DEAKIN
UNIVERSITY AUSTRALIA

Faculty of Business and Law
School of Accounting, Economics and Finance

School Working Paper

ECONOMICS SERIES

SWP 2013/3

Inflation tax in the lab: a theoretical and experimental study of competitive search equilibrium with inflation

Nejat Anbarci, Richard Dutu and Nick Feltovich

Inflation tax in the lab: a theoretical and experimental study of competitive search equilibrium with inflation

Nejat Anbarci
School of Accounting, Economics and Finance
Deakin University
70 Elgar Road, Burwood VIC 3125, Australia
nejat.anbarci@deakin.edu.au

Richard Dutu
Université de Paris Ouest
200 Avenue de la République
Nanterre 92000, France
rdutu@u-paris10.fr

Nick Feltovich*
Department of Economics
Monash University
Clayton VIC 3800, Australia
nicholas.feltovich@monash.edu

September 25, 2013

Abstract

How does the inflation tax impact on buyers' and sellers' behaviour? How strong is its effect on aggregate economic activity? To answer, we develop a model of directed search and monetary exchange with inflation. In the model, sellers post prices, which buyers observe before deciding on cash holdings that are costly due to inflation. We derive simple theoretical propositions regarding the effects of inflation in this environment. We then test the model's predictions with a laboratory experiment that closely implements the theoretical framework. Our main finding confirms that not only is the inflation tax harmful to the economy – with cash holdings, GDP and welfare all falling as inflation rises – but also that its effect is relatively larger at low rates of inflation than at higher rates. For instance, when inflation rises from 0% to 5%, GDP falls by 2.8 percent, an effect 5 to 7 times stronger than when inflation rises from 5% to 30%. Our findings lead us to conclude that the inflation tax is a monetary policy channel of primary importance, even at low inflation rates.

Journal of Economic Literature classifications: E31, E40, C90.
Keywords: money, inflation tax, directed search, posted prices, cash balances, welfare loss, frictions, experiment.

*Corresponding author. Financial support from Deakin and Monash Universities is gratefully acknowledged. We thank Gabriele Camera, James Costain, Pedro Gomis–Porqueras, Alexander Smith, Neil Wallace, Michael Woodford, Randall Wright and participants at the 2013 Chicago Fed Summer Workshop on Money, Banking, Payments and Finance, the 2013 LEEX International Conference on Theoretical and Experimental Macroeconomics, the 2013 Asia–Pacific Economic Science Association meetings and the 2012 Southern Economic Association meetings for helpful suggestions and comments.

1 Introduction

When prices rise, the real value of individuals' money holdings falls. This phenomenon is known as the *inflation tax*. It affects agents' behaviour, inducing them to adopt strategies (such as shifting consumption away from cash–intensive activities, or simply holding less money) to avoid its effects. Both the inflation tax itself and the resulting behavioural distortions can have implications for social welfare. How strong is the inflation tax's effect on aggregate economic activity? How exactly does it impact buyers' and sellers' behaviour? What is the welfare cost of the inflation tax? Those are central questions in macroeconomics that have received substantial attention in the theoretical literature.

Understanding the strength of the inflation tax is also important for formulating sound monetary policy. In classical macroeconomics (e.g., Cooley and Hansen, 1989) and money search theory (e.g., Lagos and Wright, 2005), the key mechanism through which money has real effects is the inflation tax. By contrast, in the New Keynesian framework (e.g., Galí, 2002) those effects are considered to be quantitatively small and not to capture the main sources of monetary non–neutralities at work in actual economies, i.e., nominal rigidities. Importantly, the recommendations that each of these channels entails can be quite different. If sticky prices are the primary channel, then the optimal policy requires price stability (no inflation or deflation). But if the inflation tax predominates, then pursuing deflation is optimal.

In this paper we propose a new approach to assess the impact and strength of the inflation tax. We begin by developing a model of monetary exchange suitable for experimental testing. We do so by fitting the n–buyer m–seller (n x m) price–posting model analysed by Burdett, Shi and Wright (2001), BSW hereafter, into the money search environment proposed by Lagos and Wright (2005). We then use the model to derive predictions relating to the effect of inflation on price–setting decisions by sellers, cash–holding decisions by buyers, and aggregate economic activity. Finally, we test the model's predictions by conducting a laboratory experiment that implements the model's strategic setting. We use 193 subjects, participating in a total of 2322 trading rounds, who take on the role of buyers and sellers in one of two types of price–posting market (2x2 or 2x3), and make their decisions in an environment where the inflation rate is 0%, 5% or 30%.

We anticipated that in the experiment – as in the model – inflation would work as a tax: reducing real (net of inflation) prices, cash holdings, GDP and welfare. But the experiment also allows us to test the strength of the inflation tax and assess how it compares to the theory. The goals of this study are thus twofold. First, we aim at a better understanding of the effect of inflation on price posting and cash holding decisions in a moderate–inflation environment. Second, we use the results to evaluate the strength of the inflation tax at the aggregate level, and the implications for the modelling of inflation in macroeconomics. To our knowledge, despite their importance, neither of these questions has been addressed empirically in a controlled setting.

We find that behaviour in the experiment is qualitatively in line with the theoretical predictions, with some striking quantitative results in addition. First, statistical tests easily reject the null hypothesis of no difference across our three inflation rates, showing that *the inflation tax matters*. Second, *the effect of the inflation tax is powerful*. In the 2x2 market, for example, real posted prices fall by 10.61 percent and real GDP by 13.2 percent as inflation rises from 0% to 5%, and 11.35 percent and 9.38 percent respectively when inflation jumps from 5% to 30%. Third, *a rise in inflation is relatively more consequential when initial inflation is low*. As inflation rises from 5% to 30% in the 2x2 market, each one–point increase in the inflation rate translates into a 0.5 percent drop in the transaction price and a 0.4 percent drop in real GDP. But when inflation rises from 0% to 5% we find that for each percentage–point increase in inflation, transaction prices fall by 2.6 percent and real GDP falls by 2.8 percent, an effect 5 to 7 times stronger. Results in the 2x3 market largely reinforce those findings: with the effect of a percentage–point increase in inflation between 2.5 and 3 times stronger under low inflation than under high inflation.

Inflation's effects are also seen in welfare; in both 2x2 and 2x3 markets, welfare falls by roughly 5 percent as inflation rises from 0% to 5%. Further increases in inflation have a less pronounced effect, with an inflation rise from

0% all the way to 30% leading to a welfare loss of roughly 15 percent. These losses are higher than those typically seen in the theoretical literature, though we note that differences in methodology limit their comparability.

In the end our research points to a significant effect of the inflation tax on real economic activity, perhaps greater than one may have expected, and apparent even – indeed, especially – when inflation is low. We view our findings as a reminder that the inflation tax is a channel of primary importance.

2 Other relevant work

Although there is a huge literature concerned with inflation, we will focus on a handful of theoretical and experimental papers most closely related to ours, with emphasis on papers not mentioned in the introduction. Much of the work on the effects of inflation has been conducted within theoretical macroeconomic models. This work has been very useful – for instance, in quantifying the costs of inflation to the economy. The earliest attempts by Bailey (1956) and Friedman (1969) treated real money balances as a consumption good and inflation as a tax on these balances, leading to a deadweight loss like that of an excise tax on a commodity. Following Lucas (1987), compensated measures of the costs of inflation within a general equilibrium setting (based on the increase in consumption that an individual would require to be as well off as under zero inflation) were computed, such as Cooley and Hansen (1989) using a cash–in–advance model or Lucas (2000) including money as an argument in the utility function. The welfare cost of 10% inflation was found to be as high as one percent of GDP.

Recently, Burstein and Hellwig (2008) developed a model combining nominal rigidities and the inflation tax. They found that the welfare cost of raising inflation from 2.2% to 12.2% varies widely by model parameters, from roughly zero to almost 7 percent of GDP. More importantly, and directly related to our findings, they showed that the contribution of relative price distortions to the welfare effect of inflation is negligible compared the other channel: the inflation tax (or more precisely, the opportunity cost of holding money since in addition to inflation their model also has a positive real interest rate).

Using short–cuts such as cash–in–advance constraints to introduce money, however, makes it difficult to source the effect of inflation on agents' decisions. As noted by Lucas (2000), "[these models] are not adequate to let us see how people would manage their cash holdings at very low interest rates. Perhaps for this purpose theories that take us farther on the search for foundations, such as the matching models introduced by Kiyotaki and Wright (1989), are needed" (p. 272). Since then, several papers have studied the costs of inflation using money search theory (e.g., Lagos and Wright, 2005; Craig and Rocheteau, 2011). They found that eliminating a 10% inflation rate can have a fairly large welfare benefit – as much as 4 percent of consumption in some circumstances.

While clarifying the theoretical effect of inflation on individual behaviour, none of the papers mentioned above has attempted controlled testing of the qualitative and quantitative implications of their models. Laboratory experiments – which not only give the researcher precise control over the decision making environment, but also allow exogenous manipulation of key parameters such as the inflation rate – can serve as a useful complement to theory and to empirical studies using observational data from the field. Macroeconomics was long considered beyond the reach of experimental methods, but the rise of micro foundations in macro models has made experiments increasingly feasible.[1] Studies of inflation are among the oldest examples of macroeconomics experiments (Marimon and Sunder 1993, 1994, 1995; Lim, Prescott and Sunder 1994; Bernasconi and Kirchkamp 2000), but these experiments focused on hyperinflation while our interest is in moderate inflation levels. Money search – one component of our model – has been studied in the lab, by Brown (1996) and Duffy and Ochs (1999, 2002), and more recently Camera and Casari (2011) and Duffy and Puzzello (2011). These papers concentrated on testing some of the fundamental implications of money–search theory, such as the acceptance of fiat money or the multiplicity of equilibria, rather than examining

[1]See Duffy (2008) for a detailed survey of experimental macroeconomics, and in particular his section on monetary economics.

inflation specifically. More recently, Kryvtsov and Petersen (2013) use an experimental dynamic stochastic general equilibrium (DSGE) model to examine another channel of monetary policy: expectations of future macroeconomic variables.

There is also a small experimental literature examining posted prices and directed search, the other component of our model. Cason and Noussair (2007) examine pricing in an implementation of Burdett, Shi and Wright (2001), comparing observed prices to BSW's predictions and to alternative predictions attributed to Montgomery (1991), where firms are assumed to ignore their strategic interaction with other firms. They found that the data were more consistent with BSW than Montgomery. Anbarci and Feltovich (2013) find that BSW's restriction of sellers to choosing a single price is in some sense without loss of generality, since giving them that additional flexibility to post demand–contingent prices (as in Coles and Eeckhout, 2003) has little effect on market prices, nor does it improve market efficiency.

3 The model

Our model starts with the directed–search environment from Burdett, Shi and Wright (2001). There are $n \geq 2$ buyers and $m \geq 2$ sellers. Sellers produce a homogeneous good, with cost of production 0 for the first unit, and production beyond the first unit impossible. Buyers are also identical, each with valuation $Q > 0$ for the first unit and zero for any additional unit.[2]

Sellers compete in prices in order to attract buyers. Each seller simultaneously posts a price, which is observed by all buyers. Buyers then simultaneously make their visit choices; each can visit only one seller. Trade takes place at the seller's posted price; if multiple buyers visit the same seller, one is randomly chosen to be able to buy. Buyers who aren't chosen, and sellers who aren't visited, do not trade.

3.1 BSW with money

To make the BSW model monetary, we fit it into the Lagos and Wright (2005) model of monetary exchange. Each trading period is divided into two sub–periods. In the first sub–period buyers and sellers participate in a centralised Walrasian market where they can produce and consume any quantity of a single, homogeneous consumption good, called the *general good*. Then they enter a second, frictional, market where a second good, called the *search good*, is allocated via price posting and directed search. As in Lagos and Wright, sellers use the Walrasian market to spend any money earned during the previous round of decentralised trading, while buyers use it to acquire the cash they need for the decentralised market. In our model, the Walrasian market is additionally used by sellers to post prices for the decentralised market, and by buyers to decide on their visit strategies for that market.[3]

Money comes in the form of a perfectly divisible and storable object whose value relies on its use as a medium of exchange. It is available in quantity Z_t at time t, and can be stored in any non–negative quantity z_t by any agent. (Thus z_t is an agent's *nominal* cash balance at time t.) New money is injected or withdrawn via lump–sum transfers by the central bank in the centralised market at rate τ such that $Z_{t+1} = (1+\tau)Z_t$. Only buyers receive the transfer. Inflation is forecasted perfectly and both the quantity theory and the Fisher equation apply: if the money supply increases at

[2] In the experiment, we will set specific values for these parameters, but we will keep our notation general for now.

[3] Rocheteau and Wright (2005) have an earlier price–posting model with money (their "competitive search" case). While their paper makes several important theoretical points, their model would be at best quite difficult to implement in an experiment, due to its assumption of continua of buyers and sellers. Our model, on the other hand, uses finite numbers of buyers and sellers, making experiments feasible. This feature additionally allows the matching between buyers and sellers to emerge from their individual decisions, rather than requiring an exogenous matching function. Other papers introducing price posting in money search models are Kultti and Riipinen (2003) and Julien, Kennes and King (2008). Both papers assume that goods are divisible but money is indivisible, which makes them ill–suited for studying inflation. See also Corbae, Temzelides and Wright (2003) for a model of directed matching, also with indivisible money, and with prices determined by bargaining.

rate τ, so do prices and the nominal interest rate. Denoting r the real interest rate, since $\beta = \frac{1}{1+r}$, the nominal interest rate is $i = \frac{1-\beta+\tau}{\beta}$ (from $(1+i) = (1+r)(1+\tau)$). In the experiment we will assume that the real interest rate is zero ($\beta = 1$) so that $\tau = i$, and thus shifts in the nominal interest rate simply reflect changes in expected (and actual) inflation. The price of the general good is normalised to 1 and the clearing price of money in terms of the general good is denoted ϕ_t (i.e., 1 unit of the general good costs $1/\phi_t$ units of money).

Trade in the search good takes place as in BSW (see Figure 1). Sellers simultaneously post prices for the search good during the centralised market. Buyers observe all prices, then simultaneously choose (1) which seller to visit and (2) how much cash to carry. Buyers and sellers then proceed to the decentralised market where sellers are committed to their price and buyers are constrained by their money holdings. In particular the buyer is only able to buy if he is carrying enough cash to cover the posted price; we define a *serious buyer* as one satisfying this condition. Carrying cash is costly, due to the inflation tax: carrying z units of cash incurs an inflation cost of τz irrespective of whether the buyer is able to buy.

Figure 1: Sequence of decisions associated with the decentralised market

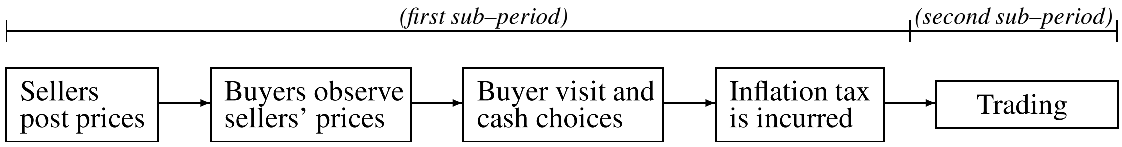

If a seller is visited by exactly one serious buyer, then they trade at the posted price. If a seller is visited by two or more serious buyers, then one is randomly chosen (with equal probability) to buy at the posted price. A seller visited by no serious buyers is unable to sell. [4]

3.2 The payoffs

Buyers have the instantaneous utility function $U_t^b = x_t + \beta u(q_t)$, where x_t is net consumption of the general good at time t, $u(q_t)$ is the utility from consuming q_t units of the search good (with $u(0) = 0$ and $u(q_t) = Q$ for $q_t \geq 1$), and $\beta \in [0, 1)$ is the discount factor between the centralised and decentralised market. Sellers' instantaneous utility function is $U_t^s = x_t - \beta c(q_t)$, where $c(q_t)$ is the cost of producing q_t units of the search good (with $c(0) = c(1) = 0$ and $c(q_t) = +\infty$ for $q_t > 1$).

