The role of automatic stabilizers in the U.S. business cycle*

Alisdair McKay — Boston University

Ricardo Reis — Columbia University

August 2012

PRELIMINARY

Abstract

Most countries have automatic rules in their tax-and-transfer systems that are partly intended to stabilize economic fluctuations. This paper measures how effective they are at lowering the volatility of U.S. economic activity. We identify seven potential stabilizers in the U.S. data and discuss four theoretical channels through which they may operate. We then present a calibrated business cycle model including all of these stabilizers, and compare the volatility of output in the data with counterfactuals where some, or all, of the stabilizers are shut down. Our first finding is that proportional taxes, like sales, property income and corporate income taxes, contribute little to stabilization. Our second finding is that a progressive personal income tax can be effective at stabilizing fluctuations but at the same time leads to significantly lower average output. Our third finding is that safety-net transfers lower the volatility of output with little cost in terms of average output, but they significantly raise the variance of aggregate consumption. Overall, we estimate that if the automatic stabilizers were scaled back in size by 2% of GDP, then U.S. output would be XX% more volatile.

*Contact: amckay@bu.edu and rreis@columbia.edu. We are grateful to Susanto Basu for useful comments. Reis is grateful to the Russell Sage Foundation's visiting scholar program for its financial support and hospitality.

1 Introduction

Many features of the fiscal rules in developed countries guarantee that, during recessions, tax revenues fall and transfer spending rise. The CBO (2011) estimates that the automatic responses of policy account for $343 of the $973 billion U.S. deficit for 2012. These automatic stabilizers, as they are usually called, provide some countercyclical fiscal stimulus. While there is strong disagreement on the efficacy of discretionary fiscal spending to fight recessions, there is more consensus about the value of automatic stabilizers.[1] This consensus is especially strong among policy circles, with for instance the IMF (2009) recommending that countries should enhance the scope of these fiscal tools as a way to reduce macroeconomic volatility. In spite of this enthusiasm, as Blanchard (2006) noted: "very little work has been done on automatic stabilization [...] in the last 20 years."

This paper examines the efficacy of automatic stabilizers in attenuating the magnitude of the business cycle. More concretely, the goal is to answer the question: by how much do the automatic stabilizers in the U.S. tax-and-transfer system lower the volatility of aggregate activity? Our approach is to use a calibrated modern business-cycle model that captures the most important channels through which automatic stabilizers can affect the business cycle.

First, the model has nominal rigidities. Therefore, aggregate demand plays a role in the business cycle, so that stabilizing after-tax income and the demand for consumption and investment can stabilize fluctuations. This Keynesian channel is the most often cited reason for why automatic stabilizers would be effective. Second, the agents in our model intertemporally optimize so incentives and relative prices matter as well. This includes the distortions in the allocation of labor and capital induced by the tax and transfer system, which may affect behavior in a way that either attenuates or accentuates fluctuations. Third, households are heterogeneous in their wealth and income and there are incomplete insurance markets. Therefore, aggregate dynamics depend on the distribution of income and wealth. Because the stabilizers redistribute resources, they can affect the business cycle. Fourth, households have a precautionary demand for savings in response to the uncertainty they face. Because the stabilizers provide social insurance, their presence changes the targets for wealth and the ability of agents to smooth out shocks.

We start in section 2 by identifying the automatic stabilizers and measuring their size in the data. We propose a new measure of stabilization as the fraction by which the volatility

[1] See Auerbach (2009) and Feldstein (2009) in the context of the 2007-09 recession, and Auerbach (2002) and Blinder (2006) for contrasting views on the merit of countercyclical fiscal policy, but agreement on the importance of automatic stabilizers.

of aggregate activity would increase if we removed some, or all, of the automatic stabilizers. This differs from the measure of "built-in flexibility" introduced by Pechman (1973), which equals the ratio of changes in taxes to changes in before-tax income, and is widely used in the public finance literature. Whereas it measures whether there are automatic stabilizers, our goal is instead to estimate whether they are effective.

Sections 3 and 4 present our quantitative business-cycle model. With complete insurance markets, the model is similar to the neoclassical-synthesis DSGE models used for business cycles, as in Christiano et al (200x), but augmented with a series of taxes affecting every decision. With incomplete insurance markets, our model is similar to the one in Krusell and Smith (1998), but including nominal rigidities and many taxes and transfers. Methodologically, we believe this is the first model to include aggregate shocks, nominal frictions and heterogeneous agents in an analysis of aggregate fluctuations.[2] A technical contribution of this paper is to use the methods developed by Reiter (2009) to numerically solve for the ergodic distributions of endogenous aggregate variables so that we can compute their second moments.

Section 5 has our findings. First, under some extreme circumstances, even though the revenue and spending from the automatic stabilizers can be very cyclical, their effect on the business cycle is zero. Therefore, even if the stabilizers are present, they may not be effective. The intuition for this result leads to a lesson that persists under more general circumstances: proportional taxes, such as those on consumption, property or corporate income are ineffective stabilizers. It is well known that these taxes are distortionary, and can have a large effect on average economic activity. But their effect on the volatility of aggregate output is small.

Second, social transfers are very effective stabilizers in that they significantly lower output volatility with a negligible effect on its average level. Because they redistribute resources away from agents who choose to work longer in response and towards agents who have a higher propensity to spend them, transfers stabilize fluctuations. However, transfers greatly raise the volatility of consumption. Because they provide social insurance against idiosyncratic shocks, they induce fewer savings, which leaves households less able to smooth out aggregate shocks.

Third, the progressivity of the income tax is potentially quite stabilizing but also leads to a significantly lower average output. This progressivity ensures that marginal tax rates are

[2]Guerrieri and Lorenzoni (2011) and Oh and Reis (2011) are important precursors, but they both solve only for one-time unexpected aggregate shocks, not for recurring aggregate dynamics.

procyclical, which is both stabilizing but also discouraging of work and savings on average.

One common finding across these results is that social insurance and redistribution are the powerful channels through which stabilizers have their effects. Stabilizing income and cash-flow, while being the most emphasized channel in policy discussions of the stabilizers, is quantitatively weak in our calibration.

It is important to emphasize from the start that none of these conclusions are normative. We stay away from discussion of welfare, in part because with heterogeneous agents and fiscal redistributions, it would require controversial assumptions on how to calculate social welfare and weigh different individuals. Instead, this paper is an exercise in positive fiscal policy, in the spirit of Summers (1981) and Auerbach and Kotlikoff (1987). Like them, we propose a model that fits the US data and then change the tax-and-transfer system within the model to make positive predictions on what would happen to the business cycle.

Literature Review

There is an old literature discussing the effectiveness of automatic stabilizers (e.g. Musgrave and Miller 1947), but very few recent papers using modern intertemporal models. Christiano (1984) uses a modern consumption model, Gali (1994) a simple RBC model, and Andres and Domnech (2006) a new Keynesian model to ask a similar question. However, they typically consider the effects of a single automatic stabilizer, the income tax, whereas we comprehensively evaluate several of them. Moreover, they assume representative agents, therefore missing out on the redistributive channels of the automatic stabilizers that we end up finding to play an important role. Christiano and Harrison (1999), Guo and Lansing (1998) and Dromel and Pintus (2007) ask whether progressive income taxes change the size of the region of determinacy of equilibrium, whereas we use a model with a unique equilibrium, and focus on the impact of a wider set of stabilizers on the volatility of endogenous variables at this equilibrium.

Cohen and Follete (2000) are closer to our paper in their goal but their model is simple and qualitative, whereas our goal is to provide quantitative answers. Huntley and Michelangeli (2011) and Kaplan and Violante (2012) are closer in terms of modeling, but they focus on the effect of discretionary tax rebates, whereas our goal is to look at the automatic features of the fiscal code.

Empirically, Auerbach and Feenberg (2000, 2008), Auerbach (2009) and Dolls et al (2010) use micro-simulations of tax systems to estimate the changes in taxes that follows a 1% increase in aggregate income. The OECD (van den Noord, 2000) and the IMF (2009a) measure automatic stabilizers using instead macro data, to assess which components of revenue and spending are strongly correlated with the business cycle. Blanchard and Perotti (2002) and Perotti (200x) use these estimates to identify the effects of fiscal policy in vector autoregressions. We take these papers' measurement of the automatic stabilizers as inputs into our study of the effectiveness of these stabilizers.

We build on recent work by Oh and Reis (2011) and Guerrierri and Lorenzoni (2011) to try to incorporate business cycles and nominal rigidities into what Storesletten et al (2010) call the "standard incomplete markets." Close to our paper in emphasizing tax and transfer programs are Alonso-Ortiz and Rogerson (2010), Floden (2001), and Horvath and Nolan (2011), but they focus only on the effects of policies on average output and employment. Our focus is on volatility instead.

Finally, relative to the recent literature on fiscal policy during the recession, this paper focuses on taxes and government transfers, as opposed to government purchases.[3] In the United States in 2011, total government purchases were 2.7 trillion dollars. Government transfers amounted to almost as much, at 2.5 trillion. Focussing on the cyclical components, during the 2007-09 recession, which saw the largest increase in total spending as a ratio of GDP since the Korean war, 3/4 of that increase was in transfers spending (Oh and Reis, 2011), with the remaining 1/4 in government purchases. Looking at the largest individual discretionary spending, of the $494 billion spent in the American Recovery and Recovery Act, only $37 billion went to purchases program. Yet, relative to the large literature estimating purchase multipliers, the literature on the stabilizing properties of tax-and-transfer system is smaller. This paper contributes to close that gap.

2 The automatic stabilizers and their role

The automatic stabilizers are sometimes presented as the fiscal rules that attenuate the business cycle. For our purposes, this confuses defining the object of our study with measuring its effectiveness. Before proceeding, we must define what are the stabilizers, discuss by which channels they may affect the business cycle, and propose an independent measure to their effectiveness.

[3]For a survey, see the symposium in the *Journal of Economic Literature*, with contributions by Parker (2010), Ramey (2010), and Taylor (2010).

