the price of good B is p_B . The value function of a type i agent who enters market B with money holdings m^i is

$$\begin{aligned} \max_{\hat{m}^i, x_B^i} W^i(m^i) &= \max_{\hat{m}^i, x_B^i} \{-x_B^i + v^i + \beta V^i(\hat{m}^i)\} \\ \text{subject to } \hat{m}^i &= p_B x_B + m^i, \end{aligned}$$

where $v^A = 0, v^B = v_0$ and \hat{m}^i is the choice of money holdings in the next market A. Substituting the budget constraint into the objective, the value function simplifies to

$$\max_{\hat{m}^i} W^i(m^i) = \max_{\hat{m}^i} \left\{ \frac{-\hat{m}^i + m^i}{p_B} + \beta V^i(\hat{m}^i) \right\}.$$

The optimal choice for \hat{m}^i solves

$$\beta \frac{\partial V^i(\hat{m}^i)}{\partial \hat{m}^i} - \frac{1}{p_B} \leq 0, \text{ with equality if } \hat{m}^i > 0.$$

As is usual, the value function $W^i(m)$ is linear in m and the choice of money holdings in the next period, \hat{m} , is independent of the current money holdings. The envelope result for both types of agents is

$$\frac{\partial W^i(m)}{\partial m} = \frac{1}{p_B}.$$

Market A Decision Problems

Agents in market A can trade good A and spend money in a competitive market at market price p_A . Type B agents, who are producers in market A, incur a linear production cost to

produce x_A units of good A. Their decision problem is

$$V^B(m) = \max_{x_A} \left\{ -x_A + \frac{(m + x_A p_A)}{p_B} + W^B(0) \right\}.$$

Notice that we used the envelope result $\frac{\partial W^i(m)}{\partial m} = \frac{1}{p_B}$. The first-order condition of type B's problem implies

$$p_A = p_B. \tag{A.1}$$

The envelope result is

$$\frac{\partial V^B(m)}{\partial \hat{m}} = \frac{1}{p_B}.$$

Type A agents, who are consumers in market A, can buy and consume x_A units of good A. Their value function in market A is

$$\begin{aligned} V^A(m) &= \max_{x_A} \left\{ u(x_A) + \frac{(m - p_A x_A)}{p_B} + W^A(0) \right\} \\ \text{subject to } p_A x_A &\leq m. \end{aligned}$$

If the cash constraint does not bind, then $u'(x_A) = p_A/p_B$, which when combined with type B's decision gives $x_A = x_A^*$. If the cash constraint binds, then $x_A = m/p_A$. In either case, the envelope condition is

$$\frac{\partial V^A(m)}{\partial m} = \frac{u'(x_A)}{p_A}.$$

Equilibrium

We now combine agents' decision problems from markets A and B to derive the equations that characterize the stationary monetary equilibrium. For type B agents, the net marginal value of carrying money to the next market A is

$$\begin{aligned} -\frac{1}{p_B} + \beta \frac{\partial V^B(\hat{m})}{\partial \hat{m}} &= -\frac{1}{p_B} + \beta \frac{1}{\hat{p}_B} \\ &= \frac{1}{p_B} [-1 + \beta] < 0. \end{aligned}$$

Expressed in words, the money carried by type B agents to market A will be idle in market A and can be used to purchase good B in the next market B. Given $\beta < 1$, holding idle balances is costly. As a result, type B agents will spend all their money in market B and enter market A with zero balances.

For type A agents, the net marginal value of carrying money to the next market A is

$$\begin{aligned} -\frac{1}{p_B} + \beta \frac{\partial V^A(\hat{m})}{\partial \hat{m}} &= -\frac{1}{p_B} + \beta \frac{u'(x_A)}{\hat{p}_A} = -\frac{1}{p_B} + \beta \frac{u'(x_A)}{\hat{p}_B} \\ &= \frac{1}{p_B} [-1 + \beta u'(x_A)]. \end{aligned}$$

Under the assumption $u'(0) = \infty$, type A agents hold a positive amount of money in market A and, in equilibrium, the net marginal benefit of carrying money is zero and the output in market A (per consumer or producer) is

$$u'(x_A) = 1/\beta. \quad (\text{A.2})$$

Given $\beta < 1$, agents will carry just enough money to spend in market A and the cash constraint will bind in that market.

