How Inflation Affects Macroeconomic Performance: An Agent-Based Computational Investigation*

Quamrul Ashraf[†]

Boris Gershman[‡]

Peter Howitt[§]

Williams College

American University

Brown University

and NBER

*Useful comments were received from two anonymous referees, workshop participants at the Lorentz Center in Leiden, and seminar participants at the OFCE in Sophia Antipolis, the Federal Reserve Board, and the National Bank of Austria. The C++ code for our model can be downloaded at http://www.econ.brown.edu/fac/Peter_Howitt/working/Inflation.zip. It was compiled and run as a 64-bit Windows application. The output files sufficient to replicate the results reported in the paper can be downloaded at http://www.econ.brown.edu/faculty/Peter_Howitt/working/Inflation_results.zip. Ashraf acknowledges research support from the Hellman Fellows Program. Howitt acknowledges research support from the NSF (SES-0617869). This paper was partly written while Ashraf was visiting Harvard Kennedy School and the Center for International Development at Harvard University.

[†]Department of Economics, Williams College, 24 Hopkins Hall Drive, Williamstown, MA 01267, USA; e-mail: Quamrul.H.Ashraf@williams.edu.

[‡]Department of Economics, American University, 4400 Massachusetts Avenue NW, Washington, DC 20016, USA; e-mail: boris.gershman@american.edu.

[§]Department of Economics, Brown University, 64 Waterman Street, Providence, RI 02912, USA; e-mail: Peter_Howitt@brown.edu. Please address academic correspondence to this author.

Running head: INFLATION AND MACROECONOMIC PERFORMANCE

Please send page proofs to the following author (in this case, different from the “corresponding” author after the article appears in print):

Quamrul Ashraf

Department of Economics

Williams College

24 Hopkins Hall Drive

Williamstown, MA 01267

USA

e-mail: Quamrul.H.Ashraf@williams.edu

Phone: 1 (401) 499-3672

Fax: 1 (413) 597-4045

Abstract: We use an agent-based computational approach to show how inflation can worsen macroeconomic performance by disrupting the mechanism of exchange in a decentralized market economy. We find that, in our model economy, increasing the trend rate of inflation above 3 percent has a substantial deleterious effect, but lowering it below 3 percent has no significant macroeconomic consequences. Our finding remains qualitatively robust to changes in parameter values and to modifications to our model that partly address the Lucas critique. Finally, we contribute a novel explanation for why cross-country regressions may fail to detect a significant negative effect of trend inflation on output even when such an effect exists in reality.

Keywords: Agent-based computational model, trend inflation, price dispersion, firm turnover

1 Introduction

Inflation-targeting central banks frequently justify their narrow focus on inflation by arguing that maintaining a low and stable rate of inflation is the best way for them to promote high employment and output. Yet, neither mainstream economic theory nor existing empirical studies offer much support for the belief that a country's real economic performance is significantly improved by reducing the trend rate of inflation, except in extreme circumstances. Indeed, Fortin (1996) and Akerlof et al. (1996) argue that reducing inflation too close to zero worsens economic performance because of downward nominal wage rigidity. Moreover, too low a rate of inflation can result in real instability because of the zero lower bound on nominal interest rates.

Given the lack of support from economic science, and because of the dangers posed by nominal wage rigidity and the zero lower bound, inflation-targeting central banks are reluctant to reduce their targets below 2 or 3 percent. However, the commonly acknowledged weakness of the microfoundations of mainstream monetary theory leaves many wondering if there might still be significant costs of inflation that theorists have not yet discovered.

The research reported here goes somewhat beyond mainstream monetary economics to investigate this issue of how economic performance might be enhanced by reducing the trend rate of inflation. The paper explores an idea that has been advanced by economists such as Heymann and Leijonhufvud (1995), namely that inflation might impede the market mechanisms responsible for coordinating economic activities (see also Ragan, 1998). Strictly speaking, this idea cannot be investigated in a conventional rational-expectations-equilibrium framework because the concept of rational-expectations equilibrium presupposes that peoples' actions and beliefs are always perfectly coordinated without the need for any explicit mechanism. We therefore exploit a less conventional methodology, namely that of agent-based computational economics (ACE).¹

More specifically, we use a modified version of the ACE model that we developed in

Ashraf et al. (2011), henceforth AGH (2011), for addressing the role of banks in an economic system. In this model economy, all trades are coordinated by a self-organizing network of trade specialists, as they are in reality by shops, brokers, middlemen, banks, realtors, and so forth. Here, we simplify the model by ignoring the existence of banks. The model is calibrated to U.S. data and simulated many times under different target inflation rates to see what difference the inflation target makes to the average level and variability of GDP, the variability of inflation, and the average rate of unemployment.

The basic reason for thinking that this model may help to account for significant costs of inflation is that it provides a propagation mechanism through which inflation might impede the market processes that coordinate economic activity. The mechanism involves the entry and exit of business firms that collectively constitute the network of trade specialists. Because price changes are not perfectly coordinated, a higher trend rate of inflation should imply, on average, more dispersion in the relative prices offered by the different specialists. Such extraneous relative-price variability introduces more volatility into the demand faced by any individual specialist, thus raising the incidence of shop failures throughout the economy. As a result, abstracting from any effects of wage rigidity and the problem of the zero lower bound on nominal interest rates, increased trend inflation should reduce average GDP by impeding the transactions process.

Our main finding from this investigation is that the trend inflation rate has a powerful adverse effect on economic performance, but only when it rises above 3 percent. For example, an increase of trend inflation from 3 to 10 percent reduces the average annual level of GDP in the median simulation by 6.8 percent. It has qualitatively similar effects on the other macroeconomic indicators that we examine.

Our paper demonstrates how an ACE approach can deal with questions not readily addressed by conventional equilibrium analyses, and also how to tackle some common objections to ACE methodology. One such objection is that simulations can be sensitive to numerical values assigned to parameters. We handle this objection, first, by calibrating our

parameter values to U.S. data and, second, by performing a sensitivity analysis, which shows that our main finding is qualitatively robust to changes in parameter values.

Another common objection to ACE is that it is subject to the “Lucas critique.” In our context, when trend inflation rises, one would expect individual behavior to adjust to this change in the environment, but a model with exogenous behavioral rules does not allow as much adaptation as a rational-expectations analysis would. To deal with this critique, we first use our sensitivity analysis to identify which behavioral parameters are most important for our results. Specifically, we find that the markup applied by a firm when setting its retail price is particularly critical – the effect of inflation is greatly magnified by having a small markup. We then modify our model to allow markups to adapt to the environment through the selection mechanism implicit in the firm exit process. We find that the economy-wide average surviving markup ends up being higher when inflation is higher, but we also find that our main result remains intact when we take this endogenous behavior of markups into consideration.