Let $W^b(z)$ and $V^b(z)$ be the value functions for a buyer holding z units of money in the centralised and frictional markets, respectively. If a buyer decides to take part in the decentralised market we have:

$$W^b(z) = \max_{x,\hat{z}} \left\{ x + \beta V^b(\hat{z}) \right\}, \quad \text{subject to} \quad \phi\hat{z} + x = \phi(z + T). \tag{1}$$

When choosing x (the net consumption of the general good) and a quantity of money to bring to the frictional market, $\hat{z}$, buyers take into account that the combined real value of these two quantities must equal the sum of the money they brought to this market, ϕz, and the amount received from the central bank, ϕT. Substituting out x yields

$$W^b(z) = \phi(z + T) + \max_{\hat{z}} \left\{ -\phi\hat{z} + \beta V^b(\hat{z}) \right\}. \tag{2}$$

[4]We assume that the seller does not change her price in response to the number of buyers visiting her. This is in keeping with Burdett, Shi and Wright's (2001) model, which assumed commitment to a single posted price, but subsequent work (Coles and Eeckhout, 2003; Virag, 2010) has loosened this restriction via a separate multi–buyer posted price or an auction. As mentioned in Section 2, Anbarci and Feltovich (2013) observed that giving sellers the extra flexibility to choose a separate multi–buyer price did not lead to different outcomes at the macro level.

If a buyer does not participate in the frictional market then $\hat{z} = 0$ and

$$W^b(z) = \phi(z + T) + \beta W^b(0). \tag{3}$$

As for sellers, they choose net consumption x in the centralised market and a nominal price p for the decentralised market, and have value function

$$W^s(z) = \max_{x,p} \{x + \beta V^s(p)\}, \quad \text{subject to} \quad x = \phi z. \tag{4}$$

We now turn to the decentralised market and characterise visit and trading probabilities. We focus on symmetric equilibria in which all sellers charge the same price and all buyers use the same mixed strategy. To facilitate comparison with BSW we use their notation and follow their exposition.

Let Φ be the probability that at least one buyer visits a particular seller when all buyers visit him with the same probability θ. Since $(1-\theta)^n$ is the probability that all n buyers go elsewhere, $\Phi = 1 - (1-\theta)^n$. Next, let Ω be the probability that a given buyer gets served when he visits this seller. Since the probability of getting served conditional on visiting this seller times the probability that this buyer visits him equals the probability that this seller serves the particular buyer, we have $\Omega\theta = \Phi/n$. Hence

$$\Omega = \frac{1 - (1-\theta)^n}{n\theta} \tag{5}$$

The Bellman equation for a buyer in the decentralised market is then

$$V^b(z) = \Omega\left\{Q + W^b_{+1}(z - p)\right\} + (1-\Omega)W^b_{+1}(z), \tag{6}$$

This equation says that with probability Ω a buyer gets served, in which case he purchases and consumes one unit of the search good, yielding utility Q. He then enters the next period's centralised market with $z - p$ units of money. With probability $1 - \Omega$ the buyer was not able to trade and proceeds to the centralised market with an unchanged amount of money. The corresponding equation for a seller is

$$V^s(p) = \Phi W^s_{+1}(p) + (1-\Phi)W^s_{+1}(0). \tag{7}$$

Now suppose that every seller is posting p, and one contemplates deviating to p^d. Let the probability that any given buyer visits the deviant be Ω^d. By (5), a buyer who visits the deviant gets served with probability

$$\Omega^d = \frac{1 - (1-\theta^d)^n}{n\theta^d}. \tag{8}$$

Since the probability that he visits each of the non–deviants is $(1-\theta^d)/(m-1)$, a buyer who visits a non–deviant gets served with probability

$$\Omega = \frac{1 - \left(1 - \frac{1-\theta^d}{m-1}\right)^n}{n\left(\frac{1-\theta^d}{m-1}\right)}. \tag{9}$$

The corresponding value function for a buyer visiting a deviant seller in the frictional market is given by

$$V^{b^d}(z^d) = \Omega\left\{Q + W^b_{+1}\left(z^d - p^d\right)\right\} + (1-\Omega)\, W^b_{+1}(z^d). \tag{10}$$

In a symmetric equilibrium of the second–stage game, buyers are indifferent between visiting the deviant seller holding p^d in cash and any other seller holding p in cash. Algebraically this means

$$-\phi p + \beta V^b(p) = -\phi p^d + \beta V^{b^d}(p^d), \tag{11}$$

where ϕp and ϕp^d are the real prices of the search good, expressed in units of the general good. Plugging (6) and (10) into (11), using $\phi_{+1}(1+\tau) = \phi$ and $z = p$, dividing by β and recalling that $i = \frac{1-\beta+\tau}{\beta}$, Equation (11) simplifies to

$$-ip + \Omega\left(\frac{Q}{\phi_{+1}} - p\right) = -ip^d + \Omega^d\left(\frac{Q}{\phi_{+1}} - p^d\right), \tag{12}$$

and given our assumption of a zero real interest rate, so that $i = \tau$, we have

$$-\tau p + \Omega\left(\frac{Q}{\phi_{+1}} - p\right) = -\tau p^d + \Omega^d\left(\frac{Q}{\phi_{+1}} - p^d\right). \tag{13}$$

As can be seen from Equation (13), inflation acts like a tax by reducing the buyer's surplus by an amount equal to the inflation rate times the amount of money carried, τp (or τp^d if he buys from the deviant).

3.3 Symmetric equilibrium

Turning to sellers, expected profit for a deviant seller is identical to that in BSW and given by

$$\pi\left(p^d, p\right) = p^d\left[1 - (1-\theta^d)^n\right]. \tag{14}$$

The first–order condition is given by

$$\frac{\partial \pi}{\partial p^d} = 1 - \left(1-\theta^d\right)^n + p^d n(1-\theta^d)^{n-1}\frac{\partial \theta^d}{\partial p^d} = 0. \tag{15}$$

Assuming $\theta^d \in (0,1)$, we differentiate (13). Inserting equilibrium condition $p^d = p$ and $\theta^d = \frac{1}{m}$ we extract $\frac{\partial \theta^d}{\partial p^d}$ which, once inserted into (15), allows us to obtain the equilibrium value of p as defined in Proposition 1 below.

Proposition 1 *The unique symmetric equilibrium has every buyer visit each seller in the frictional market with probability $\theta^* = 1/m$, and all sellers set a price p^* given by*

$$p^*(m,n,\tau) = \frac{\left[m-1-(m+n-1)(1-\frac{1}{m})^n\right]\left[1-(1-\frac{1}{m})^n\right]m^2 \cdot \frac{Q}{\phi_{+1}}}{\left[(m-1)m^2 - \left((m-1)m^2+mn\right)(1-\frac{1}{m})^n\right]\left[1-(1-\frac{1}{m})^n\right] + \tau m n^2(1-\frac{1}{m})^{n+1}} \tag{16}$$

Note that p^* is the nominal price. The real price (relative to the general good) is therefore $\phi_{+1}p^*$.

An immediate corollary of Proposition 1 is that each buyer will choose to hold p^* in cash. Also, expected nominal GDP per market – calculated as the expected total monetary value of goods traded there – is p^* multiplied by the expected number of trades in the market.[5] Since this latter quantity is

$$M(m,n) = m\left[1 - \left(1 - \frac{1}{m}\right)^n\right] \tag{17}$$

(Burdett, Shi and Wright, 2001), we have that nominal GDP is given by

$$Y^*(m,n,\tau) = m\left[1 - \left(1 - \frac{1}{m}\right)^n\right]p^*, \tag{18}$$

and so real GDP is $\phi_{+1}Y^*$.

The associated comparative statics with respect to τ are simple. Equations (16) and (18) have right–hand sides with the form $A/(B\tau + C)$ with $A, B, C > 0$ and $\tau \geq 0$. So, both p^* and Y^* are decreasing and convex in τ – as are their real analogues. That is, nominal and real prices, cash holdings and nominal and real GDP decrease as the inflation rate increases, but at a decreasing rate.

[5]We will concern ourselves with GDP per market rather than the total level of GDP, for the sake of comparability across experimental sessions with different numbers of markets. For readability, we will sometimes leave out "per market", though this is always implied.

4 The experiment

Our experiment implemented this model in a way that maintains the incentives facing buyers and sellers while automating parts that are less relevant to our research questions. We present an overview of the experimental design and procedures in this section, with additional methodological details in Appendix A for the interested reader.

4.1 Experimental design and procedures

Subjects in the experiment played 54 replications ("rounds") of a one–shot stage game that corresponds to the first period of the infinite–horizon game analysed in Section 3. The sequence of play in a round was as in Figure 1. First, sellers simultaneously post prices which are observed by buyers; second, buyers simultaneously choose both whom to visit and how much cash to hold (thus determining the size of the inflation tax they incur); third, trade takes place. In each round, the continuation (periods two and beyond) was automated, with optimal behaviour assumed, so that a seller automatically got $W^s_{+1}(z)$ in addition to her stage–game payoff – where z is her money holding at the end of the decentralised market – and similarly for buyers. See Appendix B for a demonstration that this simplified single–stage setting preserves the incentives of the infinite–horizon setting analysed in Section 3. Note that since every round of the experiment corresponds to the first period of the infinite–horizon game, ϕ_{+1} is constant throughout the experiment, and without loss of generality can be set to unity, making it unnecessary to distinguish between nominal and real prices in the analysis.

In the experiment, there were two types of market (2x2 or 2x3) and three possible inflation rates (0%, 5% or 30%). Inflation rates of 0% and 5% were chosen because they roughly bracket the actual rates seen in many developed countries in the present and recent past. Our third inflation rate of 30% is high by today's standards, at least in developed countries, but is comparable to the highest levels seen in those countries outside of hyperinflations. (E.g., inflation in the UK in 1974 was estimated at over 24%.) Each subject faced all three inflation rates (within–subject variation), within one of the two markets (between–subject variation). The 54 rounds of a session were split into three blocks of 18 rounds each, with a different inflation rate in each block – with the ordering of the inflation rates varied to control for order effects. Subjects kept the same role (buyer or seller) in all rounds, but were randomly assigned to markets in each round, so as to preserve the one–shot nature of the stage game by having subjects interact with different people from round to round. Some large sessions were partitioned into two "matching groups" that were closed with respect to matching (i.e., subjects in different matching groups were never assigned to the same market), allowing two independent observations from the same session.

The experiment was computerised, and programmed using the z–Tree experiment software package (Fischbacher, 2007). Subjects were primarily Monash University undergraduates. All interaction took place anonymously via the computer program; subjects were visually isolated and received no identifying information about other subjects (not even persistent ID numbers). Instructions were given in writing and orally, the latter in an attempt to make the rules common knowledge.[6] For the same reason, the inflation rate was announced publicly whenever it changed (before rounds 1, 19 and 37).

All subjects began each round with zero cash holdings. Buyer valuations and seller costs were set to 20 Australian dollars and zero respectively, and both prices and cash holdings were allowed to be any multiple of $0.05 between zero and $20 inclusive. A seller's profit was her posted price if she was able to sell, and zero otherwise. A buyer's profit was 20 minus the price paid if he was able to buy, or zero if not, minus the inflation tax in either case. Subjects received end–of–round feedback that included prices and visit choices in their market, and their own profits. At the end of the last round, subjects were paid, privately and individually, the sum of their profits from six randomly chosen

[6]See Appendix C for the instructions and Appendix D for sample screen–shots. Other experimental materials and the raw data are available from the corresponding author upon request.

rounds out of the 54. Total earnings averaged just under \$50 and ranged from \$10 to \$102.10, for a session that typically lasted about 90 minutes.[7]

4.2 Hypotheses

Our hypotheses are based on the implications of Proposition 1. The predicted effect of the inflation rate τ on price, based on m = 2 or 3 sellers, n = 2 buyers and a buyer valuation of $Q = 20$, is shown in Figure 2. (Recall that we have set ϕ_{+1} to 1.) Also shown is the corresponding effect on GDP per market. For the actual inflation rates used in the

Figure 2: Equilibrium posted prices (= cash holdings) and GDP

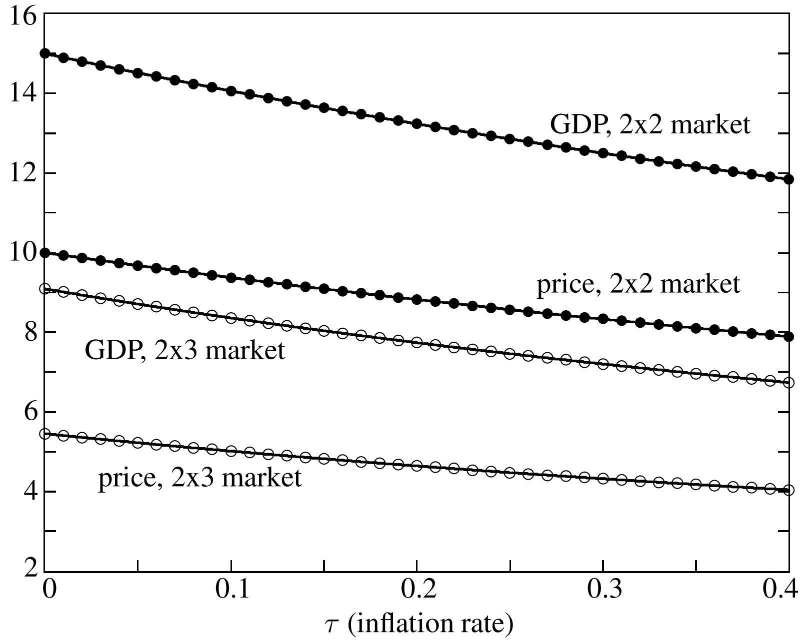

experiment ($\tau = 0$, 0.05, 0.30), the specific predictions of Proposition 1 are shown in Table 1. The table also shows the corresponding semi–elasticities of price and GDP with respect to the inflation rate; these are calculated as $\frac{\partial ln(p^*)}{\partial \tau}$ and $\frac{\partial ln(Y^*)}{\partial \tau}$, and can be interpreted as the proportion change in price or GDP associated with a 1–percentage–point change in the inflation rate. (Note that since, from Equation (18), Y^* is a constant multiplied by p^*, both variables have the same semi–elasticity.) Finally, the table shows predicted levels of efficiency, defined as quantity traded as a proportion of the maximum possible quantity (which is 2 in both markets).

As shown in the figure and the table, raising the inflation rate (ceteris paribus) leads to lower prices and lower GDP, though the size of this predicted effect decreases as inflation increases, as shown by the lower magnitudes of semi–elasticities at higher inflation rates.[8] We therefore have:

Hypothesis 1 *Holding the market constant, prices and GDP will decrease as the inflation rate increases.*

Hypothesis 2 *Holding the market constant, the magnitude of the effect of inflation on prices and GDP will decrease as the inflation rate increases.*

[7] At the time of the experiment, the Australian and US dollars were roughly at parity, while the *Economist*'s Big Mac index estimated their purchasing powers at approximately 1.08 AUD = 1 USD on 30 January 2013 (Economist, 2013).

[8] We manipulated the number of sellers in the experiment in order to determine the robustness of the effect of the inflation rate, rather than out of interest in the effect of market structure per se. We therefore do not state hypotheses concerning the number of sellers. Of course, the model clearly implies that increasing the number of sellers results in lower prices (Figure 2 and Table 1), and although we do not emphasise the corresponding results in the following sections (however, see Note 10), the interested reader can verify that the predicted effect is observed in the data.

Table 1: Theoretical predictions for the treatments used in the experiment

Market	Inflation rate	Price ($)	GDP ($)	Semi–elasticity	Efficiency
2x2	0%	10.00	15.00	–3.00	0.750
	5%	9.68	14.52	–2.90	
	30%	8.33	12.50	–2.50	
2x3	0%	5.45	9.09	–5.58	0.833
	5%	5.23	8.71	–5.35	
	30%	4.32	7.20	–4.42	

Note: semi–elasticity is proportion change in price or GDP associated with a 1–percentage–point change in the inflation rate.

5 Experimental results

We conducted fourteen experimental sessions (not including four earlier pilot sessions, with some differences in model parameters and manipulated variables, and which we do not discuss further in this paper), with a total of 193 subjects.

5.1 Market aggregates

Table 2 reports aggregate experimental data for our main variables of interest. Two measures of price are shown: the average posted price (i.e., the actual choices of sellers) and the average transaction price (those posted prices at which a unit was traded). Also shown is GDP per market. These averages are shown for each combination of market (2x2 or 2x3) and inflation rate ($\tau = 0\%$, 5% or 30%). Finally, for the reader's interest, the table shows levels of efficiency; these are very close to the theoretically predicted levels (see Table 1) and do not vary systematically with the inflation rate, so we will not discuss them further.

