# How Inflation Affects Macroeconomic Performance: An Agent-Based Computational Investigation[1]

Quamrul Ashraf and Peter Howitt
Department of Economics
Brown University

September 30, 2008

[1]Useful comments were received from participants at Leiden and Sophia Antipolis. Financial support from the NSF is gratefully acknowledged.

# 1 Introduction

Central banks around the world are increasingly focusing their attention on a single indicator of macroeconomic performance - the rate of inflation. When questioned on how this can be justified when people are concerned not only about inflation but also about the level of real activity, inflation-targeting central bankers typically respond that maintaining a low and stable rate of inflation is the best way for them to promote consistently high levels of GDP and low levels of unemployment. This answer echoes the beliefs expressed by such eminent economists as Friedman (1977) and Laidler (1997). Yet, as we document in some detail below, neither mainstream economic theory nor existing empirical studies offer much support for the belief that inflation-targeting countries perform better than others, or 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, Dickens and Perry (1996) argue that reducing the trend rate of inflation too close to zero worsens economic performance because at low rates of inflation downward nominal wage rigidity raises the natural rate of unemployment. 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 the dangers posed by wage rigidity and the zero lower bound, inflation-targeting central banks are reluctant to reduce their targets below 2 or 3 percent. But the commonly acknowledged weakness of the microfoundations of mainstream monetary theory leaves many of them wondering whether there might yet be significant gains from reducing inflation that theorists have not yet discovered.

The research reported here goes a little outside 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 several economists, such as Heymann and Leijonhufvud (1995) (see also Ragan, 1998), namely that inflation might impede the working of the market mechanisms that coordinate economic activities. Strictly speaking, this idea cannot be investigated in a conventional rational-expectations-equilibrium framework because the concept of a rational expectations equilibrium presupposes that peoples' actions and beliefs are always perfectly coordinated without the need for any explicit mechanism. Therefore we deploy a somewhat unconventional method, namely agent-based computational economics.[1]

As described by Tesfatsion (2006), agent based computational economics is a set of

[1] A survey of literature using the method in economics is provided by Tesfatsion and Judd (2006).

techniques for studying a complex adaptive system involving many interacting agents with exogenously given behavioral rules. The idea motivating the approach is that complex systems, like economies or anthills, can exhibit behavioral patterns beyond what any of the individual agents in the system can comprehend. So instead of modelling the system as if everyone's actions and beliefs were coordinated in advance with everyone else's, people are assumed to follow simple rules, whose interaction might or might not lead the system to approximate a coordinated equilibrium. The approach is used to explain system behavior by "growing" it in the computer. Once one has devised a computer program that mimics the desired characteristics of the system in question one can then use the program as a "culture dish" in which to perform experiments.

More specifically, we use a modified version of the adaptive model developed by Howitt and Clower (2000) in which all trades are coordinated by a self-organizing and self-regulating network of trade specialists, as they are in reality by shops, brokers, middlemen, banks, realtors, lawyers, accountants, employers and so forth. Howitt and Clower show that the model generates monetary exchange as an endogenous outcome. Here we modify the model to allow durable goods, fiat money and government bonds, to include a central bank that follows a Taylor rule with an explicit inflation target, and to allow explicitly for the zero lower bound on nominal interest rates. The model is calibrated to US 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 aggregate 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 does provide a propagation mechanism through which inflation might impede the market processes that coordinate economic activity. The mechanism involves business failures that disrupt the network of trade specialists. It works as follows. When established trading patterns are disturbed by the failure of some trade specialists, the suppliers (workers) of those specialists suffer a capital loss, which will be induce them to reduce their expenditures at least temporarily. This may cause other specialists to fail for lack of business, thus amplifying the effect. Eventually the failed specialists will be replaced by new entrants, but only after a turbulent shake-out period involving a large amount of business turnover, during which the economy will be producing below capacity.

Howitt (2006b) showed that because of this propagation mechanism, the basic Howitt-Clower model is capable of mimicking some of the time-series properties of US real GDP, specifically the hump-shaped impulse-response pattern. This mechanism is like the usual Keynesian multiplier process except that without the novel element of business turnover it is too weak to generate a significant effect. What we investigate in this paper is whether the same mechanism might be triggered by a high trend rate of inflation. More specifically, as in many other models, a higher trend rate of inflation should imply on average more dispersion in the relative prices offered by the different specialists, relative to what one would find in an undisturbed steady state with no inflation, because not all prices rise simultaneously. This extraneous relative-price variability should introduce more volatility into the demand faced by any individual specialist, thus raising the probability of failure and amplifying the propagation mechanism. As a result, abstracting from any effects of wage rigidity, and abstracting from the problem of the zero lower bound on nominal interest rates, increased trend inflation should reduce the mean and raise the variance of GDP in the model.

Agent based modelling raises more methodological questions that we have space to answer adequately in this paper. Nevertheless two brief remarks on methodology are in order. First, although resorting to exogenous behavioral rules sounds like a reversion to 50-year old Keynesian economics, it makes a world of difference that we are imposing these rules on individual behavior rather than on macroeconomic aggregates. From the point of view of central bank forecasting and projection analysis the main failure of old-style Keynesian models consisted in their unbelievable and often unstable long-run behavior, which arose from a failure to take into account supply constraints and stock-flow relationships. But our "bottom-up" approach, which starts by explicitly modeling all individual constraints and interactions, ensures that these constraints and relationships are always obeyed. This provides an explicit microfoundation for the aggregate patterns that emerge. Moreover, given the heterogeneity of individuals and the nonlinearity in their behavioral rules, these aggregate patterns need not bear any obvious relationship to their microeconomic counterparts, a point which is impossible to take into account using the old-style Keynesian approach.

Our second methodological remark has to do with the famous "Lucas critique." When trend inflation rises, one would expect that peoples' behavioral rules will adapt to this change in the environment, and a model with exogenous behavioral rules does not allow this kind of adaptation as fully as a rational-expectations analysis would. Our way of dealing with this issue is to do sensitivity analysis to our calibrated model, so as to identify which behavioral parameters are most responsible for our results. We then modify the analysis so as to permit greater flexibility in the choice of these parameters. Repeated simulation then allows us to see to what extent there is an individual incentive for people to change these parameters and to what extent these changes modify our results.

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. Section 4 describes the details of the program implementing the model. Section 5 discusses the workings of the model, describing what its equilibrium would look like if all shocks were suppressed, and how shocks can disturb that equilibrium by affecting shop profits and hence altering the probability of shop failure. Section 6 discusses our method for calibrating the parameters of the model. Section 7 presents the main simulation results indicating how the trend rate of inflation affects macro performance, the role of business failures in generating these results, and the role of adaptation in key parameters of the model. To anticipate this section, we find that the central bank's inflation target has a positive effect on unemployment, the output gap and the variability of output and inflation. Section 8 concludes with some discussion of the overall significance of the results, most of which we argue are at this stage methodological rather than substantive, and makes some suggestions for further research that might strengthen the substantive message of the paper.

## 2 Literature on the real effects of inflation

Many of the papers dealing theoretically with the effects of inflation on the path of real activity have focused on effects that work through the opportunity cost of holding non interest bearing money. Raising this opportunity cost can affect output through various channels, such as reducing steady-state saving behavior (Stockman 1981, Gomme 1993) or altering search intensity (Bénabou, 1988, 1992; Head and Kumar, 2005; Lagos and Rocheteau, 2005; Rocheteau and Wright, 2005). But it is difficult to make the case that these effects are empirically important given that non interest bearing money is such a small item in people's portfolios. For example, Chari, Jones and Manuelli (1995) calculate the effects of inflation on growth in a variety of endogenous growth models where the opportunity cost of holding money is the only channel through which inflation matters, and find that a 10 percentage point increase in inflation has at most a 0.027 percentage point negative effect on the steady-state growth rate.

Others have argued that non-indexation of private accounting systems can lead to distortions and inefficiencies in savings decisions and the allocation of capital under high rates of inflation, and 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). However, 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. Feldstein (1997) has argued that there would be a significant output gain (4 to 6 percentage points) to reducing inflation from 2 percent to 0 because of the nonindexation of capital income taxation. However, this effect seems more a problem to be dealt with by tax authorities than central bankers.

In New Keynesian dynamic stochastic general equilibrium models (for example, Woodford 2003), trend inflation affects output mainly by affecting the equilibrium degree of price dispersion. This is especially true in those versions that use the Calvo (1983) price-setting model. Ascari (2004) and Yun (2005) show the effect to be very large indeed, and much larger than one would guess from log-linearized approximations. However, this seems to be sensitive to two features of the Calvo model that most economists would not want to rely on, specifically the assumption that producers are constrained always to 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 which gives no upper bound to the length of time between price changes. This means that when inflation is high, a significant fraction of the economy's GDP will be produced by firms who have been waiting a long time for a price change opportunity and hence are selling at a price below marginal cost. Specifically, Howitt and Milionis (2007) show, using a standard New Keynesian Model, that when trend inflation is 10 percent, 36 percent of output is being produced by firms selling at a marginal loss, and that once one eliminates this counterintuitive prediction by replacing the Poisson process with fixed price-change intervals, or with optimally chosen state-dependent intervals, the output cost of inflation falls by a factor of 20.

The empirical literature on the effects of trend inflation on the growth rate of real output (thoroughly surveyed by Ragan, 1998) has produced little evidence to date of a significant effect, except in extreme circumstances. Nor is there much evidence that countries whose central banks adopted inflation targets starting in the early 1990s had better macroeconomic performance than those that did not. As is well know, aggregate volatility has fallen considerably in most countries over this period. Cecchetti, Flores-Lagunes and Krause (2004) estimate that for many of the 24 countries they examine this reduction in volatility indicates an improvement in monetary policy, but although many of these countries were inflation-targeters there are notable exceptions such as Japan and the United States. Levin, Natalucci and Piger (2004) find that the unconditional standard deviation of output growth is no lower among inflation-targeting countries than others.

There is also almost no direct evidence of a positive effect of trend inflation on the level of real output, as opposed to its growth rate or volatility. For example, Bullard and Keating (1995) examine the relationship between permanent shocks to inflation and permanent shocks to output in 58 different countries over a long period of time. They find that of the 25 countries in their sample that experienced permanent inflation shocks, in only 4 was there a significantly positive permanent effect on output, while in 1 there was a significant negative permanent effect.

There is however considerable direct evidence regarding at least one of the theoretical channels through which inflation might affect the level of output, namely the direct effect of inflation on relative-price dispersion which plays a major role in New Keynesian models. Vining and Elwertowski (1976) and Parks (1978) found a significant effect of inflation on dispersion, although without carefully distinguishing between expected and unexpected inflation. Lach and Tsiddon (1992) estimate that expected inflation has a large positive effect on price dispersion in Israel, even larger than the effect of unexpected inflation. The results of Binette and Martel (2005) suggest that trend inflation is the main determinant of relative-price dispersion in Canada.

# 3 The model

## 3.1 Preliminaries

Our agent-based model is a variant of the adaptive model developed by Howitt and Clower (2000), modified as by Howitt (2006b and 2007). 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 ("shopkeeper") on one side or the other of the market. Shopkeepers are the agents that match buyers and sellers, arrange terms of exchange and bear the costs of adjustment in the face of day-to-day imbalances between spending plans and productive capacity; in short, they are the real-world counterparts of the fictional auctioneer of general equilibrium theory.

We add several components to this model so as to make it considerably less stylized. In adding these new components we have tried to make the structure and its macroeconomic aggregates comparable to the baseline New Keynesian analysis (for example, Woodford, 2003) that is now commonly used to model the practice of inflation targeting. 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 in three important senses. First, we have introduced elements of search, in both goods (retail) markets and labor (wholesale) markets, whereas the canonical New Keynesian model has a Walrasian labor market and no search in the goods market.[2] Second, we assume that firms are subject to failure and that the process of replacing failed firms is a costly one, whereas the population of firms is fixed in the New Keynesian framework. Third, instead of the perfect and complete set of contingent financial markets assumed in the New Keynesian literature we assume that the only available financial instruments are government issued (non contingent) money and bonds.

We are interested in these differences from New Keynesian analysis because we think it is worth exploring the possibility that it is not just price stickiness but also coordination problems of other sorts in goods and labor markets that give rise to real effects of inflation. Howitt (2006b) showed in a related model that the degree of price flexibility, as measure by the speed of adaptation $\alpha$ does not necessarily affect the propagation mechanism of the model. This is because what causes output to fall is not the inability of prices to seek their equilibrium values but rather the disappearance of firms that make up part of the economy's transactions mechanism. The idea being explored in this project is that inflation impedes the transactions process by endangering the survival of shops.

The model is also related to the literature on the economics of search. The shopping activities described above are similar to what one finds in the class of directed-search models recently surveyed by Rogerson, Shimer and Wright (2005). Another directed search model with fiat money is analyzed by Rocheteau and Wright (2005). A directed search model with fiat money that also focuses directly on price dispersion is presented by Head and Kumar (2005). The main difference between the setup of this model and that of Head and Kumar is that we take into account that trading relationships are durable, as is typically assumed in the labor search literature, whereas in the Head-Kumar setup they are purely transitory. This is a key difference because if all trading relationships were purely transitory then the temporary failure of trading specialists would not generate the above described propagation mechanism. Moreover, the present model is complicated by the intertemporal savings and portfolio allocation decisions of households.

### 3.1.1 The role of money

An important reason for picking the Howitt-Clower model as our starting point is that in order for inflation to matter for the workings of the market mechanism, money must also matter for that mechanism, and this is a model in which money clearly does plays such a role. Specifically, Howitt and Clower (2000) show that the model contains forces that could generate monetary exchange as the outcome of an evolutionary process even

[2] See however Blanchard and Gali (2008).

if the economy started with no organizational structure at all, and that in those cases when the forces were not strong enough to generate monetary exchange, a large portion of the potential gains from trade would remain unexploited in the long run.

The model is similar in many respects to the more standard search theoretic model of monetary exchange.[3] Specifically, in both models monetary exchange can be supported in equilibrium because of a network externality arising from trading frictions. In the Kiyotaki-Wright framework that network externality comes from the fact that search time is reduced by accepting in exchange an object that others will accept. In the present framework it comes from the fixed costs of operating trade facilities, which are minimized by having the same object traded in all shops.

The underlying mechanism of exchange is most closely related to that of Starr and Stinchcombe (1998, 1999), who first introduced costs of market making into the market-game setup of Shubik (1973) and showed how these costs could generate monetary exchange as an optimal or equilibrium phenomenon. Starr and Stinchcombe however dealt exclusively with perishable commodity money and did not introduce any search or other intertemporal considerations.

### 3.1.2 Disequilibrium analysis

Most importantly our analysis differs from both New Keynesian analysis and search theory because we are investigating the out-of-equilibrium behavior of the system. The hypothesis that we are investigating is that inflation can affect real economic performance by degrading the system's ability to track a coordinated state of equilibrium. With few exceptions, the above-described analyses deal exclusively with rational-expectations equilibria. The main exception is the literature on learning in New Keynesian models (for example, Bullard and Mitra, 2002; Evans and Honkapohja, 2003; Preston, 2004) in which people form expectations about future output and inflation adaptively rather than rationally. A recent paper by Milani (2005) indicates that such learning behavior is capable of explaining a lot of the otherwise difficult to explain persistence of output and inflation in an otherwise canonical New Keynesian model.

The approach of the present paper allows one to model the coordination problem in more microeconomic detail than this New Keynesian learning theory, as it focuses on the exit and entry of the coordinating agents and allows for considerable heterogeneity in individual expectations and experiences. In a sense it is a continuation of a line of research into disequilibrium macroeconomics that reached its pinnacle in the Barro-

[3] For example, Kiyotaki and Wright, 1989. The relation of our approach to the search-theoretic approach is discussed in greater detail by Howitt and Clower (2000) and Howitt (2005).

Grossman (1976) book, which attempted to flesh out the details of what happens when people trade at prices that make their plans mutually incompatible. That line of research ran into the problem that trading out of equilibrium in one market generated rationing constraints that affected traders in other markets, in ways that were hard to rationalize and hard to model. The Howitt-Clower model provides a systematic framework for modelling those out-of equilibrium interactions through the agency that seems to manage the interactions in real-world markets, namely specialist trading firms that create and operate markets.

Moreover, agent-based computational analysis provides a way to deal with the complexities that stalled the earlier disequilibrium literature, not by sidestepping them through assumptions that limit their scope but rather by embracing them, putting them into the computer code, and using the power of modern computer technology to explore the macroeconomic patterns that emerge from seemingly chaotic microeconomic interactions.

#### 3.1.3 Inflation targeting

In designing our model, we have attempted to focus on relationships that seem particularly relevant for an economy in which the central bank is engaged in inflation targeting as it has evolved since the early 1990s. Specifically, the central bank is portrayed as setting a very short-term interest rate according to a Taylor rule that explicitly includes an inflation target and that feeds back from recent observations of inflation and GDP.

Moreover, since our computational approach allows us to model the economy's evolution from week to week we can be very explicit about the interest setting process. So instead of having a new interest setting each period, we assume that the central bank has a predetermined set of "fixed action dates" (to use the Bank of Canada's term) at which it may or may not alter the setting of its policy rate. In this version of the paper those fixed action dates occur at the end of each month, but in principle we could come even closer to current US practice by making them take place 8 times per year as in the United States.

As an aside, our approach also allows us to be explicit about the information available to the central bank when setting its policy rate. In the following we will suppose that the central bank knows exactly the value of the aggregate price level at the end of each month, and can remember those values, and likewise for the economy's level of actual output. But it would be a simple matter to introduce noisy information, because our code includes an explicit statistical agency that keeps track of that information.

We also let the current practice of inflation targeting guide us in our modeling of inflationary expectations. One of the lessons that many Central Banks have taken from their experience with inflation targeting is that it seems to have been remarkably successful in anchoring inflationary expectations. In Canada, for example, where inflation has been 2% plus or minus some small noise for more than a decade under inflation targeting, conventional measures of inflation expectations seem to be more or less fixed at 2% even during periods when inflation has risen significantly above or below its target.[4] It seems that one way in which people adapt to a successful inflation-targeting regime is by simply adopting the target as their forecast, thus economizing on whatever resources they had been using in forecasting inflation through a more sophisticated method. Accordingly, in what follows we set the inflation expectation in our wage-adjustment equation equal to the official inflation target, with no further adaptation on the part of private agents. Likewise, when people are comparing prices at different shops with prices that they remember from a week ago they adjust their remembered prices automatically by the target inflation rate without doing any more sophisticated calculation.