One question raised by our numerical results concerns the absence of conclusive evidence from existing empirical studies, which, for the most part, have failed to find a sizable and robust effect of long-run inflation on a country’s GDP. Interestingly, we show that if the baseline calibrated model were true, then there would still be a 15 percent chance that a cross-country regression of log GDP on inflation would reveal a statistically insignificant effect. The reason is that our model implies an inherent endogeneity issue, which arises from a positive Phillips-type relationship between output and inflation that tends to mask the underlying negative long-run causal effect of trend inflation on output.

The paper is organized as follows. In Section 2, we briefly survey some of the relevant literature on the real effects of inflation. Section 3 lays out the basic structure of our agent-based model and discusses its workings. Section 4 reveals our main results and those from our sensitivity analysis. In Section 5, we discuss our technique for dealing with the Lucas critique. Section 6 shows the implications of the model for cross-country regressions of output

on inflation, and Section 7 concludes.

2 Literature on the real effects of inflation

There is an extensive literature on the potential real effects and costs of inflation. One strand of the literature has focused on effects that work through the opportunity cost of holding non-interest-bearing money. Raising this opportunity cost can reduce output by lowering capital accumulation and labor supply (Stockman, 1981; Cooley and Hansen, 1989; Gomme, 1993) or by altering search intensity (Rocheteau and Wright, 2005; Lagos and Rocheteau, 2005). It is hard to make the case, however, that such effects are empirically important given that non-interest-bearing money is such a small item in people's portfolios. Accordingly, Chari et al. (1995) report a very modest impact of inflation on the steady-state growth rate of output in a variety of endogenous growth models where the opportunity cost of holding money is the only channel through which inflation is allowed to matter.

Another line of research has argued that non-indexation of private accounting systems distorts savings decisions and the allocation of capital under high rates of inflation and, moreover, that inflation diverts talent away from real productive activity and into less socially productive financial activity designed simply to cope with or profit from changes in the value of money (Leijonhufvud, 1977; Howitt, 1990; Summers, 1991). These possible costs of inflation are difficult to quantify, and we are not aware of any attempts to estimate their effects on the path of real output. On the other hand, Feldstein (1999) has argued that reducing inflation from 2 percent to 0 would raise output by 4 to 6 percentage points because of the non-indexation of capital income taxation. However, this effect seems more like a problem to be dealt with by tax authorities than central bankers.

In New Keynesian dynamic stochastic general equilibrium models, trend inflation affects output mainly through its effect on the equilibrium degree of price dispersion. This is especially true in those versions that use the Calvo (1983) price-setting model.² The rise

in price dispersion due to inflation can also have real effects by increasing the amount of resources devoted to search on the part of buyers, as emphasized in the search-theoretic literature (Benabou, 1992; Head and Kumar, 2005). It will also play a crucial role in our framework, which highlights the detrimental effect of rising price dispersion on business failures.

On the empirical side, there is virtually no robust evidence on the relationship between inflation and real outcomes. The literature on the effects of trend inflation on the growth rate of output, as surveyed in Ragan (1998), has produced little evidence of a significant impact, except in extreme circumstances. There is also almost no direct evidence of an effect of trend inflation on the level of real output as opposed to its growth rate or volatility (Bullard and Keating, 1995). Later in the paper, we provide a novel explanation for why it may be hard to detect a significant negative effect of trend inflation on output.

Despite the absence of robust reduced-form evidence on the real effects of inflation, there is considerable empirical support for the relationship between inflation and relative-price dispersion, which plays a major role not only in New Keynesian and search-theoretic models but also in our agent-based framework. For example, Parks (1978), Parsley (1996), Debelle and Lamont (1997), and Lastrapes (2006) document a strong positive association between inflation and price dispersion in U.S. data.³

3 The model

Our model is a special case of the one developed in AGH (2011). Here, for simplicity, we abstract from the banking sector. The model attempts to portray, in an admittedly crude form, the mechanism by which economic activities are coordinated in a decentralized economy. It starts from the proposition that, in reality, almost all exchanges in an advanced economy involve a specialized trader (“shop owner”) on one side or the other of the market. On this foundation, we have built a model whose structure and macroeconomic aggregates are

comparable to the canonical New Keynesian analysis of Woodford (2003). That is, prices are set by competing firms acting under monopolistic competition, the rate of interest is set by a monetary authority following a Taylor rule, and consumer demands depend, *inter alia*, on current wealth. However, it is quite different from the standard New Keynesian framework in three important respects. First, we have introduced elements of search in both goods (retail) markets and labor markets. Second, instead of having a fixed population of firms, we assume that firms are subject to failure and that the process of replacing failed firms is a costly one that disrupts established trading patterns. Third, rather than assuming complete contingent financial markets, we assume that the only available financial instruments are government-issued money and bonds.

To simplify the model and abstract from the issue of expectations formation, we assume throughout that peoples' expectations of inflation are firmly anchored by the central bank. Thus, expected inflation is always equal to the announced inflation target.⁴

3.1 The conceptual framework

There is a fixed number N of people, a fixed number n of different durable goods, and the same number n of different types of labor. Labor of type i can be used only to produce good i . Time is discrete, indexed by “weeks” $t = 1, \dots, T$. There are 48 weeks per “year.” In addition to the n goods, there are two nominal assets: fiat money and bonds. A bond is a promise by the government to pay one unit of money (“dollar”) next week.

Each person has a fixed type (i, j) , where $i \neq j$ and $i \neq j + 1 \pmod{n}$, meaning that each week he is endowed with one unit of labor of type i (his “production good”) and can eat only goods j and $j + 1 \pmod{n}$ (his primary and secondary “consumption goods”). There is exactly one person of each type. Thus, the population of the economy is $N = n(n - 2)$. In everything that follows, we set the number of goods n equal to 50, implying a population of $N = 2400$ people.

Because no person can eat his own production good, everyone must trade to consume.

Trading can take place only through facilities called “shops.” Each shop is a combined production/trading operation. There are n different types of shops. A shop of type i is capable of buying type i labor with money, converting type i labor into good i , and selling good i for money. The number of shops of each type will evolve endogenously.