In market A, each type A agent consumes x_A and each type B agent produces x_A , where x_A solves equation (A.2). The price level, which is the same for both markets according to equation (A.1), is given by the money market clearing condition

$$p_A = p_B = \frac{M}{Nx_A}. \quad (\text{A.3})$$

In market B, the amount consumed by each type B agent (which is the same as the amount of production by each type A agent) is

$$x_B = \frac{M}{Np_B} = \frac{M}{Np_A} = x_A. \quad (\text{A.4})$$

B Equilibrium for Treatment DD

In this Appendix, we check and verify that both the stationary monetary and nonmonetary equilibria in the infinite-horizon model are an equilibrium for the economy that describes the DD treatment. Going backwards, we first describe the continuation value for the agents at the end of period T , followed by the choice problem in stages A and B in period T , and then the choice problem in each stage in the first $T - 1$ periods. Notice that only the last decision stage (stage B in period T) is different from the standard Lagos and Wright (2005) set up, due to the continuation value.

Continuation Value at End of Period T

Let \bar{p}_{AT} (\bar{p}_{BT}) denote the average price of good A (B) across the first T periods, $x_{At} = \frac{1}{N} \sum_{i=1}^N x_{Ait}$ ($x_{Bt} = \frac{1}{N} \sum_{i=1}^N x_{Bit}$) be the average consumption of good A (B) per consumer in period t , and $\bar{x}_{AT+1} = \frac{1}{T+1} \sum_{t=1}^{T+1} x_{At}$ ($\bar{x}_{BT+1} = \frac{1}{T+1} \sum_{t=1}^{T+1} x_{Bt}$) be the average per capita consumption of good A (B) in the first $T + 1$ periods.

The continuation value of a type A agent holding m_{T+1} units of tokens at the end of period T is

$$\begin{aligned} V^A(m_{T+1}) = & \beta^T \left\{ u \left(\frac{m_{T+1}}{\bar{p}_{AT}} \right) - \frac{M}{N\bar{p}_{BT}} \right\} \\ & + \frac{\beta^{T+1}}{1-\beta} [u(\bar{x}_{AT+1}) - \bar{x}_{BT+1}]. \end{aligned}$$

Expressed in words, type A agents' continuation values have two components. The first line in the above equation is the utility in period $T + 1$: type A agents spend all of their money holdings in market A at the average past price \bar{p}_{AT} ; in market B, each type A agent produces an equal share of good B for type B agents (note that in market B, type B agents hold all of the tokens and spend everything).

The continuation value of a type B agent holding m_{T+1} units of tokens at the end of period T is

$$\begin{aligned} V^B(m_{T+1}) = & \beta^T \left\{ -\frac{1}{N\bar{p}_{AT}} \sum_{i=1}^N m_{AT+1}^i + \frac{1}{\bar{p}_{BT}} \left[m_{T+1} + \frac{1}{N} \sum_{i=1}^N m_{AT+1}^i \right] + v_0 \right\} \\ & + \frac{\beta^{T+1}}{1-\beta} [-\bar{x}_{AT+1} + \bar{x}_{BT+1} + v_0]. \end{aligned}$$

Expressed in words, a type B agent in period $T + 1$ produces an equal share of good A for type A agents in market A; they take this cash revenue together with their own money holdings to market B and spend everything.

From period $T + 2$ onward, assume that all agents consume and produce the average amounts consumed and spent in the past $T + 1$ periods. Notice that the term m_{T+1} only enters into the $T + 1$ part of the continuation value.

Stage B in Period T

This is the last decision stage, where agents decide on their money holdings. Type A agents solve

$$\begin{aligned} W_T(m_T) = & \max_{m_{T+1}} \beta^{T-1} \frac{m_T - m_{T+1}}{p_{BT}} \\ & + \beta^T \left\{ u \left(\frac{m_{T+1}}{\bar{p}_{AT}} \right) - \frac{M}{N \bar{p}_{BT}} \right\} \\ & + \frac{\beta^{T+1}}{1 - \beta} [u(\bar{x}_{AT+1}) - \bar{x}_{BT+1}]. \end{aligned}$$