We do however make one important exception to this rule for modelling expectations. Specifically, as discussed in more detail in section 4.5 below, we suppose that people take as given the central bank's forecasts of how the short-term interest rate will adjust over time. When it is above target, because either inflation or output is unusually high, the central bank forecasts its future setting of interest rates by forecasting the future evolution of inflation and output and feeding those forecasts into its Taylor rule. Those forecasts of future nominal interest rates, combined with the central bank's forecasts of inflation rates, imply forecasts of real short-term interest rates, which the private agents in turn use in calculating the value of their human wealth when deciding on current expenditures.

## 3.2 The conceptual framework

There are $N$ transactors and $n$ goods. In addition to these goods there are two nominal assets, money and bonds. In addition to these transactors there is a central government that issues the bonds and money, and imposes a sales tax on retail transactions. Time is discrete, indexed by "weeks" t = 1,...,T. Each bond is a promise to pay one unit of money ("dollar") next week.

Each transactor can eat only two of the goods (his "demand goods" 1 and 2) and is endowed each week with enough labor services to produce one unit of a particular good (his "supply good"), which is never one of his demand goods. A transactor whose supply

[4] See Howitt (2006a) for a discussion of related issues.

good is $i$ and whose first demand good is $j$ is said to be of type $(i, j)$, and his second demand good is $j + 1 \pmod n$. For each of the $n(n-2)$ ordered pairs $(i, j)$ of goods such that $i \neq j$ and $i \neq j + 1 \pmod n$, there is exactly one transactor of type $(i, j)$. The population of the economy is thus $N = n(n-2)$.

Because no transactor eats his own supply good, he must trade to eat. Trading is such a costly affair that it can only take place through an organized facility called a "shop." To trade with a shop a transactor must form a trading relationship with it. Each shop is capable of dealing in only one good, which it exchanges for money. Each transactor may have ongoing trading relationships with at most one shop (his "employer") that deals in his supply good, and at most one shop ("store") that deals in each of his demand goods. Each week each transactor delivers his unit endowment for sale to his employer if he has one, and visits his stores to buy demand goods from them. Goods are storable, but only by shops. Money and bonds are storable by all agents.

Each shop posts a pair of prices for the good that it deals in; a wholesale price, or "wage," $w$ and a retail price. We denote by $p$ the inverse of a shop's retail price. The shop agrees to buy all goods delivered to by each employee at the price $w$, provided that it has enough cash on hand at the time of delivery, and to sell whatever quantity is demanded by each customer at the retail price $p^{-1}$ provided that it has enough inventories on hand at the time the customer places the demand.

Every week, transactors that are without a complete set of shops search for information about possible trading relationships, through a combination of direct search and referral. The former consists in examining a randomly selected shop, and the latter consists in making inquiries with random people having the same demand or supply good. Given the results of such search activities a transactor will make his choice myopically - choosing the employer that is currently offering the highest wage and the shop that is offering the lowest retail price. In the case of a referral, however, the relevant price is not the actual price posted by the shop but rather the most recent effective price experienced by the referee, which will differ from the actual price if the shop has rationed the referee because of a shortage of goods or money. (see section 4 below). A shop will take on any given number of new customers but is not obligated to take on new employers in this process. So an employment relationship will form only if agreed to by both parties. Each trading relationship that is formed in this process is like those in standard search/matching theory - it will last until either the shop exits, the match dissolves for exogenous reasons or the transactor finds a preferable shop willing to match with him as the result of future search activities.

Shops can be opened only by transactors that innovate. More specifically, each week a random subset of transactors ("entrepreneurs") are struck by an idea for opening a shop that trades the transactor's supply good. The entrepreneur will open such a shop if he thinks he can earn a profit higher than his current estimate of permanent income (see below). In estimating the profit he can earn he is influenced by a random draw of "animal spirits" which determines his estimate of what volume of sales he can achieve, and by his estimate of the average real wage rate in the economy. The bigger his sales estimate and the lower his real wage estimate the more likely he is to enter.

Each shop posts a retail price that constitutes a fixed markup over its wage net of the sales taxes. The wholesale price, or "wage" is adjusted periodically, at fixed intervals. As in an expectations-augmented Phillips Curve, the adjustment is proportional to expected inflation. Here and elsewhere we assume that transactors always believe the inflation-targeting central bank and hence always forecast an inflation rate equal to the target rate. Given the target rate of inflation the wage adjustment is also an increasing function of a measure of the shop's excess demand for employees; ie the ratio of its target employment level (see below) to its actual labor force. Whenever a shops wage is adjusted, its retail price is automatically marked up at the same time.

Each shop has a fixed overhead cost that must be incurred each period in the form of the shop's traded good. It also aims to keep an inventory level equal to one week's sales. When inventories fall below this target level it plans to increase them by a fixed fraction of this gap each period. Its target employment level each period is thus equal to its fixed cost plus the quantity demanded last week plus the planned inventory adjustment. Each week in which a shopowner's estimate of permanent income falls below his estimate of the average market wage rate, his shop exits with a given probability.

Before trading with shops each week, transactors go to a bond market that is organized by the government, which participates in the market by agreeing to buy or sell whatever quantity private transactors wish to exchange at the government's target interest rate. This target interest rate is adjusted on the first week of each "month" (for ease of computation we assume 4 weeks in each month and 12 months in each "year") according to a Taylor Rule that responds positively to the difference between the most recent year-to-year inflation rate and the inflation target, and to the most recent observation of real GDP. (Real GDP is the sum across all firms of input minus fixed costs.)

The government adjusts its tax rate once per year in response to any discrepancy between the actual level of outstanding debt and some fixed target level.

All of the rules of behavior that we have outlined above are simple intuitive rules that require no optimization or forecasting. However there is one place where the agents are assumed to be somewhat sophisticated, and that is in formulating their demands for consumption goods and their related demands for money and bonds. Each household plans to spend a fixed fraction of current wealth on its consumption goods. If he has just one shop then all of its expenditure will be allocated towards that shop. Otherwise it will be allocated to maximize a CES current utility function of the two consumption quantities.

The sophistication enters into this simple rule in the calculation of current wealth. We assume the transactor understands that wealth consists in not just current bond and money holdings but also future income prospects. He calculates this second "human" wealth component in two stages. First he estimates "permanent income" according to a simple adaptive-expectations rule as in Friedman's original formulation of the permanent income hypothesis. Next he converts this permanent income into wealth using a capitalization factor that comes from a forecaster of interest rates. Although each transactor has a different permanent income they all use the same capitalization factor. This factor in turn is consistent with the Taylor rule being used to set interest rates, under the assumption that the output and inflation factors underlying the rule follow a fixed AR1 stochastic process. We believe that this sophisticated calculation of expected future interest rates is consistent with the philosophy of agent-based modelling because of the growing evidence that the transparency of inflation-targeting has made interest-rate forecasts relatively accurate and widely available. In any event, these forecasts are what allow monetary policy to have a non-trivial effect on demands in our model despite the fact that the monetary authority is constrained to operate only in very short-term bond markets and is assumed not to engage in deliberate interest-smoothing.

We also assume that transactors demand money according to a cash-in-advance procedure. That is, after leaving the bond market to buy and sell goods, the transactor does not know whether he will receive payment for his current sales before or after he buys his consumption goods. To ensure against being unable to pay for his consumption he always leaves the bond market with enough cash to pay for his planned consumption. In the case of a transactor who is also a shopkeeper, he plans to leave with enough to pay also for his planned wage bill, assuming that he will reach his target input this period.

# 4 Algorithmic Details

The model sketched above has been implemented as an computer program, written in the C++ programming language.[5] It represents the economy as an algorithm. The present section describes this algorithm in more detail, and enumerates the 22 free parameters

[5] The source code is posted on http://www.econ.brown.edu/fac/Peter_Howitt.

on which behavior depends. The initial conditions each week consist of a certain number of established shops, the goods they trade and their predetermined targets, a historically given configuration of ongoing trading relationships between transactors and shops, a set of targets, permanent income estimates and effective prices from last week, an array of money, bond and inventory holdings, a tax rate, an interest rate, a history of price and GDP levels, and a lagged measure of the economy-wide average wage rate. The algorithm then proceeds to generate a new set of initial conditions for the following week in nine stages, each of which represents an important component of the workings of a decentralized economy. These nine stages of weekly activities, described below, are repeated for a total of $T$ weeks.

## 4.1 Firm Entry

In the first stage of the algorithm representing weekly activities, each transactor who is not already a shopowner is considered as a potential innovator. Each one becomes an actual innovator with probability $\theta/N$. The innovator is first assigned a random realization $x$ of target sales (animal spirits) from a uniform distribution over $[1, \bar{x}+1]$, where the upper bound of the support represents the maximal target sales of a potential entrant, and a random markup $\mu$ drawn from a uniform distribution over $[0, 2\overline{\mu}]$. The innovator then estimates his weekly profit as a shopowner as $W(\mu x - F)$, where $F$ is a fixed cost parameter and $W$ is the estimated economy-wide aggregate wage rate for the current week.[6] If the innovator's profit estimate is above his latest estimate of permanent income, updated during the bond market trading stage in the previous week, the innovator becomes an entrant by opening up a shop that trades his supply good.

The key parameters introduced in this stage therefore include: $\theta$, which determines the frequency of innovations; $\bar{x}$, governing the animal spirits of potential entrants; and the markup and fixed cost parameters $\overline{\mu}$ and $F$.

## 4.2 Search and Matching

Next, the program gives each transactor an opportunity to search for possible trading relationships. This comprises both job search (for a shop (employer) that buys the transactor's supply good) as well as "store" search (for shops that sell either of his two demand goods). Each transactor engages in job search with probability $\sigma$. Job search

[6] In particular, the statistical agent computes the aggregate wage rate annually in the last week of a given year during the fiscal policy stage. Therefore, in any given week the current estimate of the aggregate wage rate is calculated as this lagged figure adjusted for the inflation that would have occurred in the intervening weeks if the annual inflation rate was at its annual target rate.

consists in asking one randomly selected "comrade" (someone with the same supply good) what his effective wage is. If it exceeds the searcher's current effective wage, then the searcher attempts to switch to the comrade's employer. The switch will be implemented if and only if the employer's current employment level is currently less than its target employment level (to be determined in section 4.9 below). If so, the searcher's former employment relationship (if any) is severed and the transactor sets his effective wage equal to the comrade's.

Store search, on the other hand, is undertaken by every transactor. This type of search comprises not just referral-based but also direct search. First the transactor asks a randomly selected "soulmate" (someone with the same two demand goods) for his effective inverse retail prices. If either is higher than the searcher's, the searcher will switch to the corresponding store and set his effective inverse retail price equal to the soulmate's. Then the transactor selects a shop at random. If the shop trades either of his demand goods and is posting an inverse retail price higher than the searcher's effective inverse retail price the searcher will switch to that store and set his effective inverse retail price equal to the store's posted price. Every time he switches, the transactor will sever any relationship with a store trading the same good.

The key parameter introduced by this stage of the weekly activities is the probability $\sigma$ that a transactor will engage in job search.

## 4.3 Bond Market Trading

After the matching process has concluded, everyone goes to the bond market. First, the government redeems all bonds, which are just one-week instruments. At this point there are no more bonds outstanding, and all the financial wealth of the transactors is in the form of money. Consider first a non-shopowner transactor. He must calculate how much of his wealth $A_i$ to reinvest in bonds $B_i$, and therefore how much to retain in the form of money $M_i$ to facilitate trading in the next stage, subject to the wealth constraint

$$A_i = M_i + \frac{B_i}{1 + r_w} \tag{1}$$

$r_w$ is the per-week nominal interest rate, set by the government's monetary policy, to be discussed in section 4.5 below.

In making this decision he must at his point plan how much $E_i$ to spend on goods this week. We assume that he plans to spend a fixed fraction $v$ of total wealth, which is the sum of financial wealth and the capitalized value of his permanent income $Y_i^p$:

$$E_i = v \cdot (A_i + V \cdot Y_i^p)$$

where $V$ is a capitalization factor, equal to the present value of a nominal income stream that grows each week at the constant weekly target rate of inflation, given the sequence of nominal interest rates that the central bank is projecting. Note that this is precisely the expenditure function that would apply if the transactor knew for certain what future incomes and interest rates would be and were choosing $E_i$ so as to maximize a standard intertemporally additive logarithmic utility function with a weekly rate of time preference $\rho_w = v/(1-v)$. We will use this interpretation of the above expenditure function when calibrating the model, and will calibrate it in terms of the annual rate of time preference $\rho$ defined by $(1+\rho) = (1+\rho_w)^{48}$.

Before using this formula to determine $E_i$, the transactor updates his estimate of permanent income according to the adaptive rule:

$$\Delta Y_i^p = \lambda_y (Y_i - Y_i^p)$$

where $Y_i$ is his actual income (effective wage) last period, and then he updates $Y_i^p$ again to adjust for estimated weekly inflation assuming that inflation is taking place each week at the target rate.

Having thus decided on planned expenditure he chooses a desired money holding $M_i = E_i$. Our idea in applying this rule here is that because there is no private credit in the economy he will need to have $E_i$ in the form of money when he visits his demand stores. But he does not know whether he will be paid his wage income before or after visiting his stores, so he carries $E_i$ out of the bond market to ensure against being unable to fulfil his expenditure plans.[7]

A shopowner transactor lays a similar plan expect that his permanent income is updated using his profit income as $Y_i$ and his desired money holding equals not just enough to pay for his planned goods expenditure but also enough to pay his target wage bill, equal to his current wage times his current target input minus one (since he does not have to pay himself).

At this point, if $M_i < A_i$ the transactor buys the quantity of bonds that satisfies the constraint (1). Otherwise he leaves the bond market holding $A_i$ units of money and no

[7] This motivation for a precautionary demand for money is similar to the "stochastic payment process" that Patinkin (1965) used to rationalize putting money in the utility function. In this case we are using it to justify what looks like a conventional cash-in-advance constraint.

bonds. The key parameters that have been introduced in this bond-market stage are the annual rate of time preference $\rho$ and the adjustment speed of permanent income $\lambda_y$.

## 4.4 Labor and Goods Market Trading

Next, for each transactor, the program simulates trading in both the labor market (i.e., with the transactor's employer) and the goods markets (i.e., with his stores). With probability 1/2 he first executes his labor market trading; otherwise he first executes his goods market trading. Labor market trading proceeds as follows. If the transactor is a shopowner (i.e., is self-employed), he simply gives his unit endowment to the shop. If he is not a shopowner then he gives trades his endowment for money with the shopowner at the posted wage, subject to the constraint that if the shopowner's money holding is less than the posted wage then they trade the amount that just exhausts that money holding, so the transactor is rationed. The transactor then sets his effective wage equal to the amount of money he has just received.

Goods market trading in the following manner. For any given transactor, his planned expenditure for the current week is set to the minimum of his current wealth, determined earlier in the bond market trading stage, and his current money holding. For each of his consumption goods, he determines his planned expenditure on that good by maximizing a two-good CES utility function

$$(c_1)^{1-1/\varepsilon} + (c_2)^{1-1/\varepsilon}$$

with an elasticity parameter of $\varepsilon$. If the transactor has established relationships with stores for both of his consumption goods he trades with both his stores, which meet his demand up to the point where their inventories are exhausted. The transactor then sets his effective inverse retail price for each good equal to the actual posted inverse retail price times the fraction of his demand that was satisfied. If the transactor has a customer relationship with only one shop, he goes through the same routine attempting to spend his entire planned expenditure on the corresponding consumption good.

At all stages of trading, the program adjusts inventories and money holdings. In goods market trading the program also makes sure to deduct the shop's tax liability, equal to the prevailing tax rate $\tau$ times the value of all executed retail transactions. Given the possibility of customer rationing by stores, the program also records the potential sales for each store that would occur in the absence of insufficient inventories, and the potential input for each store that would occur in the absence of insufficient money holdings.

The key parameter introduced by this stage of the weekly activities is the demand elasticity $\varepsilon$ determining the proportion of the weekly planned expenditures allocated to each of a transactor's consumption goods.

### 4.5 Monetary Policy

Next the program comes to the stage at which the central bank sets the interest rate$r_w$. First it checks whether this is a fixed action date, which is true every fourth week. If not, this stage is skipped and the interest rate remains unchanged. If it is a fixed action date, then the central bank calculates real GDP $y$ (the sum of each shop's input in excess of its fixed cost, over the past month) and the price level $P$ (the GDP deflator using the same data as used in calculating $y$). The bank then sets the per annum rate of interest $r$ according the rule:

$$1 + r = \max \left\{ \gamma \left(1 + \pi\right)^{\gamma_\pi} y^{\gamma_q}, 1 \right\} \tag{2}$$

where $\gamma$, $\gamma_\pi$ and $\gamma_q$ are fixed coefficients, and where $1+\pi = P/P_{-12}$ is the inflation factor over the past 12 months. This determines the weekly interest rate according to

$$1 + r_w = (1 + r)^{1/48}$$

Equation (2) is a non-linear form of a conventional Taylor rule, which respects the zero lower bound on nominal interest rates.

We interpret this policy rule as being that of an inflation targeting central bank we need to make some heroic assumptions. Specifically, we suppose that the central bank has an accurate estimate of the long-run interest rate $\rho^*$ (which in more conventional models would be approximately the rate of time preference $\rho$) and of the output level $y^*$ that will prevail on average. It then sets $\gamma$ to the value consistent with in inflation target $\pi^*$, so that when the zero lower bound is not binding we have:

$$1 + r = (1 + \rho^*) \left( \frac{1 + \pi}{1 + \pi^*} \right)^{\gamma_\pi} \left( \frac{y}{y^*} \right)^{\gamma_q}$$

A linear approximation to this rule, considered as a function of $\rho^*, \pi, \pi^*$ and $y$, around the point $\rho^* = \pi = \pi^* = 0, y = y^*$, yields the more conventional linear Taylor Rule:

$$r \simeq \rho^* + \gamma_\pi \left(\pi - \pi^*\right) + \gamma_q x$$

where $x = y/y^* - 1$ is the percentage deviation of output from its target. Thus we may interpret our coefficients $\gamma_x$ and $\gamma_q$ as the coefficients of this more conventional rule. In our calibration below we normalize the output target by capacity output $y^{cap} =$

$n(n-2-F)$, the level of GDP when everyone is working or operating a shop and there is exactly one shop per good. More precisely, the central bank is assumed to pick a target output gap: $q^* = 1 - y^*/y^{cap}$.