Table 2: Aggregate market data

	2x2 market			2x3 market		
Inflation rate (%):	0	5	30	0	5	30
Posted price ($)	11.64	10.32	9.13	7.93	6.86	5.26
Transaction price ($)	11.48	10.04	8.90	7.24	6.32	4.87
GDP ($)	16.97	14.73	13.33	12.22	10.29	7.83
Efficiency (%)	73.9	73.4	74.9	84.9	81.5	80.3

Consistent with Hypothesis 1, the table shows a negative association between the inflation rate τ and posted prices, transaction prices and GDP. Non–parametric statistical tests on the matching–group–level data verify this apparent effect.[9] Page tests reject the null hypothesis of no difference across the three inflation rates in favour of an ordered alternative hypothesis – decreasing price as τ increases – at the 0.1% level for both markets and for posted prices, transaction prices and GDP, except for GDP in the 2x2 market, where significance is at the 1% level. Additionally, pairwise Wilcoxon signed–ranks tests (for matched samples) reject the null hypothesis of no difference

[9] See Siegel and Castellan (1988) for descriptions of the non–parametric tests used in this paper, and Feltovich (2006) for critical values for the robust rank–order tests used later in this section. As noted in Section 4.1, the matching group is the smallest independent unit of aggregation, making it the appropriate unit for non–parametric tests.

in price between $\tau = 0$ and $\tau = 0.05$, and between $\tau = 0.05$ and $\tau = 0.30$, at the 1% level for both markets and for both posted and transaction prices, with only two exceptions ($p \approx 0.02$ for the difference in transaction prices between $\tau = 0$ and $\tau = 0.05$ in the 2x3 market, and $p \approx 0.04$ for the difference in GDP between $\tau = 0.05$ and $\tau = 0.30$ in the 2x2 market).[10]

Table 2 also provides suggestive evidence in favour of Hypothesis 2, as the impact of the rise in inflation from 0% to 5% on prices and GDP is relatively larger than that of the rise from 5% to 30% (i.e., accounting for the fact that the latter is five times as large an increase in the inflation rate). In Table 3, we look more closely at inflation's effects at different inflation levels, by calculating semi–elasticities based on the aggregate data found in Table 2. For each

Table 3: Semi–elasticities based on aggregate market data (percent change in market variable associated with a one–percentage–point rise in the inflation rate)

	2x2 market			2x3 market		
Inflation rate interval (%):	0–5	5–30	p–value, significance of differences	0–5	5–30	p–value, significance of differences
Posted price ($)	–2.38	–0.49	$p \approx 0.01$	–2.86	–1.06	$p \approx 0.06$
Transaction price ($)	–2.64	–0.48	$p \approx 0.01$	–2.68	–1.03	$p \approx 0.10$
GDP ($)	–2.79	–0.40	$p \approx 0.04$	–3.38	–1.09	$p \approx 0.05$

variable (posted prices, transaction prices and GDP) and market (2x2 and 2x3), we compute the percent change in the variable associated with a one–percentage–point rise in the inflation rate, over the interval from 0% to 5% and over the interval from 5% to 30%.[11] Table 3 also reports the results of Wilcoxon signed–ranks tests of differences between semi–elasticities from 0% to 5% inflation and corresponding ones from 5% to 30% inflation (again using matching–group–level data). In all six cases, inflation's impact is higher over the low interval than over the high interval, and in five of the six, the difference is significant at the 10% level or better (in the sixth case, transaction prices in the 2x3 market, the p–value is 0.102, just missing significance at the 10% level).

Figure 3 shows the time series of posted and transaction prices and GDP for each market and inflation rate. Differences in prices across inflation rates are fairly stable over time, as are the prices themselves with one exception: posted prices in the 2x3 market, where there is a steady downward trend over the first several rounds.[12] Differences in GDP across inflation rates are somewhat noisier, but the noise doesn't obscure the treatment effect, and we observe no systematic time trend in these either.

5.2 Parametric analysis of prices and GDP

We move to regressions with posted price, transaction price and GDP as the dependent variables. Our primary explanatory variables are indicators for inflation rates of 0.05 and 0.30 (with 0 as the baseline) and an indicator for the 2x3 market. To allow for time–varying effects, we include the round number (running from 1 to 18, and re–starting

[10] As one might expect, we also find that posted and transaction prices and GDP are significantly lower in the 2x3 treatment compared to 2x2, holding the inflation rate constant (robust rank–order test, $p < 0.001$ for all comparisons, except for GDP when $\tau = 0$, where $p \approx 0.004$).

[11] We compute the semi–elasticity as the value ϵ that solves $x_{\tau_2} = x_{\tau_1}(1-\epsilon)^{\tau_2-\tau_1}$, multiplied by 100 to be expressed as a percent. Note that these are interval semi–elasticities, as opposed to point semi–elasticities such as those displayed in Table 1. We computed point semi–elasticities there because we were working from the theory, and thus knew the exact formula for price as a function of the inflation rate. Here, we have no functional form to work with, so we calculate average semi–elasticities from the data.

[12] This downward trend, combined with the lack of time trend in transaction prices, suggests that many sellers in the 2x3 market initially fail to appreciate the substantial market power buyers have in this market, choose prices that would have been better suited to a market with a more equitable distribution of market power, fail to sell at these prices, and learn to choose lower prices in subsequent rounds. Other explanations are possible.

Figure 3: Time series of posted prices, transaction prices and GDP by inflation rate (2x2 and 2x3 markets)

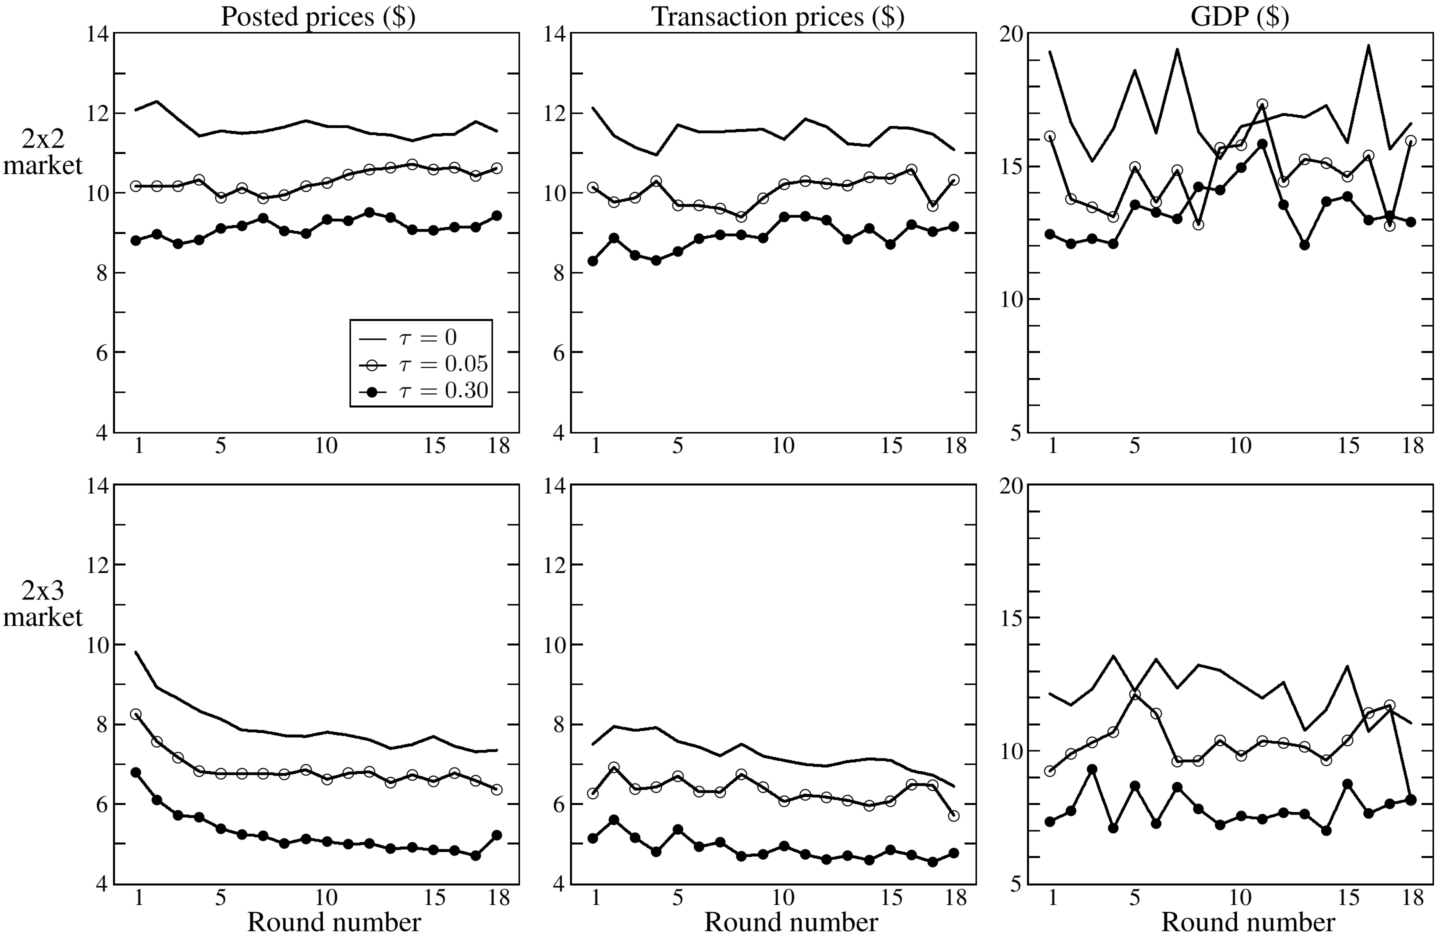

at 1 when the inflation rate changes) and its square on the right–hand side, as well as all of the two– and three–way interactions between the inflation–rate, 2x3–market, and round–number variables.

We also include a number of "nuisance" variables. To control for any order effects due to our within–subject variation of the inflation rate, we include indicator variables for the second and third inflation rates in a session (equivalent to the second and third block of 18 rounds). To control for any attempts by sellers to tacitly collude (due to our repetition of the stage game), we include the number of sellers in the entire session (which was observable to subjects) and the number of sellers in the matching group (not observable, but included in case subjects somehow managed to infer this). Finally, to control for between–subject differences in time spent thinking about their decision, we included the seller's decision time (i.e., the number of seconds from the beginning of the seller–decision stage to the time the seller entered her price). Descriptive statistics for these variables are shown in Table 4. We use Stata (version 12) to estimate panel Tobit models with endpoints 0 and 20, and with individual–seller random effects.

Table 5 reports the estimation results: marginal effects (taken at variables' means) and standard errors. These results reinforce the non–parametric test results presented earlier. Consistent with Hypothesis 1, both inflation–rate dummies have the predicted negative effect on both measures of price and on GDP, and all six of these effects are significant ($p \approx 0.007$ for the effect of the $\tau = 0.05$ dummy on GDP, $p < 0.001$ for the other five effects). Examination of the other variables indicates that prices and GDP are lower in the 2x3 market than in the 2x2 market, that prices (but not GDP) tend to decline over time, and there is evidence of order effects amongst the inflation rates (so we were correct to control for them), while the number of sellers in either the session or the matching group has no significant effect (though the effects on price do have the expected sign).

Table 4: Descriptive statistics for variables used in regressions

Variable	Mean	Std. Dev.	Min.	Max.
Posted price	8.20	3.14	0	20
Transaction price	8.09	3.12	0	19
GDP (per market)	10.14	9.26	0	38
$\tau = 0.05$ dummy	0.333	0.471	0	1
$\tau = 0.30$ dummy	0.333	0.471	0	1
2x3–market dummy	0.589	0.492	0	1
Round number	9.50	5.189	0	18
Second–inflation–rate dummy	0.333	0.471	0	1
Third–inflation–rate dummy	0.333	0.471	0	1
Number of sellers (session)	8.03	1.84	6	12
Number of sellers (group)	6.46	1.81	4	9
Seller decision time (sec.)	7.51	8.59	0	78

Table 5: Tobit results – estimated marginal effects (at variable means) and standard errors

Dependent variable:	Posted price	Transaction price	GDP
$\tau = 0.05$ dummy	-1.094^{***}	-1.123^{***}	-1.036^{***}
	(0.086)	(0.094)	(0.383)
$\tau = 0.30$ dummy	-2.514^{***}	-2.426^{***}	-2.216^{***}
	(0.086)	(0.094)	(0.375)
Significance of difference:	$p < 0.001$	$p < 0.001$	$p \approx 0.002$
2x3–market dummy	-3.881^{***}	-3.878^{***}	-6.537^{***}
	(0.357)	(0.361)	(0.480)
Round number	-0.037^{***}	-0.015^{***}	–0.001
	(0.005)	(0.005)	(0.021)
Second–inflation–rate dummy	-0.239^{***}	-0.137^{**}	–0.054
	(0.060)	(0.066)	(0.268)
Third–inflation–rate dummy	-0.156^{**}	–0.058	–0.036
	(0.061)	(0.067)	(0.269)
Number of sellers (session)	–0.007	–0.024	–0.047
	(0.081)	(0.082)	(0.105)
Number of sellers (matching group)	–0.018	–0.038	–0.156
	(0.096)	(0.096)	(0.124)
Seller decision time (sec.)	0.011^{***}	–0.001	–0.011
	(0.003)	(0.004)	(0.015)
Constant term?	Yes	Yes	Yes
Interaction effects?	Yes	Yes	Yes
N	5778	3620	5778
$\lvert ln(L) \rvert$	11713.98	6876.97	13127.25

* (**,***): *Coefficient significantly different from zero at the 10% (5%, 1%) level.*

While Table 5 shows that the inflation tax has an effect on prices and GDP, and the effect is larger for the higher inflation rate than the lower one, it does not shed any light on curvature: does the effect change more or less quickly at high inflation rates than low ones? A straight comparison of the $\tau = 0.05$ and $\tau = 0.30$ dummies finds that the latter has a significantly larger effect ($p < 0.001$ in both price regressions, $p \approx 0.0017$ for GDP), but this comparison doesn't tell the whole story, since the $\tau = 0.30$ dummy represents a change in inflation five times the size of that of the $\tau = 0.05$ dummy. In Table 6, we make a like–for–like comparison by estimating the ratio of the respective semi–elasticities: that is, the ratio between the effects of a one–percentage–point change in the inflation rate between

0% and 5% inflation (i.e., the marginal effect of the $\tau = 0.05$ dummy, divided by five) and the same change between 5% and 30% inflation (i.e., the difference in marginal effects between the $\tau = 0.30$ and $\tau = 0.05$ dummies, divided by 25). The table reports the point estimate of this ratio of semi–elasticities for posted price, transaction price, and GDP, separately for the 2x2 and 2x3 markets, and for the two markets pooled together. Also shown in the table are the corresponding 95% confidence intervals.

Table 6: Estimated ratio of marginal effects of one–percentage–point rise in inflation based on Table 5 models

	Posted price		Transaction price		GDP	
Market	Point estimate	95% C.I.	Point estimate	95% C.I.	Point estimate	95% C.I.
2x2	5.980	(3.589, 8.371)	6.149	(3.768, 8.530)	4.020	(–4.393, 12.433)
2x3	2.942	(2.027, 3.856)	3.095	(1.890, 4.300)	4.640	(–1.887, 11.167)
Significance of difference:	$p \approx 0.020$		$p \approx 0.025$		$p > 0.20$	
Pooled 2x2 and 2x3	3.856	(2.942, 4.770)	4.306	(3.161, 5.450)	4.389	(–0.767, 9.546)

Notes: Ratio defined as $(5 \cdot \beta_{\tau=0.05})/(\beta_{\tau=0.30} - \beta_{\tau=0.05})$. *A value of 1 indicates a linear effect of* τ*; larger values indicate a diminishing marginal effect.*

The results in this table are fairly striking, and consistent with those seen in Table 3 (thus supporting Hypothesis 2). The nine point estimates vary between about 3 and 6, well above the value of unity that would imply a linear effect of the inflation rate. In the case of GDP, the confidence intervals are wide enough to include 1, so that we can't reject the null hypothesis of a linear effect; however, for both price variables, 1 is well outside any of the confidence intervals, confirming that changes in the inflation rate have larger effects under low inflation than under high inflation.