The first line is the disutility from working for tokens; the second and third lines denote the continuation value described earlier. The first-order condition is

$$\frac{1}{p_{BT}} = \beta u' \left(\frac{m_{T+1}}{\bar{p}_{AT}} \right) \frac{1}{\bar{p}_{AT}}.$$

For type B agents (who consume in stage B),

$$\begin{aligned} W_T(m_T) = & \max_{m_{T+1}} \beta^{T-1} \frac{m_T - m_{T+1}}{p_{BT}} \\ & + \beta^T \left\{ -\frac{1}{N \bar{p}_{AT}} \sum_{i=1}^N m_{AT+1}^i + \frac{1}{\bar{p}_{BT}} \left[m_{T+1} + \frac{1}{N} \sum_{i=1}^N m_{AT+1}^i \right] + v_0 \right\} \\ & + \frac{\beta^{T+1}}{1 - \beta} [-\bar{x}_{AT+1} + \bar{x}_{BT+1} + v_0]. \end{aligned}$$

The first-order condition is

$$\frac{1}{p_{BT}} \geq \beta \frac{1}{\bar{p}_{BT}}.$$

For both types, we have

$$W'_T(m_T) = \beta^{T-1} \frac{1}{p_{BT}}.$$

Stage A in Period T

In period T , type B agents solve

$$V_T(\hat{m}) = \max_{x_{AT}} -\beta^{T-1}x_{AT} + W_T(\hat{m} + x_{AT}p_{AT}).$$

The first-order condition is

$$\beta^{T-1} = p_{AT}\beta^{T-1}\frac{1}{p_{BT}} \rightarrow p_{AT} = p_{BT} = p_T.$$

Type A agents solve

$$\begin{aligned} V_T(\hat{m}) &= \max_m \beta^{T-1}u(x_{AT}) + W_T(\hat{m} - x_{AT}p_{AT}) \\ \text{s.t. } x_{AT} &\leq \hat{m}/p_{AT}. \end{aligned}$$

If the cash constraint binds, then $x_T = \hat{m}/p_{AT}$ and

$$V'_T(\hat{m}) = \beta^{T-1}u'(x_{AT})/p_{AT}.$$

If the cash constraint does not bind, then $u'(x_{AT}) = \frac{p_{AT}}{p_{BT}} = 1$.

Stage B in Periods Before T

In the stage B in the periods before T , agents solve

$$W_t(m_t) = \max_{m_{t+1}} \beta^{t-1} \frac{m_t - m_{t+1}}{p_{Bt}} + V_{t+1}(m_{t+1}),$$

which implies

$$\begin{aligned} \beta^{t-1} \frac{1}{p_{Bt}} &\geq V'_{t+1}(m_{t+1}), \\ W'_t(m_t) &= \beta^{t-1} \frac{1}{p_{Bt}}. \end{aligned}$$

Stage A in Periods Before T

In the stage A of the periods before T , type B agents (who are producers of good A) solve

$$V_t(m_t) = \max_{x_{At}} -\beta^{t-1}x_{At} + W_t(m_t + x_{At}p_{At}),$$

which implies

$$\begin{aligned}\beta^{t-1} &= p_{At}\beta^{t-1}\frac{1}{p_{Bt}} \rightarrow p_{At} = p_{Bt} = p_t, \\ V'_t(m_t) &= W'_t(m_t + x_{At}p_{At}).\end{aligned}$$

Type A agents solve

$$\begin{aligned}V_t(m_t) &= \max_{q_{At}} \beta^{t-1}u(x_{At}) + W_t(m_t - x_{At}p_{At}) \\ s.t. \ x_{At} &\leq m_t/p_{At}.\end{aligned}$$

If the cash constraint binds, then $x_{At} = m_t/p_{At}$ and

$$V'_t(m_t) = \beta^{t-1}u'(x_{At})/p_{At}.$$

If the cash constraint does not bind, then $u'(x_{At}) = 1$.