At this point the central bank also announces its projections of future inflation and future nominal interest rates. It does this by assuming that $\pi$ and $y$ will approach their long-run targets in yearly movements according to the non-linear autoregressive processes:

$$\frac{1+\pi_{+12}}{1+\pi^*} = \left(\frac{1+\pi}{1+\pi^*}\right)^{\lambda_\pi}$$
$$\frac{y_{+12}}{y} = \left(\frac{y}{y^*}\right)^{\lambda_q}$$

where $\lambda_\pi$ and $\lambda_q$ are fixed autocorrelation coefficients. These are the projections that are used in computing the capitalization factor used by transactors when formulating their expenditure plans.

The key parameters introduced in this section are the Taylor-rule coefficients $\gamma_\pi$ and $\gamma_q$, the autocorrelation coefficients $\lambda_\pi$ and $\lambda_q$, the target output gap $q^*$, the target real interest rate $\rho^*$ and the target inflation rate $\pi^*$.

## 4.6 Match Breakups

The program next simulates the random breakup of established trading relationships in the economy. In particular, each transactor in the economy who does not own a shop is subjected to a probability $\delta$ of "quitting" the economy, which entails the unconditional severance of all current trading relationships by the transactor with his employer as well as his consumption stores.

## 4.7 Fiscal Policy

Next comes the stage where the retail sales tax rate $\tau$ is adjusted. This happens only once a year, in the last week of the year. In all other weeks this stage is bypassed and $\tau$ remains unchanged. When deciding on the new tax rate the government first calculates the size of the government debt (normalized by the price level) relative to capacity GDP $y^{cap}$. It then sets the tax rate equal to a value $\tau^*$ which is the value that would leave the debt/GDP ratio undisturbed in the unshocked equilibrium to be described in section 5.1 below, plus an adjustment factor that is proportional to the difference between the actual target debt to GDP ratio $b^*$. The factor of proportionality in this adjustment factor is $\lambda_\tau$.

Thus there are two new key parameters introduced in this section, the target Debt-GDP ratio $b^*$ and the fiscal adjustment speed $\lambda_\tau$.

## 4.8 Firm Exit

Next, the program gives each unprofitable shopowner an opportunity to close down his shop with probability $\phi$. To be specific, the condition of unprofitability is met if a shopowner's permanent income estimate is below what he believes he could earn as a non-shopowner transactor, which is simply the estimated economy-wide aggregate wage rate $W$ for the current week. Once a shop exits, all trading relationships (with both employees and customers) are severed.

## 4.9 Updating of Targets and Prices

In the final stage of weekly activities, the program first updates each shop's employment (input) target, and then attempts to calculate its corresponding offer prices. In updating the target employment of a shop, the program first determines the planned inventory adjustment by a shopowner who aims to reduce a proportion $\lambda_I$ of the current deviation of his shop's inventory from its potential sales estimate based on the current week's goods trading activities. The parameter $\lambda_I$ therefore captures the "speed" with which a shop's actual inventory would converge to its sales in the absence of future shocks to the system. A shop's employment target is then calculated to meet the sum of its fixed cost, its potential sales experienced that week and its planned inventory adjustment.

The program then proceeds to update the shop's wages and retail prices but does so only if the current week happens to be a "price-updating" week for that shop. This is condition is met if either $\Delta$ weeks have elapsed since the shop's most recent price update or if the current week is the shop's very first price-updating week, which is a random realization from the discrete uniform distribution over $[1, \Delta]$ assigned when the shop first opens. The random assignment of initial updating weeks implies that price setting will be "staggered", and the fraction $1/\Delta$ of all prices will change in an average week.

Given that the current week is indeed a price-updating week for the shop being considered, its offer price updates are performed in the following manner. First, a desired annual real wage adjustment figure is calculated as a proportion $\beta$ of the percentage deviation in the shop's target employment from its actual employment, based on all current relationships with employees regardless of whether they were paid or not.[8] The

[8] In computing this expression we actually use not actual employment but the minimum of actual employment and the shop's fixed cost $F$, so as to avoid division by zero when actual employment falls

parameter $\beta$ hence indexes the degree of wage and price flexibility in the economy. This annual wage adjustment figure is then rescaled to obtain a corresponding adjustment for the length of the contract period, taking into account the shopowner's projected estimate of inflation during this period as derived from the weekly target inflation rate. This rescaled wage adjustment figure is then applied to the current wage offered by the shop and its retail price (net of taxes) is correspondingly altered to retain its markup over the wage.

The key parameters introduced in this final stage of weekly activities therefore include the length of the contract period, $\Delta$, and the speeds of inventory and wage adjustment, $\lambda_I$ and $\beta$.

### 4.10 Simulating the model

The entire run of the algorithm over T weeks is then repeated for R different runs, where a run always starts off in the flexible-price, no-shock equilibrium of the economy (see below). Each run, however, is unique in the initial seeding of the computer's random number generator for that run. This allows us to exploit the randomness that is built into the system by ultimately enabling us to examine the average performance of the system across different realizations of the economy's behavior over T weeks.

## 5 The workings of the model

### 5.1 Equilibrium with price flexibility and no shocks

As the preceding discussion has made clear, all shocks in this economy are individual shocks. Unlike in the standard New Keynesian framework, we have postulated no exogenous shock process impinging on aggregate productivity, price adjustment, aggregate demand, monetary policy or fiscal policy. 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. However if we do turn off all shocks, there is a deterministic equilibrium that the economy would stay in if left undisturbed by breakups and entry if the inflation target $\pi^*$ were equal to zero and the output gap were also equal to zero. Moreover, if the contract period $\Delta$ were equal to one week, the economy would remain in this equilibrium for any positive

to zero.

rate of inflation. This equilibrium will serve as an initial position for all the experiments we perform below, and a brief description of it helps to illustrate the workings of the model.

The equilibrium is one in which all the potential gains from trade have been exhausted. Each transactor is matched with one employer and two stores. There are $n$ shops, one trading in each of the $n$ goods. Each shop begins each week with actual, potential and target input all equal to $n-2$, which is the number of suppliers of each good, and with actual and potential sales equal to inventory holdings equal to its actual output: $n-2-F$. So the economy's total output equals full capacity: $y^* = n(n-2-F)$.

Each shop begins each week with a wage rate equal to $W = (1+\pi_w^*) W_{\_}$, which is the same for all shops, where $W_{\_}$ was the common wage rate last week, and with a retail price equal to $P = W * (1+\mu)/(1-\tau)$ where the tax rate $\tau$ equals:

$$\tau^* = 1 - (1+\pi_w^*)(1-48\rho_w b^*) \Big/ \left(1 - \pi_w^* \frac{n-3}{(n-2-F)(1+\mu)}\right)$$

The initial outstanding stock of bonds is

$$B = b^*(1+r_w) 48 y^* P_{-}$$

where $P_{\_} = P/(1+\pi_w^*)$ is last week's price level and where the weekly interest rate $r_w$ is given by:

$$1 + r_w = (1+\rho)^{1/48}(1+\pi_w^*).$$

The money supply at the start of the week is

$$M = W_{\_} * (N-n) + (1-\tau) P_{\_} y^*$$

which is the sum of all wage receipts of non-shopowners, and all sales receipts (ex taxes) of shopowners, from last period.

Each agent starts the period with an effective supply price equal to $W_{\_}$ and and both effective demand prices equal to $P_{\_}$. The owner of each shop starts with a permanent income equal to last period's profit: $(1-\tau) P_{\_}(n-2-F) - W_{\_}(n-3)$ and with money holding equal to last period's revenue: $(1-\tau) P_{\_}(n-2-F)$. Each of the non-shopowner transactors begins with money holding equal to permanent income, which in turn is equal to last period's wage income $W_{\_}$.

The aggregate bond supply $B$ is assumed to be initially distributed across agents in proportion to their initial money holdings.

The initial history is one in which the output gap has been equal to zero for the past

12 months and inflation has equaled its target rate for the past 12 months.

It is straightforward to verify that if prices are changed every period this configuration will repeat itself indefinitely, except that all the nominal maginitudes – money and bond holdings, prices and permanent incomes – will rise at the weekly target inflation rate $\pi_w^*$.

### 5.1.1 Inflation in the unshocked flexible-price equilibrium

The model has been set up so that inflation has no almost no effect on this unshocked flexible-price equilibrium. An increase in the inflation target $\pi^*$ will leave unchanged the aggregate level of GDP per week, which will still equal its capacity level $y^{cap} = n(n-2-F)$ because everyone will continue to maintain his employment relationship indefinitely and aggregate demand will remain equal to $y^{cap}$. There will be a minor effect on the equilibrium tax rate $\tau$ because with higher inflation the government is now collecting more seigniorage and hence can maintain its target real debt level with a lower sales tax. Because of this reduction in the sales tax, the real wage $W/P = (1-\tau)/(1+\mu)$ of all non-shopowner transactors will rise a little, and (since aggregate income remains equal to $y^{cap}$) the income of each shopowner will fall a little. Real money balances at the start of each period:

$$\begin{aligned} M/P &= \left(W_{\_} * (N-n) + (1-\tau) P_{\_} y^{cap}\right) / P \\ &= \left((W/P) * (N-n) + (1-\tau) y^{cap}\right) / (1+\pi_w^*) \end{aligned}$$

will also change a little, although in which direction is not clear. However, aggregate output and employment will be unaffected, the real rate of interest will remain equal to the rate of time preference $\rho$, and the volatility of output and inflation will remain equal to zero.

To be more exact, under the parameter values assigned by the calibration exercise below, an increase in inflation from 0 to 9 percent per annum will reduce the equilibrium tax rate from 3.92 percent to 3.59 percent, raise the real wage rate from 0.835 to 0.838, reduce each shopkeeper's income from 4.73 to 4.60, and leaves real money balances per capita unchanged at 1.7.

**The effects of staggered price setting** When we leave the shocks turned off (by having no entry, exit or job breakup) but introduce staggered price setting we introduce some dispersion of relative prices, since not all prices will increase at the same time. However, this by itself does not have an effect on output, employment or the average rate of inflation, except to the extent that shopowners for some reason start running low on money holdings. That is, as long as they can continue to pay their employees we will have full employment of all transactors, with one shop active in each market, and the economy will continue to operate at full capacity.

In fact it turns out that when we simulate the economy with the parameter values discussed below, some shops do run out of money holdings, because of the randomness of price changes. But the economy continues to track very closely to full capacity output, full employment and zero volatility.

## 5.2 Shop profits

As we have seen, a shop that is earning profits lower than the aggregate wage rate is at risk of failing. Shop failure will be a key component in what follows, so a brief word is in order concerning what determines a shop's profit rate. By definition, a shop paying a wage rate $w$, with an input of $x$ and sales equal to $y$ will earn a profit equal to

$$\pi = ((1+\mu)\, y - x)\, w$$

In the long run, (or once it exhausts its inventories) its input will equal sales plus fixed cost:

$$y = x + F \tag{3}$$

so we have:

$$\pi = (\mu y - F)\, w \tag{4}$$

and its profit will be reduced by having a small markup or a large fixed cost. Profit will be positive if and only if sales are at least equal to the critical level:

$$y^c = \frac{F}{\mu} \tag{5}$$

which is increasing in $F$ and decreasing in $\mu$. In the unshocked equilibrium examined in the previous section we saw that each firm's sales will equal $n-2-F$. It follows that in this equilibrium all shops will be closer to the verge of failure the larger is $F$ and the smaller is $\mu$.

Once we introduce shocks to the system that reduce a firm's sales and/or input, anything that increases the variability of sales will make it more likely that a given shop will fall below the critical level needed to maintain profitability. This is because input will always be tending towards the level $y - F$. For a firm that employs more than this amount will have a target input greater than actual input only for as long as its inventories are below target, but having $x > y - F$ means inventories will be rising, and they will continue to rise beyond the point where they are equal to the desired stock, which is just $y$ itself. At this point the shop's target input will fall below actual (assuming the shopowner has enough cash to pay all its workers), and the firm will cease hiring. Exogenous quits will then cause $x$ to fall down to the level $y - F$. Indeed there will then be some overshooting because during the adjustment process the excess of actual inventories over desired inventories has been rising, so that target input will actually be less than $y - F$. Likewise a firm that employs $x < y + F$ will soon run out of inventories, in which case it will have to turn customers away and sell only the amount $x - F$.

With a reasonably fast adjustment speed of inventories, the firm's inventories will thus tend to track closely the level determined by sales according to equation (3), in which case its profit will tend to track the level given by equation (4), and accordingly anything that reduces its sales below the critical level (5) will quickly put the shop in danger of failure. Such an event might be a reduction in the price of some rival shop, which draws reduces this shop's sales, or a rash of quits. The bigger is the fixed cost $F$ and the smaller is the markup $\mu$, the more likely is such an event to trigger failure of the shop.

In particular, as we shall see, inflation will tend to raise the incidence of those events that bring at least some shops closer to the verge of failure, and this is a reason why inflation can have a deleterious effect on macro performance in the model.

## 5.3 The effects of shocks

When we turn the shocks on by allowing entry, exit and job breakup, the economy ceases to track full employment and full capacity, and the level of output becomes volatile, even though, to repeat, these shocks are purely individual-specific. The two shocks that disturb what would otherwise be an unchanging equilibrium are random job breakups and random innovation. Because unemployed transactors do not always find re-employment immediately, the fact that jobs are continually breaking up implies that the economy will almost always be operating at less than full capacity. The economy has a natural rate of unemployment greater than zero.

Innovation is also continually disturbing each market. New entrants compete with incumbents for employees and customers. The more successful they are in this competition the more likely are incumbents to fail, thus adding to the job loss attributable to exogenous breakups. This is because of the fixed costs, which imply that once a firm's sales and/or employment have fallen below a critical level it cannot break even and hence will face a positive probability of exiting.

Although these shocks are individual in nature, they can induce aggregate volatility because there is just a finite number of agents. For example, in the unshocked steady state examined above there is just one shop per market. If that one shop fails then approximately 1/n of the economy's labor force will move into unemployment. Moreover, job failure can spread through the economy via the multiplier process examined in detail by Howitt (2006b) in a related model. That is, those who lose their employment in one market because of a shop failure will demand less in other markets as their permanent income falls, thus putting other shops at greater risk of failure. Thus any constellation of shocks in some week that causes the failure of a shop that was employing a large number of transactors also threatens to cause a cumulative rise in unemployment that could last for a considerable time (in Howitt, 2006b the response of GDP to such a shock was very close to that predicted by the AR(2) process fit to US data by Chari, Kehoe and McGrattan (2001).

### 5.3.1 The interaction of inflation and shocks

The effects of these shocks will tend to be larger when inflation is higher, simply because inflation introduces an extra source of dispersion in relative prices, given that individual prices are changed at random times, and by a larger percentage amount the higher is the trend inflation rate. This variability in relative prices induces variability in each shop's sales and hence in its profitability, and this widening of the distribution of profitability ends up putting more firms near the edge of failure, and hence more firms likely to pushed into failure by a random shock to innovation or match breakup. In order to see in more detail how this might give rise to a real effect of inflation we need first to calibrate the model and then simulate it under different target inflation rates.

# 6 Calibration

Although the model has many agents we have imposed a great deal of ex ante symmetry. We have done this partly so that we can fully characterize the unshocked equilibrium that serves as a reference point, which will facilitate our analysis of what is generating our experimental results, and also partly so that we can keep the number of parameters small enough to calibrate them to US economic data. This section describes our calibration procedure.

There are a total of 22 parameters, which we have categorized as shop parameters, transactor parameters and government parameters. These are listed in the following three tables, along with their assigned values.

7 Shop parameters:

| | | |
|---|---|---|
| $\overline{x}$ | Maximal target sales | 84 |
| $\Delta$ | Length of contract period (in weeks) | 16 |
| $\overline{\mu}$ | Average markup over variable cost | 0.172 |
| $\phi$ | Failure rate of unprofitable shops (weekly) | 0.025 |
| $F$ | Fixed cost | 4.0 |
| $\lambda_I$ | Inventory adjustment speed (weekly) | 0.159 |
| $\beta$ | Wage adjustment coefficient (annual) | 0.3 |

6 Transactor parameters

| | | |
|---|---|---|
| $\varepsilon$ | Demand elasticity parameter | 10.0 |
| $\lambda_y$ | Permanent income adjustment speed (weekly) | 0.0185 |
| $\rho$ | Rate of time preference (annual) | 0.04 |
| $\theta$ | Frequency of innovation (weekly) | 10 |
| $\delta$ | Quit rate (per week) | 0.0032 |
| $\sigma$ | Search propensity | 0.41 |

9 Government parameters

| | | |
|---|---|---|
| $b^*$ | Target Debt-GDP ratio | 0.328 |
| $\lambda_\tau$ | Fiscal adjustment speed (annual) | 0.054 |
| $\lambda_\pi$ | Inflation autocorrelation factor (annual) | 0.286 |
| $\lambda_q$ | Output gap autocorrelation factor | 0.71 |
| $\gamma_\pi$ | Inflation coefficient in Taylor rule | 1.5 |
| $\gamma_q$ | Output gap coefficient in Taylor rule | 0.5 |
| $q^*$ | Target output gap in Taylor rule | 0.0725 |
| $\rho^*$ | Target real interest rate in Taylor rule | .0347 |
| $\pi^*$ | Target inflation in Taylor rule | 0.03 |

The calibration of parameters took place at three different levels. At the first level, one subset of parameter values was chosen to match empirical counterparts in the US data. At the second level, the values of other parameters were chosen so as to be internally consistent with average simulation outcomes. At the third level the values of the remaining parameters, for which we could find no convenient empirical counterparts, were chosen to make the average simulation outcomes match certain properties of the US data. We first describe the former level.

## 6.1 First level of calibration

### 6.1.1 Shop parameters

Bils and Klenow (2004) report that about 1/4 of prices are changed in an average month, indicating an average pricing interval of 4 months, or 1/3 of a year. Accordingly we set the length of contract period $\Delta$ to 16, which in our model is 1/3 of a year.

Estimates of the degree of returns to scale in the US economy vary from 0 to about 30 percent. This is commonly measured as the ratio of average to marginal cost (minus unity). In our model the typical firm in a steady state with input equal to $x$ and sales equal to $x - F$ would thus have a degree of returns to scale equal to

$$\frac{AC}{MC} = \frac{Wx/(x-F)}{W} = \frac{1}{1-F/x}$$

If the economy was operating with a 5 percent average unemployment rate then the typical firm would have

$$x = 0.95 \cdot (n-2) = 45.6$$

so by setting the fixed cost $F$ equal to 4 we get a typical degree of returns to scale equal to 9.6 percent.

The inventory adjustment speed $\lambda_I = 0.159$ corresponds to the estimate by Durlauf and Maccini (1995) of a monthly adjustment speed equal to approximately 0.5 ( $= 1 - (1-0.159)^4$ ).