To trade with a shop, a person must form a trading relationship with it. Each person may have a trading relationship with at most one shop that deals in his production good, in which case that shop is his “employer” and he is one of its “employees,” and at most one shop that deals in each of his consumption goods, in which case that shop is one of his “stores” and he is one of its “customers.” Each person’s trading relationships will evolve endogenously.

Each shop of type i has a single owner whose production good is i . Operating the shop entails a fixed overhead cost of F units of type i labor per week and a variable cost of one unit of type i labor per unit of good i produced. All trade takes place at prices that are posted in advance by the shop. Specifically, each shop posts a retail price and a weekly wage rate, each of which may be adjusted periodically.

There is no depreciation or other physical storage cost. Goods produced but not sold in a week are kept in inventory. Former shop owners continue to hold their “legacy inventories” until these are sold in special “firesale” markets.

The government does not purchase goods or labor but it does issue money and bonds, and it services the interest on bonds through a sales tax on every retail transaction. It adjusts the ad valorem tax rate τ once per year. The central bank pegs the interest rate i on bonds by open market operations. It adjusts this rate every 4 weeks according to a Taylor rule. In addition, the central bank is engaged in the forecasting of future GDP, inflation, and interest rates.

3.2 Protocol and behavioral rules

Each week our model economy proceeds through the same sequence of actions consisting of nine stages: 1) entry; 2) search and matching; 3) bond market trading; 4) labor and goods market trading; 5) monetary policy; 6) match breakups; 7) fiscal policy; 8) exit; and 9) wage and price setting. We believe that each of these 9 stages represents a process that is of central importance to the functioning of a market economy. Our underlying strategy is to model each process as simply as possible, that is, with unsophisticated yet plausible behavioral rules that maintain the spirit of ACE modeling, while attempting to be roughly consistent with mainstream macroeconomics. The rest of this sub-section sketches the protocol and behavioral rules of the model. We refer the reader to AGH (2011) for the exhaustive details of the algorithm that implements the model, further discussion of the behavioral rules, and the way to incorporate a banking sector in the model.

Entry. In the first stage, each person who is not already a shop owner (that is, has another employer or is unemployed) becomes a potential entrant (“entrepreneur”) with a fixed probability θ/N , where the parameter θ represents the supply of entrepreneurship or the weekly frequency of innovation. Each entrepreneur goes through the following process to decide whether to enter.

First, the entrepreneur formulates a business plan for his prospective shop, which consists of a wage rate w and a sales target y^{trg} . He sets the initial wage rate equal to the previous week’s (publicly known) economy-wide average wage, adjusted for expected inflation over the prospective wage contract period of Δ weeks, where Δ is a fixed parameter. He then picks his sales target y^{trg} from a uniform distribution over $[1, n]$ and plans to post the retail price $p = (1 + \mu)w/(1 - \tau)$, which is a fixed percentage markup μ over after-tax marginal cost.⁵ These choices imply a profit flow Π that the firm would earn if it opened and if its sales target were realized.

Given this business plan, the entrepreneur will allow the opportunity for entry to lapse

if his financial wealth A (the sum of his money and bond holdings) is not enough to pay for the fixed cost of operating the shop during the first month. He will also allow the entry opportunity to lapse if his prospective profit Π is less than his currently estimated permanent income, which is discussed further below.

The entrepreneur also conducts “market research” before deciding whether to enter. In particular, he sends invitations to two people, one a potential employee and the other a potential customer, to form trading relationships with him if the shop opens. The potential employee accepts if the shop’s wage w exceeds his current “effective wage,” and the potential customer accepts if the shop’s planned retail price p is less than his current “effective price.”⁶ If either invitee declines, the entrepreneur allows the entry opportunity to lapse.

If he passes all these tests, the entrepreneur opens a shop with an “input target” equal to $x^{\text{trg}} = y^{\text{trg}} + F + \lambda_I (y^{\text{trg}} - I)$, where λ_I is the weekly inventory adjustment speed and I is current inventories, which, for a new shop, are just equal to the entrant’s legacy inventories. Implicit in this equation is a desired inventory level equal to one week’s sales.

Search and matching. Next, each person is given an opportunity to search for possible trading relationships. Each person who is not a shop owner engages in job search with probability σ . Job search consists in asking one “comrade” (someone with the same production good) what his effective wage is. If this exceeds the searcher’s current effective wage, then the searcher will ask the comrade’s employer for a job. The employer will take on the new worker if his last period’s labor input does not exceed his current input target.

Store search is undertaken by every person. The searcher first asks a “soulmate” (someone with the same consumption goods) for his effective retail prices, and then asks a randomly chosen shop if it has a posted price for one of the searcher’s consumption goods. For each of his consumption goods, the searcher decides to switch, to either the soulmate’s shop or to the randomly located shop, or decides to stay with his current shop, depending on which choice yields the lowest effective price.

Bond market trading. Next, everyone goes to the bond market. First, the government redeems all bonds. At this point, all the financial wealth A of any transactor is in the form of money. Then, each person decides on this week's planned consumption expenditures E . He first adjusts his permanent income Y^p according to the adaptive rule $Y^p = \bar{Y}^p + \lambda_p (Y - \bar{Y}^p)$, where Y is actual income (either wages or profit) and $\bar{Y}^p$ is permanent income from the previous period, and λ_p is the weekly permanent income adjustment speed. He then adjusts permanent income for expected weekly inflation. Finally, he sets planned consumption equal to a fixed fraction v of total wealth.⁷ Total wealth is A plus any legacy inventories plus the expected present value of his estimated permanent income, $V \cdot Y^p$, where V is a capitalization factor equal to the present value of a perpetual weekly (real) income stream of one dollar in today's prices. The capitalization factor is calculated by the central bank, based on its projected inflation and nominal interest rates, and is announced to the public on a monthly basis.⁸

People then adjust their portfolios. Someone who does not own a shop retains E in the form of money and invests $A - E$ in bonds. However, if $E > A$, he sets his money holdings equal to A . A shop owner goes through a similar procedure except that his desired money holding is the amount needed to cover not only his planned consumption expenditure but also his target wage bill, which is his current posted wage rate times his input target x^{trg} .