Equilibrium

We can summarize the equilibrium as follows. Define $p_t \equiv p_{At} = p_{Bt}$. For type A agents before period T , the equilibrium condition is $\beta^{t-1}\frac{1}{p_t} = V'_{t+1}(m_{t+1}) = \beta^t u'(x_{At+1})/p_{t+1}$, or

$$\frac{1}{p_t} = \frac{\beta u'(x_{At+1})}{p_{t+1}}.$$

For type B agents before period T , the condition is $\beta^{t-1}\frac{1}{p_t} \geq V'_{t+1}(m_{t+1}) = \beta^t \frac{1}{p_{t+1}}$, or

$$\frac{p_{t+1}}{p_t} \geq \beta.$$

For type A agents in period T

$$\frac{1}{p_T} = \beta u'(x_{AT+1}) \frac{1}{\bar{p}_T}.$$

For type B agents in period T

$$\frac{1}{p_T} \geq \beta \frac{1}{\bar{p}_T}.$$

If we focus on an equilibrium path where $p_{t+1} \geq p_t$, then only type A agents hold tokens in market A. Imposing money market clearing: i.e., $x_{At} = M/p_t$, the conditions above become

$$\begin{aligned} x_{At} &= \beta x_{AT+1} u'(x_{At+1}), \text{ for } t \leq T-1; \\ \frac{x_{At}}{x_{At+1}} &\geq \beta, \text{ for } t \leq T-1; \\ x_{AT} &= \beta u'(x_{AT+1}) \frac{1}{\bar{p}_T} = \beta u'(x_{AT+1}) \frac{T}{\sum_{t=1}^T \frac{1}{x_{At}}}; \\ x_{AT} &\geq \beta \frac{T}{\sum_{t=1}^T \frac{1}{x_{At}}}. \end{aligned}$$

It is straightforward to check that both the monetary equilibrium in the infinite-horizon model and the nonmonetary equilibrium are both in equilibrium in the environment where agents make decisions for T periods with discounting and then receive a continuation value as described in the implementation of treatment DD.

C Spending ratios and fraction of tokens by buyers

Table C.1: Regression of Aggregate Token Spending Ratio on Treatment Dummies

Variables	Market A Spending Ratio	Market B Spending Ratio
BRT	-1.273 (1.954)	-2.803 (1.813)
DD	-0.423 (1.653)	-0.384 (1.612)
Constant	74.010*** (1.364)	78.588*** (1.235)
Observations	428	428
R-squared	0.001	0.007

Notes. (1) Standard errors are clustered at the subject level. (2) * p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01. (3) In each market, the spending ratio is calculated as aggregate token spending as a percentage of the sum of tokens held by buyers.

Table C.2: Regression of % of money held by buyers on Treatment Dummies

Variables	% of money held by buyers in market A	% of money by held buyers in market B
BRT	-2.923** (1.447)	-1.137 (1.699)
DD	-0.703 (1.281)	-0.746 (1.437)
Constant	85.684*** (0.944)	78.197*** (1.156)
Observations	428	428
R-squared	0.011	0.001

Notes. (1) Robust standard errors are in parentheses. (2) * p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01.

D Time Trends: Last Sequence and Period 10 Effects

E Time Path of Output and Price by Session

Table D.1: Time Trends in Treatment RT: Last Sequence and Period 10 Effects

	Market A			Market B		
	All seq	Non-last seq	Last seq	All seq	Non-last seq	Last seq
Inflation						
Period	0.023** (0.121)	-0.010 (0.015)	0.040*** (0.010)	0.026*** (0.007)	0.006 (0.012)	0.036*** (0.009)
Period10	0.033 (0.090)	0.045 (0.123)	0.080 (0.105)	0.079 (0.115)	0.164 (0.165)	0.052 (0.157)
Observations	134	79	55	134	79	55
R-squared	0.386	0.314	0.451	0.458	0.197	0.459
Output						
Period	-0.209*** (0.059)	-0.151 (0.107)	-0.244*** (0.070)	-0.139** (0.064)	-0.107 (0.100)	-0.155* (0.081)
Period10	-0.341 (0.649)	0.523 (1.156)	-1.063 (0.494)	-0.132 (0.790)	-0.427 (1.421)	0.025 (0.895)
Observations	134	79	55	134	79	55
R-squared	0.428	0.316	0.527	0.669	0.649	0.609
Spending ratios						
Period	0.479* (0.252)	-0.386 (0.513)	0.936*** (0.250)	0.246 (0.244)	-0.117 (0.336)	0.431 (0.313)
Period10	-2.402 (2.244)	-0.403 (4.625)	-2.363 (2.090)	0.413 (3.006)	2.638 (3.639)	-0.571 (4.397)
Observations	134	79	55	134	79	55
R-squared	0.706 (3.109)	0.687 (4.535)	0.774 (4.455)	0.723 (1.963)	0.820 (4.339)	0.539 (1.210)

Notes. (1) Robust standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01. (3) The regression also includes a constant, and the session-sequence dummies and their coefficients are omitted.