Roberts (1995) estimated aggregate expectations-augmented Phillips relations with a coefficient on detrended output between 0.25 and 0.334 using annual data. A linear approximation to our wage-adjustment equation yields a same relation if we assume that actual/capacity output ratio is proportional to the target/potential input ratio. Accordingly we chose $\beta = 0.3$ to lie near the mid point of Roberts's range of estimates.

### 6.1.2 Transactor parameters

Carroll (2001) recently confirmed Friedman's (1957) estimate that the marginal propensity to consume out of transitory income is about 1/3 in annual data. Setting the weekly adjustment speed $\lambda_y$ equal to 0.0185 matches this figure, in the sense that if income were to rise by one unit for 48 weeks straight, then the average effect on consumption over these weeks implied by our consumption function (holding financial wealth constant) would be 1/3.

We set the annual rate of time preference $\rho$ equal to 0.04 as is standard in the real business cycle literature.

### 6.1.3 Government parameters

The target Debt-GDP ratio $b^*$ was set equal to 0.328 because this is the average ratio of federal marketable debt to GDP in the US on average between 1969 and 2005. The fiscal adjustment speed $\lambda_\tau$ was estimated at 0.054 by Bohn (1998). The autocorrelation factors $\lambda_\pi$ and $\lambda_q$ were taken from estimates of a univariate AR(1) process on CPI inflation and on (linearly-detrended) log per-capita GDP using annual data for the US over the 1984-2006 period. The coefficients $\gamma_\pi$ and $\gamma_q$ are Taylor's original specification in an interest-rate rule that is a linear approximation to ours around a point where the inflation, target inflation and the output gap are all equal to zero. In the calibration exercise to be described immediately below we took the inflation target $\pi^*$ to equal 3%, which is the average in the US over the period from 1984-2006.

## 6.2 Second level of calibration

### 6.2.1 Government targets - Finessing Wicksell

The two government targets - the target real interest rate $\rho^*$ and the target output gap $q^*$ were chosen to be internally consistent. That is, they were chosen so as to make the average real rate equal $\rho^*$ and the average gap equal to $q^*$ over all years of all simulation runs. This procedure effectively assumes away many of the pitfalls of interest-rate control that economists from Wicksell to Friedman and up through Orphanides have warned of - the danger that no-one knows the economy's natural rate of interest or the steady-state level of the output gap, and hence that controlling the rate of interest to the neglect of the money supply risks aggravating volatility or even engendering a cumulative inflation.

Our motive in doing this is to focus on instability problems that arise from the trading process, not from these well-known pitfalls. In future work we intend to introduce learning on the part of the central bank in order to take a fresh analytical look at these old issues.

### 6.2.2 Markups and the Lucas Critique

In early trials with the model we found a very large negative effect of target inflation on all our indicators of economic performance. But this was under the assumption that all shops applied the same markup, which was invariant to the trend inflation rate. Awareness of the Lucas critique prompted us to revise the model in favor of the current assumption, namely that each shop picks its markup at birth. This variant allows the economy-wide average markup to respond endogenously to the policy environment, through the evolutionary selection mechanism implicit in the exit process. As we will see, average markups do indeed respond, being systematically higher when trend inflation is higher. Moreover, this endogenous response of markups tends to mitigate the large effects of inflation that we found when there was a single uniform markup.

This still leaves us with the difficult question of deciding the mean of the distribution from which markups are drawn. We decided to let that be determined the same way that the values of the government targets $\rho^*$ and $q^*$ were determined - by internal consistency. We find that with a target inflation rate of 3 percent the ex post average markup across all years and runs is equal to the mean of the distribution from which each firm's markup is drawn when that mean equals 0.172.

## 6.3 Third level of calibration

This leaves six parameters still to be determined, $\overline{x}, \phi, \theta, \delta, \sigma$ and $\varepsilon$ . These were chosen by searching (manually) for parameter values that would generate average values of six different indicator variables, on average over 10,000 simulations of 40 years each, that approximated their counterparts in US data.

The six indicator variables are listed below, along with their data values and the values in our fully calibrated model averaged over 10,000 simulations:

| Variable | Data | Model |
|---|---|---|
| Unemployment rate | 5.0 | 5.0 |
| Volatility of output gap | 1.59 | 1.74 |
| Volatility of inflation | 1.16 | 1.46 |
| Average markup | 10 to 20 | 0.17 |
| Exit rate | 46.4 | 47.1 |
| Job loss rate | 0.69 | 0.69 |

All numbers are expressed as percentage points. The unemployment rate is the average over the period from 1984 to 2006. The volatility of the output gap is the standard deviation of linearly-detrended log per-capita GDP over the same period and the volatility of inflation is the standard deviation of annual US CPI inflation over the same period. Golosov and Lucas (2007) indicate that estimates of the percentage markup vary between 10 and 20 percent. The exit rate is the fraction of all firms found operating in a given industry in one census year that are found still operating in that industry in the next census year (five years later), which Dunne, Roberts and Samuelson (1988) report to be

46.4 percent. The job-loss rate is the weekly rate of job loss that would give rise to the number reported by Hall (1995), namely that 71.8% of people working at a given date have been working continuously for the same employer for the past year.[9]

These variables were chosen because it seems they would represent the effects of the remaining variables to be calibrated. That is, an increase in the innovation rate $\theta$ or in animal spirits $\overline{x}$ might raise volatility by raising the fraction of potential entrants that disrupt the status quo in a given industry, or it might lower volatility by helping to replace failed shops before the multiplier process has gained full force; an increase in the demand elasticity parameter $\varepsilon$ should reduce the average markup; an increase in the failure rate $\phi$ should raise the exit rate; an increase in the quit rate $\delta$ should increase the job loss rate and an increase in the search propensity $\sigma$ should reduce the average unemployment rate.

One of the exciting aspects of doing this kind of agent-based modeling is that it when it comes to matching the data on some variable that was collected in a way that makes it different from any of the parameters in the model, one is able to introduce into the code a data collector that acts like the agency that gathered the data in questions. We did this in particular with the exit rate, by introducing a census taker who identified all the firms that were active in the final week of each 5th year, and were able to identify which of them was active in the previous census.

For the other variables what we did was much more mundane. The unemployment rate in the model was the average across all weeks and across all simulations of the number of transactors having a positive effective supply price, as a fraction of the total number of transactors. The volatility of the output gap was computed during each simulation from annual data on total output (itself compiled from weekly observations of all transactions), and then averaged across all simulations. The volatility of inflation was computed similarly. The average markup was computed as the average across all weeks and all runs of the sales-weighted average markup across all shops. And the job loss rate was the average across all runs and weeks of the number of transactors losing their employer, through an exogenous quit or the exit of the employer.

As the above table shows, we were only partially successful in mimicking the data with these remaining 6 parameters. Interestingly enough the variable for which we went to the greatest trouble to measure as it had been done in the data, the exit rate, is one that we were able to match almost exactly. The unemployment rate also came out to be almost exactly the same as in the data. However both of our variance measures came out too high.

[9] That is, $(1-\alpha)^{48} = 0.718$ if $\alpha = 0.00688$.

### 6.3.1 Volatility and individual shocks

We suspect that the excessive volatility of the model arises partly because we do not have any interest smoothing in our Taylor Rule, something that we plan to introduce in the next round of simulations. Another reason for excess volatility is that we have too few goods. In the undisturbed equilibrium there are only 50 active shops, one for each good, so an individual shop failure can cause roughly 2 percent of the economy's GDP to disappear. If there were more goods and hence more shops, then individual shocks would have less impact on the aggregate economy. In the infinite limit we would expect a law of large numbers to result in zero volatility. So with the right number of goods we could no doubt match any given degree of volatility.

Rather than going this route we gave up on matching volatility more closely. The main reason[10] is that we think individual shocks do matter for aggregate volatility, but because of something that does not exist in our model, namely fat tails in the size distribution of firms (see Gabaix 2005). In the current version of the model the symmetry of types and the limited potential for substitution in demand guarantee that we will not get this kind of size distribution. One of the items on the top of our agenda is to modify the model to allow a more realistic size distribution, and to see if these modifications also deliver a realistic degree of aggregate volatility.

## 6.4 Approximating a steady state

Although each of the data points described in the above table is an average over 10,000 40-year (1920-week) runs, we actually allowed each run to proceed for 60 years (2880 weeks), and ignored the first 20 years of data when computing the averages. This was in order to give the economy time to approximate a steady state, having been started in the above-described initial position which was a steady state only under the assumption of no shocks and weekly price adjustment by all firms. As Figure 1 below shows, 20 years was indeed enough for the cross-run average values of the real interest rate, the output gap, the inflation rate and the markup to become more or less constant.

INSERT FIGURE 1 HERE

[10] Another reason is that we are trying to match US data over the period since the beginning of the "great moderation." The model's volatility is not excessive in comparison to the US economy since 1969, over which period the standard deviation of the detrended gap and of inflation were 2.94 and 3.10 respectively.

# 7 Main results

To investigate the effects of the central bank's inflation target on different indicators of macro performance, we ran a baseline experiment, in which we used the parameter values calibrated as above, and for each value of 10 different values of the inflation target $\pi^* = 0, ..., 9$ we ran 5,000 simulations of 60 years each, again counting only the last 40 years. For each inflation target we computed the cross-simulation average of the following 4 statistics, each of which can serve as a (negative) indicator of macroeconomic performance: the average (across the 40 years) output gap and unemployment rate, the volatility (standard deviation across the 40 years) of the output gap and of annual inflation. The results are illustrated in Figure 2a below.

INSERT FIGURE 2a HERE

As this Figure illustrates, we find that as the target rate of inflation rises from 0 to 9 percent per annum, expected macroeconomic performance deteriorates somewhat according to all the indicator variables except output volatility, which is hardly affected. The output gap for example rises by 6.5 percent (from 7.215% of capacity when $\pi^* = 0$ to 7.684% when $\pi^* = .09$). unemployment rises by 4.25 percent and inflation volatility rises by 9.7 percent.

The increase in these three indicators was not however monotonic. Instead, in all cases the economy appeared to perform better under a one-percent inflation target than under a zero target, but then performance deteriorated as the target was raised above one percent. In this sense the model seems to indicate that one percent is the optimal inflation target.

Before discussing the economic significance of these results we first address two questions that are raised by our computational procedure. The first question is whether 5,000 runs under each inflation target was enough to establish statistically significant results. Figure 2b shows the 95% confidence interval bands for the average output gap. For example, for $\pi^* = 0$ we have a sample of size 5,000 on the 40-year average gap; the mean of the sample is 7.215% and the 95% confidence interval goes from 7.195% to 7.235%. Inspection of this figure reveals that the U-shaped behavior of the average output gap with a minimum at a one percent target is indeed likely to be a property of the model rather than a statistical fluke.

INSERT FIGURE 2b HERE

The second question relates to the behavior of average markets and the Lucas critique. Figure 2c indicates that markups were higher on average the higher the inflation target. For each inflation target and each year of the simulation it shows the average across all runs of the sales-weighted average markup of all firms during that year. As discussed above, the rise in endogenous markups tends to mitigate much of the deleterious effect of higher inflation on economic performance.

INSERT FIGURE 2c HERE

To assess the economic significance of the results we first examine the likelihood that raising the inflation target above one percent will actually raise the average output gap. Figure 2d shows the interquartile range of the average gap across all runs for each inflation target, as well as the mean, median, 5th percentile and 95th percentile. It shows that there is enough variation from simulation to simulation that the interquartile ranges overlap considerably, suggesting that according to the model there is a significant probability that raising the inflation target even from 0 to 9 percent will actually reduce the average output gap.

INSERT FIGURE 2d HERE

To investigate this issue more closely we then restricted out comparison of average gaps to runs with the same initial seed of the random number generator. The results are shown in Figure 2e, which shows for each inflation target what fraction of the 5,000 seeds initiated a run in which the average output gap under that inflation target exceeded the average gap under a one percent target. This fraction is our estimate of the likelihood that the output gap will be higher than under a one percent target. The figure also indicates the 95% confidence interval band around that estimated likelihood. Thus the model predicts a 72 to 74 percent likelihood a nine percent target will generate a larger gap than a one percent target.

INSERT FIGURE 2e HERE

Another measure of the significance of these results is estimated cost of deviating from the gap-minimizing one-percent target. This can be measured in two different directions. What we call "right volatility" is the amount by which the gap rises going from its minimum to its value under a nine percent target. The amount by which it rises going from its minimum to its value under a zero percent target is what we call left volatility.

Thus right volatility is the cost of going all the way to nine percent whereas left volatility is the cost of going in the other direction and eliminating trend inflation altogether.

According to our baseline experiment, right volatility is equal to 0.52% of capacity GDP, indicating that raising trend inflation from its gap-minimizing value to nine percent would cost just over one half of one percent of capacity output. This is not a huge cost but it is a permanent one. So it is not negligible. For example, at a discount rate of 4%, the present value of this cost is 13% of capacity GDP. On the other hand, left volatility, which measures the cost of going all the way to zero inflation, is only 0.05 percent, an order of magnitude smaller than right volatility.

## 7.1 Interpreting the results

Our interpretation of the results is that raising the inflation target inflation intensifies the multiplier mechanism of Howitt (2006b) involving cumulative shop failures, whereas lowering it all the way to zero exacerbates the well-known problem of the zero lower bound on the nominal interest rate. A one percent target minimizes the sum of these two effects on the output gap.

Some evidence for this interpretation is provided by the behavior of price dispersion, failure rates and the incidence of binding ZLB under different target inflation rates. As discussed above, inflation should raise price dispersion and hence raise the dispersion of profitability, thus intensifying this multiplier process by making shops more vulnerable to failure. Indeed, we have kept track of the dispersion of prices, as measured by variance of the sales-weighted log of retail prices across all active shops, averaged across all weeks and all runs for each target inflation rate. As shown in Figure 2f below, this weighted measure of price dispersion does indeed rise monotonically as trend inflation rises from 0 to nine percent. This increase in price dispersion is matched by a monotonic increase in the average number of jobs lost due to shop failure per week, expressed as a fraction of employment. Both of these indicators of the multiplier mechanism are plotted in Figure 2f. The same figure shows that the incidence of a binding ZLB, measured by the fraction of fixed action dates at which it occurs, which is 16 percent when the inflation target is zero, falls monotonically and reaches almost zero by the time the inflation target has risen to five percent.

INSERT FIGURE 2f HERE

More direct evidence of the role played by the multiplier mechanism is provided by redoing the simulation with entry and exit turned off. That is, each local monopolist lasts forever, although there are still random job breakups, which can be interpreted as people randomly relocating. The results of this simulation are shown in Figure 3 below, which shows that with no entry or exit, volatility is almost zero while the average output gap unemployment become independent of the target inflation rate.

INSERT FIGURE 3 HERE

Another piece of evidence is provided by redoing the simulation with prices being reset every week. This takes away the main channel through which inflation raises price dispersion and thus intensifies the cumulative process of shop failures. When we do this by setting the contract period $\Delta = 1$, all indicators except for the volatility of inflation become independent of the target inflation rate. The results are shown in Figure 4 below.

INSERT FIGURE 4 HERE

Finally, in order to check that it is indeed the ZLB that is responsible for the deterioration that occurs going from a one percent target to zero, we redid the baseline experiment suppressing the ZLB in the code. The result as shown in Figure 5 below was to eliminate this deterioration except for a tiny rise in unemployment.

## 7.2 Sensitivity analysis

To check that these results were not a fluke induced by the specifics of our calibration, we conducted a sensitivity analysis with respect to eight key parameters (in the next revision we will do a more complete analysis). The parameters are listed in the following table. For each of these parameters, we redid the Baseline experiment reported in Figure 1 above but with the parameter set 25% above its baseline value (other parameters remaining at baseline) and then again with the parameter set 25% below its baseline value.

Sensitivity analysis

| | |
|---|---|
| $\overline{x}$ | Animal spirits |
| $\delta$ | Quit rate |
| $\varepsilon$ | Demand elasticity |
| $\phi$ | Failure rate |
| $F$ | Fixed cost |
| $\theta$ | Innovation rate |
| $\sigma$ | Search intensity |
| $\lambda_w$ | Wage flexibility |

The detailed results are indicated in Figures A1 ~A8 in the Appendix, and we can summarize them as follows. Qualitatively, it is still the case that for almost all parameter values, and according to all the macroeconomic indicators except for output volatility (which did not respond to the inflation target in the baseline experiment), macro performance was either monotonically increasing or U-shaped in the target inflation rate. (The one exception was search propensity; when it was reduced by 25%, all the indicators except inflation volatility improved slightly but monotonically when target inflation rose above five percent.) In all cases, the output gap was minimized at a target value of either 0, 1 or 2 percent. Figure 8b below shows that for all parameter values the gap-minimizing inflation target was zero for either the high or low value and above zero for the other value.

INSERT FIGURE 8b HERE

The sensitivity analysis does more than confirm the qualitative message of the baseline experiment reported above. It also helps to understand the experiment by identifying which parameters matter the most, both for the overall level of macro performance and also for the deterioration of performance resulting from high inflation. That is, in some cases the increasing relationship between, say the output gap and the inflation target is shifted up or down a lot by a change in the parameter and in some cases its slope is raised or lowered by the parameter change.

Figures 8a below shows a bar for each of the eight key parameters listed above, showing the direct effect of that parameter on the average output gap when the inflation target is 3%. The total length of the bar is the absolute difference between the average output gap when that parameter is 25% above its baseline value and when it is 25% below. The green portion of the bar represents the output gap when the parameter is 25% below baseline, and the orange portion represents the gap when the parameter is 25% above. On the vertical axis we indicate the size of the output gap as a proportion of capacity GDP. The red line shows the average gap in the baseline calibration. So, for example, when the animal spirits variable $\overline{x}$ is set 25% above its baseline value (at 105), the average output gap over the 5,000 simulations with $\pi^* = 0.03$ rises from 7.3% to 9.5% of capacity GDP, and when $\overline{x}$ is set 25% below baseline (at 63) the average gap falls to 4.4% of capacity GDP. The figure indicates that the four most important parameters in terms of the average size of the output gap are animal spirits, fixed cost, innovation frequency and wage adjustment speed.

INSERT FIGURE 8a HERE

The bar for each parameter in Figure 8c indicates right volatility, our measure of the cost of a 9% inflation target. The blue bar shows right volatility when the parameter is set 25% below baseline and the red bar shows 25% above. So according to this figure, the four parameters that are most important in determining the cost of inflation, as measured by the total length of the red and blue bars, are search propensity, innovation frequency, fixed cost and wage adjustment speed. Figure 8d does the same for left volatility, our measure of the cost of reducing the target all the way to zero.