Labor and goods market trading. This stage starts with trade in the firesale markets, where inventories can be purchased at a discounted price from former shop owners. Next, each person in turn engages in labor market trading (with his employer) and goods market trading (with his stores). Labor market trading proceeds as follows. If the person has an employer with positive money holdings, he offers to trade his endowment in exchange for the effective wage $w^{\text{eff}} = \min\{w, M\}$, where w is the employer's posted wage and M is the employer's money holdings just prior to this trade. The employer accepts the offer unless the shop's labor input already exceeds its input target and the ratio of inventory to sales

target exceeds a critical threshold value $IS > 1$.

Goods market trading happens in the following manner. First, the customer learns the price p_s currently posted by each of his two stores ($s = 1, 2$). Then, when a person visits a store with positive inventory I , he can place an order for some amount c_s , subject to a cash-in-advance constraint. The store then sells the person an amount $c_s^{\text{eff}} = \min\{c_s, I\}$. The person's effective price for that good is defined as $p_s^{\text{eff}} = p_s c_s / c_s^{\text{eff}}$. However, if $I = 0$ or the person does not have a store for that good, the person's effective price is set to infinity.

Each customer chooses his desired consumption bundle (c_1, c_2) to maximize his utility function $u(c_1, c_2) = c_1^{\varepsilon/(\varepsilon+1)} + c_2^{\varepsilon/(\varepsilon+1)}$ with a "demand parameter" $\varepsilon > 0$, subject to his budget constraint. If he has established relationships with stores for both of his consumption goods, the budget constraint is $p_1 c_1 + p_2 c_2 = E$. However, if he has a relationship with only one store, he spends the entire amount E on that store's good.

Monetary policy. Every fourth week, the central bank sets the rate of interest i according to the Taylor rule, $\ln(1+i) = \max\{\ln(1+i^*) + \gamma_\pi [\ln(1+\pi) - \ln(1+\pi^*)] + \gamma_y [y - \tilde{y}], 0\}$, where γ_π and γ_y are fixed coefficients, $1+\pi$ is the inflation factor over the past 12 months, π^* is the fixed inflation target, y is the current 3-months moving average for the weekly average log GDP, $\tilde{y}$ is the central bank's evolving estimate of weekly log potential output, and $i^* \equiv r^* + \pi^*$, where r^* is the evolving target for the long-run real interest rate.

Since the central bank cannot be presumed to know the economy's natural interest rate or potential output, it must estimate them adaptively. Accordingly, it adjusts r^* up or down depending on whether the current inflation rate exceeds or falls short of the inflation target π^* , with an adjustment speed η_r . It also estimates $\tilde{y}$ as the long-run expected value of y under an AR(1) process whose parameters are re-estimated right after r^* is adjusted. As mentioned previously, the central bank also calculates the capitalization factor required by people for planning their consumption. To this end, it estimates an AR(1) process for inflation and feeds the predictions of its estimated AR(1) processes for y and π into its Taylor

rule to obtain the nominal interest rate projections.⁹

Match breakups. Established trading relationships may break up randomly. In particular, each person in the economy who does not own a shop is subjected to a probability δ of quitting the labor and goods markets, which entails the unconditional severance of all current trading relationships by the person with his employer and his stores.

Fiscal policy. Next, once a year, and in the last week of the year, the retail sales tax rate τ is adjusted. The government sets the tax rate equal to a value τ^* , which is the value that would leave the debt-to-GDP ratio undisturbed in the no-shock equilibrium outlined in the next sub-section, plus an adjustment factor that is proportional to the difference between the actual and the target debt-to-GDP ratio b^* , with a fiscal adjustment coefficient λ_τ .

Exit. Now, each shop has an opportunity to exit. First, each shop exits for exogenous reasons with probability δ . Second, each remaining shop can choose voluntarily to exit if the shop owner is “hopeless,” that is, if his financial wealth A is not enough to pay for the coming week’s fixed overhead cost. Third, each unprofitable shop exits with probability ϕ . In computing profitability, the shop owner takes into account the opportunity cost of his labor services, that could be earning a wage, and the interest-opportunity cost of maintaining the shop’s inventory. Once a shop exits for any reason, all of its trading relationships (with both employees and customers) are dissolved.

Wage and price setting. Finally, each shop has an opportunity to update its posted wage and retail price. The shop first sets its sales target equal to the current week’s actual sales and adjusts its input target accordingly. Then, it proceeds to update its wage but only if the previous contract period of Δ weeks has just ended. Its wage adjustment obeys a firm-level expectations-augmented Phillips curve, that is, $w = \bar{w} \cdot (1 + \beta \cdot x) \cdot (1 + \pi^*)$, where $\bar{w}$ is the pre-existing wage, π^* is the inflation rate expected over the next contract period, β is a fixed adjustment parameter, and x is a measure of the shop’s excess demand for labor, namely the average percentage deviation of its input target from the amount of labor it has been

offered in the trading process over the previous contract period.

The shop has the opportunity each week to change its retail price. Following Blinder et al. (1998) we assume that price adjustments are mainly due to changes in costs. Specifically, the shop will update its price, so as to maintain a constant markup, whenever there is a change in its marginal cost, $w/(1 - \tau)$, because of a change in either its wage w (which happens every Δ weeks) or in the tax rate τ (which happens once a year). Motivated by the fact that most price changes are temporary in nature (Bils and Klenow, 2004), we also assume that the shop can make price adjustments based on its inventory-to-sales ratio. In particular, when that ratio exceeds the fixed upper threshold $IS > 1$, the price is reduced by a factor $\delta_p > 1$, and when the ratio falls below the lower threshold $1/IS$, the price is increased by this same factor. Unless there is a change in marginal cost or a crossing of the inventory-to-sales ratio beyond either threshold, the shop will leave its price unchanged.

3.3 The workings of the model

As the preceding discussion has made clear, all shocks in this economy are individual shocks. Nevertheless, the individual shocks that cause matches to break up and shops to enter or leave particular markets do have aggregate consequences because there is only a finite number of agents. So, in general, the economy will not settle down to a deterministic steady state unless we turn off these shocks. If we do turn off all shocks, there is a deterministic equilibrium that the economy would stay in if left undisturbed by entry and breakups, with wages being adjusted each week ($\Delta = 1$). This equilibrium will serve as an initial position for all the experiments we perform on the model below.

The equilibrium is one in which all the potential gains from trade are realized. Each person is matched with one employer and two stores. There are n shops, one trading in each of the n goods. To preserve symmetry across goods, we suppose that each good is the primary consumption good for exactly one shop owner. Each shop begins every week with actual and target input both equal to $n - 2$, which is the number of suppliers of each

good, and with actual and target sales equal to inventory holdings equal to actual output, $n - 2 - F$. Thus, the economy's total output equals full capacity, $y^* = n(n - 2 - F)$.