5.3 Buyer behaviour

We move to an examination of buyers' behaviour. Buyers make two inter–connected decisions: how much cash to hold, and which seller to visit. Both of these decisions are worthy of study not only for their own sake, but because of their role in determining transaction prices and GDP (both of which are affected by buyers' visit and cash–holding decisions), and indirectly in shaping the incentives sellers face when choosing their posted prices. We look at cash holdings here, and at visit decisions in Section 5.5.

Figure 4 shows, for both markets, how average cash holdings change with the inflation rate. The left panel shows a money demand curve for each market (at each inflation rate, the average amount of cash held by all buyers). These

Figure 4: Buyer behaviour – demand for money

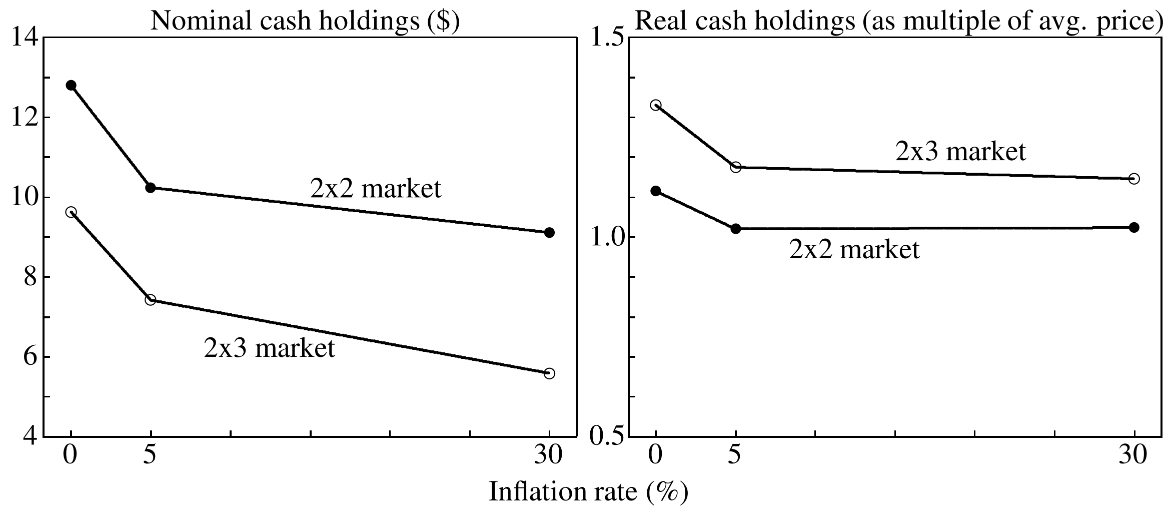

curves replicate patterns found in real data (see, e.g., Lucas, 2000, Figures 2 and 3, or Lagos and Wright, 2005, Figure 2), with cash holdings declining as the inflation rate increases. The right panel shows those same average cash holdings, normalised for each combination of market and inflation rate by dividing by the associated mean transaction price. Even in real terms, money demand tends to decrease as inflation increases (though the difference from 5% to 30% inflation is close to zero in both markets).[13]

5.4 Welfare cost of inflation

We have already documented that a rise in inflation hurts the aggregate economy via falling GDP (see Table 2). In this section we take advantage of the framework we have built and the data we collected in the experiment to compute an estimate of the welfare costs of inflation in our economy.

In our experiment, whenever a transaction occurs, the buyer receives a consumer surplus equal to the difference between his valuation for the good and the seller's posted price, less the amount of the inflation tax incurred in order to participate in the market. The seller receives a producer surplus equal to her posted price (since the cost of production is zero). The natural measure of absolute welfare is then the sum of consumer and producer surplus per market, and consequently welfare loss is the difference in total surplus between zero inflation and a given positive inflation rate.[14] Table 7 summarises the findings.

Table 7: Total surplus (consumer + producer) per market, and welfare loss from inflation

	2x2 market	2x3 market
Total surplus, 0% inflation	\$29.55	\$33.76
Total surplus, 5% inflation	\$28.32	\$31.85
Total surplus, 30% inflation	\$24.48	\$28.76
Welfare loss, 0%-5%	4.2%	5.7%
Welfare loss, 0%-30%	17.2%	14.8%

On average, raising inflation from 0% to 5% is associated with a 5 percent decrease in total surplus, with the loss somewhat higher in the 2x3 market than in the 2x2 market, although the difference is not significant (robust rank–order test, $p > 0.20$). Further increases in inflation lead to additional welfare losses, though the rate of increase slows as inflation rises; the total decrease in surplus as inflation rises from 0% to 30% averages roughly 15 percent, and is significantly larger in the 2x2 market than in the 2x3 market (robust rank–order test, $p \approx 0.05$).

For low levels of inflation, we thus find welfare losses to be higher than previous measures using field data have found. (By comparison, the largest effect mentioned in Section 2 is a 7 percent welfare loss from a 10–percentage–point rise in inflation, in one of Burstein and Hellwig's (2008) cases.) They are not directly comparable, however. A first obvious difference is that our experimental approach gives us access to information such as exact buyer valuations and seller costs, allowing their use in our calculations. Such information would be at best difficult, and at worst impossible, to obtain from field data.[15] Second, our measure deducts the inflation tax in full from total surplus. In a

[13]The negative overall relationship between inflation and money demand continues to hold if we instead put nominal or real cash holdings *net of the price paid* on the vertical axis (figure available from the corresponding author).

[14]Our welfare measure also equals total subject profit, which is the standard measure of well–being used in analysis of economics experiments.

[15]Hence, welfare–loss measures using field data tend to be *compensated*; i.e., they don't measure welfare directly (since they cannot), but instead measure the amount of some other variable, such as income or consumption, that agents would have to gain in order to offset a rise in inflation. Our measure can also be thought of as a compensated measure, since it also represents the change in consumption of the general good that would offset the effect of the inflation tax.

general equilibrium model, however, the inflation tax contributes to the gross income of some lending institutions and therefore cannot be subtracted in full.

5.5 Buyer visit choices

In Figure 5, we examine the other component of buyer behaviour – the choice of which seller to visit. We begin by noting that in a given market and round, the profile of seller prices and inflation rate gives rise to a symmetric subgame played between the two buyers. Each buyer has m+1 pure strategies: one for visiting each of the m sellers, and another strategy we might call "stay home". (Even though buyers in our experiment are required to visit a seller, they can "stay home" by choosing to hold zero cash.) Staying home yields a certain payoff of zero; visiting a seller yields a expected payoff that depends on that seller's price and the inflation rate. This game often, but not always, has multiple Nash equilibria; however, it is easy to show that there is always a unique *symmetric* Nash equilibrium. This symmetric equilibrium is the one used in Section 3.3 to find the equilibrium in seller price choices, and is the one selected by Burdett, Shi and Wright (2001) and others; because buyers in the model have no external information on which to coordinate on an asymmetric equilibrium, the symmetric equilibrium is eminently reasonable.

From this symmetric equilibrium, we construct a *reliability diagram*, showing how closely the predicted and actual probabilities of visiting a seller correspond. This is known as the *calibration* (Yates, 1982) of mixed–strategy equilibrium as a predictor of buyer visit choices.[16] The reliability diagram is constructed as follows. First, for each buyer and round, the predicted probability of visiting each seller is computed, and the associated actual visit choice is recorded.[17] Second, for a given seller number and for each of thirteen intervals of predicted probability, the average of all the predicted probabilities lying in the interval is calculated, as is the frequency of actual visits to that seller in those occurrences.[18] Then, a circle is plotted at the ordered pair (average predicted probability, actual frequency), with an area proportional to the number of occurrences. As an example, one of the intervals we used was the singleton $\{0.5\}$. If there were 100 cases where the predicted probability of visiting Seller 1 was 0.5, and the buyers actually visited Seller 1 in 47 of those cases, a circle would be plotted at (0.5, 0.47), with area proportional to 100.

Figure 5 is the result of this process for visits to Seller 1 (using a different seller number has no qualitative impact), with circles plotted separately for 2x2 and 2x3 markets, for each of the three inflation rates, and for each of the thirteen intervals. Also shown are OLS trend lines for each market and inflation rate, along with the 45–degree line (where predicted and observed probability are equal, and hence calibration is perfect).

Two aspects of buyer visit behaviour are apparent. First, calibration varies between the two markets: buyers in the 2x3 market are remarkably well calibrated, with actual visit frequencies very close to the corresponding predicted probabilities in all three panels, while buyers in the 2x2 market tend to visit a seller too often when the predicted probability is low, and too seldom when it is high. Since predicted probability is based primarily on the seller's price relative to the other seller price(s) – along with the inflation tax – this result means that buyers are insufficiently price–elastic compared to the theory in the 2x2 market, while they have roughly the right level of price sensitivity in the 2x3 market.[19]

[16]For example, if calibration is high, then in all cases where mixed–strategy equilibrium predicts Seller 1 is visited with probability 0.4, buyers should actually have chosen to visit that seller four–tenths of the time; and when the predicted probability is 0.7, she should have been visited seven–tenths of the time by any given buyer; and so on.

[17]As it turns out, in every observation in the experiment, sellers' prices were such that staying home was always strictly dominated by visiting at least one of the sellers, so the predicted probability of staying home was always zero.

[18]Our intervals are $\{0\}$, $(0, 0.1]$, $(0.1, 0.2]$, $(0.2, 0.3]$, $(0.3, 0.4]$, $(0.4, 0.5)$, $\{0.5\}$, $(0.5, 0.6]$, $(0.6, 0.7]$, $(0.7, 0.8]$, $(0.8, 0.9]$, $(0.9, 1)$, $\{1\}$.

[19]Either risk aversion or loss aversion would imply less price elasticity than under expected–payoff maximisation. However, the results in Figure 5 are not solely due to risk aversion, as risk aversion also implies *lower* price elasticity in the 2x3 market than in the 2x2 market, rather than the higher sensitivity that we observe. Loss aversion, on the other hand, does imply higher price elasticity in the 2x3 market than in the 2x2 market, and therefore can explain the *qualitative* relationships we observed. It has less success in characterising behaviour *quantitatively*, as extremely high levels of loss aversion are required to match the amount of price inelasticity observed in the experiment: a loss–aversion

Figure 5: Buyer behaviour – symmetric Nash equilibrium probability of visiting Seller 1 versus actual frequency of visits to Seller 1 *(area of a circle is proportional to the number of observations it represents)*

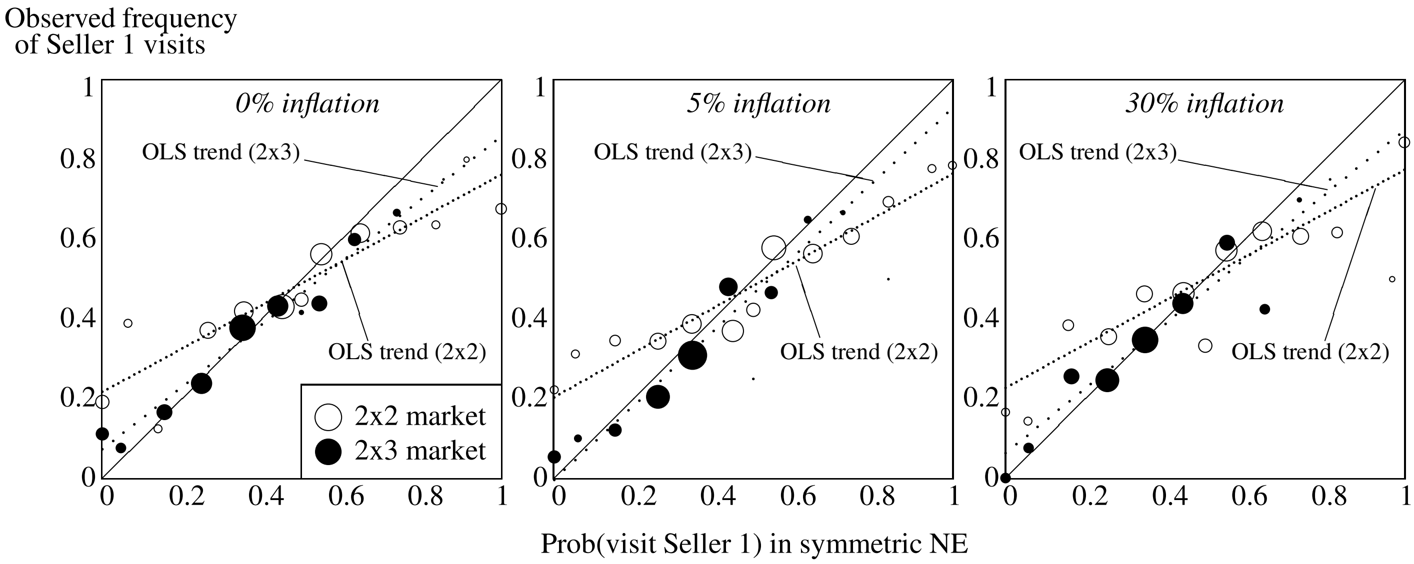

Second, there are no economically relevant differences in calibration across inflation rates within either the 2x2 or 2x3 market. This apparent lack of difference is given further support by a panel probit regression, with Seller 1 visit as the dependent variable (again, using a different seller number doesn't change the conclusions), and with the predicted probability, a dummy for the 2x3 market, and dummies for inflation rates of 5% and 30% on the right–hand side, along with all interactions and a constant term. The results are shown in Table 8.

Table 8: Marginal effects of selected factors on observed frequency of Seller 1 visits (panel probit)

Average marginal effects (std. errors)		Marginal effects of predicted probability at particular variable values				
		Inflation	2x2 market		2x3 market	
		rate	Point estimate	95% CI	Point estimate	95% CI
Predicted probability	0.703***					
	(0.033)					
$\tau = 0.05$ dummy	–0.014					
	(0.017)	0%	0.557	(0.420, 0.694)	0.805	(0.641, 0.969)
$\tau = 0.30$ dummy	0.006					
	(0.017)	5%	0.558	(0.412, 0.704)	0.989	(0.817, 1.160)
2x3–market dummy	−0.043**					
	(0.018)	30%	0.555	(0.412, 0.698)	0.828	(0.655, 1.001)

Notes: $N = 4644$, $|LL| = 2888.21$. ** (**,***): Coefficient significantly different from zero at the 10% (5%, 1%) level.*

The left column shows that the marginal effect of the predicted probability is positive but significantly less than one ($p < 0.001$), implying that buyers do respond to prices in their visit choices, but they are less price–elastic than they should be, as Figure 5 illustrated. The right side of the table shows marginal effects of the predicted probability for each combination of market and inflation rate: clearly, responsiveness is higher in the 2x3 market than in the 2x2 market, and these differences are significant at each inflation rate ($p < 0.001$ at 0% inflation, $p \approx 0.024$ at 5%, $p \approx 0.019$ at 30%). On the other hand, there is no difference in responsiveness across the three inflation rates within a market ($p \approx 0.99$ in the 2x2 market, $p \approx 0.26$ in the 2x3 market).[20] As always, we must be careful in drawing any

parameter well above 12, as compared to values of 2 or 3 typically estimated from individual decision–making tasks. See Appendix E for illustrations of the effects of risk and loss aversion on predicted buyer visit behaviour.

[20]Pairwise tests between any two inflation rates within a market also yield no significant differences.

positive conclusion from a failure to reject null hypotheses, but based on these high p–values, we are fairly confident that there actually is no difference in buyer' price elasticity across inflation rates.

It is worth commenting on this lack of difference across inflation rates. As mentioned already, responsiveness to predicted probability essentially means responsiveness to prices and the inflation rate. Irrespective of the overall level of responsiveness observed in the experiment, or the difference between the 2x2 and 2x3 markets, a systematic difference in responsiveness across inflation rates within a market would suggest that buyers weren't appropriately accounting for the effect of the inflation tax. If buyers ignored the inflation tax, responsiveness would decrease as inflation rises, while if they focus too much on the inflation tax, responsiveness would increase with inflation. The fact that we see neither a rise nor a fall in responsiveness within either market suggests that buyers are – on average – correctly incorporating the inflation tax in with prices when making their visit decisions. It also implies that our comparative–static theoretical predictions for the effect of the inflation rate on *seller prices* could reasonably be expected to prevail in the experiment – as we have seen they do.[21]

6 Discussion

We examine the effects of the inflation tax with a theoretical and experimental analysis. Our theoretical model inserts Burdett, Shi and Wright's (2001) posted–price directed–search model into Lagos and Wright's (2005) money–search model. Sellers in a frictional market independently post prices, which are observed by buyers who then independently decide (a) which seller to visit, and (b) how much cash to hold. Holding cash is necessary in order to buy the seller's item, but is costly because of inflation. We show the model implies that rises in the inflation rate are associated with decreases in prices and GDP, at a rate that diminishes with inflation.