INSERT FIGURES 8c AND 8d HERE

# 8 Conclusion

In summary, the agent-based computational model that we have laid out here shows how inflation can affect the mechanism of exchange in a decentralized market economy The effect works through the increased dispersion of relative prices caused by staggered price setting, and the attendant increased failure rate of the specialist trading enterprises that create and organize markets.

Although this effect does not seem to be huge it does seem to be robust to sizeable changes in parameter values in the model. Moreover, its size could grow as we make the model less stylized by making agents less sophisticated in adjusting for estimated inflation when making price comparisons or by adding elements of reality not yet taken into account, such as private financial institutions that allocate credit using behavioral rules that are ridden with inflation illusion (see Howitt, 1990).

Because of the highly stylized nature of the model, we see our contribution at this stage as being more methodological than substantive. More specifically, the fact that our computational analysis has been able to uncover effects of inflation that remain invisible to more conventional approaches, the fact that it can be calibrated to data generating through surveys by directly modelling the behavior of the surveyor, and especially the fact that it is capable of modelling the behavior of individuals in rich microeconomic detail, all suggest that this approach is worth pursuing further.

In future work we plan to shore up the substantive message of the paper by working on several problems to which we have drawn attention in the current analysis. In particular, we will study the endogenous evolution of behavioral parameters other than markups, especially contract periods and the degree of wage flexibility. We also plan to introduce such important elements of reality as downward wage inflexibility and interest-smoothing by the central bank. We will also explore alternative network structures instead of limiting ourselves to the simple symmetric configuration of tastes and endowments in the present model, with the goal of making the model consistent with the observed size distribution of firms.

# References

Akerlof, George A., William T. Dickens, and George L. Perry. "The Macroeconomics of Low Inflation." *Brooking Papers on Economic Activity* I (1996): 1-59.

Ascari, Guido. "Staggered Prices and Trend Inflation: Some Nuisances." *Review of Economic Dynamics* 7 (July 2004): 642-67.

Barro, Robert J., and Herschel I. Grossman. *Money, Employment and Inflation.* New York: Cambridge University Press, 1976.

Bénabou, Roland. "Search, Price Setting and Inflation." *Review of Economic Studies* 55 (July 1988): 353-76.

Bénabou, Roland. "Inflation and Efficiency in Search Markets." *Review of Economic Studies* 59 (1992): 299-329.

Bils, Mark, and Peter J. Klenow. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (October 2004): 947-85.

Binette, André, and Sylvain Martel. "Inflation and Relative Price Dispersion in Canada: An Empirical Assessment." Bank of Canada Working Paper 2005-28, October 2005.

Blanchard, Olivier, and Jordi Gali. "Labor Markets and Monetary Policy: A New-Keynesian Model with Unemployment." NBER Working Paper 13897, March 2008.

Bohn, Henning. "The Behavior of US Public Debt and Deficits." *Quarterly Journal of Economics* 113 (August 1996): 949-63.

Bullard, James, and John W. Keating. "The Long-Run Relationship between Inflation and Output in Postwar Economies." *Journal of Monetary Economics* 36 (December 1995): 477-96.

Bullard, James, and Kaushik Mitra. "Learning about Monetary Policy Rules." *Journal of Monetary Economics* 49 (September 2002): 1105-29.

Calvo, Guillermo A. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (September 1983): 383-98.

Carroll, Christopher D. "A Theory of the Consumption Function, with and without Liquidity Constraints." *Journal of Economic Perspectives* 15 (Summer 2001): 23-45.

Cecchetti, Stephen G., Alfonso Flores-Lagunes, and Stefan Krause. "Has Monetary Policy Become More Efficient? A Cross Country Analysis." NBER Working Paper No. 10973 2004.

Chari, V. V., Larry E. Jones, and Rodolfo E. Manuelli. "The Growth Effects of Monetary Policy." *Federal Reserve Bank of Minneapolis Quarterly Review* 19 (Fall 1995): 18-32.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?" *Econometrica* 68 (September 2000): 1151-79.

Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. "Patterns of Firm Entry and exit in US Manufacturing Industries." *RAND Journal of Economics* 19 (Winter 1988): 495-515.

Durlauf, Steven N., and Louis J. Maccini. "Measuring Noise in Inventory Models." *Journal of Monetary Economics* 36 (December 1995): 65-89.

Evans, George, and Seppo Honkapohja. "Adaptive Learning and Monetary Policy Design." *Journal of Money, Credit and Banking* 36 (December 2003, Part 2): 1045-72.

Feldstein, Martin. "Capital Income Taxes and the Benefit of Price Stability." NBER Working Paper No. 6200 1997.

Fortin, Pierre. "The Great Canadian Slump." *Canadian Journal of Economics* 29 (November 1996): 761-87.

Friedman, Milton. *A Theory of the Consumption Function.* Princeton, NJ: Princeton University Press, 1957.

Friedman, Milton. "Nobel Lecture: Inflation and Unemployment." *Journal of Political Economy* 85 (June 1977): 451-72.

Gabaix, Xavier. "The Granular Origins of Aggregate Fluctuations." Unpublished, MIT 2005.

Golosov, Mikhail, and Robert E. Jr Lucas. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (April 2007): 171-99.

Gomme, Paul. "Money and Growth Revisited: Measuring the Costs of Inflation in an Endogenous Growth Model." *Journal of Monetary Economics* 32 (August 1993): 51-77.

Hall, Robert E. "Lost Jobs." *Brooking Papers on Economic Activity* (1995): 221-73.

Head, Allen, and Alok Kumar. "Inflation, Output and Welfare." *International Economic Review* 46 (May 2005): 495-522.

Heymann, Daniel, and Axel Leijonhufvud. *High Inflation.* Oxford: Clarendon Press, 1995.

Howitt, Peter. "Zero Inflation as a Long-Run Target for Monetary Policy." In *Zero Inflation: The Goal of Price Stability*, edited by Richard G. Lipsey, 66-108. Toronto: C.D. Howe Institute, 1990.

Howitt, Peter. "Beyond Search: Money in Organized Exchange." *International Economic Review* 46 (May 2005): 405-29.

Howitt, Peter. "Discussion." In *Issues in Inflation Targeting: Proceedings of a Conference held by the Bank of Canada, April 2005*, 359-66. Ottawa: Bank of Canada, 2006a.

Howitt, Peter. "The Microfoundations of the Keynesian Multiplier Process." *Journal of Economic Interaction and Coordination* 1 (May 2006b): 33-44.

Howitt, Peter. "Macroeconomics with Intelligent Autonomous Agents." Unpublished, Brown University 2007.

Howitt, Peter, and Robert Clower. "The Emergence of Economic Organization." *Journal of Economic Behavior and Organization* 41 (January 2000): 55- 84.

Howitt, Peter, and Petros Milionis. "On the Robustness of the Calvo Price Setting Model." Unpublished, Brown University, June 25, 2007.

Kiyotaki, Nobuhiro, and Randall Wright. "On Money as a Medium of Exchange." *Journal of Political Economy* 97 (August 1989): 927-54.

Lach, Saul, and Daniel Tsiddon. "The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Price Data." *Journal of Political Economy* 100 (April 1992): 349-89.

Lagos, Ricardo, and Guillaume Rocheteau. "Inflation, Output, and Welfare." *International Economic Review* 46 (May 2005): 495-522.

Laidler, David. *Money and Macroeconomics.* Cheltenham, UK: Elgar, 1997.

Leijonhufvud, Axel. "Costs and Consequences of Inflation." In *The Microeconomic Foundations of Macroeconomics*, edited by G. C. Harcourt. London: Macmillan, 1977.

Levin, Andrew T., Fabio Massimo Natalucci, and Jeremy Piger. "Explicit Inflation Objectives and Macroeconomic Outcomes." European Central Bank Working Paper 383, August 2004.

Milani, Fabio. "Expectations, Learning and Macroeconomic Persistence." Unpublished, UC Irvine, August 2005.

Parks, Richard W. "Inflation and Relative Price Variability." *Journal of Political Economy* 86 (February 1978): 79-95.

Patinkin, Don. *Money, Interest, and Prices.* 2nd ed. New York: Harper and Row, 1965.

Preston, Bruce. "Adaptive Learning and the Use of Forecasts in Monetary Policy." unpublished, Columbia University 2004.

Ragan, Christopher. "On The Believable Benefits of Low Inflation." Bank of Canada Working Paper 98-15, September 1998.

Roberts, John M. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking* 27 (November (Part 1) 1995): 975-84.

Rocheteau, Guillaume, and Randall Wright. "Money in Search Equilibrium, in Competitive Equilibrium, and in Competitive Search Equilibrium." *Econometrica* 73 (January 2005): 175-202.

Rogerson, Richard, Robert Shimer, and Randall Wright. "Search-Theoretic Models of the Labor Market: A Survey." *Journal of Economic Literature* 43 (December 2005): 959-88.

Shubik, Martin. "Commodity Money, Oligopoly, Credit and Bankruptcy in a General Equilibrium Model." *Western Economic Journal* 11 (March 1973): 24-38.

Starr, Ross M., and Maxwell B. Stinchcombe. "Monetary Equilibrium with Pairwise Trade and Transaction Costs." Unpublished, University of California San Diego, September 3, 1998.

Starr, Ross M., and Maxwell B. Stinchcombe. "Exchange in a Network of Trading Posts." In *Markets, Information and Uncertainty: Essays in Economic Theory in Honor of Kenneth J. Arrow*, edited by Graciela Chichilnisky, 217-34. Cambridge UK: Cambridge University Press, 1999.

Stockman, Alan C. "Anticipated Inflation and the Capital Stock in a Cash-in- Advance Economy." *Journal of Monetary Economics* 8 (October 1981): 387-93.

Summers, Lawrence. "How Should Long-Term Monetary Policy be Determined?" *Journal of Money, Credit and Banking* 23 (August 1991, Part 2): 625-31.

Tesfatsion, Leigh. "Agent Based Computational Economics: A Constructive Approach to Economic Theory." In *Handbook of Computational Economics*, edited by Leigh Tesfatsion and Kenneth L. Judd, 831-80. Vol. 2. Amsterdam: North-Holland, 2006.

Tesfatsion, Leigh, and Kenneth L. Judd, eds. *Handbook of Computational Economics.* Vol. 2. Amsterdam: North-Holland, 2006.

Vining, Daniel R. Jr, and Thomas C. Elwertowski. "The Relationship between Relative Prices and the General Price Level." *American Economic Review* 66 (September 1976): 699-708.

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

Yun, Tack. "Optimal Monetary Policy with Relative Price Distortions." *American Economic Review* 95 (March 2005): 89-109.

**FIGURE 1: TIME PATHS OF MACROECONOMIC VARIABLES IN THE CALIBRATION EXERCISE**

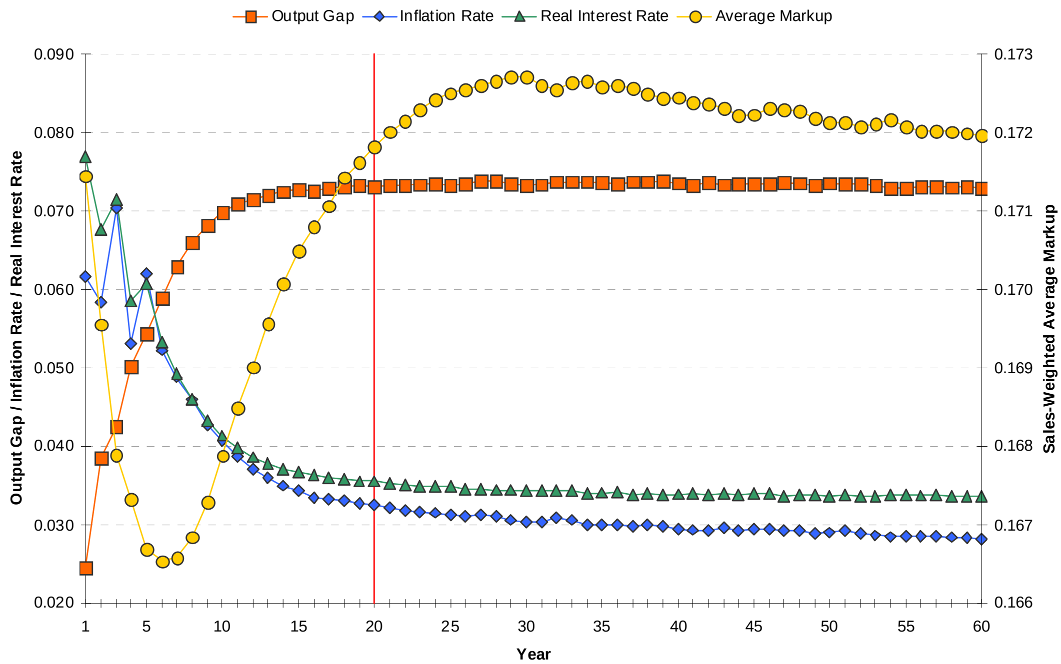


**FIGURE 2a: THE BASELINE EXPERIMENT**
Zero-Target Values: Output Gap=7.215%, Unemployment Rate=5.062%, Gap Vol.=1.750%, Inflation Vol.=1.478%

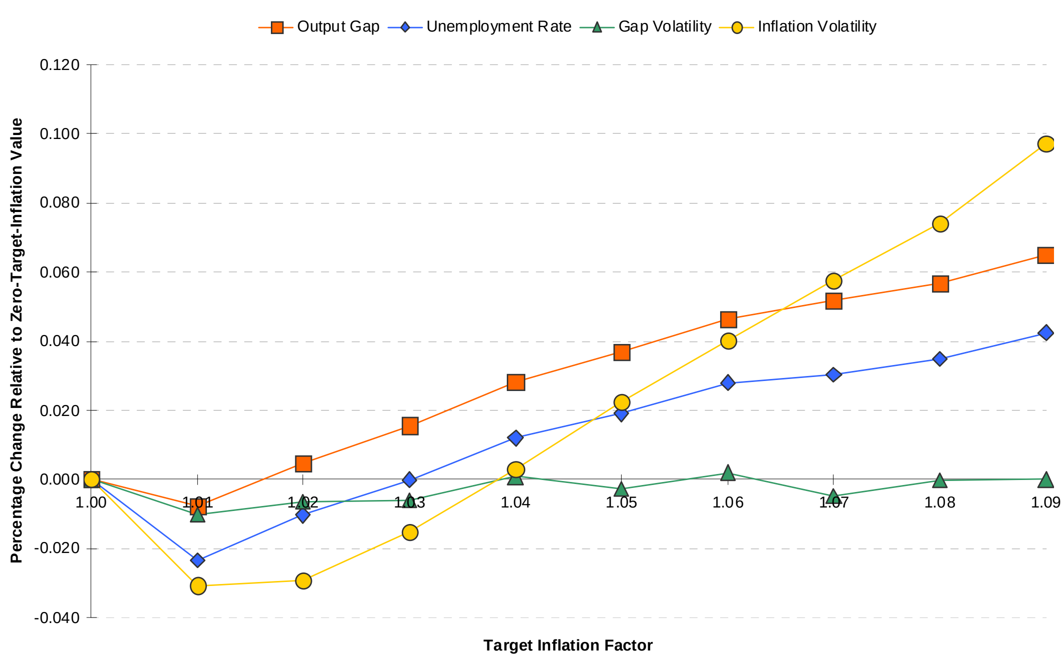

**FIGURE 2b: MEAN OUTPUT GAP IN THE BASELINE EXPERIMENT WITH 95% CONFIDENCE INTERVAL BAND**
Sample Size (Runs) Per Target = 5000; Error Bars depict 95% Confidence Intervals

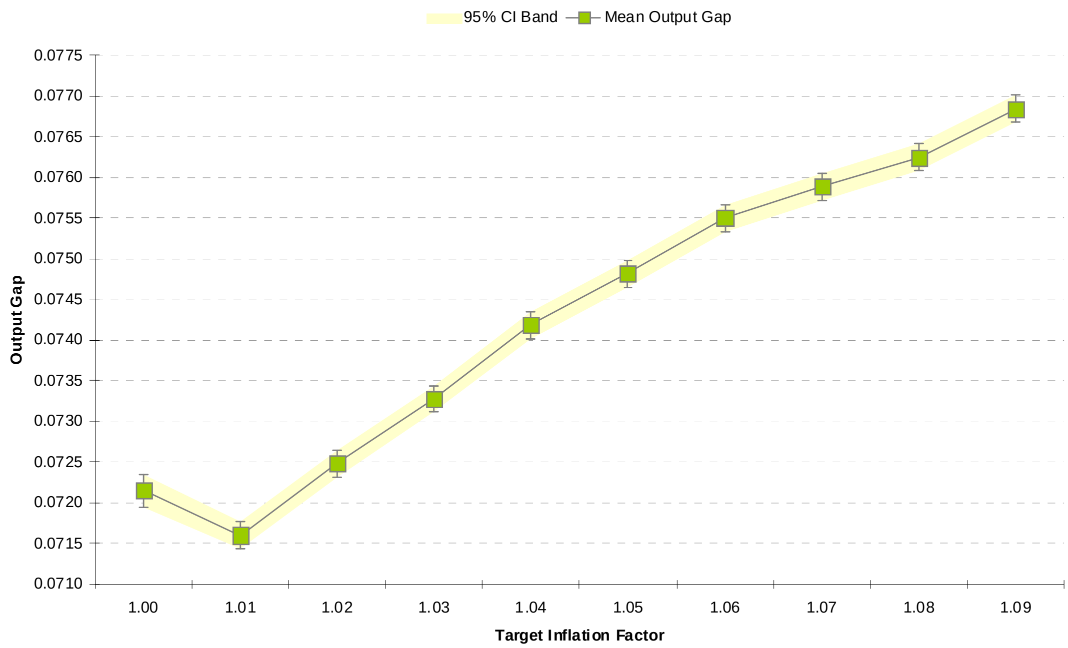


**FIGURE 2c: TIME PATH OF THE AVERAGE MARKUP UNDER VARIOUS TARGET INFLATION FACTORS**

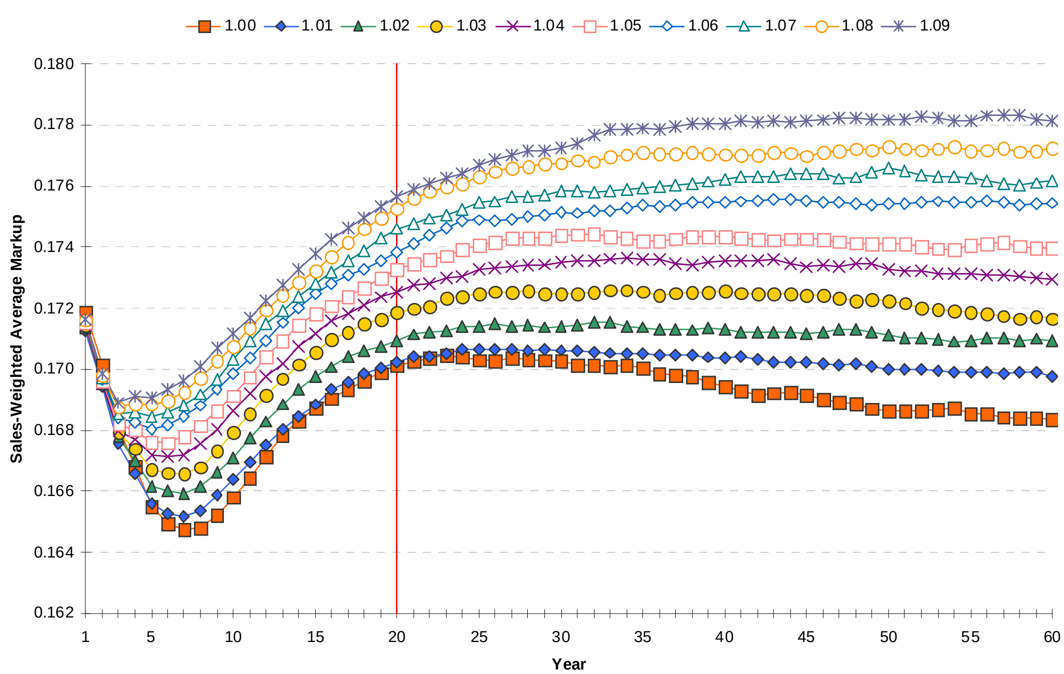

**FIGURE 2d: STATISTICAL DISPERSION OF THE OUTPUT GAP IN THE BASELINE EXPERIMENT**
Sample Size (Runs) Per Target = 5000; Error Bars Stretch Below to the 5th Percentile and Above to the 95th Percentile