In AGH (2011), we characterize the rest of this equilibrium, which will repeat itself indefinitely, with all nominal magnitudes – money and bond holdings, actual and effective wages and prices, and permanent incomes – rising each week at the constant target rate of inflation, provided that the fixed cost of operation F is small enough that shops always pass the profitability test during the exit stage.

Inflation has almost no effect in this flexible-price equilibrium without any shocks. There will be a minor effect on the equilibrium tax rate τ because, with higher inflation, the government is collecting more seigniorage and, hence, can maintain its target real debt level with a lower sales tax. However, aggregate output and employment will be unaffected, the real rate of interest will remain equal to the rate of time preference, and the volatility of output and inflation will remain equal to zero.

When the shocks are turned on, the economy ceases to track full employment and full capacity. In particular, GDP falls whenever a shop that was satisfying some consumers goes out of business or a customer loses a store because of a random breakup. GDP also falls whenever a new shop enters and diverts workers from old shops that were satisfying some customers, because some of these workers' efforts will now be used up in defraying the fixed cost of the new shop rather than producing goods that can be consumed by customers of the old shop.

These events that reduce GDP are constantly being offset to some degree by the entry of new shops that are able to satisfy customers in markets where there had previously been no viable shop, and by the exit of shops that were using up fixed costs but not producing enough to satisfy their customers. Thus, both entry and exit are critical to the system's ability to approximate full-capacity utilization. However, although entry of new shops is useful in markets where there are no incumbents, or where the incumbents are not hiring all the potential workers because of layoffs or because of financial problems that prevent

them from meeting their payroll, entry can be harmful (to the level of real GDP) in cases where incumbent shops were hiring most of the potential workers and satisfying most of the potential customers. Likewise, although exit is important in cases where the shop has ceased to play an active intermediation role, whether because of financial difficulties, a surfeit of inventories, or a too high markup, exit can also be very harmful in cases where the incumbent was previously doing well, because it can cascade across markets causing a cumulative loss of output (Howitt, 2006).

The effects of the shocks will tend to be larger when inflation is higher because inflation introduces an extra source of dispersion in relative prices, given that individual prices are changed at random times, and it does so by a larger percentage amount the higher is the trend inflation rate.¹⁰ This dispersion in relative prices induces variability in each shop's sales and, hence, in profitability, and the concomitant widening of the distribution of profitability ends up putting more firms near the edge of failure, making it more likely for firms to fail due to a random shock to innovation or a match breakup. The bigger is the fixed cost and the smaller is the markup, the more likely is such a shock to trigger the failure of the shop. Thus, a higher target rate of inflation intensifies the propagation of shocks through the failure of trading enterprises. To explore this mechanism in more detail, however, we need first to calibrate the model and then simulate it under different target inflation rates.

4 Main results

The 25 parameters of this model were calibrated to U.S. data using the same procedure as in AGH (2011).¹¹ In order to assess the effect of higher inflation on macro performance, we simulated the model while varying the central bank's target rate of inflation. Specifically, we ran an experiment in which, for each integer-valued inflation target from 0 to 10, we created 10,000 artificial 60-year histories of the economy by running the model starting in a neighborhood of the no-shock equilibrium. To eliminate transient effects due to initial

FIGURE 1: Inflation and macroeconomic performance

conditions, we then threw out the first 20 years of each history and considered only the remaining 40 years. For each set of simulations (runs), we recorded the cross-run median of the 40-year average value of each of 5 different macroeconomic indicators: the log output gap, the unemployment rate, the annual inflation rate, the volatility of output, and the volatility of inflation.¹² Hereafter, we refer to this statistic for a given indicator as the “median” value of that indicator.

The results of this experiment are reported in Figure 1. The most salient finding is a monotonic deterioration in performance of all 5 indicators for target inflation rates between 3 and 10 percent per year. Over this range, the median output gap increased by over 7 percent while median unemployment rose by over 6 percentage points. Median output volatility nearly doubled while median inflation volatility rose by about 40 percent. In addition, over the whole range, the median inflation rate rose approximately point for point with the inflation target, indicating that the central bank does typically hit its long-run target in the model.

During some simulations, our model economy crashed, in the sense that the log output gap became infinite in finite time. Further, we observed that the incidence of these crashed runs increased monotonically with the target inflation rate. Specifically, while we did not

FIGURE 2: Statistical significance of the main results

observe any crashes for low inflation targets, the fraction of crashed runs increased from 0.07 percent at the 4 percent target to 17.77 percent at the 10 percent target. Given that we omitted crashed runs from the sample when calculating the cross-run medians, the adverse effect of inflation on economic performance in our model is presumably higher than what is reported in Figure 1.

For low rates of target inflation, that is, between 0 and 3 percent, the results shown in Figure 1 indicate little or no deterioration as inflation rises. Indeed, it seems that the economy performs slightly better at 3 percent inflation than at zero inflation, at least in terms of output and unemployment, although both measures of volatility are virtually constant over this range.¹³ Exactly why there is this small beneficial effect of inflation at very low positive rates is not something we have managed to answer. There is apparently more going on than the zero lower bound on nominal interest rates since, when we suppress that lower bound by allowing the nominal interest rate to become negative, we get almost identical results in terms of our macro indicators between 0 and 3 percent inflation.¹⁴

The deterioration in performance when inflation rises above 3 percent is significant not just economically but also statistically. Figure 2 shows, for each inflation rate, the 95 percent confidence intervals for the output gap and the unemployment rate based on our simulation

sample. The confidence interval bands are narrow, and their bounds are rising monotonically for inflation rate targets between 3 and 10 percent. Similar results were found for all our other macro indicators. Thus, it seems certain that the monotonic deterioration we have found is indeed implied by the model and is not just a result of sampling error from too small a number of simulations.

To check the robustness of our main findings to variations in parameter values, we conducted a sensitivity analysis. We did this one parameter at a time and for all the 25 parameters of the model, except for the target inflation rate π^* . Specifically, for each of these 24 parameters, we redid our experiment assuming that the parameter value was first 25 percent higher than in the baseline calibration and then 25 percent lower. In the case of 5 parameters, the model crashed more than 90 percent of the time at one of these extremes when the inflation target was 10 percent, so we reduced the variation to plus or minus 10 percent. In one of those cases, namely the fixed cost F , there was still more than a 90 percent incidence of crashes, so we reduced the variation to plus or minus 5 percent. Our confidence in our results was increased when we found, without exception, that the monotonic increase in the output gap, unemployment, and output volatility, as the inflation target increases from 3 to 10 percent, survived the variations in the parameter values. The same was largely true of the monotonic increase in inflation volatility.