2.1 What are automatic stabilizers?

An automatic stabilizer is a rule in the fiscal system that leads to significant automatic adjustments in government revenues and outlays relative to total output in response to business-cycle fluctuations, partly with the intent of attenuating these fluctuations. In the words of Musgrave and Miller (1947), they are the built-in-flexibility in the tax-and-transfer system that ensure that in recessions taxes fall and spending rises. While this definition is broad, it does exclude some government policies.

First, it focuses on fiscal stabilizers. There are many other dimensions of public policy, notably monetary policy, that have features aimed at stabilizing real activity. Our focus is solely on the rules in the tax code and government spending programs.

Second, it excludes discretionary changes in policy. Most recessions come with "stimulus packages" of some form or another. There is already a tremendous amount of research on their impact. But, as **?** put it: "The advantage of automatic stabilization is precisely that it is automatic. It is not vulnerable to the perversities that arise when a discretionary "stimulus package" (or "cooling-off package") is up for grabs in a democratic government."

Third, while the automatic stabilizers are a component of a fiscal policy rule, they do not include all of the other systematic responses of fiscal policy to the state of the economy. To give one example, receiving benefits when unemployed is an automatic feature of unemployment insurance, while the decision by policymakers to extend the duration of unemployment benefits in most recessions is not.[4] There is still no consensus in the literature on what are the policy rules followed by U.S. policymakers, or even on how to best estimate them. In contrast, measuring automatic stabilizers is easier, because it requires reading and interpreting the written laws and regulations.

Fourth, we focus on the automatic fiscal rules that, either by initial design or by subsequent research, have been singled out as potentially contributing to mitigate output fluctuations. There are more government programs than a lifetime of research could study.

Given these restrictions on the rules that we will consider, we turn to the actual components of the U.S. budget to identify the stabilizers.

[4] To give another example, this one from monetary policy, the Taylor rule may be a systematic policy rule, but it is not automatic: there is no written rule that tries to enforce it on the actions of the Federal Reserve.

Figure 1: Ratio of change in taxes to change in gross income, Auerbach (2009)

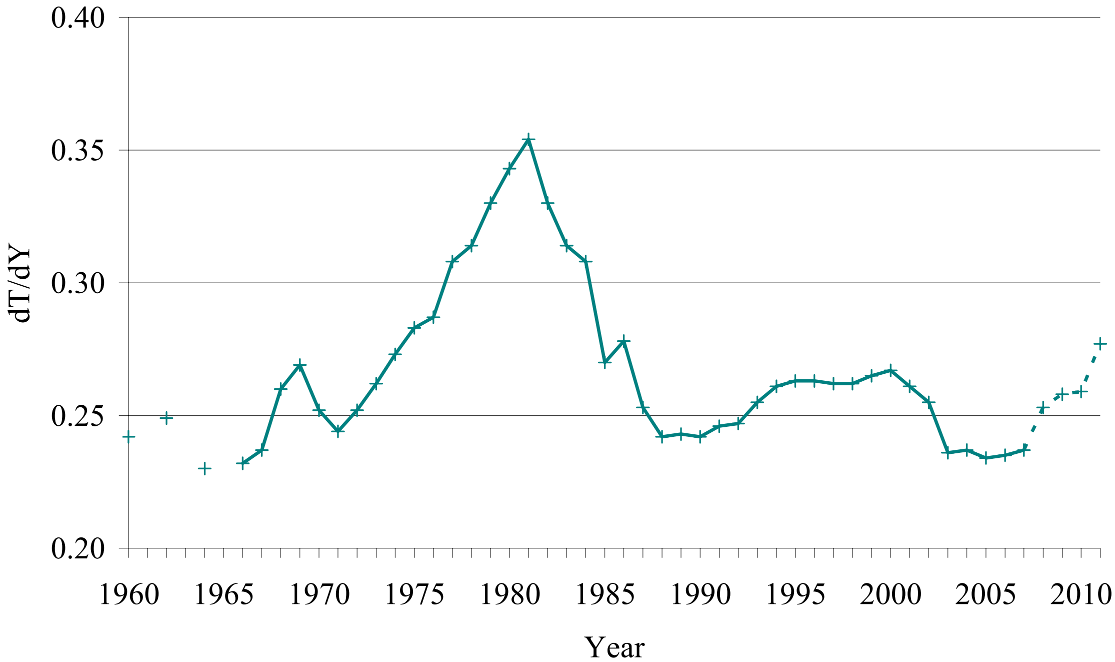

2.2 Automatic stabilizers in the United States

The classic automatic stabilizer is the *personal income tax* system. Because it is progressive in the United States, its revenue fall by more than income during a recession. Because they lower the variance of after-tax income, it is often argued that personal income taxes stabilize private spending. Figure 1 shows an estimate of their automatic component due to Auerbach and Feenberg (2000). Using micro-simulations based on the TAXSIM program, they asked by how much would a 1% increase in a typical household's income affect the amount of income taxes the household pays. The figure shows that a significant fraction of extra income goes into taxes, although this fraction has become less sensitive to the business cycle in the last decade.

While they are the most studied, personal income taxes are not the only stabilizer. Table 1 shows the main components of spending and revenue in the United States. The data comes from the National Income and Product Accounts (NIPA), so as to focus on the consolidated government flow of funds, across the different levels of government. The numbers are an average between 1988 and 2007, because earlier data would average significant changes in the structure of government, and the more recent years include large discretionary changes.[5]

[5]For instance, including the data form the start of the 1980s would imply averaging over a time with a

Table 1: Automatic stabilizers in the U.S. budget, average 1988-2007

Revenues		**Outlays**	
Progressive income taxes		*Transfers*	
Personal Income Taxes	11.22%	Unemployment benefits	0.29%
		Safety net programs	0.91%
Proportional taxes		Supplemental nutrition assistance	0.20%
Corporate Income Taxes	2.59%	Family assistance under PRWORA	0.17%
Property Taxes	2.75%	Security income to the disabled	0.35%
Sales and excise taxes	3.80%	Others	0.19%
Budget deficits		*Budget deficits*	
Public deficit	0.92%	Government purchases	15.14%
		Net interest income	2.25%
Out of the model		*Out of the model*	
Payroll taxes	6.26%	Retirement-related transfers	7.27%
Customs taxes	0.21%	Health benefits (non-retirement)	1.74%
Licenses, fines, fees	1.79%	Others (esp. rest of the world)	1.92%
Sum	29.52%	Sum	29.52%

We wanted a long enough sample to capture a few business cycles, but short enough to not mix very different fiscal regimes. The appendix describes how we aggregated the components of the government budget into the categories in the table.

Beyond personal income taxes, we consider three more stabilizers on the revenue side. Corporate income tax revenues vary by more than aggregate output because corporate profits are more volatile than national income, and it has been argued they may stabilize economic activity by lowering the volatility of corporate investment and dividends. Property taxes likewise vary with property prices and affect residential investment. Sales and excise taxes are rarely studied as automatic stabilizers, but we include them as they lower the variance of after-tax income needed to sustain a fixed real quantity of consumption. Because all of these three taxes have, approximately, a fixed statutory rate, we will refer to them as a group as *proportional taxes*.[6]

On the spending side, we consider two stabilizers working through *transfers*. The first, and most studied, is unemployment benefits, which greatly increase in every recession as the number of unemployed goes up. The second are safety-net programs, providing minimum support to poor households. Its main three components are food stamps, cash assistance

much higher corporate income tax rate, no earned income tax credit, and significantly lower spending on health care.

[6]Average corporate income taxes are progressive mostly as result of recurrent changes in investment tax credits during recessions that are not automatic. The rules for corporate income taxes or investment tax credits have few automatic features that would seem to vary over the business cycle.

to the very poor, and transfers to to the disabled. Most of the recipients of these three programs are out of the labor force, and their numbers increase during recessions.

A seventh stabilizer is the *budget deficit*, or the automatic constraint imposed by the government budget constraint. The previous stabilizers imply a rise in expenses and a fall in revenues during recessions, but there is no rule forcing these programs to pay for themselves during booms. There is also no automatic rule for government purchases, so these are typically not categorized as automatic stabilizers.[7] We will consider different rules for how deficits are paid and how fast to measure the impact of the deficit and the debt on volatility.

The last rows of the table include the fiscal programs that we will exclude. Some include licenses and fines, which have no obvious stabilization role. Others include international trade, like customs taxes and transfer to the rest of the world, which we leave out because we will consider a closed-economy model. Either way, these do not account for a large share of the government budget.

The main omissions are retirement, both in its expenses and the payroll taxes that finance it, as well as health benefits, which mostly are due to Medicare (for the elderly) and Medicaid (for the poor). These are large categories of the government budget, that we exclude from our study for two complementary reasons. First, to follow the convention as they are also excluded from the literature that measures structural deficits, stripping the government budget from the automatic stabilizers. Even the increase in medical assistance to the poor during recessions is questionable: for instance, in 2007-09 the proportional increase in spending with Medicaid was as high as that with Medicare. Second, we wanted our model to retain the core of conventional business-cycle models that are known to provide a satisfactory fit to the data. They typically ignore the life-cycle considerations that dominate choices of retirement and health spending, and so do we. This is a high priority for future work.

2.3 Channels for stabilization

The literature so far has proposed four possible channels by which automatic stabilizers can attenuate the business cycle.

First, is the *disposable income* channel, emphasized especially in Keynesian models and that dominates much of the policy discussion around stabilizers. The argument is that if after-tax income is less volatile than pre-tax income, then consumption and investment will

[7]See IMF, OECD, and Perotti for arguments why purchases should be excluded from the stabilizers, and Darby and Melitz (200x) for a rare dissenting view.

also be more stable. As long as aggregate demand determines output, then this will stabilize production. All four of the tax stabilizers discussed in the previous section make after-tax income less volatile than pre-tax income. For instance, transfers provide a minimum amount of income when pre-tax income has fallen to zero as a result of losing a job or leaving the labor force. This channel requires that disposable income has an effect on aggregate demand, and in turn that aggregate demand affects the business cycle. With rational forward-looking agents, under complete markets, changes in disposable income have almost no effect on consumption, which is driven by movements in permanent income. Moreover, with flexible prices, aggregate demand affects prices but not output. We include this channel in our model by assuming that households face liquidity constraints and that firms face nominal rigidities in setting prices.