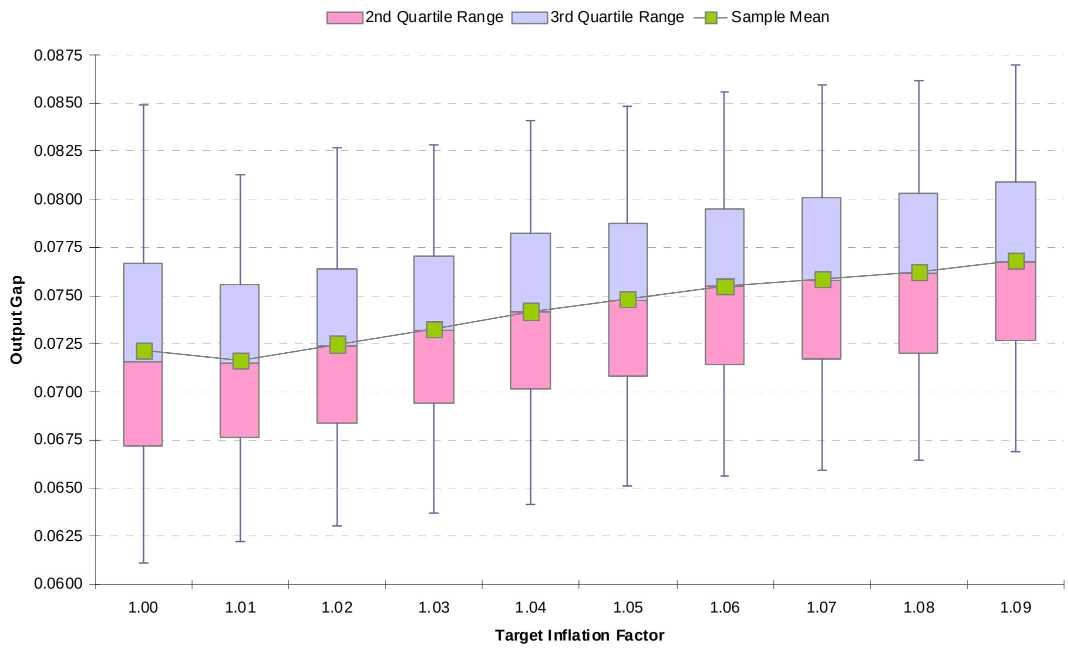


**FIGURE 2e: LIKELIHOOD THAT OUTPUT GAP IS LARGER THAN THAT AT 1% TARGET INFLATION RATE**
Sample Size (Runs) Per Target = 5000; Error Bars depict 95% Confidence Intervals

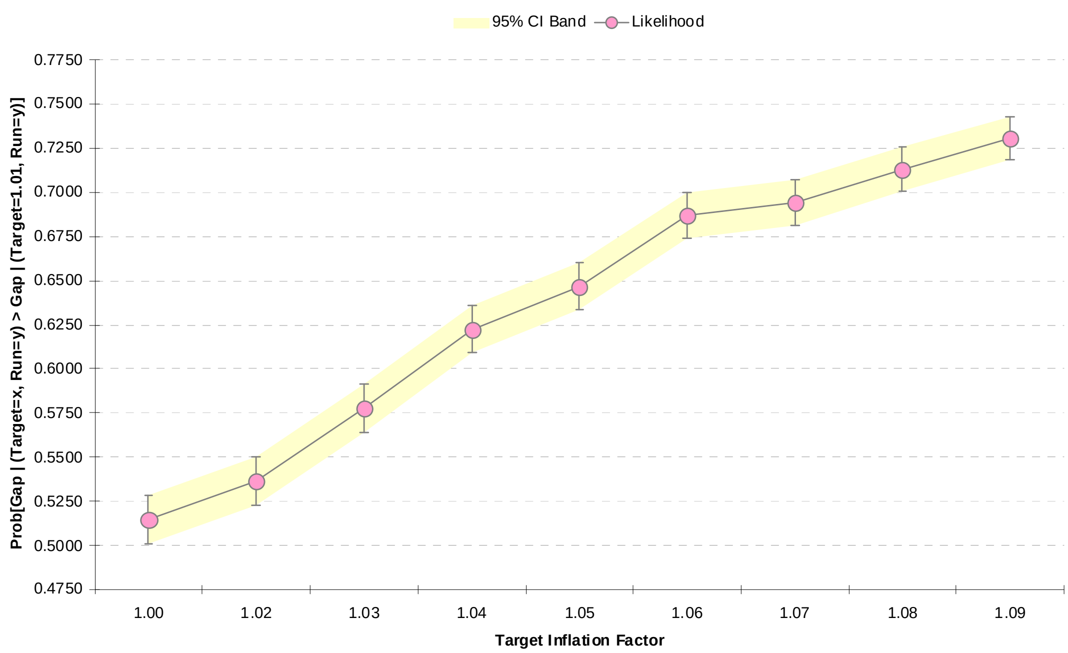

**FIGURE 2f: PRICE DISPERSION, FIRM FAILURES, AND HITS OF THE ZERO LOWER BOUND**
Zero-Target Values: Price Dispersion=5.418%, Firm Failures=0.320%, Zero Lower Bound Hits=15.795%

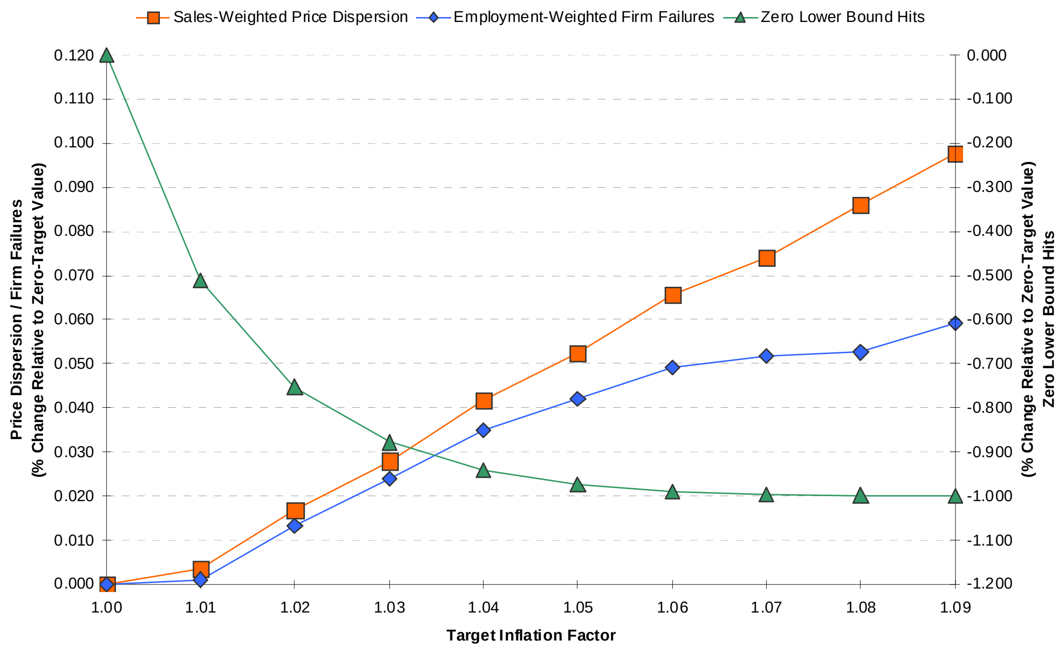


**FIGURE 3: THE BASELINE EXPERIMENT WITHOUT FIRM TURNOVER**
Zero-Target Values: Output Gap=0.497%, Unemployment Rate=0.456%, Gap Vol.=0.047%, Inflation Vol.=0.344%

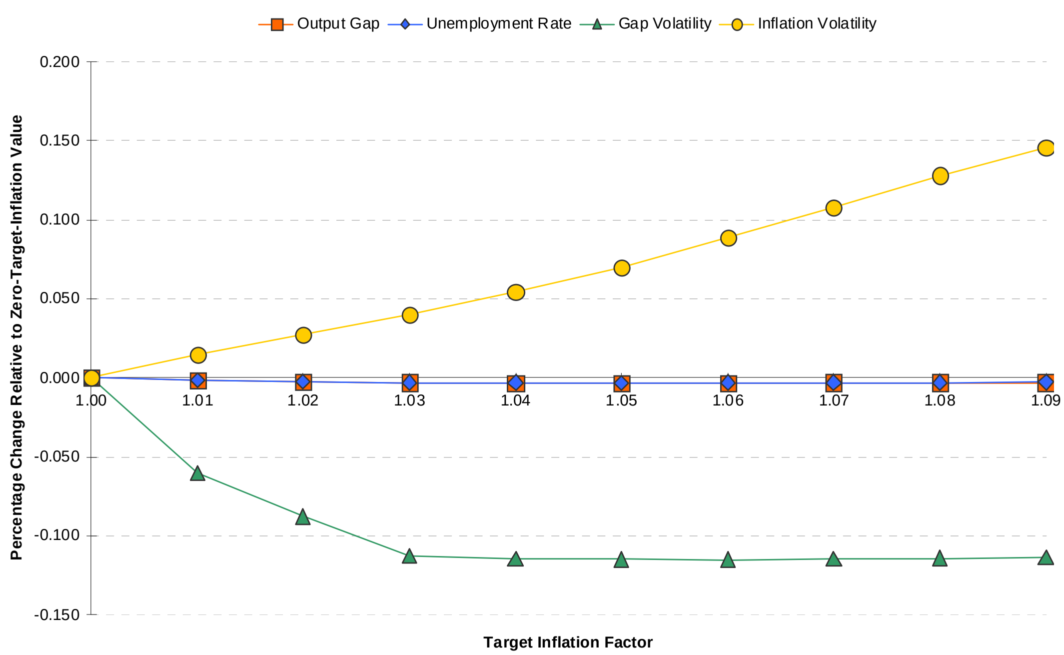

**FIGURE 4: THE BASELINE EXPERIMENT WITH FLEXIBLE PRICES**
Zero-Target Values: Output Gap=6.254%, Unemployment Rate=4.549%, Gap Vol.=1.600%, Inflation Vol.=1.037%

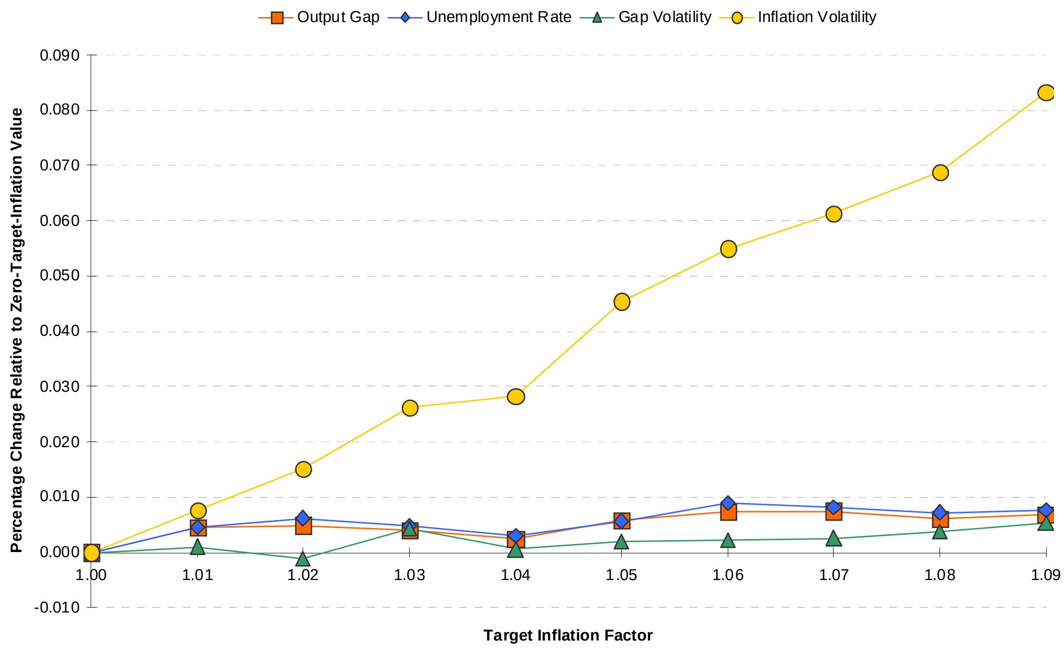


**FIGURE 5: THE BASELINE EXPERIMENT WITHOUT THE ZERO LOWER BOUND**
Zero-Target Values: Output Gap=7.073%, Unemployment Rate=4.948%, Gap Vol.=1.718%, Inflation Vol.=1.344%