This sensitivity analysis also helped us to understand which of the parameters in the model are the most critical ones for our main results. Specifically, for each parameter, and for both high and low values, we calculated the difference between the cross-run median output gaps observed under 3 and 10 percent inflation targets. This difference was the largest in the case of the high value for the markup parameter μ , even though we were only considering variations of plus or minus 10 percent for this parameter. This finding will be used in Section 5 below where we address the Lucas critique.

Our interpretation of the deterioration of macro performance when trend inflation rises above 3 percent focuses on the costs of maintaining the trading facilities that are necessary

FIGURE 3: Identifying the mechanism behind the main results

for economic transactions. Specifically, given that price changes are not perfectly coordinated across shops, an increase in inflation should create an artificial increase in the dispersion of relative prices, that is, an increase unrelated to tastes for consuming and technology for producing the wares of different sellers. This will clearly add to the volatility and unpredictability of demand faced by any seller, making it more likely that some sellers will be at risk of failure because of their inability to cover the fixed overhead cost. The result should be a reduction in the number of shops, an increase in the number of potential transactions that go unconsummated for lack of a trading facility, and, hence, a fall in the level of economic activity and a rise in the rate of unemployment.

This interpretation is consistent with the facts shown in Figure 3. Specifically, the left panel of the figure shows that, as the inflation target rises from 3 to 10 percent, there is indeed a monotonic increase in both sales-weighted price and profit dispersion across sellers. This panel also shows that the annual average number of shops in the median run declines monotonically over this range. While this decline is only on the order of 10 percent, it is economically significant in an economy that needs at least 50 shops to function at peak efficiency and has only a 20 percent margin of redundant shops (60 shops in total) when inflation is at 3 percent. The right panel of Figure 3 shows that, with target inflation increasing from 3 to 10 percent, both the job loss rate and the duration of unemployment

FIGURE 4: Breakdown of the main results under price flexibility and no turnover

increase substantially and monotonically, which is what one would expect if the main effect of increased inflation was working through a higher incidence of business failures and a smaller number of shops offering employment.

We sought further corroboration of our interpretation by running a number of related experiments. First, we suppressed the source of price variability by forcing all wages and prices to be reset every week. As shown in the left panel of Figure 4, this did not eliminate unemployment or the output gap, which are still non-zero because of other frictions in the model, but it did cause both indicators to be virtually independent of inflation, suggesting that our main results are indeed attributable to the effects of uncoordinated price changes. Next, we suppressed all turnover in firms, thereby eliminating job loss as a result of shop failures. There was still a monotonic increase in unemployment and the output gap for trend inflation rates higher than 3 percent, but, as shown in the right panel of Figure 4, the increase in each case was much smaller in magnitude than in our main experiment, suggesting that much of the deterioration in the main experiment did indeed come from people losing their jobs when shops failed. Finally, we combined the two previous experiments by suppressing price variability and turnover simultaneously. As a result, both the unemployment rate and the output gap fell to about 1 percent for all rates of inflation, suggesting that these two factors taken together are at the root of our main results.

5 Lucas critique

Our approach thus far leaves us open to the Lucas critique, namely that our predicted effects of a change in the policy environment may be undone by endogenous changes in behavior due to individual or social learning on the one hand and Darwinian selection on the other. In this section, we discuss how one might address such a critique in an agent-based framework, and we illustrate our strategy with respect to the behavioral parameter identified by our sensitivity analysis as being the most critical for our main results – the markup.

In particular, our analysis begs the question of whether an increase in the trend rate of inflation, which affects the economy by putting more firms at risk of failure, might not naturally induce shops to set a larger markup, thereby delivering them more profits in any given situation and, hence, alleviating the increased failure risk. We address this concern here by modifying our model such that shops in a given simulation, rather than entering with identical markup coefficients, now draw their own respective markups from a uniform distribution that remains fixed throughout the simulation. This variant of our model essentially permits the economy-wide average markup in a given simulation to change over time due to the evolutionary selection mechanism that is implicit in the exit process. Furthermore, the long-run outcome of this selection process will be different depending on the policy environment (target inflation rate) of the simulation.

However, in pursuing this strategy, we are left with the difficult task of choosing, for each target inflation rate in our main experiment, the mean of the distribution from which markups are drawn. We decided to let these policy-environment-specific means be determined by internal consistency. Specifically, for each target inflation rate, we used an iterative procedure to select an “equilibrium” value for the mean of this distribution, such that, when new entrants draw from this distribution, the median markup (that is, the median across simulations of the 40-year mean of the economy-wide average markup) ends up being equal to that value. Having thus determined the means of the policy-environment-specific distri-

FIGURE 5: The main results under endogenous markups

butions from which markups are drawn, we then reconducted our main experiment using this alternative version of our model.

The results of our exercise are shown in Figure 5. As is evident from the figure, surviving markups do indeed respond to changes in the policy environment, being systematically higher on average the higher the target inflation rate. Specifically, the median markup rises monotonically from 13 to 20 percent as the target inflation rate increases from 3 to 10 percent. Interestingly, however, as shown in the figure, allowing for this endogenous response of markups does not seem to mitigate the deleterious effects of inflation uncovered by our baseline analysis under fixed markups. For instance, with endogenous markups, as the target inflation rate rises from 3 to 10 percent, the median output gap increases by 7.2 percent while median unemployment rises by 7.7 percentage points. We conclude that, at least with respect to endogenous variations in markups, our estimates of the macroeconomic costs of inflation are less vulnerable to the Lucas critique than one might think.

6 Cross-country regressions

As already discussed, existing empirical studies have generally failed to uncover a systematic relationship between long-run inflation and the level or growth of GDP at the country level. In attempting to address the discrepancy between this finding and our numerical results, we proceeded to quantify the likelihood that, under the null hypothesis that our baseline calibrated model were the true data generating process, a cross-country regression of GDP on the observed rate of inflation would yield a statistically insignificant effect.