The second channel works through *marginal incentives*, especially on labor supply. If the previous channel focusses on aggregate demand, this one works through aggregate supply. The intertemporal response of labor supply and investment to changes in marginal returns is the key driving force behind real business cycles. We include it by having an elastic labor supply in our model. This channel works especially through the progressivity of the personal income tax. In recessions, households move to lower tax brackets, which increases the relative return to working. The progressive income tax therefore stabilizes labor supply by encouraging intertemporal substation of labor from booms to recessions. A less studied example comes from property and corporate income taxes, which lower the variance of the after-tax return to investments.

The third channel is *redistribution*, and it interacts with the previous two. Both the progressive personal income tax and, especially, the transfer payments, imply a redistribution from higher-income to lower-income households. As discussed in Blinder (1974), this may raise aggregate demand if those that receive the funds have higher propensities to spend them than those who give the funds, and through nominal rigidities this may raise output in recessions. Redistribution may also work through labor supply, as in Oh and Reis (2011), if the recipients of transfer payments are at a corner solution with respect to their choice of hours to work, whereas those being taxed to fund the program, work more to offset the negative income effect. We include this channel by having incomplete insurance markets, so that the distribution of after-tax income affects economic aggregates.

Finally, we consider a *social insurance* channel working through precautionary savings. The automatic stabilizers provide insurance to households by lowering the taxes they pay and increasing the transfers they receive when they get hit by a bad shock. On the one hand, this reduces income and wealth inequality, while on the other hand it reduces the desire for precautionary savings, lowers aggregate savings and may increase pre-tax inequality (Floden, 2001, Alonso-Ortiz and rogerson, 2009). With incomplete markets, the wealth distribution will affect aggregate output. For instance, the social insurance provided by the stabilizers will likely lead agents to save less and become liquidity constrained more often, while at the same time making their spending choices less sensitive to hitting the liquidity constraint.

2.4 How to measure the effectiveness of automatic stabilizers?

At the macroeconomic level, the automatic stabilizers are effective if the variance of aggregate variables is lower in their presence. That is, letting $Y(\cdot, \tau)$ be a measure of real activity, then each element of the vector τ measure the strength of each stabilization program. We let $\tau = 1$ correspond to the status quo, and lower τ towards zero as we shrink the size of each automatic stabilizers in terms of its size in the budget. Our measure of effectiveness is the stabilization coefficient:

$$\mathbf{S} = \frac{Var(Y(\cdot, \tau))}{Var(Y(\cdot, 1))} - 1.$$

The measured $\mathbf{S}$ is the fraction by which the volatility of aggregate activity would increase if the stabilizers were decreased by τ. In the denominator is the status quo represented by our model calibrated to mimic the U.S. business cycle, while the counterfactuals in the numerator consist of shutting off different automatic stabilizers.

Some work in public finance, of which Dolls et al (2010) is a recent example, starts from the measures of the stabilizers in figure 1 and then makes behavioral assumptions on how demand changes with income for different households and how this affects output. Namely, they assume that households with certain certain characteristics (e.g., low financial wealth or no home) increase consumption one-to-one with income, while the marginal propensity of the other households is zero, and that aggregate demand equals output. This provides a different measure of the effectiveness of the stabilizers.[8]

This work measures exclusively the disposable income channel of stabilization. Moreover, it assumes extreme behavioral responses of consumption and aggregate output in the short run, while shutting off their dynamic effect especially in the long-run adjustment of prices and the wealth distribution. Finally, it does not take into account the general-equilibrium effect that changes in disposable income will have on rates of return, wages and prices in the economy. To include all of these effects and to assess how large they are, one needs a fully

[8]Devereux (2008) is a recent example of the same approach but applied to corporate income taxes.

specified model of, not just consumers, but all agents and markets. In short, one needs a business-cycle model. The next section provides one.

3 A business-cycle model with automatic stabilizers

Following the discussion of the channels by which automatic stabilizers may matter, we need a model that includes liquidity constraints, incomplete insurance markets, nominal rigidities, elastic labor supply, and precautionary savings. The model must also have room for the seven stabilizers that we want to study. And finally, we would like it to be close to business-cycle models that are known to capture the main features of the U.S. business cycle. The model that follows is the simplest we could write—and it is already quite complicated—while satisfying these three requirements.

Time is discrete, starting at date 0, and all agents live forever. The population has a fixed measure of $1 + \nu$ households.[9] Of these, a measure 1 refers to participants in the stock market, or capitalists, while the remaining ν refers to other households. The main difference between them is that capitalists are more patient. As a result, they end up accumulating all of the capital stock and owning all of the shares in firms. Following Krusell and Smith (1998), having heterogeneous discount factors allows us to math the very skewed wealth distribution that we observe in the data. Linking it to participation in financial markets matches the well-known fact since Mankiw and Zeldes (1989) that most U.S. household do not own any equity.

On the side of firms, there is a measure 1 of monopolistic intermediate-goods firms, a representative final-goods firm, and another representative capital-goods firm. Some of these agents could be centralized into a single household and a single firm without changing the predictions of the model. We keep them separate to ease the presentation, and introduce one automatic stabilizer with each type of agent.

The notation for the automatic stabilizers is that $\bar{\tau}$ are taxes collected, τ are tax rates, and T are transfers.

3.1 Capitalists and the personal income tax

The stock-owners are all identical ex ante in period 0 and share risks perfectly. We assume they have access to financial markets where all idiosyncratic risks can be insured, but this is

[9] Because we assume balanced-growth preferences, it would be straightforward to include population and economic growth.

not a strong assumption since they enjoy significant wealth and would be close to self-insuring even without financial assets.

We can then talk of a representative stock-owner, whose preferences are:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln(c_t) - \psi_1 \frac{n_t^{1+\psi_2}}{1+\psi_2}, \right] \tag{1}$$

where c_t is consumption and n_t are hours worked. These preferences ensure that there is a balanced-growth path in our economy and are consistent with the survey on the responses of labor supply to taxes in Chetty (2011).

The representative stock-owner budget constraint is:

$$\hat{p}_t c_t + b_{t+1} - b_t = p_t \left[x_t - \bar{\tau}^x(x_t) \right] + T_t^e. \tag{2}$$

The left-hand side has the uses of funds: consumption at the after-tax price $\hat{p}_t$ plus savings in riskless bonds b_t in nominal units. The right-hand side has real after-tax income, where x_t is the pre-tax income and $\bar{\tau}^x(x_t)$ are personal income taxes. The pre-tax price of consumption goods is p_t. The T_t^e are lump-sum transfers, which we will calibrate to zero as in the data, but will be useful later to discuss counterfactuals.

The real income of the stock owner is:

$$x_t = (i_t/p_t) b_t + w_t \bar{s} n_t + d_t. \tag{3}$$

It equals the the sum of the returns on bonds at nominal rate i_t, wage income, and dividends d_t from all the firms in the economy. The wage is the product of the average wage in the economy, w_t, and the agent's productivity $\bar{s}$. This productivity could be an average of individual-specific productivities of the measure 1 of stock-owners, since these idiosyncratic draws are perfectly insured within capitalists.

The first automatic stabilizer in the model is the *personal income tax* system. It satisfies:

$$\bar{\tau}^x(x) = \int_0^x \tau^x(x') dx', \tag{4}$$

where $\tau^x : \Re^+ \to [0, 1]$ is the marginal tax rate that varies with the tax base, which equals real income. The system is progressive because $\tau^x(\cdot)$ is weakly increasing.

3.2 Other households and transfers

Other households are indexed by $i \in [0, \nu]$, so that an individual variable, say consumption, will be denoted by $c_t(i)$. They have the same period felicity function as capitalists, but they are potentially more impatient $\hat{\beta} \leq \beta$, as we discussed earlier. Just like owners, individual households choose consumption, hours of work and bond holdings $\{c_t(i), n_t(i), b_{t+1}(i)\}$ to maximize:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \hat{\beta}^t \left[\ln(c_t(i)) - \psi_1 \frac{n_t(i)^{1+\psi_2}}{1+\psi_2} \right]. \tag{5}$$

Also like owners, households can borrow using government bonds, and pay personal income taxes, so their budget constraint and real income are:

$$\hat{p}_t c_t(i) + b_{t+1}(i) - b_t(i) = p_t \left[x_t(i) - \bar{\tau}^x(x_t(i)) \right] + p_t T_t^s(i), \tag{6}$$

$$x_t(i) = (i_t/p_t) b_t(i) + s_t(i) w_t n_t(i) + T_t^u(i). \tag{7}$$

There are two further constraints on the household choices. They also applied to capitalists, but will only bind for non stock owners. First is a borrowing constraint, $b_{t+1}(i) \geq 0$, which is equal to the natural debt limit if they cannot borrow against future government transfers. Second is a constraint that hours worked must be non-negative, $n_t(i) \geq 0$, which will bind if the household receives a sufficiently bad wage offer and chooses to not work.

Unlike stock owners, households face two sources of idiosyncratic risk regarding their labor income: on their labor-force status, $e_t(i)$, and on their skill, $s_t(i)$. If the household is employed, then $e_t(i) = 2$, and she can choose how many hours to work, since if $e_t(i) \neq 2$, then $n_t(i) = 0$ is an extra constraint. While working, her labor income is $s_t(i) w_t n_t(i)$. The shocks $s_t(i)$ captures shocks to the worker's skill, her productivity at the job, or the wage offer she receives. They generate a cross-sectional distribution of labor income.

With some probability, the worker loses her job, in which case $e_t(i) = 1$ and labor income is zero. However, now the household collects unemployment benefits $T_t^u(i)$, which are taxable in the United States. Once unemployed, the household can either find a job with some probability, or exhaust her benefits and qualify for poverty benefits. This is the last state, and for lack of better terms, we refer to their members as the needy, the poor, or the long-term unemployed. If $e_t(i) = 0$, labor income is zero but the household collects food stamps and other safety-net transfers, $T_t^n(i)$, which are non-taxable. Households escape poverty with some probability at which they find a job.