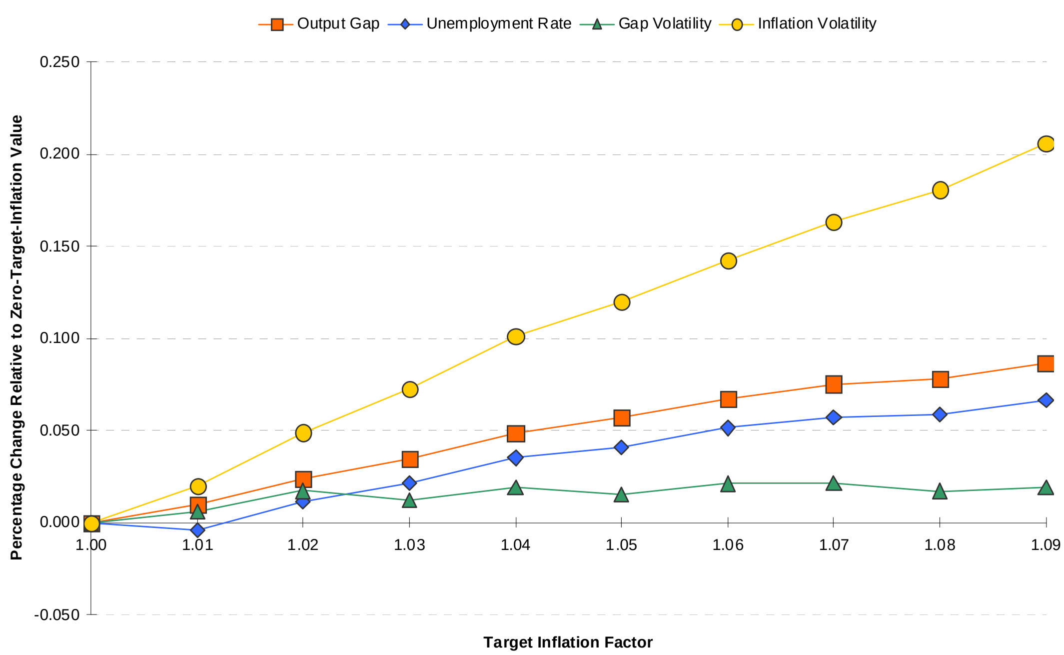

**FIGURE 6a: SENSITIVITY OF THE OUTPUT GAP AT 3% TARGET INFLATION RATE**

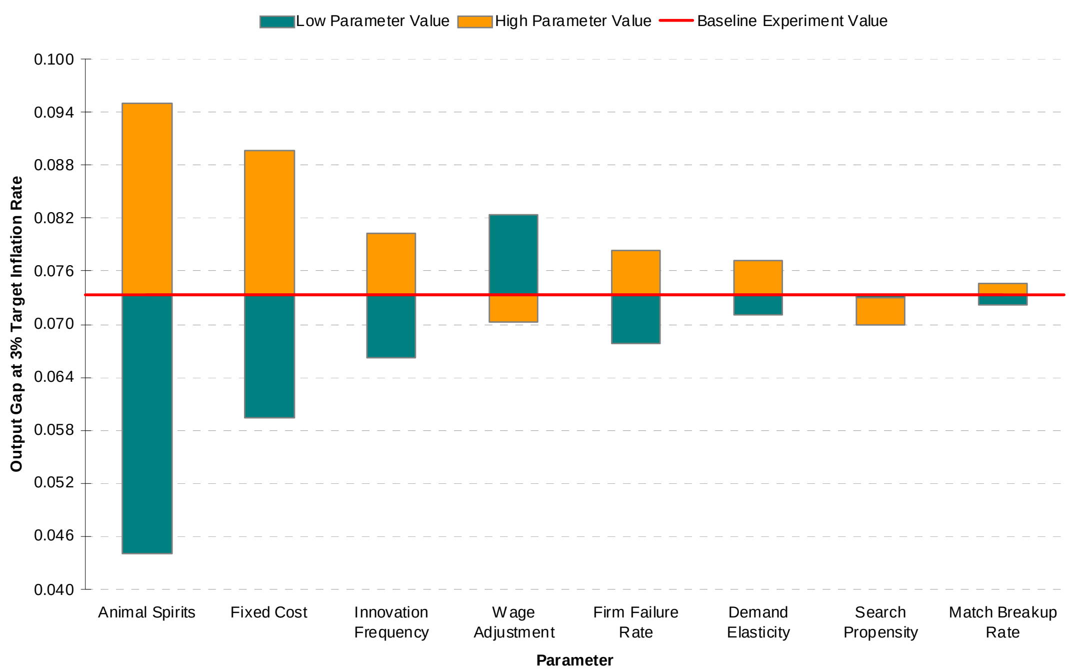


**FIGURE 6b: SENSITIVITY OF THE OUTPUT GAP-MINIMIZING TARGET INFLATION FACTOR**

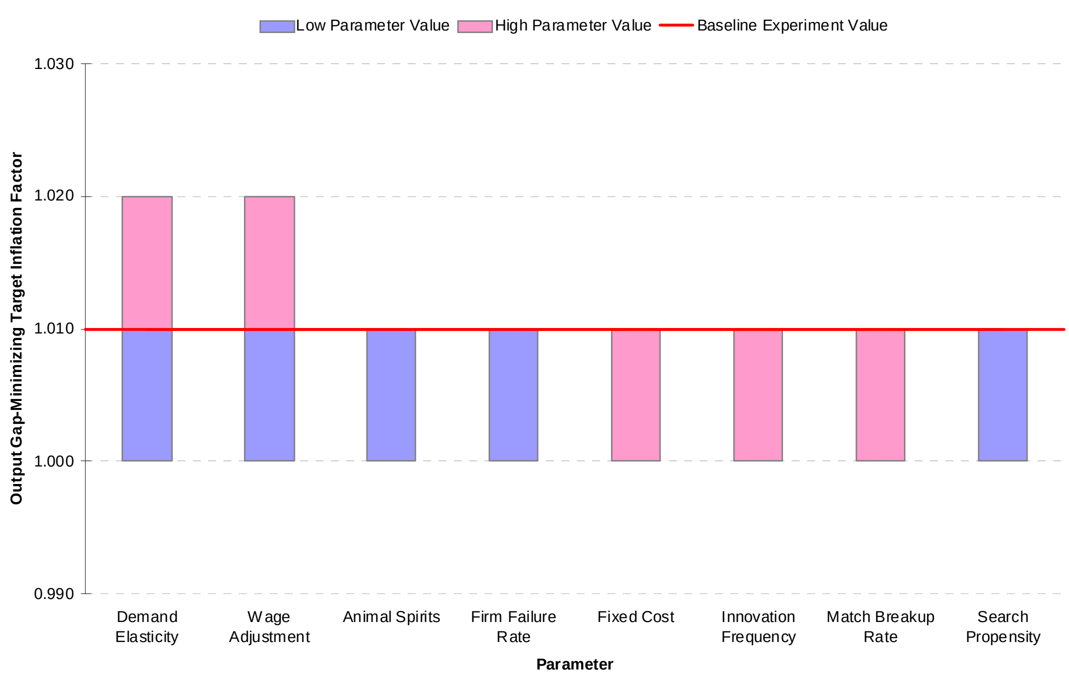

**FIGURE 6c: SENSITIVITY OF THE RIGHT AMPLITUDE OF THE OUTPUT GAP-TARGET INFLATION PROFILE**

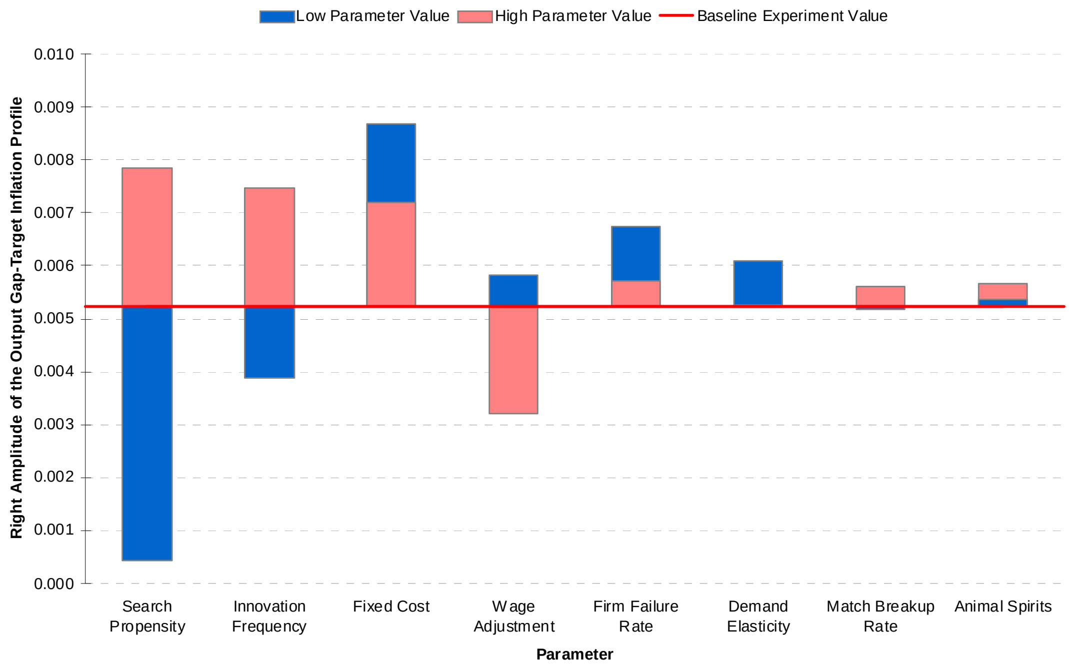


**FIGURE 6d: SENSITIVITY OF THE LEFT AMPLITUDE OF THE OUTPUT GAP-TARGET INFLATION PROFILE**

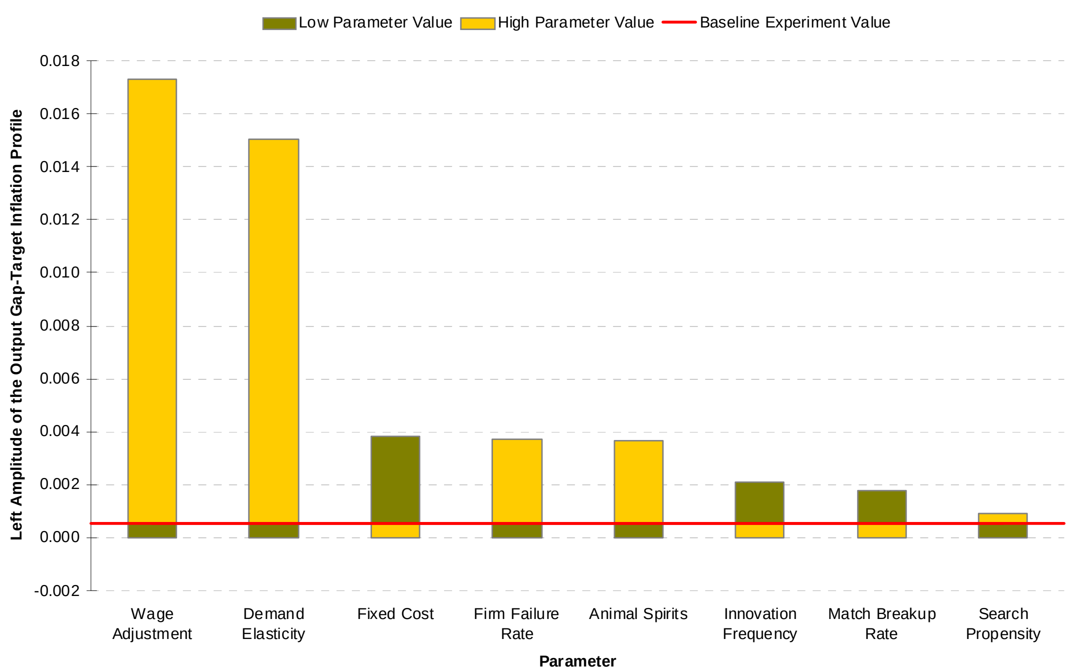

**FIGURE A.1a: EXPERIMENT WITH LOW ANIMAL SPIRITS [-25% OF BASELINE]**
Zero-Target Values: Output Gap=4.129%, Unemployment Rate=2.909%, Gap Vol.=1.360%, Inflation Vol.=1.217%

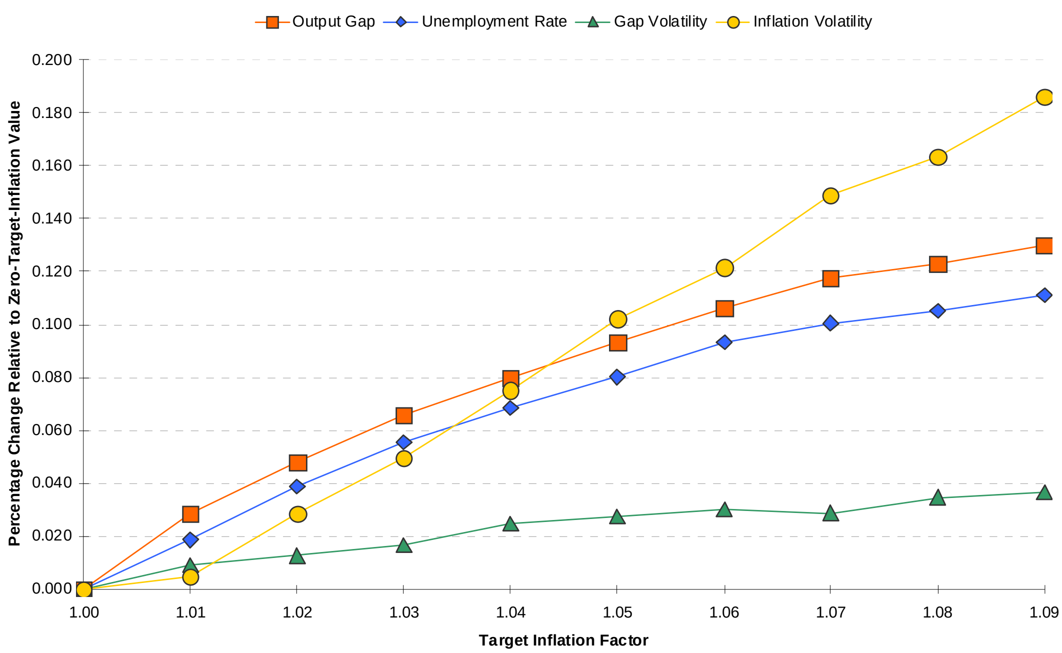


**FIGURE A.1b: EXPERIMENT WITH HIGH ANIMAL SPIRITS [+25% OF BASELINE]**
Zero-Target Values: Output Gap=9.716%, Unemployment Rate=6.829%, Gap Vol.=2.039%, Inflation Vol.=1.718%

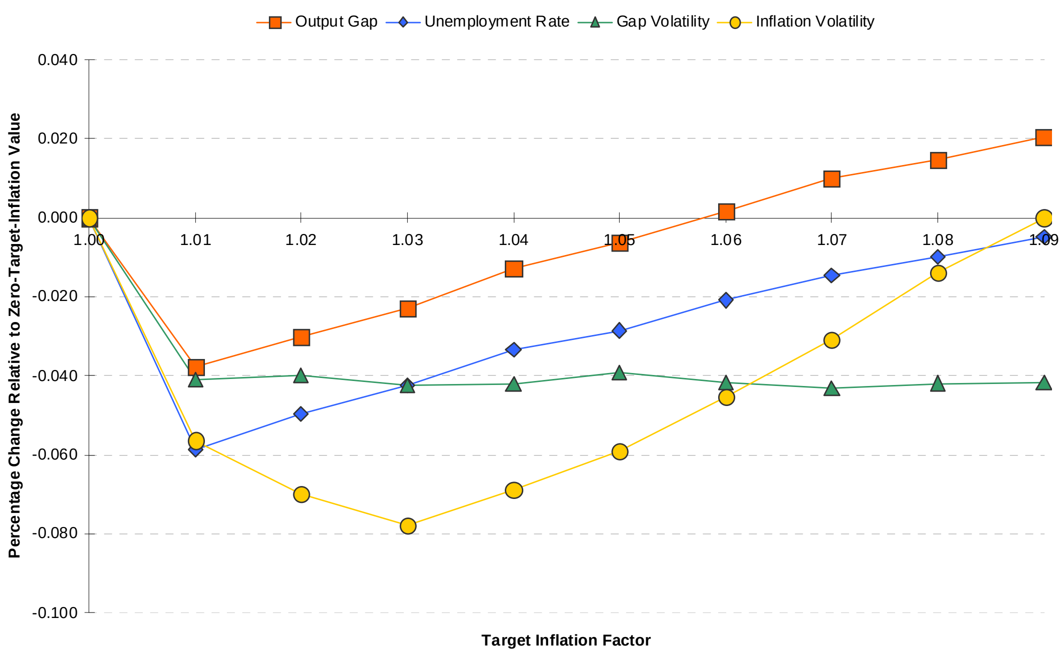

**FIGURE A.2a: EXPERIMENT WITH LOW MATCH BREAKUP RATE [-25% OF BASELINE]**
Zero-Target Values: Output Gap=7.219%, Unemployment Rate=5.056%, Gap Vol.=1.766%, Inflation Vol.=1.485%


**FIGURE A.2b: EXPERIMENT WITH HIGH MATCH BREAKUP RATE [+25% OF BASELINE]**
Zero-Target Values: Output Gap=7.227%, Unemployment Rate=5.086%, Gap Vol.=1.732%, Inflation Vol.=1.465%

**FIGURE A.3a: EXPERIMENT WITH LOW DEMAND ELASTICITY [-25% OF BASELINE]**
Zero-Target Values: Output Gap=6.775%, Unemployment Rate=4.755%, Gap Vol.=1.680%, Inflation Vol.=1.345%

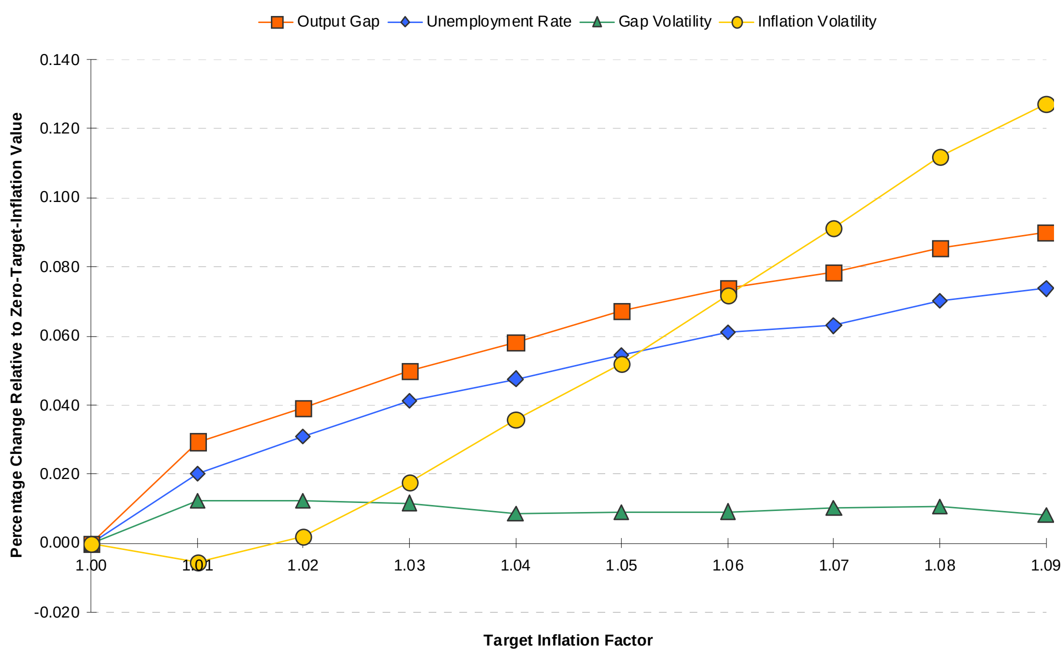


**FIGURE A.3b: EXPERIMENT WITH HIGH DEMAND ELASTICITY [+25% OF BASELINE]**
Zero-Target Values: Output Gap=9.137%, Unemployment Rate=6.494%, Gap Vol.=2.045%, Inflation Vol.=1.935%

**FIGURE A.4a: EXPERIMENT WITH LOW FIRM FAILURE RATE [-25% OF BASELINE]**
Zero-Target Values: Output Gap=6.459%, Unemployment Rate=4.248%, Gap Vol.=1.582%, Inflation Vol.=1.328%

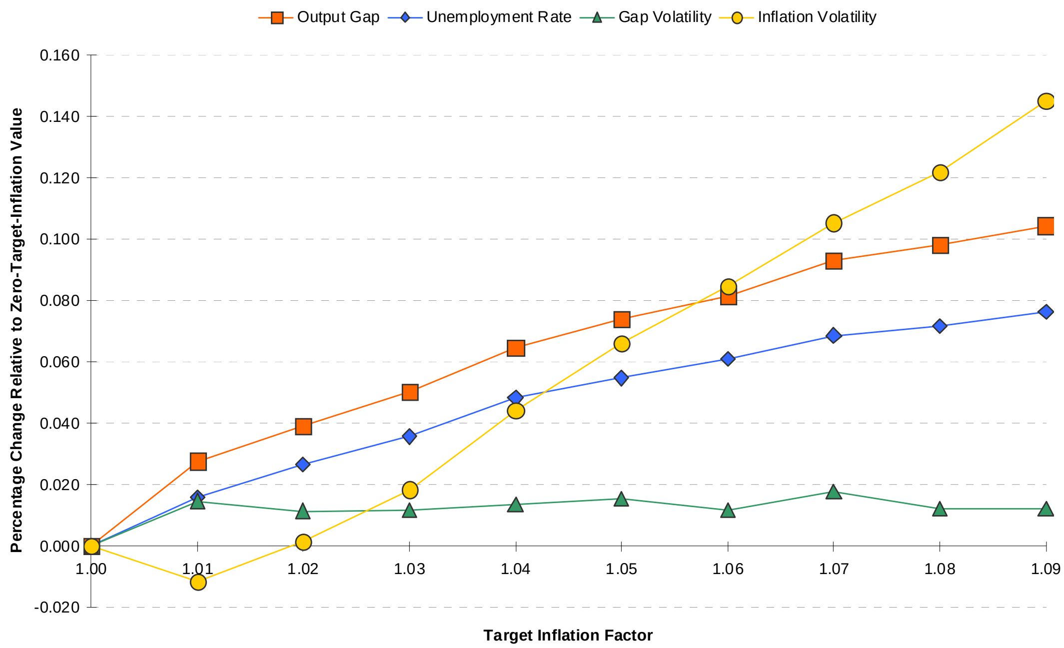


**FIGURE A.4b: EXPERIMENT WITH HIGH FIRM FAILURE RATE [+25% OF BASELINE]**
Zero-Target Values: Output Gap=8.025%, Unemployment Rate=5.906%, Gap Vol.=1.950%, Inflation Vol.=1.635%