We test the model's predictions with an experiment using three inflation rates. Our results, which are quite stark compared to many lab experiments, are largely consistent with the model in both first–order effects (higher inflation leads to lower real prices, GDP and welfare) and second–order effects: the magnitude of the effect of a one–percentage–point rise in the inflation rate between 0% and 5% inflation on a given statistic varies from 2.5 to 7 times the corresponding effect between 5% and 30% inflation. The effects we find persist as subjects become more experienced, with no economically meaningful variation across replications of the setting. Buyer behaviour also largely supports the model's comparative–static predictions.

Our results indicate that the inflation tax is a channel of primary importance, with even fairly low levels of inflation leading to significant changes in individual behaviour and market aggregates. This leads to an important implication. Although high levels of inflation are universally viewed to be harmful, it is also conventionally accepted that, if the inflation rate could be kept in the low single digits, and as long as changes could be predicted with some degree of accuracy, a society could live fairly easily under such a regime. Indeed, low positive levels of inflation are often viewed as beneficial in developed societies. Our results suggest that positive inflation – however low – entails non–negligible costs, which must be weighed against any benefits.[22]

We believe we have taken a small but worthwhile step toward quantifying the effects of the inflation tax in a controlled setting. We encourage other experimental work in this area. The clarity of our results makes us confident of their robustness to changes in experimental procedures and parameters, but future research might look at larger markets (more buyers and sellers), markets with upward–sloping supply and downward–sloping demand (by inducing heterogeneous costs and valuations), and at alternative pricing protocols, such as random matching with bilateral

[21] For example, if buyer responsiveness to predicted probability had alternatively been systematically lower as inflation increased, simultaneous best–response by sellers might have implied no effect of inflation on prices, or even an increase with inflation.

[22] Dutu, Huangfu and Julien (2011) show in another setting that low levels of inflation can have significant effects. They add inflation to Coles and Eeckhout's (2003) model of demand–contingent price posting and directed search, and find that under even an arbitrarily small positive inflation rate, Coles and Eeckhout's indeterminacy result disappears in favour of a unique equilibrium.

bargaining (as Lagos and Wright (2005) did theoretically). Another avenue for future research would allow for multiple channels through which money can affect the economy (e.g., both the inflation tax and expectations of future macroeconomic variables), thus allowing for direct comparisons between channels.

References

Abdellaoui, M., H. Bleichrodt and O. L'Haridon (2008), "A tractable method to measure utility and loss aversion under prospect theory", *Journal of Risk and Uncertainty* 36, pp. 245–266.

Anbarci, N. and N. Feltovich (2013), "Directed search, coordination failure and seller profits: An experimental comparison of posted pricing with single and multiple prices", *International Economic Review* 54, pp. 873–884.

Andersen, S., J.C. Cox, G.W. Harrison, M. Lau, E.E. Rutström and V. Sadiraj (2011), "Asset integration and attitudes to risk: theory and evidence", working paper 2011–10, Durham University.

Bailey, M. (1956), "The welfare cost of inflationary finance", *Journal of Political Economy* 64, pp. 93–110.

Beetsma, R.M.W.J. and P.C. Schotman (2001), "Measuring risk attitudes in a natural experiment: data from the television game show Lingo", *Economic Journal* 111, pp. 821–848.

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, pp. 315–343.

Brown, P.M. (1996), "Experimental evidence on money as a medium of exchange", *Journal of Economic Dynamics and Control* 20, pp. 583–600.

Burdett, K., S. Shi and R. Wright (2001), "Pricing and matching with frictions", *Journal of Political Economy* 109, pp. 1060–1085.

Burstein, A. and C. Hellwig (2008), "Welfare costs of inflation in a menu cost model", *American Economic Review* 98, pp. 438–443.

Camera, G. and M. Casari (2011), "The coordination value of monetary exchange: experimental evidence", Working paper wp754, Dipartimento Scienze Economiche, Universita' di Bologna.

Camerer, C. (2005), "Three cheers—psychological, theoretical, empirical—for loss aversion", *Journal of Marketing Research* 42, pp. 129–133.

Cason, T.N. and C. Noussair (2007), "A market with frictions in the matching process: an experimental study", *International Economic Review* 48, pp. 665–691.

Coles, M.G. and J. Eeckhout (2003), "Indeterminacy and directed search", *Journal of Economic Theory* 111, pp. 265–276.

Cooley, T.F. and G.D. Hansen (1989), "The inflation tax in a real business cycle model", *American Economic Review* 79, pp. 733–748.

Corbae, D., T. Temzelides and R. Wright (2003), "Directed matching and monetary exchange", *Econometrica* 71, pp. 731756.

Craig, B. and G. Rocheteau (2008), "State–dependent pricing, inflation, and welfare in search economies", *European Economic Review* 52, pp. 441–468.

Dave, C., C.C. Eckel, C.A. Johnson and C. Rojas (2010), "Eliciting risk preferences: When is simple better?" *Journal of Risk and Uncertainty* 41, pp. 219–243.

Deck, C., J. Lee and J. Reyes (2008), "Risk attitudes in large stake gambles: evidence from a game show", *Applied Economics* 40, pp. 41–52.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp and G.G. Wagner (2011), "Individual risk attitudes: new evidence from a large, representative, experimentally–validated survey", *Journal of the European Economic Association* 9, pp. 522–550.

Duffy, J. (2008), "Macroeconomics: a survey of laboratory research", mimeo, Department of Economics, University of Pittsburgh.

Duffy, J. and J. Ochs (1999), "Emergence of money as a medium of exchange: an experimental study", *American Economic Review* 89, pp. 847–877.

Duffy, J. and D. Puzzello (2013), "Gift exchange versus monetary exchange", mimeo, Department of Economics, University of Pittsburgh.

Dutu, R., S. Huangfu and B. Julien (2011), "Contingent prices and money", *International Economic Review* 52, p. 1291–1308.

Economist (2013), "Big Mac index", 2 February.

Feltovich, N. (2005), "Critical values for the robust rank–order test", *Communications in Statistics – Simulation and Computation* 34, pp. 525–547.

Fischbacher, U. (2007), "z–Tree: Zurich toolbox for ready–made economic experiments", *Experimental Economics* 10, pp. 171–178.

Friedman, M. (1969), *The Optimum Quantity of Money and Other Essays*. Chicago: Aldine.

Galí, J. (2002), "New perspective on monetary policy, inflation and the business cycle", NBER working paper 8767.

Greiner, B. (2004), "An online recruitment system for economic experiments", in Kremer, K. and V. Macho eds., *Forschung und wissenschaftliches Rechnen*. GWDG Bericht 63. Göttingen: Gesellschaft für Wissenschaftliche Datenverarbeitung, pp. 79–93.

Harrison, G.W. and E.E. Rutström (2008), "Risk aversion in the laboratory", in Cox, J.C. and G.W. Harrison eds., *Research in Experimental Economics, v. 12*. Bingley, UK: Emerald Group Publishing Ltd., pp. 41–196.

Julien, D., J. Kennes and I. King (2008), "Bidding for Money", *Journal of Economic Theory* 142, pp. 196–217.

Kahneman, D. and A. Tversky (1979), "Prospect theory: an analysis of decision making under uncertainty", *Econometrica* 47, pp. 273–297.

Kiyotaki, N. and R. Wright (1989), "On money as a medium of exchange", *Journal of Political Economy* 97, pp. 927–954.

Kryvtsov, O. and L. Petersen (2013), "Expectations and monetary policy: experimental evidence", working paper, Simon Fraser University.

Kultti, K. and T. Riipinen (2003), “Multilateral and bilateral meetings with production heterogeneity”, *Finnish Economic Papers* 16, pp. 27–37.

Lagos, R. and R. Wright (2005), “A unified framework for monetary theory and policy analysis”, *Journal of Political Economy* 113, pp. 463–484.

Lim, S., E. Prescott and S. Sunder (1994), “Stationary solution to the overlapping generations model of fiat money: experimental evidence”, *Empirical Economics* 19, pp. 255–277.

Lucas, R.E. Jr. (1987), *Models of Business Cycles*, New York: Basil Blackwell.

Lucas, R.E., Jr. (2000), “Inflation and welfare”, *Econometrica* 68, pp. 247–274.

Marimon, R. and S. Sunder (1993), “Indeterminacy of equilibria in a hyperinflationary world: experimental evidence”, *Econometrica* 61, pp. 1073–1107.

Marimon, R. and S. Sunder (1994), “Expectations and learning under alternative monetary regimes: an experimental approach”, *Economic Theory* 4, pp. 131–162.

Marimon, R. and S. Sunder (1995), “Does a constant money growth rule help stabilize inflation? Experimental evidence”, *Carnegie Rochester Conference Series on Public Policy*, pp. 803–836.

Montgomery, J. (1991), “Equilibrium wage dispersion and interindustry wage differentials”, *Quarterly Journal of Economics* 106, pp. 163–179.

Rocheteau, G. and R. Wright (2005), “Money in search equilibrium, in competitive equilibrium, and in competitive search equilibrium”, *Econometrica* 73, pp. 175–202.

Rotemberg, J. and M. Woodford (1997), “An optimization–based econometric framework for the evaluation of monetary policy”, *NBER Macroeconomics Annual* 12, pp. 297–346.

Sidrauski, M. (1967), “Rational choice and patterns of growth in a monetary economy”, *American Economic Review* 57, pp. 534–544.

Siegel, S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York.

Tversky, A. and D. Kahneman (1991), “Loss aversion in riskless choice: a reference–dependent model”, *Quarterly Journal of Economics* 106, pp. 1039–1061.

Virág, G. (2010), “Competing auctions: finite markets and convergence”, *Theoretical Economics* 5, pp. 241–274.

Yates, J.F. (1982), “External correspondence: decompositions of the mean probability score”, *Organizational Behavior and Human Performance* 30, pp. 132–156.

A Additional information about experimental procedures

The experiment has a 3x2 factorial design, with the inflation rate varied within–subjects over the values 0%, 5% and 30%, and the market (2x2 or 2x3) varied between–subjects. To reduce and control for order effects, we varied the ordering of the inflation rates between–subjects, using three of the six possible orderings (see Table 9). We chose the 2x2 market because it is the smallest market where both buyers and sellers face uncertainty about whether they will be able to trade. (Smallness is valuable in our experiment, since small markets allow us to collect more independent observations with a given budget for subject payments.) We chose the 2x3 market because it is one of the two next–smallest markets, the other being the 3x2 market. Both the 2x3 and the 3x2 markets have the added benefit of equilibria well away from the equal–split norm, as opposed to the 2x2 market whose equilibrium price of 10 when $\tau = 0$ gives buyer and seller equal profits. However, the 3x2 market would have led to frequent rounds of negative profits for each buyer, and – since the likely high prices give buyers less opportunity to earn large positive profits in other rounds to offset the losses – a real possibility of overall negative payments to subjects. As usual in experiments, negative payments could not credibly be enforced, thus providing incentives for risk–seeking behaviour, especially by those subjects incurring early losses, hoping to get back above zero payments and knowing that further losses would be costless. Thus, we chose the 2x3 market over the 3x2 market.

Each session lasted for 54 rounds, split into three blocks of 18 rounds each, and with subjects facing a different inflation rate in each block. In a given round, all subjects in a session faced the same inflation rate. There were a total

Table 9: Treatment and session information

Session	Market	Ordering of	Number of		
		inflation rates	subjects	markets	matching groups
5	2x3	5–30–0	10	2	1
6	2x2	0–5–30	16	4	2
7	2x2	5–30–0	16	4	2
8	2x3	0–5–30	15	3	1
9	2x3	5–30–0	10	2	1
10	2x2	30–0–5	16	4	2
11	2x3	30–0–5	15	3	1
12	2x2	30–0–5	12	3	1
13	2x3	5–30–0	15	3	1
14	2x3	30–0–5	20	4	2
15	2x2	0–5–30	16	4	1
16	2x2	5–30–0	12	3	1
17	2x3	0–5–30	10	2	1
18	2x3	5–30–0	10	2	1

Note: 4 pilot sessions not analysed in this paper.

of fourteen experimental sessions (not including four pilot sessions, with some differences in procedures and which we leave out of our data set), conducted between August 2012 and January 2013. Session size varied from two to four times the size of a market (8–16 for the 2x2 market and 10–20 for the 2x3 market). There were 193 subjects in all.

Subjects remained in the same role (buyer or seller) in all rounds, but the composition of the markets (containing two buyers and either two or three sellers) was randomly drawn in each round, so that a given subject was matched with different people from round to round; this was done primarily to lessen the likelihood of repeated–game effects like reputation building or dynamic collusion. Some larger sessions were partitioned into two independent "matching groups", each double the size of a market. Matching groups were closed to interaction (i.e., individual markets were always subsets of a matching group). There was a total of 18 matching groups: 3 for each combination of market and inflation–rate ordering. No mention was made to subjects of the existence of these partitions; subjects were told only that they would be randomly assigned to markets in each round, not that some matchings never occurred.

The experimental sessions took place at MonLEE, Monash University's experimental economics lab. Subjects were primarily undergraduate students from Monash University, and were recruited using the ORSEE web–based recruiting system (Greiner, 2004). No one took part more than once. The experiment was run on networked personal computers, and was programmed using the z–Tree experiment software package (Fischbacher, 2007); some sample screen–shots are shown in Appendix D. Subjects were visually isolated, and were asked not to communicate with other subjects except via the computer program.

At the beginning of a session, subjects were seated at desks with computers, and given written instructions (a sample is provided in Appendix C). The instructions were also read aloud to the subjects, in an attempt to make the rules common knowledge. Additionally, before the first round, a public announcement was made of the initial inflation rate, and it was announced that additional public statements would be made whenever the rate changed (and were made after rounds 18 and 36).

Each round began with both buyers and seller being reminded of the inflation rate, which was framed as an "interest rate" to subjects. Firms were prompted at this time to choose their prices, which could be any multiple of 0.05 Australian dollars, between zero and \$20 inclusive. The restriction to [0, 20] reflects the fact that choosing a price outside this interval is weakly dominated; we simplified the decision situation by making such seller choices impossible rather than merely undesirable. With \$20 as the maximum price, there is no need for a buyer to hold more than \$20. The restriction to multiples of \$0.05 is because Australia no longer circulates coins with denominations smaller than 5 cents.

After the sellers had entered their prices, buyers observed these prices and were prompted to choose which firm to visit. Within a round, the firms in a market were labelled as "Seller 1", "Seller 2", and in the 2x3 market, "Seller 3", so that buyers could make clear which one they wanted to visit. These labels were chosen randomly and i.i.d. in each round, preserving anonymity and ensuring that labels could not be used as a coordination device across rounds. Buyers also chose their cash holdings at this time; these could also be any multiple of \$0.05 between zero and \$20 inclusive. (Since \$20 is the maximum price, there is no benefit from holding more than \$20.) Once all buyers had made both of these decisions, the round ended and subjects received feedback. Firms were informed of both own and rival prices, how many buyers visited them, the quantity sold and profit. Buyers were informed of all prices, which firm each buyer visited, the quantity bought and profit.

At the end of the last round, subjects were paid, privately and individually. For each subject, two rounds from each block of 18 were randomly chosen, and the subject was paid in Australian dollars his/her earnings in those six rounds, plus a \$10 show–up fee. Subjects' total earnings averaged just under \$50 and ranged from \$10 (for one subject) to \$102.10, for a session that typically lasted about 90 minutes.

B Equivalence of the infinite–horizon model and one–shot experimental round

The model (Section 3) involved an infinite–horizon dynamic game, while the experiment (Section 4) involved stationary repetition of a one–stage game. In what follows, we show that these seemingly different settings are actually equivalent, in that the respective payoff functions the induce – holding constant other agents' behaviour and assuming optimal continuation behaviour in the dynamic game – differ at most by an affine transformation.