There are two new automatic stabilizers at play in the household problem. First, the household can collect unemployment benefits, $T_t^u(i)$ which equal:

$$T^u(e_t(i), s_t(i)) = \min\left\{\bar{T}^u s_t(i), \bar{T}^u \bar{s}^u\right\} \text{ if } e_t(i) = 1 \text{ and zero otherwise.} \tag{8}$$

Making the benefits depend on the current skill-level captures the link between unemployment benefits and previous earnings, and relies on the persistence of s_t^h to achieve this. As is approximately the case in the U.S. law, we keep this relation linear with slope $\bar{T}^u$ and a maximum cap $\bar{s}^u$.

The second stabilizer are safety-net payments $T_t^s(i)$, which equal:

$$T^s(e_t(i)) = \bar{T}^s \text{ if } e_t(i) = 0 \text{ and zero otherwise.}$$

We assume that these transfers are lump-sum, providing a minimum living standard. In the data, these transfers are mean-tested, but because in our model these families only receive interest income from holding bonds, when we modified the model to put a maximum income cap to be eligible to these benefits, we found that almost no household ever hits this cap. For simplicity, we keep the transfer lump-sum.

3.3 Final goods' producers and the sales tax

A competitive sector for final goods combines intermediate goods according to the production function:

$$y_t = \left(\int_0^1 y_t(j)^{1/\mu} dj\right)^{\mu}. \tag{9}$$

where $y_t(j)$ is the input of the j^{th} intermediate input. The representative firm in this sector takes as given the final-goods pre-tax price p_t, and pays $p_t(j)$ for each of its inputs. Cost minimization together with zero profits imply that:

$$y_t(j) = \left(\frac{p_t(j)}{p_t}\right)^{\mu/(1-\mu)} y_t, \tag{10}$$

$$p_t = \left(\int_0^1 p_t(j)^{1/(1-\mu)} dj\right)^{1-\mu}. \tag{11}$$

On top of the price p_t, there is a sales tax τ^c so the after-tax price of the goods is:

$$\hat{p}_t = (1 + \tau^c) p_t. \tag{12}$$

This consumption tax is our next automatic stabilizer, as it makes actual consumption of goods a $1/(1+\tau^c)$ fraction of pre-tax spending on them.

3.4 Intermediate goods and corporate income taxes

Each variety j is produced by a monopolist firm using a production function:

$$y_t(j) = a_t k_t(j)^\alpha l_t(j)^{1-\alpha}. \tag{13}$$

where a_t is productivity, $k_t(j)$ is capital used, and $l_t(j)$ is effective labor. In the labor market, if l_t is the total amount of effective labor, then:

$$\int_0^1 l_t(j)dj = \int_0^\nu e_t(i)s_t(i)n_t(i)dh + \bar{s}n_t. \tag{14}$$

The demand for labor on the left-hand side comes from the intermediate firms. The supply on the right-hand side comes from employed households, adjusted for their productivity, and stock owners.

The firm maximizes after-tax nominal profits:

$$d_t(j) = \left(1-\tau^k\right)\left[p_t(j)y_t(j)/p_t - w_t l_t(j) - (r_t+\delta)k_t(j) - \xi\right], \tag{15}$$

taking into account the demand function in (10). The firm's costs are the wage bill to workers, the rental of capital at rate r_t plus depreciation of a share δ of the capital used, and a fixed cost ξ. The maximized profits are rebated every period to the capitalists as dividends.

Intermediate firms set prices subject to nominal rigidities a la Calvo (1983) with probability of price revision θ. Since they are owned by the capitalists, they use their discount factor $\lambda_{t,t+s}$ to choose price $p_t(j)^*$ at a revision date with the aim of maximizing expected future profits:

$$\mathbb{E}\left[\sum_{s=0}^{\infty}(1-\theta)^s \lambda_{t,t+s} d_{t+s}(j)\right] \quad \text{subject to:} \quad p_{t+s}(j) = p_t(j)^* \tag{16}$$

The new automatic stabilizer is the corporate income tax, which is a flat rate τ^k over corporate profits. In the U.S. data, dividends and capital gains pay different taxes. While this distinction is important to understand the capital structure of firms and the choice of retaining earning, it is immaterial for the simple firms that we just described.

3.5 Capital-goods firms and property income taxes

A representative firm owns the capital stock and rents it to the intermediate-goods firm taking r_t as given. If k_t denotes the capital held by this firm, then in the market for capital:

$$k_t = \int_0^1 k_t(j)dj. \tag{17}$$

This firm invests in new capital $\Delta k_{t+1} = k_{t+1} - k_t$ subject to adjustment costs to maximize after-tax profits:

$$d_t^k = (1-\tau^k)r_t k_t - \Delta k_{t+1} - \frac{\zeta}{2}\left(\frac{\Delta k_{t+1}}{k_t}\right)^2 k_t, \tag{18}$$

The value of this firm, which owns the capital stock is then given the recursion:

$$v_t = d_t^k - \tau^p v_t + \mathbb{E}_t\left[\lambda_{t,t+1} v_{t+1}\right].$$

The new automatic stabilizer, the property tax, is a fixed tax rate τ^p that applies to the value of the only property in the noel, the capital stock. A few steps of algebra show the conventional results from the q-theory of investment:

$$v_t = q_t k_t, \tag{19}$$

$$q_t = 1 + \zeta\left(\frac{\Delta k_{t+1}}{k_t}\right). \tag{20}$$

Because, from the second equation, the price of the capital stock is procyclical, so will property values, making the property tax a potential automatic stabilizer.

Finally, note that total dividends sent to stock-owners, d_t, come from every intermediate firm and the capital-goods firm:

$$d_t = \int_0^1 d_t(j)dj + d_t^k - \tau^p q_t k_t. \tag{21}$$

We do not include investment tax credits. They are small the data and, when used to attenuate the business cycle, they have been enacted as part of stimulus packages and not as automatic rules.

3.6 The government and budget deficits

The government budget constraint is:

$$\begin{aligned}
&\tau^c \left(\int_0^\nu c_t(i)di + c_t\right) + \tau^p q_t k_t + \int_0^\nu \bar{\tau}^x(x_t(i))di + \bar{\tau}^x(x_t) + \\
&\quad \tau^k \left[\int_0^1 \hat{d}^i(j)di + r_t k_t\right] - \int_0^\nu (T_t^u(i) + T_t^s(i)di \\
&\quad = g_t + (i_t/p_t)B_t - (B_{t+1} - B_t)/p_t + T_t^e.
\end{aligned} \tag{22}$$

On the left-hand side are all of the automatic stabilizers discussed so far: sales taxes, property taxes and personal income taxes in the first line, and corporate income taxes and transfers in the second line. On the right-hand side are government purchases, g_t and government bonds B_t. Because government bonds are the only asset in positive net supply to the households, the market for bonds will clear when:

$$B_t = \int_0^\nu b_t(i)di + b_t. \tag{23}$$

In steady state, the stabilizers on the left-hand side imply a positive surplus, which is offset by steady-state government purchases $\bar{g}$. Since we set transfers to the entrepreneurs to zero, the budget constraint then determines a steady state amount of debt $\bar{B}$, which is consistent with the government not being able to run a Ponzi scheme.

Outside of the steady state, as outlays rise and revenues fall during recessions, the left-hand side of equation (23) increases, and so must the right-hand-side. This is the last stabilizer that we consider: the automatic increase in the budget deficit during recessions.

The debt that results must be paid over time. In our baseline, we consider a simple fiscal rule where debt is paid via a lump-sum tax on capitalists:

$$T_t^e = -\gamma \log\left(\frac{B_t/p_t}{\bar{B}}\right). \tag{24}$$

The parameter $\gamma > 0$ measure the speed at which the deficits from recessions are paid over time. If γ is close to infinity, then the deficits caused by a recessions are paid right away the following period; if γ is close to zero, they take arbitrarily long to get paid. We have the tax on stock-owners adjusting because it is the fiscal tool that interferes the least with the other stabilizers, affecting neither marginal returns like the distortionary tax rates or having an important effects on the wealth and income distribution as transfers to households. We will consider an alternative later.

3.7 Shocks and business cycles

Monetary policy follows a conventional Taylor rule:

$$i_t = \bar{i} + \phi_p \Delta \log(p_t) + \phi_y \log(y_t/y) + \varepsilon_t \tag{25}$$

with $\phi_p > 1$ and $\phi_y \geq 0$.[10]

Two aggregate shocks hit the economy: technology, $\log(a_t)$, and monetary policy, ε_t. Therefore, both aggregate-demand and aggregate-supply shocks may drive business cycles. We assume that both shocks follow independent AR(1)s for simplicity.

The idiosyncratic shocks to households, $e_t(i)$ and $s_t(i)$ are first-order Markov processes. Moreover, the transition matrix of labor-force status, the three-by-three matrix $\Pi_{e,e'}$, depends on a linear combination of the two aggregate shocks. This way, we let unemployment vary with the business cycle to match Okun's law. This approach to modeling unemployment is clearly reduced-form and subject to the Lucas critique. However, our model of the business cycle is already sufficiently complicated that endogenizing the extensive margin of labor supply is challenging. At the same time, recall that workers choose how many hours to work. Therefore, our model has an endogenous intensive margin of labor supply.

3.8 Equilibrium and volatility

An equilibrium in this economy is a collection of aggregate quantities $(y_t, k_t, d_t, v_t, c_t, n_t, b_{t+1}, x_t)$; aggregate prices $(p_t, \hat{p}_t, w_t, q_t)$; individual consumer decision rules $(c_t(b, s, e), n_t(b, s, e))$; a distribution of households over assets, skill levels, and employment statuses; individual firm variables $(y_t(j), p_t(j), k_t(j), l_t(j), d_t(j), d_t^k)$; and government choices (B_t, i_t, g_t) such that:
(i) owners maximize expression (1) subject to the budget constraint in equations (2)-(3),
(ii) the household decision rules maximize expression (5) subject to their budget constraint in equations (6)-(7),
(iii) the distribution of households over assets and skill and employment levels evolves in a manner consistent with the decision rules and the exogenous idiosyncratic shocks,
(iv) final-goods firms behave optimally according to equations (10)-(12),
(v) intermediate-goods firms maximize expression (16) subject to equations (10), (13), (15),
(vi) capital-goods firms maximize expression (18) so their value is in equations (19)-(20),
(vii) fiscal policy respects equation (22) and follows the rule in equation (24) while monetary

[10] Including interest-rate smoothing had a small quantitative effect on the results (details available form the authors), so we leave it out to save on one more parameter to keep track of and calibrate.