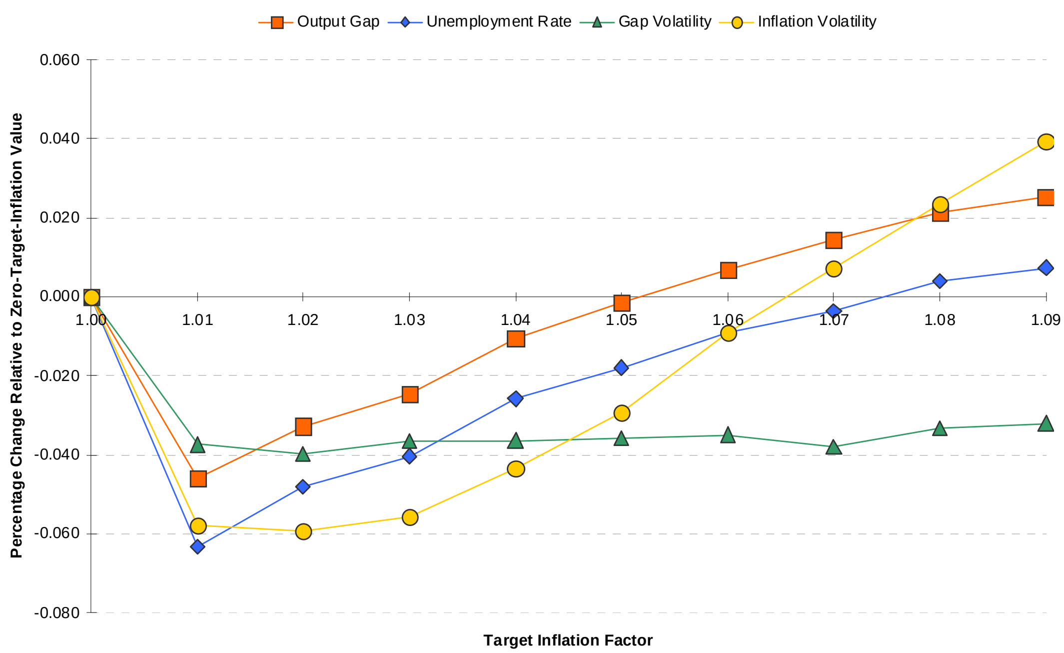

**FIGURE A.5a: EXPERIMENT WITH LOW FIXED COST [-25% OF BASELINE]**
Zero-Target Values: Output Gap=6.072%, Unemployment Rate=4.267%, Gap Vol.=1.543%, Inflation Vol.=1.349%

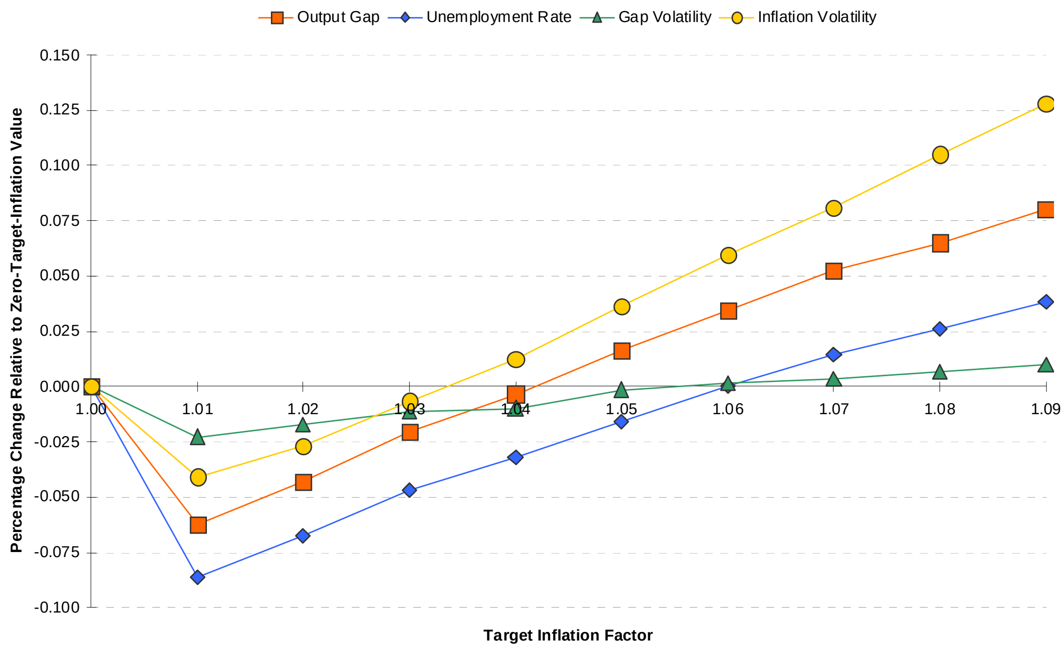


**FIGURE A.5b: EXPERIMENT WITH HIGH FIXED COST [+25% OF BASELINE]**
Zero-Target Values: Output Gap=8.763%, Unemployment Rate=6.292%, Gap Vol.=2.022%, Inflation Vol.=1.627%

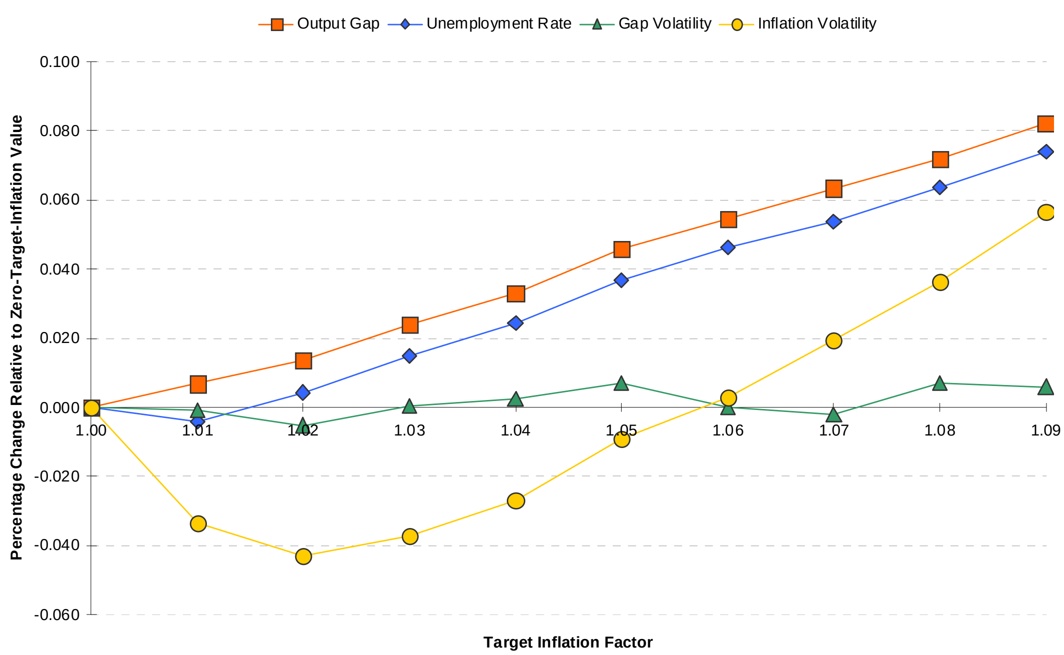

**FIGURE A.6a: EXPERIMENT WITH LOW INNOVATION FREQUENCY [-25% OF BASELINE]**
Zero-Target Values: Output Gap=6.661%, Unemployment Rate=4.850%, Gap Vol.=1.776%, Inflation Vol.=1.530%

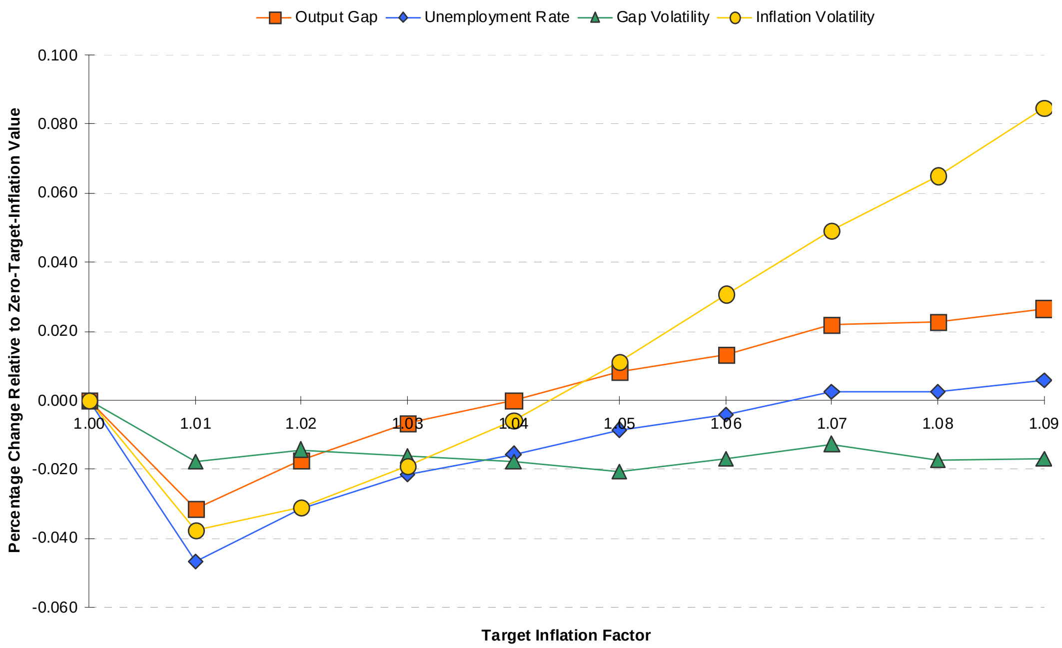


**FIGURE A.6b: EXPERIMENT WITH HIGH INNOVATION FREQUENCY [+25% OF BASELINE]**
Zero-Target Values: Output Gap=7.737%, Unemployment Rate=5.262%, Gap Vol.=1.739%, Inflation Vol.=1.443%

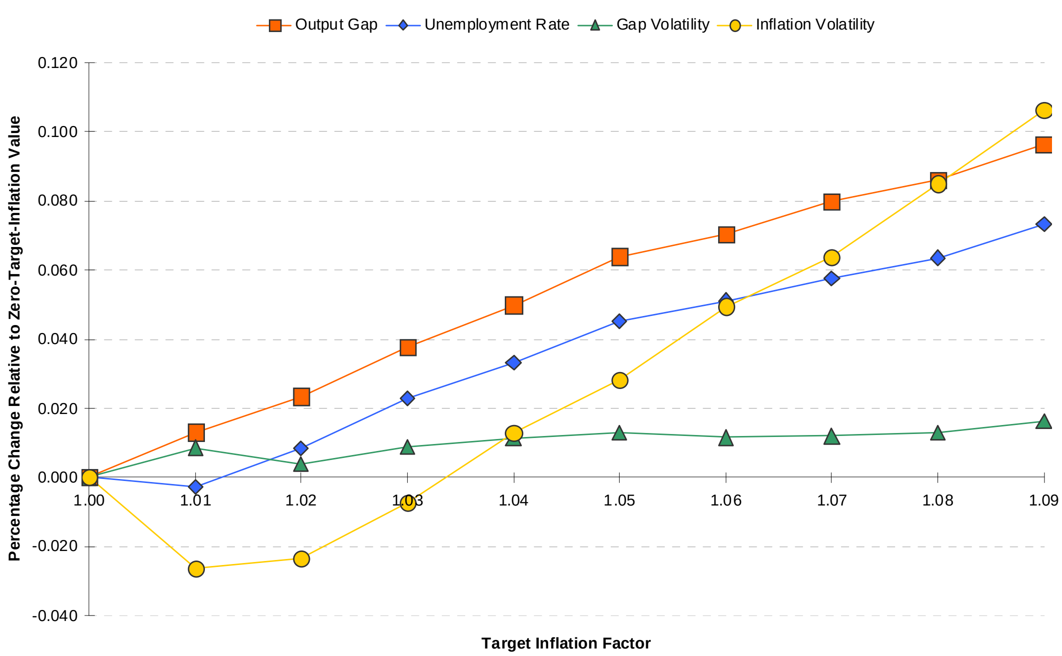

**FIGURE A.7a: EXPERIMENT WITH LOW SEARCH PROPENSITY [-25% OF BASELINE]**
Zero-Target Values: Output Gap=7.194%, Unemployment Rate=4.914%, Gap Vol.=1.726%, Inflation Vol.=1.477%

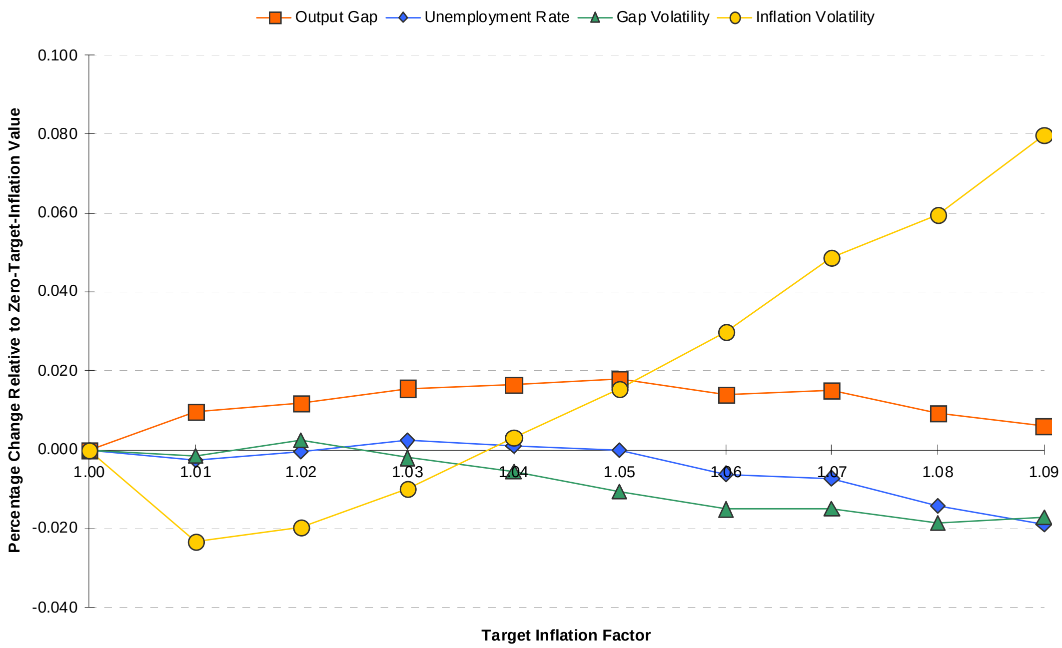


**FIGURE A.7b: EXPERIMENT WITH HIGH SEARCH PROPENSITY [+25% OF BASELINE]**
Zero-Target Values: Output Gap=6.848%, Unemployment Rate=4.857%, Gap Vol.=1.707%, Inflation Vol.=1.406%

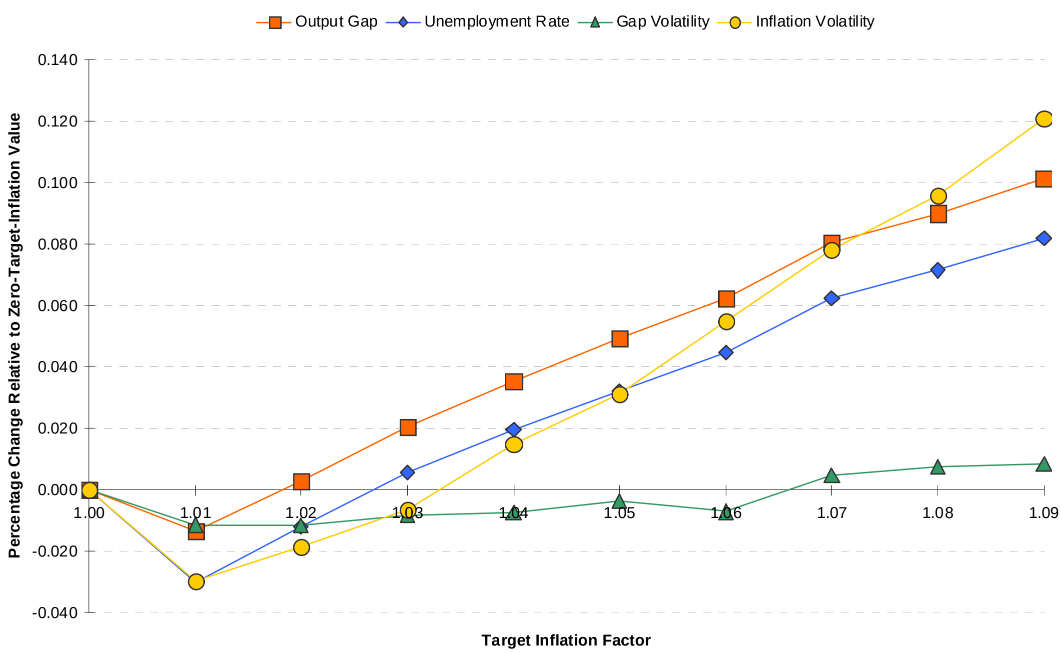

**FIGURE A.8a: EXPERIMENT WITH LOW WAGE ADJUSTMENT SPEED [-25% OF BASELINE]**
Zero-Target Values: Output Gap=7.971%, Unemployment Rate=5.786%, Gap Vol.=1.879%, Inflation Vol.=1.212%


**FIGURE A.8b: EXPERIMENT WITH HIGH WAGE ADJUSTMENT SPEED [+25% OF BASELINE]**
Zero-Target Values: Output Gap=8.745%, Unemployment Rate=6.107%, Gap Vol.=1.924%, Inflation Vol.=2.462%