Table D.2: Time Trends in BRT Treatment: Last Sequence and Period 10 Effects

	Market A			Market B		
	All seq	Non-last seq	Last seq	All seq	Non-last seq	Last seq
Inflation						
Period	0.018** (0.009)	-0.023 (0.015)	0.016 (0.012)	0.035*** (0.008)	0.051*** (0.014)	0.023** (0.009)
Period10	0.260 (0.078)	0.093 (0.076)	0.549*** (0.107)	0.310** (0.130)	0.260 (0.189)	0.308** (0.139)
Observations	134	83	51	134	83	51
R-squared	0.697	0.643	0.754	0.716	0.670	0.818
Output						
Period	-0.167** (0.070)	-0.143 (0.117)	-0.195** (0.086)	-0.067 (0.051)	-0.123 (0.097)	-0.019 (0.048)
Period10	-1.023* (0.562)	-0.482 (0.636)	-2.219** (0.994)	-1.176* (0.623)	-1.172 (0.899)	-0.823 (0.768)
Observations	134	83	51	134	83	51
R-squared	0.621	0.435	0.764	0.708	0.534	0.896
Spending ratios						
Period	0.633** (0.280)	1.396*** (0.428)	-0.057 (0.350)	0.739** (0.327)	0.639 (0.562)	0.854** (0.396)
Period10	9.141*** (2.713)	2.883 (3.278)	16.246*** (4.635)	0.846 (3.046)	-1.232 (4.146)	5.488 (4.339)
Observations	134	83	51	134	83	51
R-squared	0.691	0.760	0.565	0.515	0.425	0.636

Notes. (1) Robust standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01. (3) The regression also includes a constant, and the session-sequence dummies and their coefficients are omitted.

Table D.3: Time Trends in DD Treatment: Last Sequence and Period 10 Effects

	Market A			Market B		
	All seq	Non-last seq	Last seq	All seq	Non-last seq	Last seq
Inflation						
Period	0.025*** (0.010)	0.031*** (0.012)	0.007 (0.014)	0.033*** (0.010)	0.039** (0.013)	0.015 (0.019)
Period10	0.175* (0.090)	0.155 (0.105)	0.234 (0.180)	0.517*** (0.128)	0.553*** (0.130)	0.410** (0.172)
Observations	160	120	40	160	120	40
R-squared	0.629	0.630	0.600	0.500	0.527	0.348
Output						
Period	-0.243*** (0.072)	-0.238** (0.092)	-0.257*** (0.077)	-0.164** (0.074)	-0.138 (0.091)	-0.240** (0.116)
Period10	0.081 (0.488)	0.067 (0.616)	0.126 (0.651)	-2.025*** (0.708)	-2.432*** (0.889)	0.807 (0.859)
Observations	160	120	40	160	120	40
R-squared	0.639	0.605	0.752	0.546	0.555	0.476
Spending ratios						
Period	0.935*** (0.313)	1.200*** (0.361)	0.140 (0.640)	-0.745** (0.357)	-0.354 (0.423)	-1.915*** (0.611)
Period10	8.353*** (3.912)	7.831** (3.272)	9.918 (6.306)	5.102 (3.254)	4.035 (3.825)	8.303 (5.905)
Observations	160	120	40	160	120	40
R-squared	0.453	0.472	0.366	0.470	0.506	0.392

Notes. (1) Robust standard errors are in parentheses. (2)* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01. (3) The regression also includes a constant, and the session-sequence dummies and their coefficients are omitted.