Before beginning the proofs, we describe in some additional detail the way the infinite–horizon model is implemented in the experiment. Figure 6 shows a pictorial version of this implementation. The top of the figure (part (a)) shows a representation of the infinite–horizon model analysed in Section 3, including the expressions for buyer and seller payoffs in each period. Inflation operates in the usual way within this setting, with the price of the general good (and in equilibrium, the search good) rising by a factor of $(1+\tau)$ each period.

Figure 6: Diagram of how infinite–horizon model is implemented in experiment

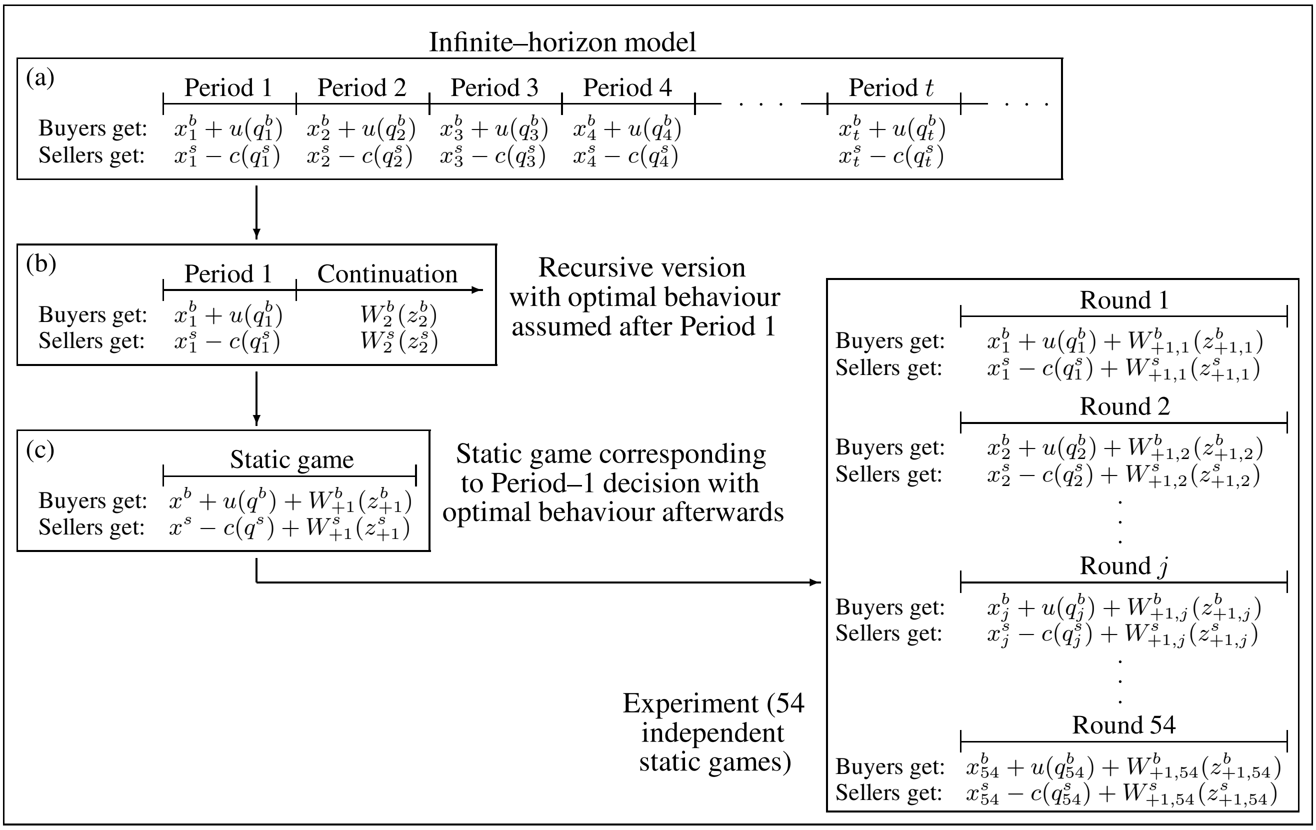

In solving the infinite–horizon game, we typically do not work with the infinite series of payoffs, but rather with its recursive analogue (b). Here, optimal behaviour in all future periods is assumed – the result of which contributes $W_2^b(z_2^b)$ to payoffs for a buyer, or $W_2^s(z_2^s)$ for a seller – and so we focus on the first–period decision (which affects W_2 by determining z). Given this, we can think of the entire setting as if it were a static game as in (c), with the first–period and continuation payoffs combined. We show in the proofs below that an individual round of the experiment (part (d)) corresponds to one of these static games. Thus the 54 rounds of an experimental session can be thought of as 54 independent replications of the entire infinite–horizon model, and in particular, *not a 54–period subset of a single infinite–horizon model*. Among other things, this means that even when inflation is positive, prices need not rise from round to round in the experiment.

We now move to the proofs, first for sellers, then for buyers.

Sellers

Consider a seller at the beginning of a period (i.e., in the Walrasian market) of the infinite–horizon model, initially holding z units of money. Suppose she chooses price p, and let $\overline{W}^s$ be the resulting payoff function, assuming optimal continuation behaviour in all future periods, but not necessarily optimal behaviour in the current period. (This last point is our reason for using $\overline{W}^s$ instead of W^s.) We have

$$\begin{aligned}
\overline{W}^s(z,p) &= \phi z + \beta V^s(p) \\
&= \phi z + \beta \Phi W^s_{+1}(p) + \beta(1-\Phi)[W^s_{+1}(0)] \\
&= \phi z + \beta \Phi[\phi_{+1} p + \text{Max}_{\tilde{p}}\{\beta V^s_{+1}(\tilde{p})\}] + \beta(1-\Phi)\text{Max}_{\tilde{p}}\{\beta V^s_{+1}(\tilde{p})\} \\
&= \beta \Phi \phi_{+1} p + \phi z + \beta^2 \text{Max}_{\tilde{p}}\{V^s_{+1}(\tilde{p})\}.
\end{aligned}$$

Since the real interest rate is assumed to be zero, we have $\beta = 1$, and therefore

$$\overline{W}^s(z,p) = \Phi\phi_{+1}p + \phi z + \text{Max}_{\tilde{p}} V^s_{+1}(\tilde{p}) = \Phi\phi_{+1}p + W^s(z). \tag{19}$$

On the right–hand side of (19), the second term doesn't depend on choice variables – only the initial money balance z. Also, the first term is the payoff function for sellers in a round of the experiment: the probability of selling a unit multiplied by the profit conditional on selling a unit. Thus, the seller payoffs in the infinite–horizon model and those in the experimental stage game differ only by a constant. □

Buyers

For a buyer initially with z units of money, and choosing to hold $\hat{z}$ in the Walrasian market, let $\overline{W}^b$ be the resulting payoff function (as previously, assuming optimal behaviour in future but perhaps not in the present). Then,

$$\begin{aligned}
\overline{W}^b(z,\hat{z}) &= \phi(z+T) - \phi\hat{z} + \beta V^b(\hat{z}) \\
&= \phi(z+T) - \phi\hat{z} + \beta[\Omega Q + \Omega W^b_{+1}(\hat{z}-p) + (1-\Omega)W^b_{+1}(\hat{z})].
\end{aligned}$$

Note that

$$\begin{aligned}
W^b_{+1}(\hat{z}-p) &= \phi_{+1}(\hat{z}-p+T) + \text{Max}_{\tilde{z}}\{-\phi_{+1}\tilde{z} + \beta V^b(\tilde{z})\} \quad \text{and} \\
W^b_{+1}(\hat{z}) &= \phi_{+1}(\hat{z}+T) + \text{Max}_{\tilde{z}}\{-\phi_{+1}\tilde{z} + \beta V^b(\tilde{z})\},
\end{aligned}$$

so that

$$\begin{aligned}
\overline{W}^b(z,\hat{z}) &= \phi(z+T) - \phi\hat{z} + \beta\Omega Q + \beta\Omega[\phi_{+1}(\hat{z}-p+T) + \text{Max}_{\tilde{z}}\{-\phi_{+1}\tilde{z} + \beta V^b(\tilde{z})\}] \\
&\quad + \beta(1-\Omega)[\phi_{+1}(\hat{z}+T) + \text{Max}_{\tilde{z}}\{-\phi_{+1}\tilde{z} + \beta V^b(\tilde{z})\}] \\
&= \phi(z+T) - \phi\hat{z} + \beta\Omega[Q - \phi_{+1}p] + \beta\phi_{+1}(\hat{z}+T) + \beta\text{Max}_{\tilde{z}}\{-\phi_{+1}\tilde{z} + \beta V^b(\tilde{z})\} \\
&= \phi(z+T) - \phi\hat{z} + \beta\Omega[Q - \phi_{+1}p] + \beta\phi_{+1}(\hat{z}-z) + \beta\phi_{+1}(z+T) + \beta\text{Max}_{\tilde{z}}\{-\phi_{+1}\tilde{z} + \beta V^b(\tilde{z})\} \\
&= \phi(z+T) - \phi\hat{z} + \beta\Omega[Q - \phi_{+1}p] + \beta\phi_{+1}(\hat{z}-z) + \beta W^b_{+1}(z).
\end{aligned}$$

Since $\beta = 1$ in our experiment, we have

$$\begin{aligned}
\overline{W}^b(z,\hat{z}) &= \Omega(Q - \phi_{+1}p) - (\phi - \phi_{+1})\hat{z} + (\phi - \phi_{+1})z + \phi T + W^b_{+1}(z) \\
&= \Omega(Q - \phi_{+1}p) - \tau\hat{z} + \left[\tau z + \phi T + W^b_{+1}(z)\right].
\end{aligned}$$

The terms inside the square brackets don't depend on either the visit choice (which affects p and Ω) or the cash holding choice ($\hat{z}$). The two remaining terms are the payoff function for buyers in the experiment: the first is the consumer surplus times the probability of being able to buy, and the second is the amount of the inflation tax. Thus, the buyer payoffs in the infinite–horizon model and those in the experimental stage game differ only by a constant. □

C Instructions from the experiment

Below are the instructions from the 2x2 treatment; the instructions from the 2x3 treatment are nearly identical and available from the corresponding author. Horizontal lines indicate page breaks.

Instructions

You are about to participate in a decision making experiment. Please read these instructions carefully, as the money you earn may depend on how well you understand them. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants during the experiment.

This experiment is made up of 54 rounds. Each round consists of a simple computerised market game. Before the first round, you are assigned a role: buyer or seller. ***You will remain in the same role throughout the experiment.***

In each round, the participants in this session are divided into "markets": groups of four containing a total of ***two buyers*** and ***two sellers***. ***The other people in your market will change from round to round.*** You will not be told the identity of the people in your market, nor will they be told yours – even after the session ends.

The market game: In each round, a seller can produce up to ***one*** unit of a good, at a cost of ***$0***. A buyer can buy up to ***one*** unit of the good, which is ***resold to the experimenter*** at the end of the round for ***$20***. It is not possible to buy or sell more than one unit in a round. Sellers begin a round by choosing the prices of their goods, which must be entered as ***multiples of 0.05, between 0 and 20*** inclusive (without the dollar sign).

After all sellers have chosen prices, each buyer observes the prices of each of the sellers in their market, then chooses which seller to visit. If the two buyers in a market visit ***different*** sellers, then ***both*** buyers have the opportunity to purchase their seller's item at that seller's price. If both visit the ***same*** seller, then since the seller can only produce one unit, ***only one*** buyer will have the opportunity to purchase the item at that seller's price.

In order to be able to purchase an item, ***a buyer must be carrying enough cash to cover its price***. Buyers begin each round with ***no cash***, but they can choose to ***borrow*** from the bank when they are choosing which seller to visit. The amount they choose to borrow can be ***any multiple of 0.05, between 0 and 20*** inclusive. The bank charges ***interest*** on any amount borrowed. The total interest charged in a round is equal to the ***interest rate*** multiplied by the amount borrowed.

Examples:
- If the interest rate is 10% (= 0.10) and the amount borrowed is $15.00, then the interest charge is 0.10 * $15.00 = $1.50.
- If the interest rate is 20% (= 0.20) and the amount borrowed is $5.00, then the interest charge is 0.20 * $5.00 = $1.00.
- If the interest rate is 40% (= 0.40) and the amount borrowed is $0.00, then the interest charge is 0.40 * $0.00 = $0.00.

The interest rate is shown on everyone's screens at the beginning of each round, and it may be ***zero or positive***. It is the same for all buyers in a round, but it may change from round to round. ***When the interest rate changes, an announcement will be made to all participants.*** Interest charges are automatically deducted from the buyer's profit, so there is no need to borrow to pay interest. Sellers have no reason to borrow, so do not pay interest.

Buying and selling: If you are a ***seller***, then you are ***able*** to sell your item as long as ***at least one*** buyer (a) visits you, and (b) has enough cash to pay the price you chose. If no buyer visits you, or if the buyers who did visit did not borrow enough to pay your price, then you are ***unable*** to sell.

If you are a ***buyer***, then:
- If you and the other buyer chose ***different sellers***, then you are able to buy your seller's item as long as you borrowed enough money to pay the price.
- If you and the other buyer chose the ***same seller***, and the other buyer ***did not borrow enough*** to pay the seller's price, then you are able to buy as long as you have enough money.
- If you and the other buyer chose the ***same seller***, and ***both of you*** have enough money to pay the price, then each of you has a ***50%*** chance of being able to buy the seller's item at that price. One of you is chosen ***randomly*** by the computer to buy; the other buyer will be ***unable*** to buy.
- If you ***did not borrow enough money*** to pay your seller's price, then you will be ***unable*** to buy.

Profits: Your profit for the round depends on the round's result.
- If you are a ***seller*** and you are ***able to sell***, your profit is the selling price.
- If you are a ***seller*** and you are ***unable to sell***, your profit is zero.
- If you are a ***buyer*** and you are ***able to buy***, your profit is $20.00 minus the price you paid, minus the amount of interest charged (if any).
- If you are a ***buyer*** and you are ***unable to buy***, your profit is zero minus the amount of interest charged (if any).

Sequence of play in a round:
(1) The computer randomly forms markets made up of two buyers and two sellers, and displays the current interest rate on everyone's screen.
(2) Sellers choose their prices.
(3) Buyers observe the sellers' prices, then each buyer chooses which seller to visit and how much to borrow from the bank.
(4) The round ends. If you are a seller, you are informed of: each seller's price, how many buyers visited you, quantity sold and profit for the round. If you are a buyer, you are informed of: each seller's price, which seller each buyer visited, your quantity bought and profit for the round.
After this, you go on to the next round.

Payments: At the end of the experiment, ***six*** rounds will be chosen randomly for each participant. You will be paid your total profit from those rounds, plus an additional $10 for completing the session. Payments are made privately and in cash at the end of the session.

D Sample screen-shots from the experiment

Below are sample screen-shots from the experiment. They were taken during a test-run of the program; in the actual experiment the top line of each screen-shot would read “1 of 54” rather than “1 of 6”. Otherwise, these screen-shots are typical of those seen by subjects in the 2x2 market with a 0% inflation rate. (Screen-shots from other cells are available from the corresponding author upon request.)