Figure 2: Steady-state capital and household bond holdings

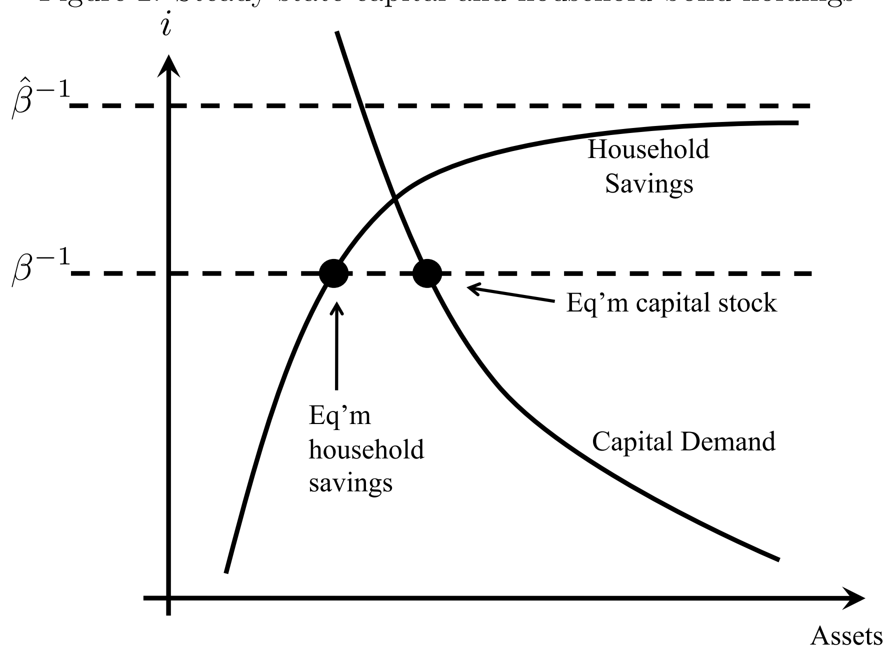

policy follows the rule in (25),
(viii) markets clear for labor in equation (14), for capital in equation (17), for dividends in equation (21) and for bonds in equation (23).

We evaluate the mean and variance of aggregate endogenous variables at the ergodic distribution at the equilibrium in this economy.

4 Properties of the model

Our model is not easy since to solve it, one must keep track not only of the aggregate variables, but also of the distribution of assets across agents and the distribution of prices across firms. At the same time, the model has familiar foundations laid out in this section.

4.1 The ergodic distribution

Figure 2 uses a simple diagram, akin to Aiyagari (199x), to describe the steady state of the model without aggregate shocks.

The downward-sloping curve is the demand for capital, with slope determined by diminishing marginal returns. The demand of stock owners for assets is perfectly elastic at their time-preference rate just as in the neoclassical growth model. Because they are the sole hold-

Figure 3: Optimal savings policies

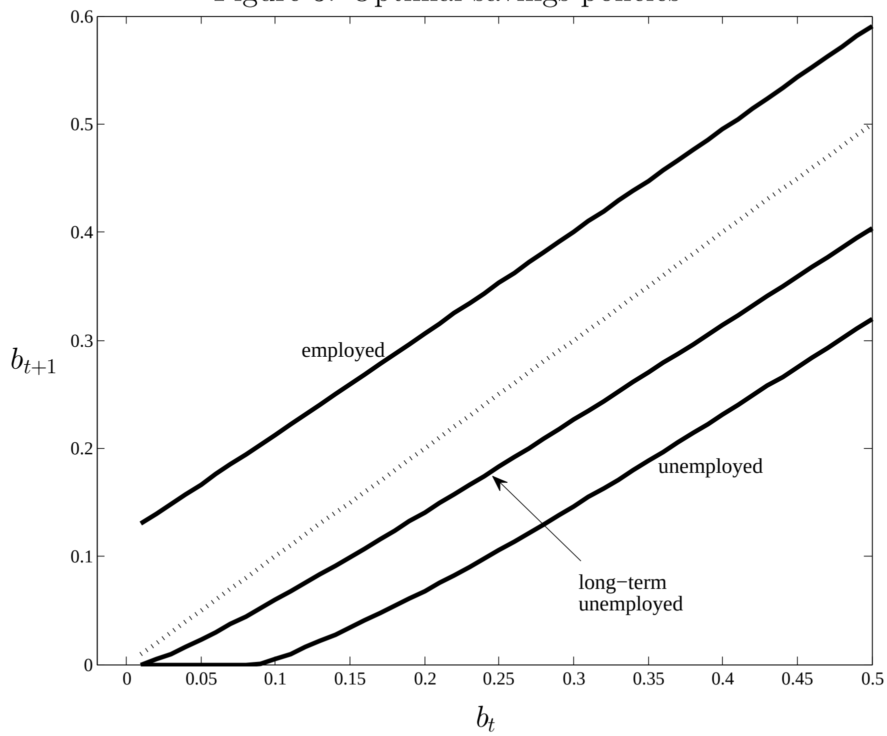

ers of capital, the equilibrium capital stock in the model is determined by the intersection of these two curves, and is affected by taxes on capital income, like the personal or corporate income taxes. If households were also fully insured their demand for assets would be the horizontal line going through their time-preference rate, but because of the idiosyncratic risk they face, they have a precautionary demand for assets. Therefore, they are willing to hold bonds even at lower interest rates. Their asset demand is given by the upward-sloping curve. Because in the steady state without aggregate shocks, bonds and capital must yield the same return, equilibrium bond holdings by households are given by the point to the left of the equilibrium capital stock. The difference between the total amount of government bonds outstanding and those held by households gives the bond holdings of stock owners.

Figure 3 shows the optimal savings decisions of households at each of their e_t states. When a household is employed, they save so the savings policy is above the 45^o line, while when they do not have a job, they run down their assets. Households who lost their job and fewer assets will run down their wealth to zero and stay there until they regain employment, leading to the horizontal segment along the horizontal axis in their savings policies.

Figure 4 shows the ergodic wealth and income distribution for households. Three features of these distributions will play a role in our results. First, that households below the poverty

Figure 4: The ergodic wealth distribution, with and without transfers

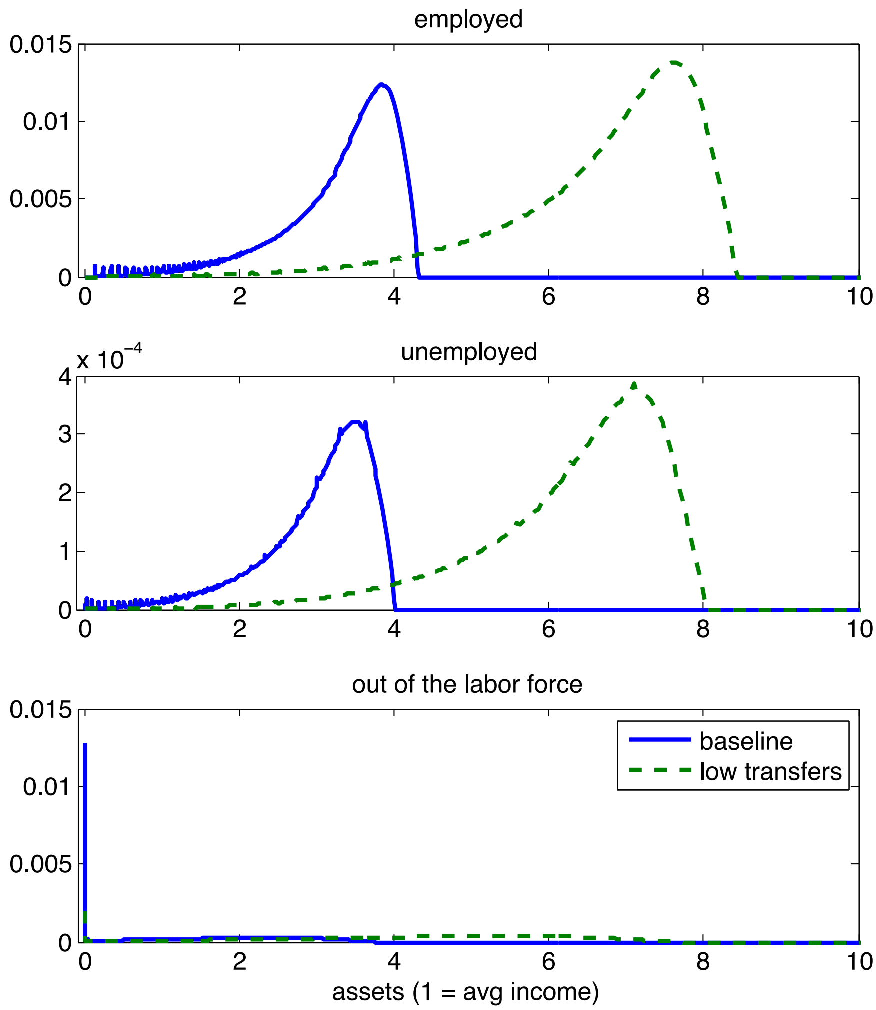

line have essentially no assets. Given the borrowing constraint they face, they live hand to mouth. Second, that employed households are wealthier than the unemployed. When a recession comes, and more households lose their jobs, they will draw down their wealth to smooth out hard times. Third, the figure shows a counterfactual wealth distribution if the two transfer programs are significantly cut. Because not being employed now comes with higher income risk, households save more, which raises their wealth in all states. This large impact of the stabilizers on the wealth distribution will play an important role in our results.