To accomplish this task, we needed to generate a large number of artificial cross-country samples using our model. We did this by focusing on a sample of 63 countries, and for each country in this sample, we set the target rate of inflation in our baseline calibrated model equal to the country's actual trend inflation rate during the 1960–2009 time period before simulating the model 5000 times over.¹⁵ For each country, this yielded 5000 artificial histories of the 40-year average values of the log output gap and the annual inflation rate, thereby allowing us to construct just as many artificial cross-country samples of these two variables. For each artificial cross-country sample, we then ran a regression of the log output gap on the annual inflation rate and recorded the t -statistic associated with the relevant slope coefficient. The left panel of Figure 6 depicts the histogram of these t -statistics across 5000 artificial cross-country regressions along with a kernel density estimate of the distribution. Our analysis of this distribution suggests that, in a world described fully by our baseline calibrated model, the likelihood that one would find a statistically insignificant (at the 5% level) result from running a cross-country regression of GDP on inflation is about 15 percent.

While this likelihood may not seem large, it is certainly significant considering the fact that a cross-country regression in a world described by our model will not be plagued by the same set of endogeneity issues afflicting such regressions in real life. We therefore sought to find an explanation for why, in some states of our model world, a cross-country regression would fail to detect a significant effect of inflation on output. In the course of our inves-

FIGURE 6: Cross-country regression results

tigations, we decided to look at the relationship between the 40-year average values of the annual inflation rate and the log output gap across country-specific artificial histories, that is, across simulations with the same target inflation rate. For any target inflation rate, what we found was a systematic Phillips-type relationship like the one shown in the right panel of Figure 6 for the country with the median trend inflation rate in our 63-country sample. We hypothesize that this negative Phillips-type relationship between the output gap and actual inflation might be responsible for partially masking the positive causal effect of trend inflation on the output gap in our model.

7 Conclusion

In this paper, we use an ACE model to show how inflation can disrupt the mechanism of exchange in a decentralized market economy and, thus, have an adverse effect on macroeconomic performance. The effect works through the increased dispersion of relative prices caused by staggered price setting, which ultimately raises the incidence of failures amongst the specialist trading enterprises crucial for creating and organizing the network of markets that coordinate economic activities.

We find that, in our model economy, the trend rate of inflation has a substantial dele-

terious effect when it rises above 3 percent. Our analysis, however, provides no support for the notion that reducing inflation targets below 3 percent can yield any gains in output or employment. Our main finding remains qualitatively robust to changes in parameter values and to modifications made to our model to partly address the Lucas critique. Finally, we contribute a novel explanation for why it may be difficult to detect a significant negative effect of trend inflation on output in a cross-country context.

Notes

¹See Tesfatsion and Judd (2006) for a survey. Examples of recent macroeconomic applications of ACE include Deissenberg et al. (2008), Delli Gatti et al. (2008), Dosi et al. (2010), and Ashraf et al. (2011).

²Ascari (2004) and Yun (2005) show this effect to be large. However, this result seems to be sensitive to the two features of the Calvo model that most economists would not want to rely on, namely the assumptions that producers are constrained to always satisfy whatever demand is forthcoming at their posted prices and that the opportunity to change one's price is delivered by a state-independent Poisson process providing no upper bound to the length of time between price changes (Howitt and Milionis, 2007).

³Similar findings have also been reported for other nations such as Israel (Lach and Tsiddon, 1992).

⁴Our model almost invariably generates a mean rate of inflation that is very close to the announced target, so expectations are correct on average. We believe that fixed expectations help stabilize our model as they do in standard old Keynesian frameworks. While rule-of-thumb expectations would admittedly be more realistic, we are ultimately interested in isolating the costs of inflation that may arise even in the absence of destabilizing expectations.

⁵Specifically, the entrepreneur's sales target is initially determined by "animal spirits." If he does enter, however, the sales target will be adjusted to actual market conditions every

week.

⁶An agent’s effective wage and effective price are defined in the “labor and goods market trading” stage of the algorithm.

⁷In AGH (2011), we point out that such an expenditure function would apply if the person knew for certain his future incomes and interest rates, and if he were choosing E to maximize an intertemporally additive logarithmic utility function with a weekly rate of time preference $\rho_w = v/(1 - v)$, subject to the usual lifetime budget constraint. We use this interpretation of the expenditure function and calibrate it in terms of the annual rate of time preference $\rho = (1 + \rho_w)^{48} - 1$.

⁸An alternative would have been to allow agents to calculate the capitalization factor on their own, based on their inflation expectations (anchored to the central bank’s target inflation rate) and the central bank’s nominal interest rate projections. This modification to our model yielded results that are almost identical to the main findings reported in the paper.

⁹We assume that the central bank begins to adjust all its estimates and targets after T_{cb} years of simulation.

¹⁰In the context of our model, the facts that (i) inflation expectations are anchored to the central bank’s target, (ii) wage updating follows a firm-level expectations-augmented Phillips curve, (iii) prices are always a fixed percentage markup over marginal costs (except for temporary price changes), and (iv) at any point in time, there is a distribution across firms in terms of how far along they are in their respective wage contract periods, together imply that a higher target rate of inflation will yield a greater percentage difference in prices between firms that have just updated their wages and those that are yet to do so.

¹¹The calibrated parameter values are reported in an extended version of this paper, available online via the authors’ websites. In AGH (2011), there were 33 parameters but 7 of them were banking parameters, which do not apply to the current simplified model, and another was a shop setup cost, which we had to suppress because there are no banks in the

present model to finance this setup cost. Further, unlike AGH (2011), the markup parameter here was held fixed across shops and calibrated to 13 percent, which falls within the 10 to 20 percent range of markup estimates cited by Golosov and Lucas (2007). Due to the omission of the banking sector and the exogeneity of the markup parameter in our model here, the calibrated values of some of the other parameters have small but noticeable differences with their counterparts in AGH (2011), despite the similarity in the calibration procedure.

¹²The output gap is defined as the ratio of capacity to actual GDP. For each run, the volatilities of output and inflation are measured as standard deviations of the annual time series of log GDP and year-to-year percentage inflation respectively.

¹³We also ran an experiment to check the performance of our model economy under negative trend inflation rates up to the -3 percent inflation target, and we found a monotonic deterioration in macroeconomic performance as the target became more negative. Specifically, going from 0 to -3 percent for the inflation target, the median output gap increased by about 1.6 percent and median unemployment rose by just over 1 percentage point.