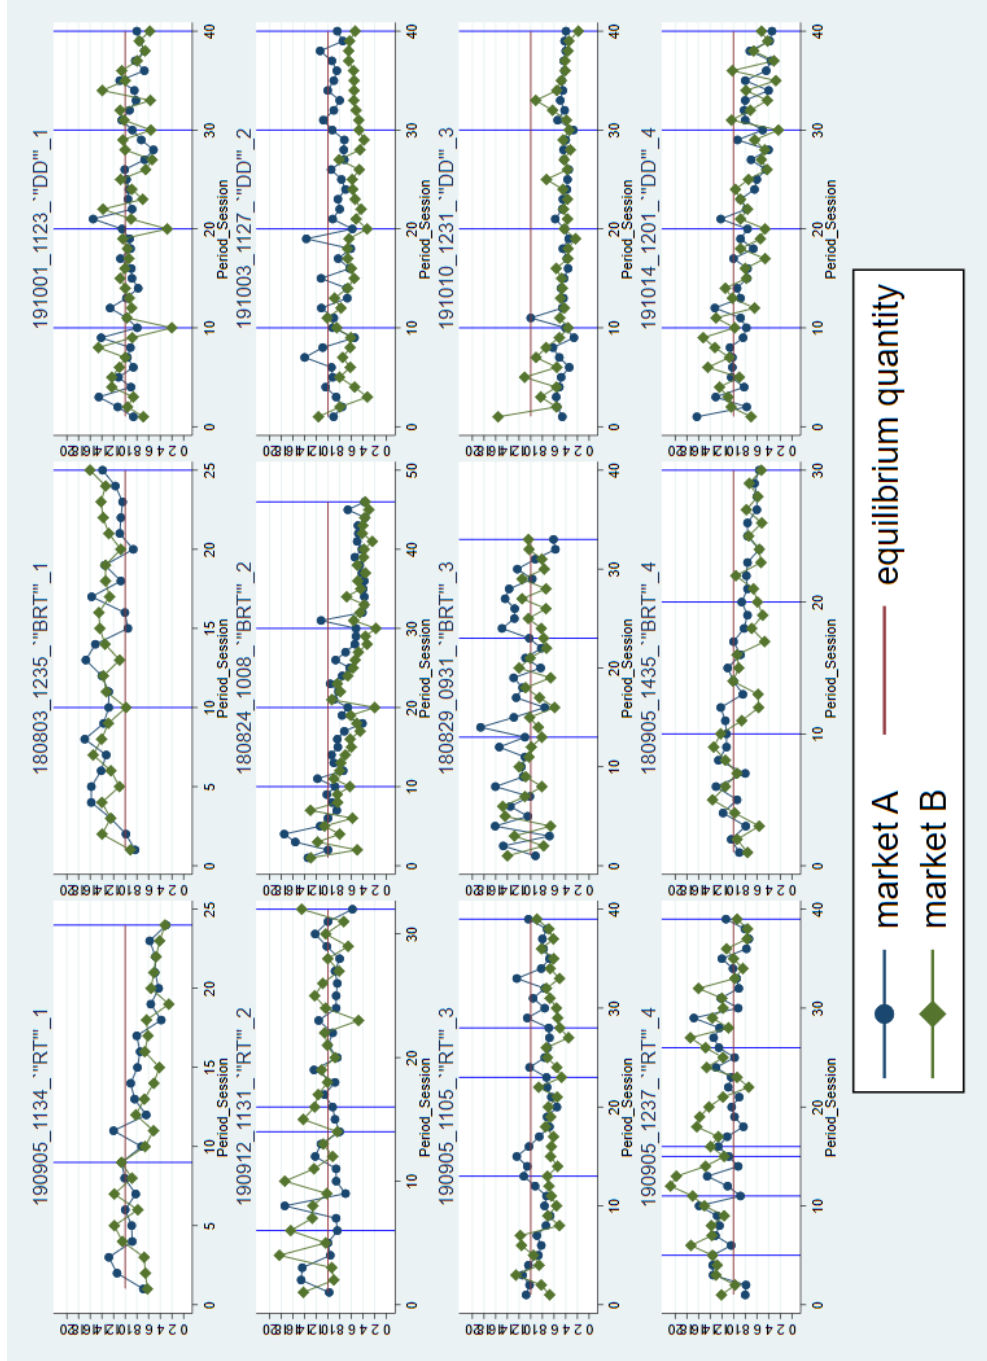


Figure E.1: Average Output in Markets A and B by Session




Figure E.2: Price Levels in Markets A and B by Session