4.2 Solution algorithm

The distribution of wealth across households is a state variable of the model, so the solution algorithm has to keep track of the dynamics of this distribution. One candidate is the Krusell-Smith (1998) algorithm, which summarizes the distribution of wealth with a few moments of the distribution. We opt instead for the solution algorithm developed by Reiter (2008, 2009) because this method can more easily be applied to models with a rich structure at the aggregate level. The Reiter algorithm approximates the distribution with a histogram that has a large number of bins. The mass of households in each bin then becomes a state variable of the model. Similarly, Reiter's algorithm approximates the household decision rules with a discrete approximate (e.g. a spline). In this way, the model is converted from one that has infinite-dimensional objects to one that has a large, but finite, number of variables. Using standard techniques, one can find the stationary competitive equilibrium of this economy in which there is idiosyncratic uncertainty, but no aggregate shocks. Reiter's (2008) method then calls for linearizing the model with respect to aggregate shocks and solving for the dynamics of the economy as a perturbation around the stationary equilibrium without aggregate shocks using existing methods (e.g. Sims, 2001). The resulting solution is non-linear with respect to idiosyncratic variables, but linear with respect to aggregate variables.

This approach works well for small versions of our model (e.g. with fewer than four discrete types of households). Increasing the number of household types leads the system to grow to a size for which the application of linear rational expectation solution methods are not feasible. Therefore, we make use of Reiter's (2009) method of compressing the system using model reduction techniques. As Reiter (2009) explains, this compression comes with virtually no loss of accuracy relative to the larger linearized system because many dimensions of the state space are not needed.[11] We can verify this claim for versions of our model for which it is possible to solve the full linear system and the reduced system. Additional details regarding the solution of the model appear in the appendix.

[11] There are two intuitive reasons that a dimension of the state space can be eliminated without loss of accuracy: 1) the system never varies along that dimension or 2) variation along that dimension is not relevant for the variables of interest. See Antoulas (2005) for a discussion of model reduction in a general context and see Reiter (2009) for their application to forward-looking economic systems.

4.3 Calibrating of the model

We calibrate as many parameters as possible to the properties of the automatic stabilizers in the data. For the government spending and revenues, we use table 1, which recall averaged over the period 1988-2007. For macroeconomic aggregates though, we average over a longer period, starting in 1960 and using quarterly data, so that we can include more recessions in the sample and periods outside th Great moderation and do not underestimate the amplitude of the business cycle.

For the three proportional taxes, we use a parameter related to preferences or technology to match the tax base in the NIPA accounts, and choose the tax rate to match the average revenue reported in table 1, or in the words of Mendoza, Razin and Tesar (1994, JME), the average effective tax rates. The top panel of table 2 shows the parameters set and the respective targets. Our average tax rates are not far from the statutory marginal tax rates for sales and corporate income.

For the personal income tax, we followed Auerbach and Feenberg and simulated TAXSIM, including federal and state taxes, for a typical household. We averaged the tax rates across states weighted by population, and across years between 1988 and 2007. We then fit a cubic function of income to the resulting scheduled, and splined it with a flat line above a retain level of income so that the fitted function would be non-decreasing. The result is in figure 2. The cubic-linear schedule approximated the actual taxes well, and its smoothness is numerically useful. We then added an intercept to this schedule to fit the effective average tax rate. This way, we made sure we fitted both the marginal tax rates (via TAXSIM) and the average tax rates (via the intercept).

Panel B calibrates the parameters on government spending. Both parameters governing transfer payments are set to equate the average outlaid from these programs. The speed at which deficits are paid comes from the autocorrelation of budget deficits in the data.

The first source of idiosyncratic risk facing household was on whether to be employed, unemployed, or below the poverty line. We calibrate the steady-state transition probabilities across these states in panel C by using the flows and average value of receiving unemployment benefits or safety-net benefits. What is relevant for our question is not whether households are unemployed per se, but whether they are receiving benefits from the government. Therefore, we keep the discipline of focussing on government outlays to calibrate the parameters. For the cyclical changes in these probabilities we use the standard deviation of the flows into the states in the sample since 1960.

The second source of risk is labor income, and it is calibrated in panel D. We used the

Table 2: Calibration of the parameters

Symbol	Parameter	Value	Target (Source)
Panel A. Tax bases and rates			
τ^c	Tax rate on consumption	0.054	Avg. revenue from sales taxes (Table 1)
β	Discount factor of stock owners	0.989	Consumption-income ratio = 0.689 (NIPA)
τ^p	Tax rate on property	0.003	Avg. revenue from property taxes (Table 1)
α	Coefficient on labor in production	0.296	Capital income share = 0.36 (NIPA)
τ^k	Tax rate on corporate income	0.282	Avg. revenue from corporate income tax (Table 1)
ξ	Fixed costs of production	1.32	Corporate profits / GDP = 9.13% (NIPA)
μ	Desired gross markup	1.1	Avg. U.S. markup (Basu, Fernald, 1997)
Panel B. Government outlays and debt			
$\bar{T}^u$	Unemployment benefits	0.185	Avg. outlays on unemp. benefits (Table 1)
$\bar{T}^s$	Safety-net transfers	0.169	Avg. outlays on safety-net benefits (Table 1)
G/Y	Steady-state purchases / output	0.130	Avg. outlays on purchases (Table 1)
γ	Fiscal adjustment speed	2.2	Autocorrel. net public savings / GDP = 0.966 (NIPA)
B/Y	Steady-state debt / output	1.66	Avg. interest expenses (Table 1)
Panel C. Labor-force status			
π_{eu}	Steady-state transition prob. E-U	0.026	Avg. insured unemp. rate = 0.023 (BLS)
π_{ue}	Steady-state transition prob. U-E	0.571	Avg. UE flow quarterly = 0.813 (Shimer, 2007)
π_{up}	Steady-state transition prob. U-P	0.297	Avg. SNAP ratio = 0.077 (USDA)
π_{pu}	Steady-state transition prob. P-E	0.087	SNAP exit hazard = 0.03 monthly (Mabli et al., 2011)
π^y_{eu}	Cyclical transition prob. E-U	-1.75	St. dev. of unemp. rate = 0.009 (BLS)
π^y_{ue}	Cyclical transition prob. U-E	9.70	St. dev. of UE flows = 0.053 (Shimer)
π^y_{up}	Cyclical transition prob. U-P	0.00	St. dev. of SNAP ratio = 0.020 (USDA)
Panel D. Income and wealth distribution			
ν	Non-participants / stock owners	4	
β^h	Discount factor of households	0.983	Wealth of top 20% by wealth
$\bar{s}$	Skill level of stock owners	4.66	Income of top 20% by wealth (SCF)
$E(s)$	Mean of non-participants skill	1.08	Avg. income in economy normalized to 1
Panel E. Business-cycle parameters			
θ	Calvo price stickiness	0.286	Avg. price spell duration = 3.5 (Klenow, Malin, 2011)
ψ_1	Labor supply	21.6	Avg. hours worked = 0.31 (Cooley, Prescott, 1995)
ψ_2	Labor supply	2	Frisch elasticity = 1/2 (Chetty, 2011)
δ	Depreciation rate	0.114	Annual depreciation expenses / GDP = 0.046 (NIPA)
ζ	Adjustment costs for investment	15.0	Corr. of Y and C = 0.88 (NIPA)
ρ_z	Autocorrelation productivity shock	0.880	Autocorrel. of log GDP = 0.864 (NIPA)
σ_z	St. dev. of productivity shock	0.004	St. dev. of log GDP = 1.539 (NIPA)
ρ_m	Autocorrelation monetary shock	0.500	Largest AR for inflation = 0.85 (Pivetta, Reis, 2006)
σ_m	St. dev. of monetary shock	0.005	Share of output variance due to shock = 0.2
ϕ_p	Interest-rate rule on inflation	1.55	St. dev. of inflation = 0.638 (NIPA)
ϕ_y	Interest-rate rule on output	0.010	Correl. of inflation with log Y = 0.198 (NIPA)

Figure 5: Personal income tax rate from TAXSIM

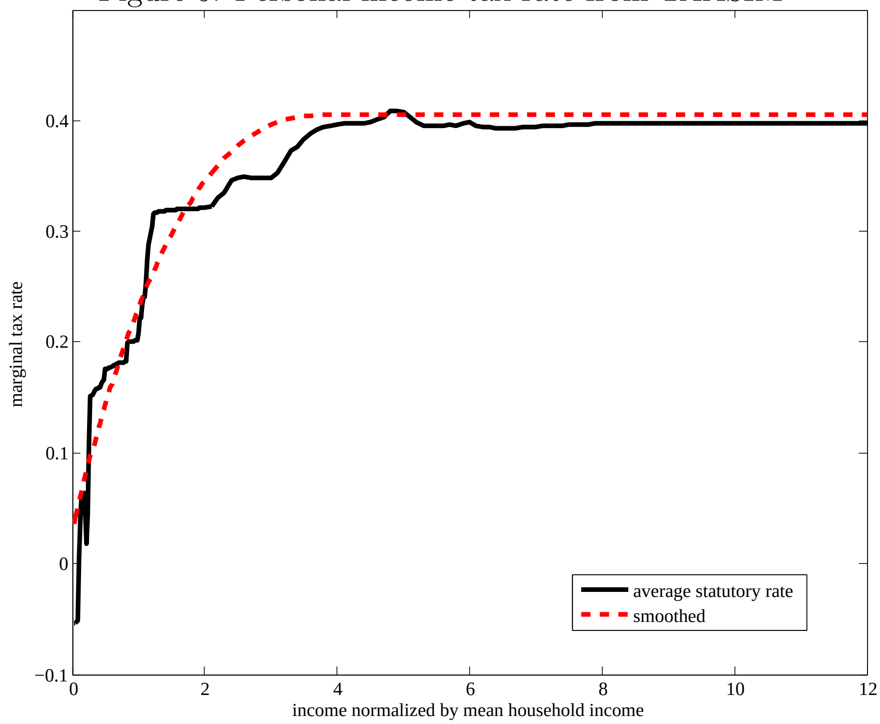

Survey of Consumer Finances to calculate that 0.84% of the wealth is held by the top 20% in the United States. We then picked the discount factor of the households to match this. Omitted from the table is the calibration of the Markov transition matrix for the s shocks, which we include in the appendix. Following XXX, we used a 3-point grid in the Panel Study for Income Dynamics, and divided annual earnings by annual hours to estimate the transition probabilities.