¹⁴Of course, negative nominal interest rates are not consistent with our assumption that the cash-in-advance constraint is binding on households, but we retained that assumption in any event just to see what difference it made when the central bank was freed from the zero lower bound constraint.

¹⁵Our sample of 63 countries is based on that used by Levine and Zervos (1993) in their influential paper that finds no relationship between inflation and GDP growth. To be consistent with the range of inflation targets used in our main experiment, we excluded those countries with historical trend inflation rates above 10 percent. We also updated the Levine-Zervos analysis with data for the 1960–2009 time horizon and verified that their baseline cross-country regression of GDP growth on inflation continues to find no statistically significant relationship.

References

- Akerlof, George A., William T. Dickens, and George L. Perry (1996) The macroeconomics of low inflation. *Brookings Papers on Economic Activity* 1, 1–76.
- Ascari, Guido (2004) Staggered prices and trend inflation: Some nuisances. *Review of Economic Dynamics* 7(3), 642–667.
- Ashraf, Quamrul, Boris Gershman, and Peter Howitt (2011) Banks, market organization, and macroeconomic performance: An agent-based computational analysis. NBER working paper 17102.
- Benabou, Roland (1992) Inflation and efficiency in search markets. *Review of Economic Studies* 59(2), 299–329.
- Bils, Mark and Peter J. Klenow (2004) Some evidence on the importance of sticky prices. *Journal of Political Economy* 112(5), 947–985.
- Blinder, Alan S., Elie R. D. Canetti, David E. Lebow, and Jeremy B. Rudd (1998) *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.
- Bullard, James and John W. Keating (1995) The long-run relationship between inflation and output in postwar economies. *Journal of Monetary Economics* 36(3), 477–496.
- Calvo, Guillermo A. (1983) Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3), 383–398.
- Chari, V. V., Larry E. Jones, and Rodolfo E. Manuelli (1995) The growth effects of monetary policy. *Federal Reserve Bank of Minneapolis Quarterly Review* 19(4), 18–32.
- Cooley, Thomas F. and Gary D. Hansen (1989) The inflation tax in a real business cycle model. *American Economic Review* 79(4), 733–748.
- Debelle, Guy and Owen Lamont (1997) Relative price variability and inflation: Evidence from U.S. cities. *Journal of Political Economy* 105(1), 132–152.
- Deissenberg, Christophe, Sander van der Hoog, and Herbert Dawid (2008) EURACE: A

- massively parallel agent-based model of the European economy. *Applied Mathematics and Computation* 204(2), 541–552.
- Delli Gatti, Domenico, Edoardo Gaffeo, Mauro Gallegati, Gianfranco Giulioni, and Antonio Palestrini (2008) *Emergent Macroeconomics: An Agent-Based Approach to Business Fluctuations*. Milan: Springer-Verlag Italia.
- Dosi, Giovanni, Giorgio Fagiolo, and Andrea Roventini (2010) Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control* 34(9), 1748–1767.
- Feldstein, Martin (1999) Capital income taxes and the benefit of price stability. In Martin Feldstein (ed.), *The Costs and Benefits of Price Stability*, pp. 9–46. Chicago: University of Chicago Press.
- Fortin, Pierre (1996) The Great Canadian Slump. *Canadian Journal of Economics* 29(4), 761–787.
- Golosov, Mikhail and Robert E. Lucas Jr (2007) Menu costs and Phillips curves. *Journal of Political Economy* 115(2), 171–199.
- Gomme, Paul (1993) Money and growth revisited: Measuring the costs of inflation in an endogenous growth model. *Journal of Monetary Economics* 32(1), 51–77.
- Head, Allen and Alok Kumar (2005) Price dispersion, inflation, and welfare. *International Economic Review* 46(2), 533–572.
- Heymann, Daniel and Axel Leijonhufvud (1995) *High Inflation: The Arne Ryde Memorial Lectures*. New York: Clarendon Press.
- Howitt, Peter (1990) Zero inflation as a long-run target for monetary policy. In Richard G. Lipsey (ed.), *Zero Inflation: The Goal of Price Stability*, pp. 66–108. Toronto: C.D. Howe Institute.
- Howitt, Peter (2006) The microfoundations of the Keynesian multiplier process. *Journal of Economic Interaction and Coordination* 1(1), 33–44.
- Howitt, Peter and Petros Milionis (2007) On the robustness of the Calvo price-setting model.

- Mimeo, Brown University.
- Lach, Saul and Daniel Tsiddon (1992) The behavior of prices and inflation: An empirical analysis of disaggregated price data. *Journal of Political Economy* 100(2), 349–389.
- Lagos, Ricardo and Guillaume Rocheteau (2005) Inflation, output, and welfare. *International Economic Review* 46(2), 495–522.
- Lastrapes, William D. (2006) Inflation and the distribution of relative prices: The role of productivity and money supply shocks. *Journal of Money, Credit and Banking* 38(8), 2159–2198.
- Leijonhufvud, Axel (1977) Costs and consequences of inflation. In G. C. Harcourt (ed.), *The Microeconomic Foundations of Macroeconomics*. London: Macmillan Press.
- Levine, Ross and Sara J. Zervos (1993) What we have learned about policy and growth from cross-country regressions? *American Economic Review Papers and Proceedings* 83(2), 426–430.
- Parks, Richard W. (1978) Inflation and relative price variability. *Journal of Political Economy* 86(1), 79–95.
- Parsley, David C. (1996) Inflation and relative price variability in the short and long run: New evidence from the United States. *Journal of Money, Credit and Banking* 28(3), 323–341.
- Ragan, Christopher (1998) On the believable benefits of low inflation. Bank of Canada working paper 98-15.
- Rocheteau, Guillaume and Randall Wright (2005) Money in search equilibrium, in competitive equilibrium, and in competitive search equilibrium. *Econometrica* 73(1), 175–202.
- Stockman, Alan C. (1981) Anticipated inflation and the capital stock in a cash-in-advance economy. *Journal of Monetary Economics* 8(3), 387–393.
- Summers, Lawrence (1991) How should long-term monetary policy be determined? *Journal of Money, Credit and Banking* 23(3), 625–631.
- Tesfatsion, Leigh and Kenneth L. Judd, eds. (2006) *Handbook of Computational Economics*.

Volume 2. Amsterdam: Elsevier North-Holland.

Woodford, Michael (2003) *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.

Yun, Tack (2005) Optimal monetary policy with relative price distortions. *American Economic Review* 95(1), 89–109.