Seller decision screen:

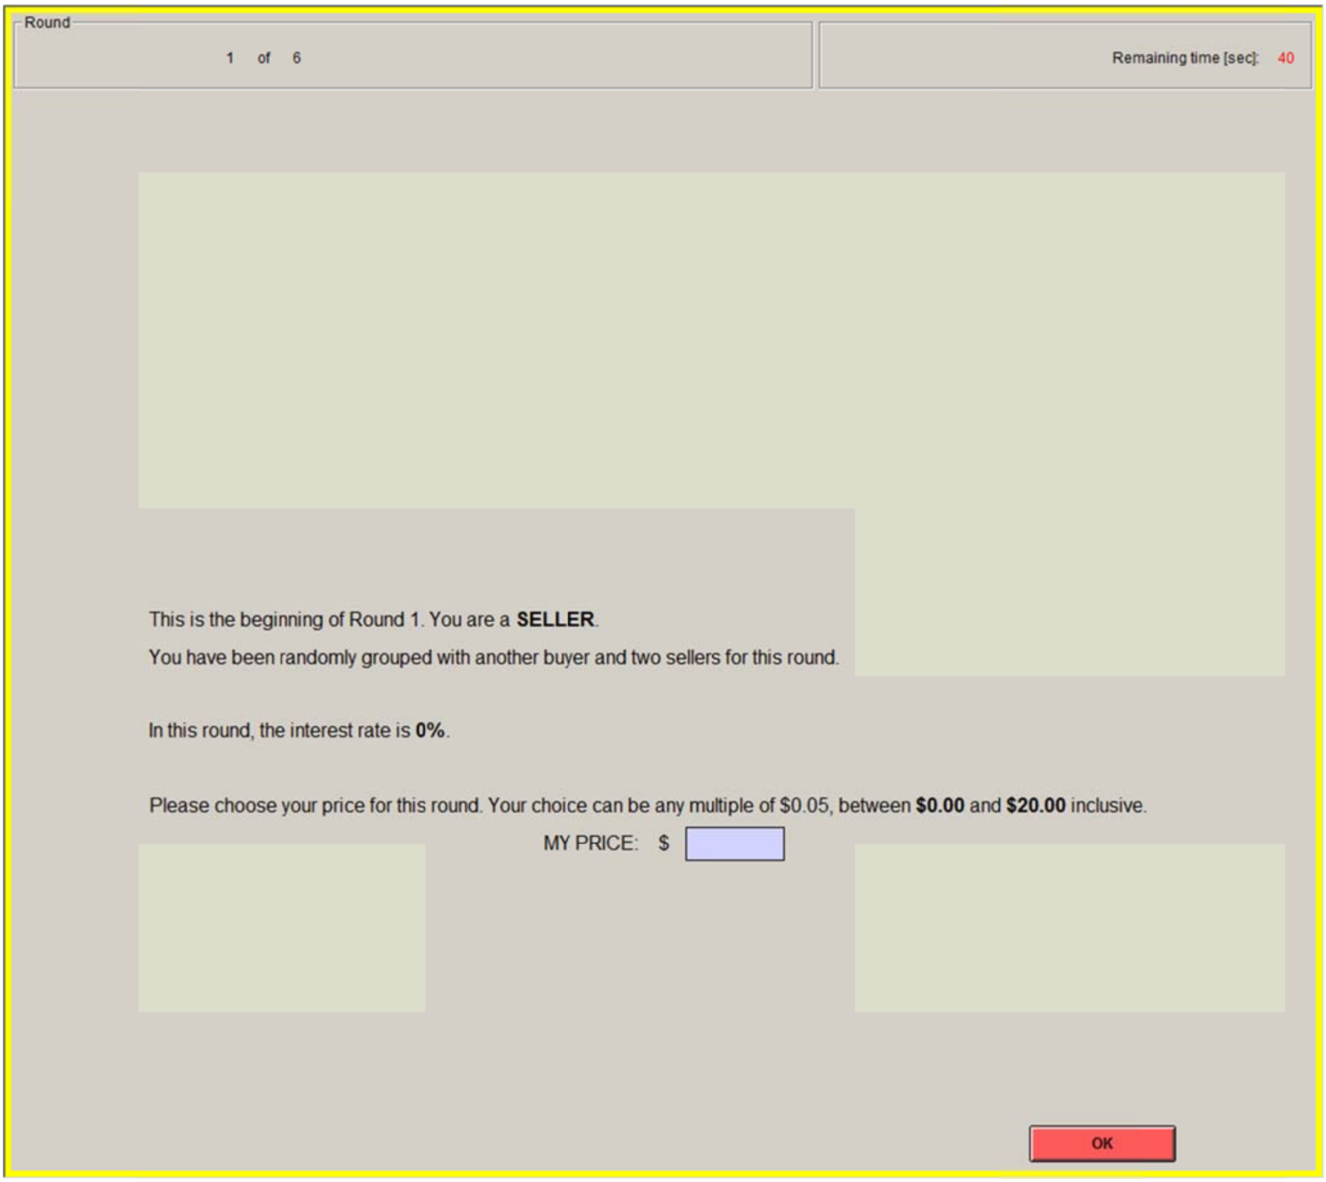

Buyer decision screen:

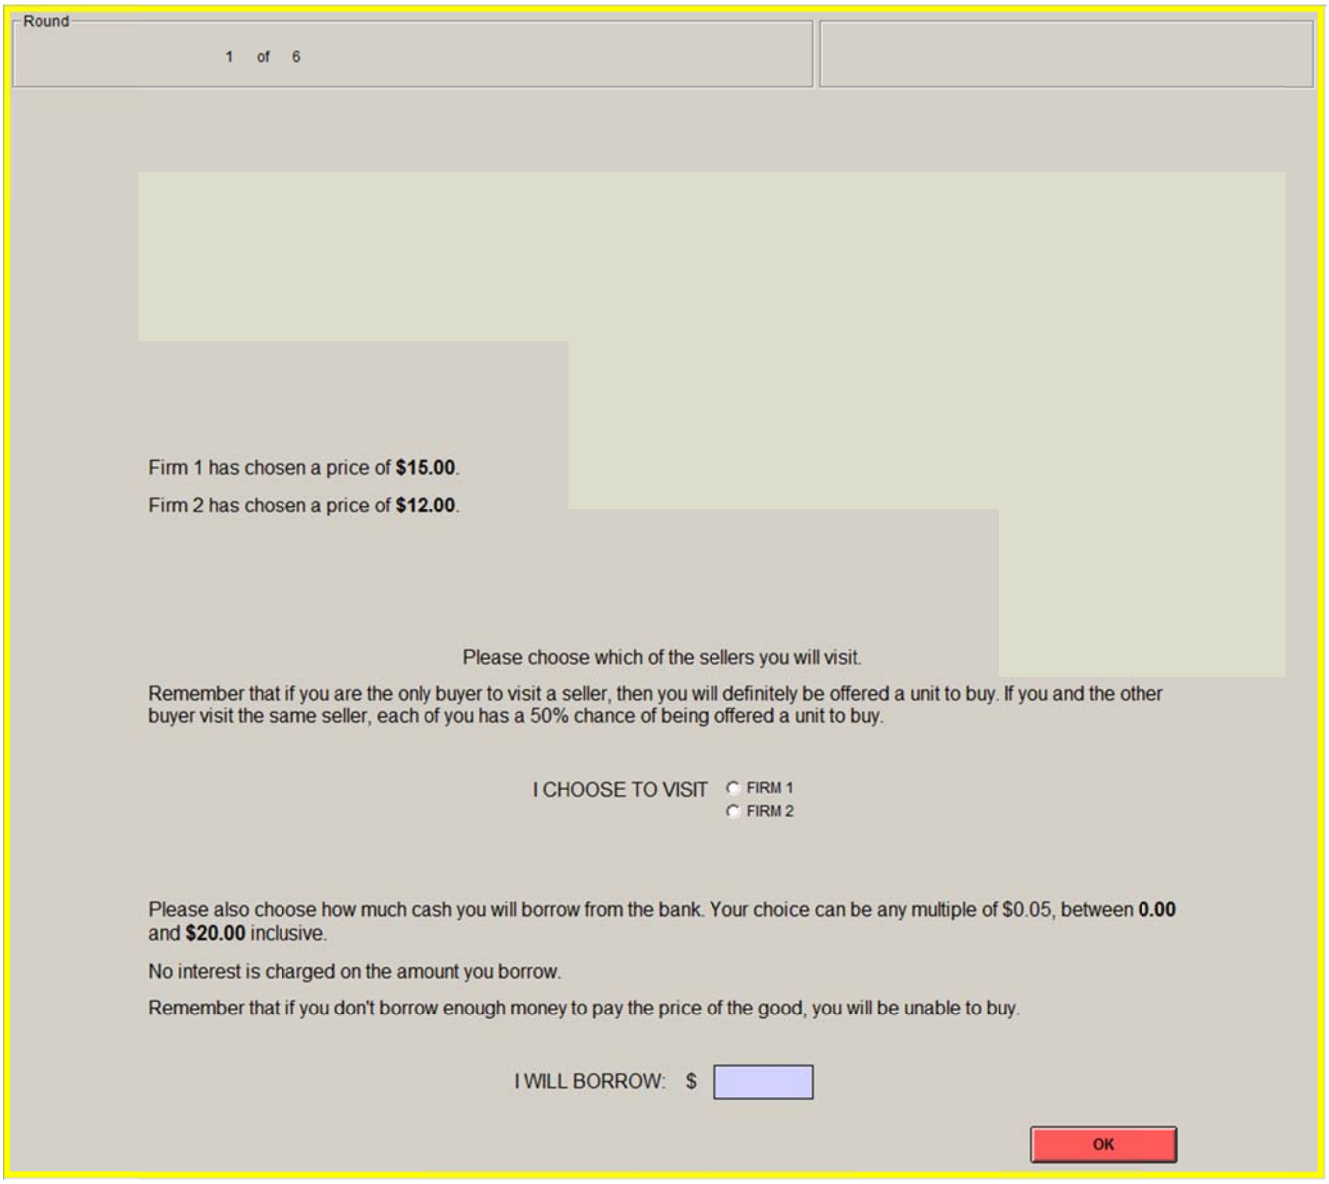

Seller feedback screen:

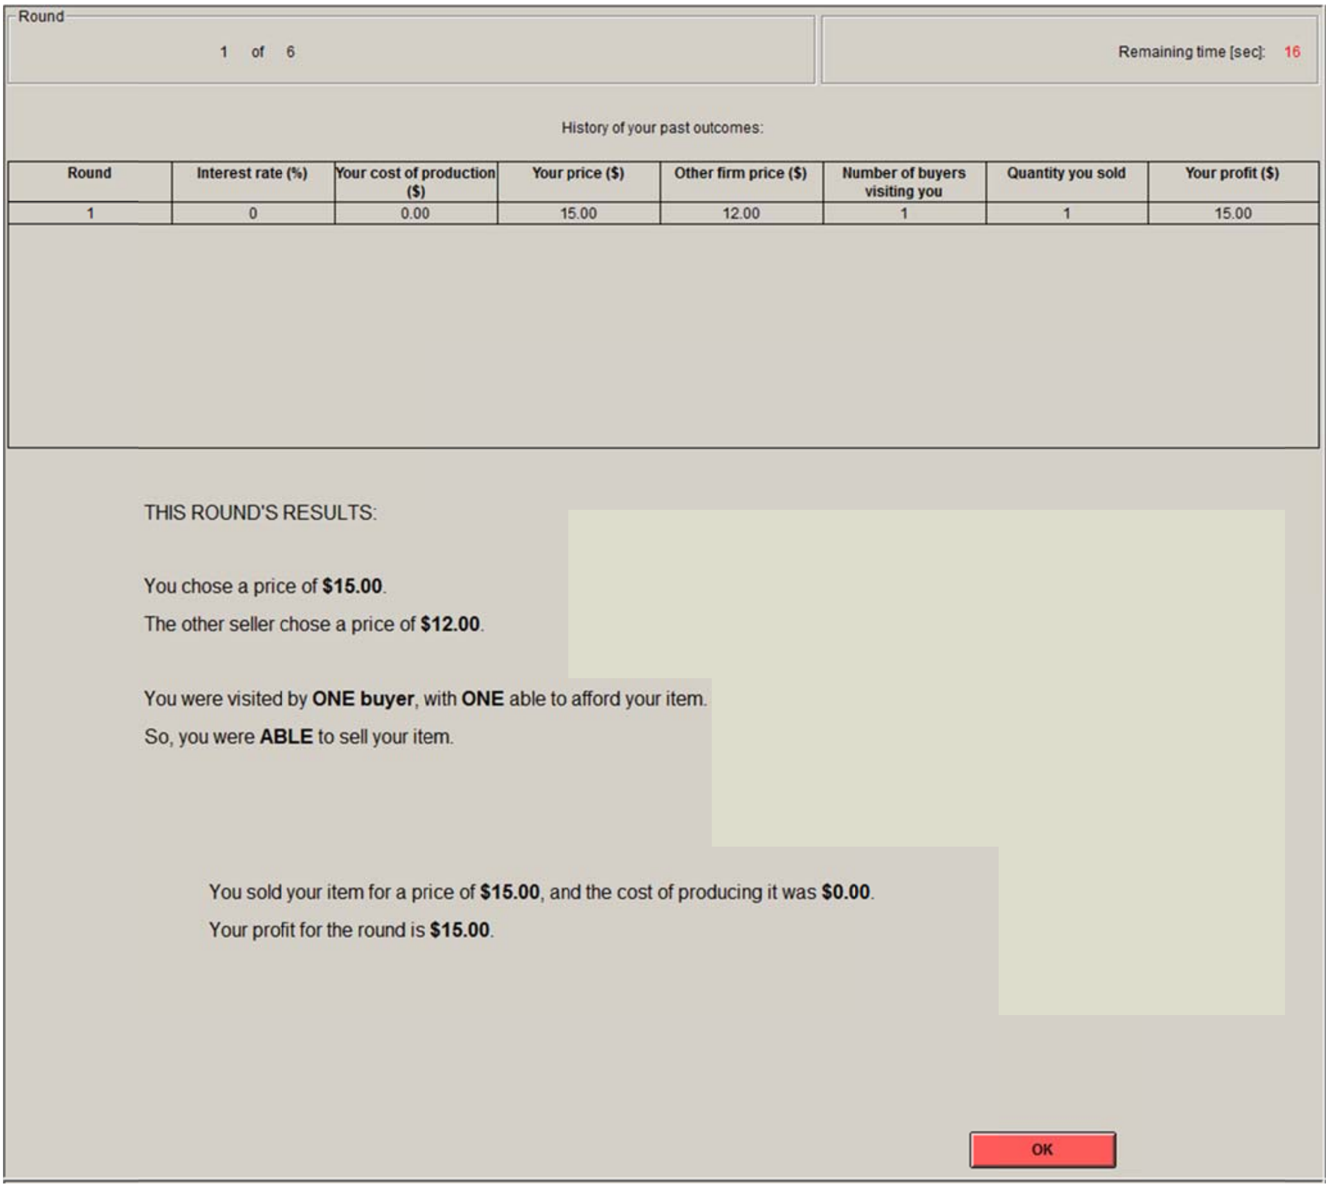

Buyer feedback screen:

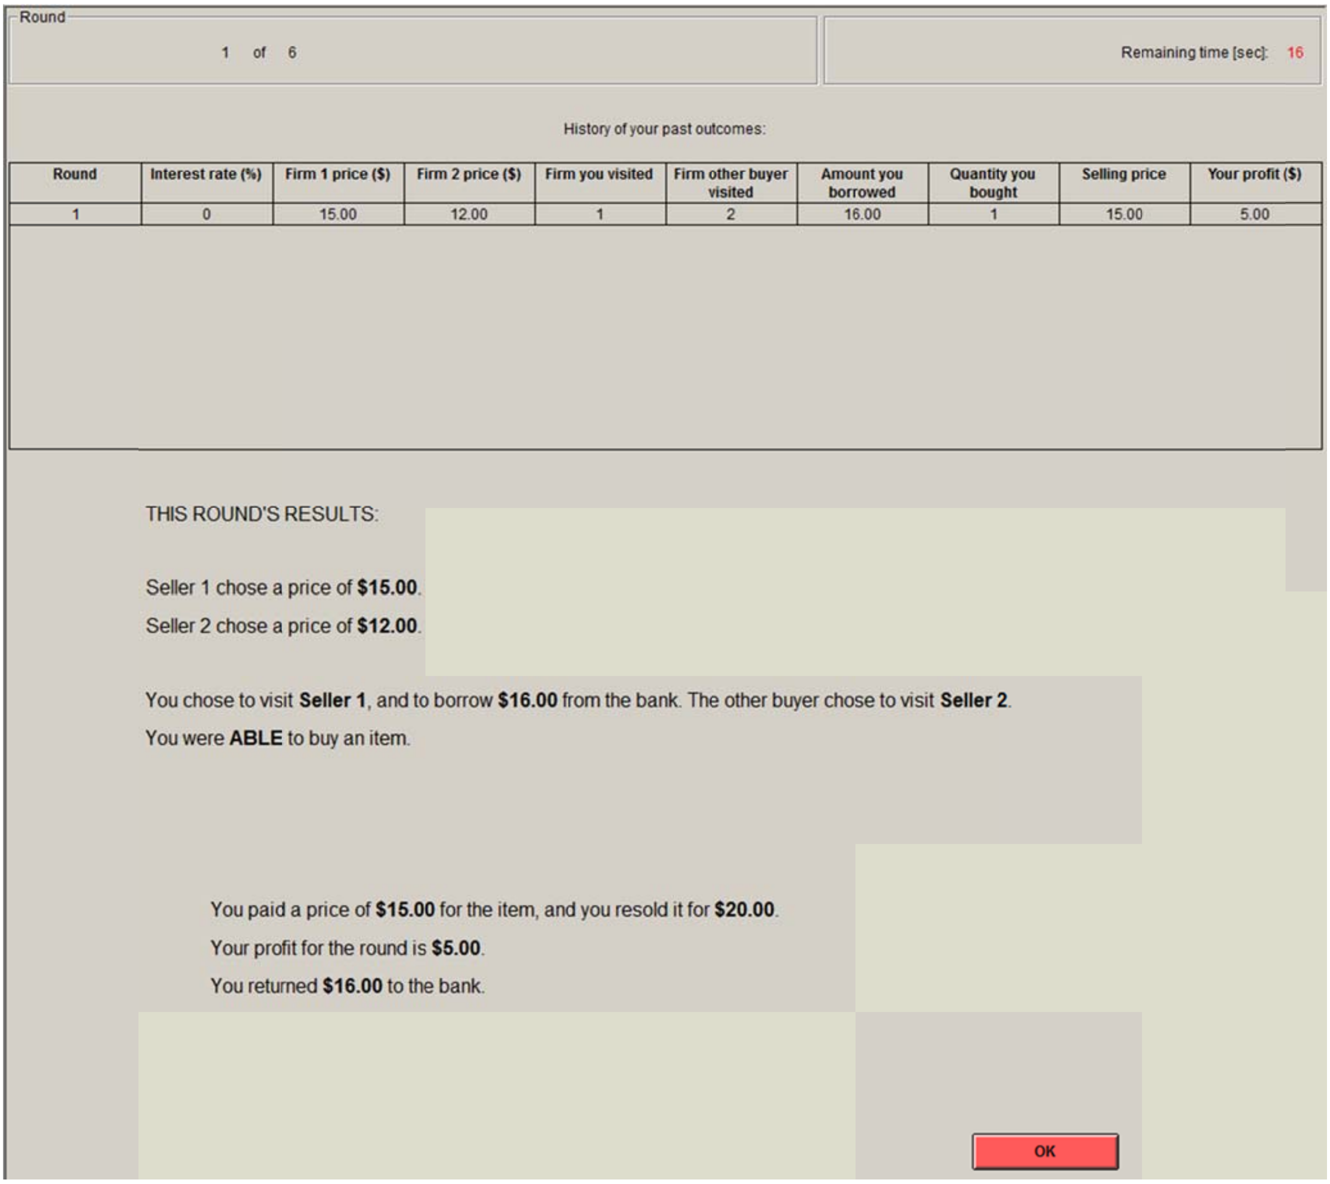

Round

1 of 6

Remaining time [sec]: 16

History of your past outcomes:

Round	Interest rate (%)	Firm 1 price ($)	Firm 2 price ($)	Firm you visited	Firm other buyer visited	Amount you borrowed	Quantity you bought	Selling price	Your profit ($)
1	0	15.00	12.00	1	2	16.00	1	15.00	5.00

THIS ROUND'S RESULTS:

Seller 1 chose a price of **$15.00**.

Seller 2 chose a price of **$12.00**.

You chose to visit **Seller 1**, and to borrow **$16.00** from the bank. The other buyer chose to visit **Seller 2**.

You were **ABLE** to buy an item.

You paid a price of **$15.00** for the item, and you resold it for **$20.00**.

Your profit for the round is **$5.00**.

You returned **$16.00** to the bank.

OK

E Can risk aversion or loss aversion explain buyer visit choices?

Two features apparent in Figure 5 (showing predicted and observed buyer visit choices) are (1) fairly substantial under–responsiveness to price in the 2x2 market, and (2) slight under–responsiveness in the 2x3 market. We look here at two potential explanations for these results: risk aversion and loss aversion.

Both risk aversion and loss aversion, from an intuitive standpoint, seem capable of explaining under–responsiveness to price. Consider the two buyers in a 2x2 market, and suppose they face a pair of prices low enough that staying home is dominated, and unequal but close enough together that the symmetric equilibrium in visit choices involves mixed strategies. In these circumstances, each buyer must choose between two lotteries. For example, for Buyer 1, visiting Seller 1 yields a large prize (the profit from buying from Seller 1) or a small prize (the non–positive profit from being unable to buy), with the latter's probability equal to the probability of Buyer 2 also visiting Seller 1, times one–half (the probability Buyer 2 is randomly chosen to buy in case both visit the same seller).

It is easy to show that under these conditions, and with any utility function that is increasing in money payment, the probability of visiting the low–priced seller will be greater than that of visiting the high–priced seller. This means that visiting the low–price seller means a higher potential profit (if the buyer is able to buy), but a lower probability of getting that profit. The high–price seller offers a lower potential profit, but a higher chance of getting it. Under risk aversion (or more precisely, decreasing marginal utility of money), the relative benefit of the higher potential profit from the low–price seller decreases, so that – other things equal – the high–price seller becomes more attractive, compared to the risk neutral case. Then the probability of choosing the high–price seller must adjust upwards to keep buyers indifferent between them in a symmetric equilibrium.

Now consider a buyer in the same situation who is loss averse (Kahneman and Tversky, 1979) – that is, he dislikes losses more than he likes equal–sized gains – but otherwise does not avoid risks. Under a positive inflation rate, the profit from being unable to buy will be strictly negative (rather than nil under zero inflation). Faced with the trade–off we've discussed (low probability of a high profit versus high probability of a low profit), a loss–averse buyer will be more sensitive to the higher probability of a loss when visiting the low–price seller, so – other things equal – the high–price seller becomes more attractive, compared to the loss neutral (i.e., expected–profit maximising) case. Again, the probability of choosing the high–price seller must adjust upwards.

In Figure 7, we illustrate these intuitive arguments with the use of reliability diagrams similar to the ones in Figure 5. Like the earlier figure, these diagrams concern buyer visit probabilities: in particular, the probability of visiting Seller 1. (The results are nearly identical for the other sellers.) Unlike the earlier figure, though, we are not comparing theoretical probabilities with observed frequencies. Instead, we compare theoretical probabilities under risk neutrality with theoretical probabilities under a particular model of risk aversion. That makes these diagrams directly comparable to those in Figure 5, in the sense that if all of the real buyers had the utility function assumed in one of the diagrams in Figure 7, they would show the same degree of responsiveness to the predicted probability in both figures.

The top–left panel of Figure 7 shows the OLS trend lines that would obtain in the 2x2 and 2x3 markets (pooling over inflation rates) if all buyers had the utility function $u(x) = \frac{1}{1-\alpha}(10 + x)^{1-\alpha}$, where x is money profit and $\alpha = 0.25$; this is essentially a constant–relative–risk–aversion utility function with coefficient α, except for the addition of 10 to profit.[23] The next four panels use the same functional form for utility, but different values of α: 0.5, 1 (i.e., $u(x) = ln(10 + x)$), 2 and 4; the range from 0 (risk neutrality) to 4 covers the values typically estimated from the lab and the field (e.g., Beetsma and Schotman, 2001; Deck et al., 2008; Harrison and Rutström, 2008; Dave et al.,

[23] The minimum profit a subject can earn in a round is –6, so adding 10 guarantees that utility is defined. Under standard expected–utility theory, the subject's wealth should be in the utility function instead of a constant 10; however, Andersen et al. (2011) report evidence that subjects fail to integrate their experimental income with their wealth outside the lab, and thus in a sense act as if their outside wealth is much lower than it actually is.

2010; Dohmen et al., 2011). The final panel shows the corresponding trend lines for the observed data (the same data as in Figure 5, but pooled over inflation rates).

Figure 7: Buyer visit probabilities – risk neutral vs. CRRA with parameter α (OLS trends, pooled inflation rates)

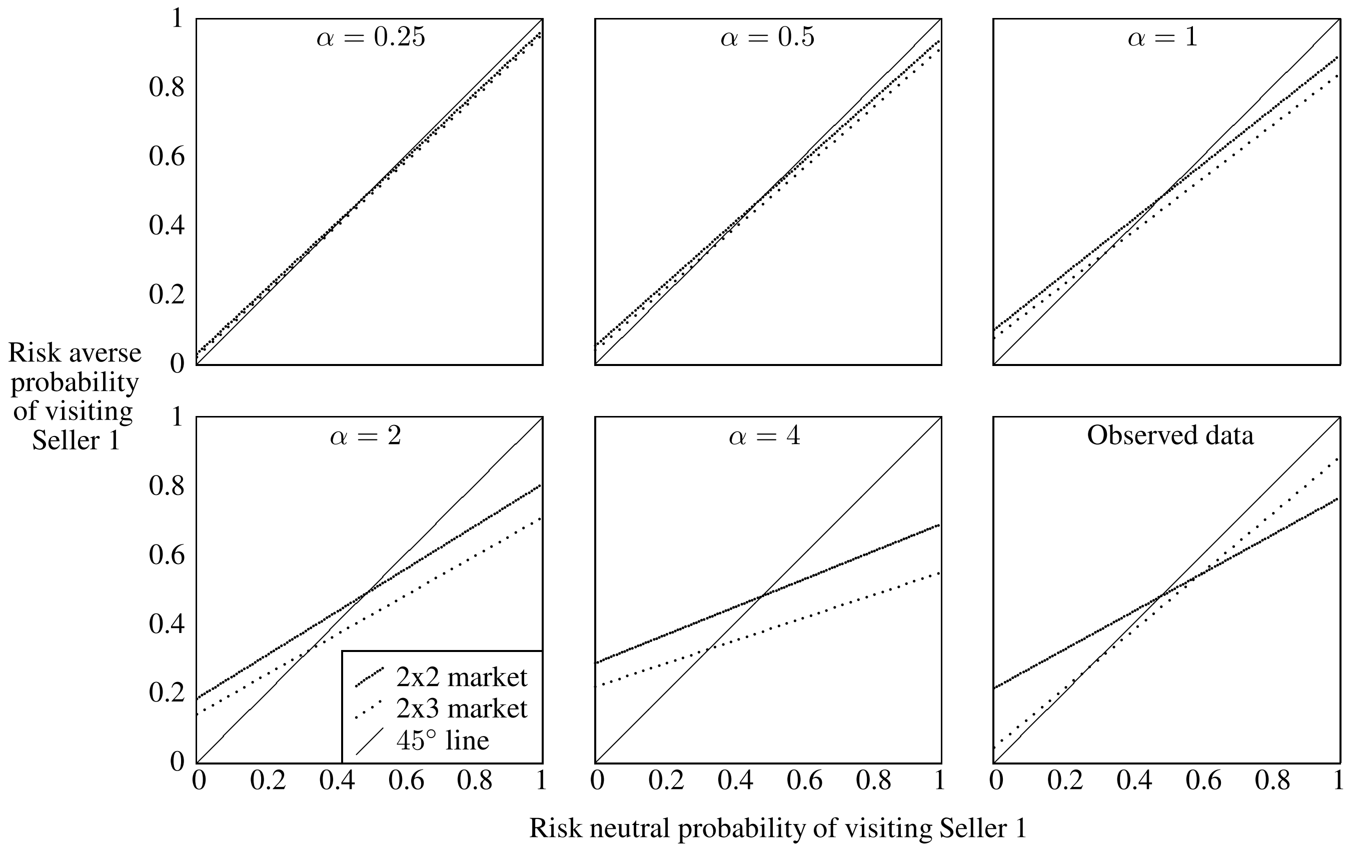

Consistent with intuition, risk aversion leads to less responsiveness to predicted probability, and equivalently less price elasticity; moreover, the effect gets larger as α increases. However, for a given α, the effect is actually slightly *larger* in the 2x3 market than in the 2x2 market, whereas the real subjects showed substantially *less* under–sensitivity to price in the 2x3 market. Thus, risk aversion has at best mixed success in characterising the price under–sensitivity we observe in the experiment.

We move to loss aversion. Figure 8 is similar to Figure 7, but these reliability diagrams are based on subjects who are loss averse. Specifically, they are assumed to have a linear utility–of–money function away from the origin, but a possible kink at the origin, so that the slope for negative values is $\beta \geq 1$ times that for positive values.[24] The top–left panel sets $\beta = 3$, a value in the neighbourhood of those commonly estimated from individual decisions (e.g., Tversky and Kahneman, 1991; Camerer, 2005; Abdellaoui, Bleichrodt and L'Haridon, 2008). The next two panels use higher values of β (6 and 12), while the last panel again shows the observed data.

As with risk aversion, loss aversion has mixed success in characterising the price under–sensitivity we observe in the experiment. As our intuition suggested, loss aversion leads to lower price elasticity, and it decreases further as β increases. Moreover, holding β constant, the size of the effect is larger in the 2x2 market than in the 2x3 market, as we had seen in the experiment. However, while loss aversion captures the qualitative effects seen in real buyer behaviour, it performs poorly in a quantitative way. Even extreme levels of loss aversion ($\beta = 12$) entail a degree of under–responsiveness substantially less than what was actually observed. Thus, while loss aversion has more success

[24]Loss aversion is one part of Kahneman and Tversky's prospect theory; other parts include diminishing marginal sensitivity to both gains and losses, and non–linear weighting of probabilities. Though the parts of prospect theory are often taken together, there is no logical reason why they need to be, and it is certainly true that loss aversion on its own does not entail any of the other parts. We assume in this exercise that subjects are loss averse, but we leave out the other parts of prospect theory.

Figure 8: Buyer visit probabilities – loss neutral vs. loss averse with parameter β (OLS trends, pooled inflation rates)

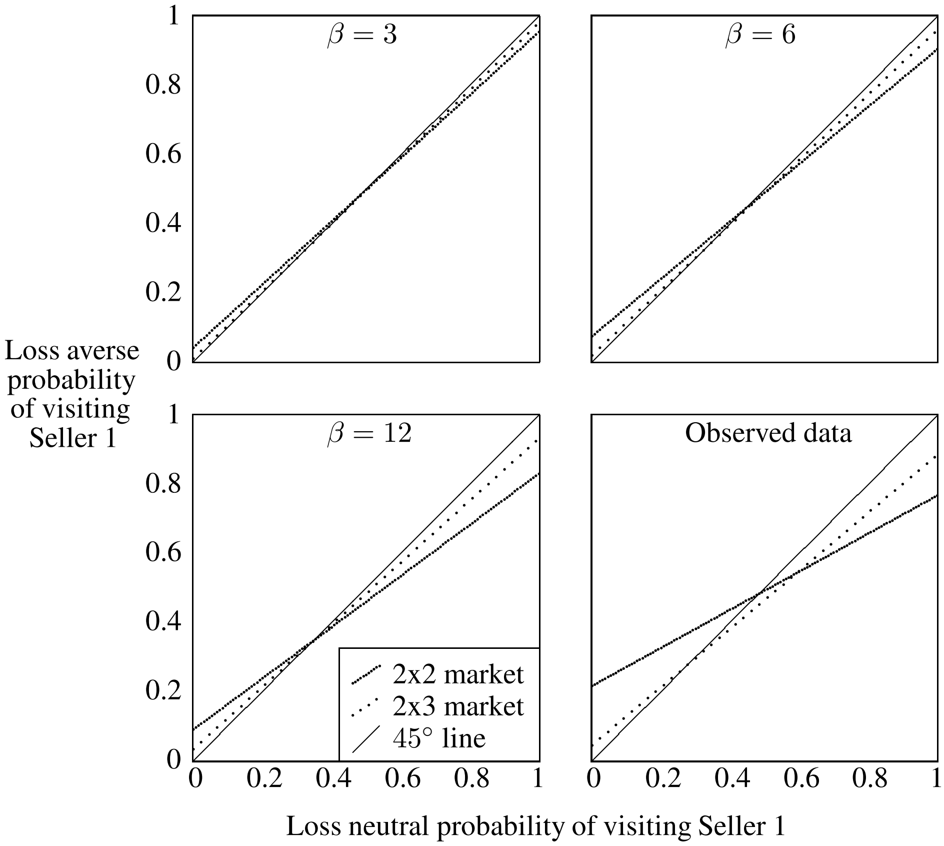

than risk aversion in explaining these results, it still falls a bit short.