Finally, panel E has other parameters that affect the aggregate dynamics of the economy. Most are standard, but two deserve some explanation. First, the Frisch elasticity of labor supply plays an important in most intertemporal business cycle models. Consistent with our focus on taxes and spending, we use the recent survey by Chetty (2011) on the response of hours worked to several tax and benefit changes. Second, we choose the variance of monetary shocks so that a variance decomposition of output attributes them 20% of aggregate fluctuations. There is great uncertainty on the empirical estimates of the sources of business cycles, but this number is not out of line with at least some of the estimates reported in Christiano et al (1999).

Figure 6: Impulse responses to the two aggregate shocks

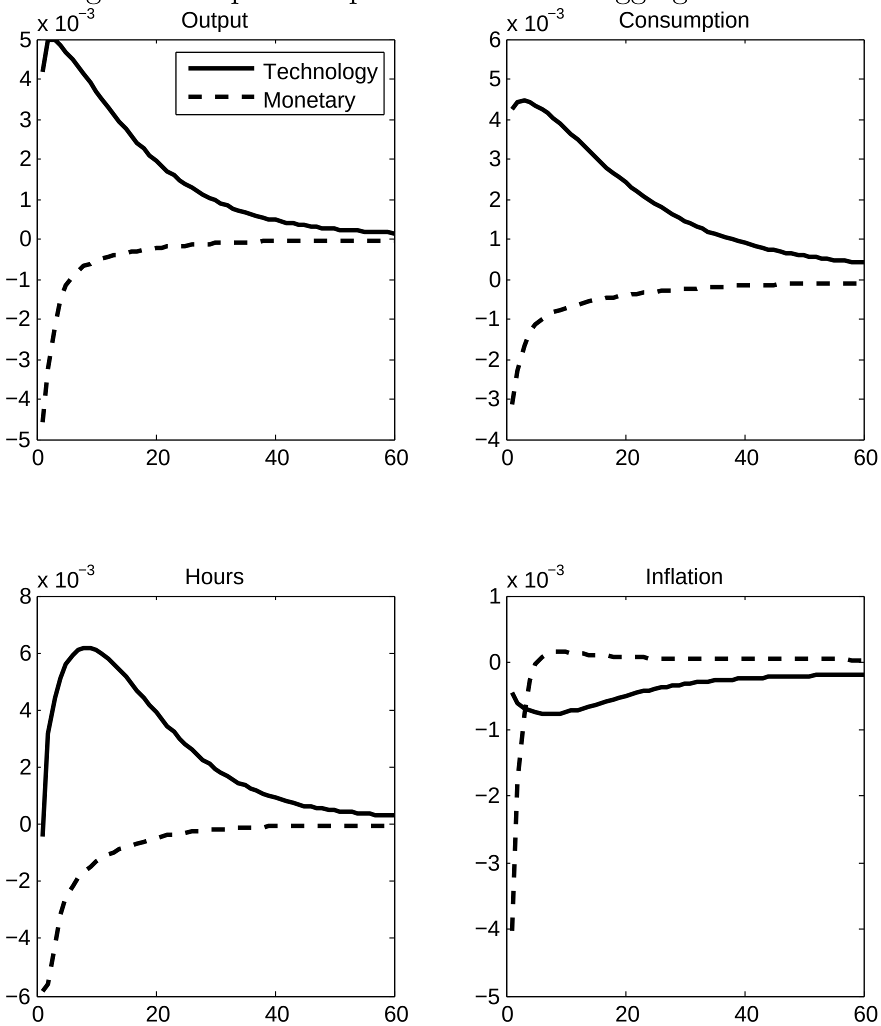

4.4 Impulse responses to shocks

Before we can use our model to make counterfactual predictions, we must verify that it provides a reasonable description of U.S. business cycles. Figure 5 shows impulse responses to a positive technology shock and a contractionary monetary shock of one standard deviation.

The model generates the positive co-movement of output and consumption, as well as the persistent responses of both variables to shocks that have been emphasized in the literature. Hours have a hump-shaped repines to a technology shock: at first they rise because higher TFP raises the marginal product of labor. Then, while TFP falls as the shock dissipates, the increase in the capital stock further raises the marginal product of working so hours keep rising. At a point, investment is no longer strong enough, so the fall in TFP dominates, and hours start converging back to zero. Finally, turning to inflation, as is well-known in new Keynesian models, both shocks lower inflation, but the simple Calvo model implies a fairly short-lived response.

In sum, the aggregate dynamics of our model resembles those in the standard new neoclassical synthesis model of Goodfriend and King (1997) and Clarida Gali and Gertler (1999) that has been widely used to study business cycles in the past decade.

4.5 Complete markets and the neoclassical synthesis

While our model is complicated, it has a familiar foundation. If prices were flexible and there were no stock owners, our model would be close to that in Krusell and Smith (1997), augmented with many taxes and transfers. This has become the standard model of incomplete markets (Storesletten, Violante, 2010, ARE)

With complete markets, households can diversify their two sources of idiosyncratic, to their labor income an their labor-force states. The following assumption eliminates this risk:

Assumption 1. *Households and capitalists trade a full set of Arrow securities, so they are fully insured.*

It will not come as a surprise that with complete markets there is a representative agent in this economy. The problem she solves is familiar:

Proposition 1. *Under assumption 1, there is a representative agent with preferences:*

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \log(c_t) - (1+E_t)\,\psi_1 \frac{n_t^{1+\psi_2}}{1+\psi_2} \right\},$$

and constraints:

$$\begin{aligned} \hat{p}_t c_t + b_{t+1} - b_t &= p_t\left[x_t - \bar{\tau}(x_t)\right] + T_t^n \\ x_t &= \frac{i_t}{p_t} b_t + w_t s_t (1+E_t) n_t + d_t + T_t^u \\ s_t &= \left[\frac{1}{1+E_t}\bar{s}_t^{1+1/\psi_2} + \frac{E_t}{1+E_t}\int_0^{\nu} s_{i,t}^{1+1/\psi_2} di\right]^{\frac{1}{1+1/\psi_2}}, \end{aligned}$$

where $1+E_t$ is total employment, including capital-owners and households.

The proof is in the appendix. With the exception of the exogenous shocks to employment, the problem of this representative agent is fairly standard. Moreover, on the firm side, optimal behavior by the goods-producing firms leads to a new Keynesian Phillips curve, while optimal behavior by the capital-goods firm produces a familiar IS equation. Therefore, with complete markets, our model is of the standard neoclassical synthesis variety (Goodfriend and King, 1997, Woodford, 2003) that has been intensively used to study business cycles over the past decade. We will use it to study the effectiveness of automatic stabilizers when distributional issues are set to the side.

5 The effectiveness of automatic stabilizers

We start with a simple but extreme case.

Assumption 2. *The following set of conditions holds:*

1. *The personal income tax is proportional, so $\tau^x(\cdot)$ is constant.*
2. *The probability of being employed $\Pi_{e,e'}$ is constant over time.*
3. *The Calvo probability of price adjustment $\omega = 1$, so prices are flexible.*
4. *There are infinite adjustments costs, $\gamma \to +\infty$, so capital is fixed.*

These strong assumptions shut down the aggregate demand channel, since prices are flexible, and the redistribution and social insurance channels, since all households are perfectly insured against idiosyncratic shocks so the wealth distribution does not matter for aggregate dynamics. Still, there remains the effect of automatic stabilizers on marginal incentives to work, save and consume.

The appendix proves the following result:

Proposition 2. *If assumptions 1 and 2 hold, then the variance of the log of output is equal to the variance of the log of productivity. Therefore, $\mathbf{S} = 0$ and the automatic stabilizers are ineffective.*

The steps of the proof provide some intuition for the result. With flexible prices, there is an aggregate Cobb-Douglas production function, so if the capital stock and employment are fixed, then the proposition will be true as long as the labor supply is fixed. Equating the marginal rate of substitution between consumption and leisure for households to their after-tax wage gives the standard labor supply condition:

$$n_t(i) = \left(\frac{(1 - \tau^x) s_t(i) w_t \hat{p}_t}{\psi_1 c_t(i) p_t} \right)^{1/\psi_2}$$

Table 3: The effect of proportional taxes on the business cycle

	Representative agent		Full model	
	variance	average	variance	average
output	0.0044	0.0159	-0.0029	0.0164
hours	0.0073	0.0007	0.0046	-0.0002
consumption	-0.0290	0.0147	-0.0106	0.0153

Note: Proportional change caused by cutting the stabilizer

Perfect insurance implies that consumption is equated across households. But then, our balanced-growth preferences and technologies imply that c_t/w_t is fixed over time, so the condition above, once aggregated over all households, gives a constant labor supply.

While this result and the assumptions supporting it are extreme, it serves a useful purpose. Note that the estimates of the size of the stabilizer following the Pechman (1973) approach would be large in this economy. Measuring the ratio of the changes in taxes over the change in output over time in a simulation of this economy would produce a plot similar to the one in figure 1 above. And yet, the stabilizers in this economy are completely ineffective. An economy may have high measured built-in flexibility while not being flexible at all.

5.1 The effectiveness of proportional taxes

Assumption 2 imposed no restrictions on proportional taxes, yet their effect on volatility was nil. Table 3 considers the effect on the variance and average level of output, hours, and consumption of the following experiment. We cut the tax rates τ^C, τ^P and τ^K by 10%, and replaced the lost revenue of 0.6% of GDP by a lump-sum tax on the entrepreneurs.

The first pair of columns has the effects under assumption 1, so there are complete markets, and the next pair with the full model. The effects of proportional taxes on the volatility of the business cycle in both cases are quantitatively small. At the same time, they lead to a significantly lower output and consumption on average.

Intuitively, a higher tax rate on consumption lowers the returns from working and so lowers labor supply and output on average. However, because the tax rate is the same after an aggregate shock, it does not induce any intertemporal substitution of hours worked, nor does it change the share of disposable income available in booms versus recessions. Likewise the taxes on cop orate and property income may discourage savings and affect the average capital stock, but they do not do so differentially across different stages of the business cycle and so they have a negligible effect on volatility.

5.2 The effectiveness of transfers

To evaluate the effectiveness of our two transfer programs, unemployment and safety-net benefits, we reduced spending on both by 0.6% of GDP, the same amount we used for proportional taxes. Again, we replaced the fall in outlays with a lump-sum transfer to stock owners. The impact on the full model is in table 4.

Transfers have a very small effect on the average level of output and consumption, yet they have a large effect on their volatility. Reducing transfer payments would raise output volatility by as much as 13%, while lowering the volatility of consumption by 22%. The main channel at work seems to be redistribution. In a recession, the households without a job receive higher transfers. These have no direct effect on their labor supply of hours worked, since they do not have a job in the first place. However, they are funded by higher taxes on the stock owners, who raise their hours worked in response to the reduction in their wealth. This stabilizes hours worked and output.

The reason behind the higher volatility of consumption is in figure 4. With transfers, households are better insured against the major idiosyncratic shocks they face, so they accumulate fewer assets. But then, when aggregate shocks hit, they are less able to smooth them out. Therefore, while the variance of household consumption increases by 91% when transfers are reduced, because there is less social insurance against idiosyncratic shocks, the variance of aggregate consumption falls as there is less smoothing of aggregate shocks.

To confirm that it is the precautionary channel that is behind these changes in volatility, we performed an additional experiment. We lowered the households' discount factor at the same time that we reduced transfers, so that the aggregate assets of the household did not change. This is not a valid policy experiment, since we are changing not just policy but also preferences, but it serves to highlight the role of precautionary savings. Now, when we lower transfers and the discount factor, both the volatility of output and consumption rise substantially. The higher volatility of output is no longer offset by precautionary savings.

Table 4 also considers an alternative economy, inspired in the savers-spenders model (Mankiw, 2001). In this model, we replaced household's optimal savings function in figure 3, with the assumption that they live hand-to-mouth, consuming all of their after-tax income at every date. Now there are no precautionary savings, so eliminating the public insurance

Table 4: The effect of transfers on the business cycle

	Full model		Hand-to-mouth	
	variance	average	variance	average
output	0.0719	0.0016	-0.0107	-0.0043
hours	0.1613	-0.0134	0.0082	-0.0017
consumption	-0.1345	0.0023	0.0586	-0.0056

Note: Proportional change caused by cutting the stabilizer

provided by transfers, raises the volatility of consumption. Moreover, the volatility of output now only slightly goes up without transfers.

The savers-spenders economy maximizes the disposable-income channel that is most often mentioned in support of automatic stabilizers. Every dollar given to households is spent, raising output because of sticky prices. Yet, we see that, quantitatively, this effect accounts for little of the stabilizing effects of transfers in our economy. Rather, is is the redistribution channel that is most at work.

5.3 The effectiveness of progressive income taxes

The next experiment we consider replaces the progressive personal income tax with a proportional, or flat, tax that raises the same revenue in steady state. Table 5 has the results.

In the representative-agent economy, progressive income taxes have a modest effect on the volatility of output. While the increase in marginal tax rates during booms and their decline in recessions, is stabilizing in theory, the level of progressivity in the current U.S. tax system is modest, as we saw in figure 2. As a result, this effect is quantitatively small.

Progressive income taxes stabilize consumption because capital-goods firms retain dividends in expansions to avoid the higher taxes, and distribute them in recessions. This stabilizes after-tax income for stock owners, and stabilizes their consumption while making investment on capital more volatile. Moving to a flat tax would raise the average level of economic activity significantly, output by 4% and consumption by 5%.

With incomplete markets, the effect on average is slightly larger, but the effect of progressive income taxes on volatility is much higher. Again, social insurance and precautionary swings are at work. Removing the progressivity of the personal income tax increases after-tax income risk inducing households to save more. Households therefore hold more bonds,

Table 5: The effect of progressive taxes on the business cycle

	Representative agent		Full model	
	variance	average	variance	average
output	-0.0224	0.0369	0.0691	0.0413
hours	-0.0129	0.0369	0.0109	0.0416
consumption	0.0831	0.0486	0.0065	0.0543

Note: Proportional change caused by cutting the stabilizer

so stock owners own a larger share of their wealth in the capital stock. They are therefore more responsive in their investment choices to to changes in marginal returns on saving, and output is more volatile.

Finally, note that progressive income taxes make consumption more volatile, the opposite of what happened with complete markets. The reason is again redistribution and the mechanism is the same as with transfers. Removing social insurance against idiosyncratic shocks increases the volatility of household consumption by 70%. This encourages households to accumulate more assets, which then allows them to smooth out better aggregate shocks in their consumption choices.

We also repeated this experiment in the savers-spenders economy and found that there that removing the progressivity of personal income taxes changed output volatility by only 0.5%, similar to the representative agent economy. This further reinforces our conclusion that the disposable-income channel is not what is behind the strong stabilizing effect of progressive income taxes.

5.4 The effectiveness of budget deficits

To assess the role of the budget deficit, we conducted two final experiments. First, we increased γ to infinity, so that the government balanced its budget every period. Table 6 shows that this had almost no effect on the volatility of any of the endogenous variables. While Ricardian equivalence does not hold in our economy, changing the time profile of the taxes on stock owners has a small quantitative effect.

In our baseline model, g_t/y_t is fixed over time. In our second experiment, we replaced our fiscal rule by an alternative where deficits are paid for by cutting government purchases

Table 6: The effect of budget deficits on the business cycle

	Balanced-budget		Purchases adjust	
	variance	average	variance	average
output	-0.0008	-0.0000	0.1471	-0.0000
hours	-0.0001	0.0000	0.0866	0.0001
consumption	-0.0078	-0.0000	-0.5548	0.0002

Note: Proportional change from altering the fiscal adjustment rule

over GDP:

$$\log\left(\frac{g_t/y_t}{g/y}\right) = -\gamma \log\left(\frac{B_t/p_t}{\bar{B}}\right). \tag{26}$$

The last two columns in table 6 show that this rule makes output and hours more volatile, and consumption less volatile. After a positive aggregate shock, public debt falls, which induces government purchases to rise. This lowers private wealth, increasing labor supply and therefore amplifying the shock. The same wealth effect stabilizes consumption.

To conclude, changing the timing of deficits has little effect on the economy, but the way in which these deficits are financed can have a significant effect on volatility.

5.5 The overall effectiveness of the automatic stabilizers

All the previous experiments together, cutting proportional taxes, transfers, flat income tax, and balance budget every period. Because cut taxes and transfers both by same amount of GDP, do not need the tax on stock owners.

XXX

6 Conclusion

As Alan Blinder (2004) put it in his survey of the use of fiscal policy for stabilizing the business cycle "virtually every contemporary discussion of stabilization policy by economists—whether it is abstract or concrete, theoretical or practical—is about monetary policy, not fiscal policy." Robert Solow had, in a similar context, stated that "Serious discussion of fiscal policy has almost disappeared."

The recession of 2007-09 drastically changed this state of affairs. During these two years, the ratio of government spending to GDP increased by more than it had in the previous 50 years. Feldstein (2009) writing on "Rethinking the role of fiscal policy, "and Blanchard et al (2010) on "Rethinking macroeconomic policy" both emphasized the use of fiscal policy as a countercyclical instrument as one of the main areas in need of new work.

This paper contributed to this effort by assessing the role of automatic stabilizers as a countercyclical tool. We constructed a business cycle model with many of the stabilizers and calibrated it to replicate the U.S. data. The model had some interesting features in its own right. First, it nested both the standard incomplete markets model, as well as the standard new-Keynesian business cycle model. Second, it matched the first and second moments of U.S. business cycles, as well as the broad features of the U.S. wealth and income distributions. Third, solving it required using new methods that may be useful for other model that combine nominal rigidities and incomplete markets.

We found that proportional taxes, like the sales tax, the property tax, and the corporate income tax have negligible effect on the volatility of economic aggregates. The progressivity of the personal income tax and transfer payments to the unemployed and those on food stamps have been quite effective stabilizers, contributing to a lower variance of output by 15% and 13% respectively. However, the progressivity of the income tax also least to significantly lower average output. Transfer payments, in turn, have a negligible effect on average output, but because they lower precautionary savings, they raise the variance of consumption substantially.

We also found that the traditional Keynesian channel used to support automatic stabilizers is quantitatively weak. While raising the disposable income of consumers during recessions increases aggregate demand and output, this has a small effect over the business cycle. We found that a more important channel for stabilization was redistributing resources from richer agents, that have lower marginal propensities to consume and change their labor supply as their after-tax wealth changes, towards poorer agents, with higher marginal propensities to consume and are without a job so cannot decrease hours worked any further. At the same time, because this redistribution provides social insurance against idiosyncratic shocks, households hold fewer assets to self-insure, which raises the volatility of consumption in response to aggregate shocks.

We have refrained from making welfare judgments. Our economy has many dimensions of heterogeneity so that any normative point would depend crucially on how to weight the welfare of different agents. Moreover, it was important to understand first the positive properties of the model. Future work may take up the challenge of looking at optimal policy.

A second limitation of our work is that each of the automatic stabilizers that we considered is more complex than our description and distorts behavior in more ways than the ones we modeled. We did so in part because we wanted our model to be sufficiently simple so that we could understand the channels through which the stabilizers might work, and in part because of computational limitations. To obtain sharper quantitative estimate of the role of the stabilizers, it would be desirable to include the findings form the rich micro literatures that study each of these government programs in isolation. Perhaps the main point of this paper is that to assess automatic stabilizers requires having a fully articulated business-cycle model, so that we can move beyond simply stating the disposable-income channel, and consider both other channels as well as quantify their relevance. Our hope is that as computational constraints diminish, we can keep this macroeconomic approach of solving for general equilibrium, while being able to consider the richness of the micro data.

XX: References to be added

A From the NIPA tables to table 1

xxx

B Calibration of the idiosyncratic skill shocks

xxx

C Further details on the solution algorithm

xxx