



WORKING PAPERS

RESEARCH DEPARTMENT

**WORKING PAPER NO. 03-20
ON THE WELFARE GAINS OF ELIMINATING
A SMALL LIKELIHOOD OF ECONOMIC CRISES:
A CASE FOR STABILIZATION POLICIES?**

Satyajit Chatterjee
Federal Reserve Bank of Philadelphia

Dean Corbae
University of Texas

Revised September 2003

FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • (215) 574-6428 • www.phil.frb.org

WORKING PAPER NO. 03-20

On the Welfare Gains of Eliminating a Small Likelihood of Economic Crises: A Case for Stabilization Policies?¹

Satyajit Chatterjee

Dean Corbae

Federal Reserve Bank of Philadelphia

University of Texas

December 1999 (Revised, September 2003)

¹We are very grateful to Hide Ichimura and Narayana Kocherlakota for help on this project and to an anonymous referee for suggesting several improvements. We also thank seminar participants at Arizona State, Boston University/Boston College, Duke, the Federal Reserve Banks of Cleveland and Minneapolis, Paris-I, Penn State, Texas, USC, the Wharton School, and conference participants at the Econometric Society, Midwest Macro, NBER Summer Institute, SED, and the Denver FR System Meeting for their input. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or of the Federal Reserve System.

Abstract

In this paper we estimate the potential benefit of policies that eliminate a small likelihood of economic crises. We define an economic crisis as a Depression-style collapse of economic activity. For the U.S., based on the observed frequency of Depression-like events, we estimate the likelihood of encountering a depression to be about once every 83 years. Even for this small probability of moving into a Depression-like state, the welfare gain from setting it to zero can range between 1 and 7 percent of annual consumption, in perpetuity. These estimates are large in comparison to welfare costs typically found for microeconomic distortions and suggest that there may be a net benefit to policies directed toward preventing economic instability.

1 Introduction

Since Robert Lucas's (1987) provocative argument, the existence of sizable welfare gains from the pursuit of stabilization policies has become a matter of debate. Lucas argued that the welfare gains from eliminating postwar variability in aggregate consumption was something on the order of one-tenth of 1 percent of annual U.S. consumption.

A number of papers have since explored the robustness of Lucas's finding. One group of authors maintained Lucas's assumption of complete markets but pursued the implications of alternative and less restrictive preference specifications on the magnitude of the welfare gain. An incomplete list includes Obstfeld (1994), Pemberton (1996), Dolmas (1998), Alvarez and Jermann (2000), and Tallarini (2000). These studies have obtained larger welfare costs of business cycles, but Otrok (2001) argues that when preference parameters are chosen to be consistent with business-cycle behavior, the welfare costs of business cycles are of the same order of magnitude as those obtained by Lucas.¹ A second group maintained standard preference assumptions but pursued the implications of incomplete markets for welfare calculations. The seminal paper in this group is Imrohorglu (1989), who assumed that the risk of unemployment was uninsurable. While she obtained larger welfare costs, Atkeson and Phelan (1994) and Krusell and Smith (1999) pointed out that if stabilization policies simply remove the correlation in the timing of individual episodes of unemployment, the effect of stabilization policies on welfare is much smaller and might even be perverse.²

We revisit the question of the welfare gains of stabilization policies by esti-

¹Ramey and Ramey (1995) have shown that, across countries, the variance of the innovations to real GDP growth tends to be negatively correlated with real GDP growth, suggesting there may be a growth benefit to reduction in business-cycle volatility. Barlevy (2000), Jones, Manuelli and Stacchetti (1999), Matheron and Maury (2000), and Portier and Puch (1999) have investigated the effects of reduced cyclical volatility on capital accumulation, but there is no consensus yet on whether these effects entail sizable welfare gains.

²One study that does claim to find sizable welfare gains from reduction of uninsured cyclical wage risk is Beaudry and Pages (1999).

imating the potential welfare gain from pursuing policies that prevent the occurrence of a Depression-style collapse of economic activity. In altering the scope of the welfare calculation from a reduction in volatility to a reduction in the likelihood of economic crises, we are moving the focus of the debate toward what we think is the ultimate goal of real-world stabilization policies. In the U.S., the various laws aimed at stabilizing the financial system (such as federal insurance of bank deposits and the concomitant regulation of commercial banks), automatic stabilizers (such as the state-run system of unemployment insurance), and the authority to conduct discretionary monetary and fiscal policy (the Employment Act of 1946) came into being during the Depression years or shortly thereafter. Aside from providing partial relief from unemployment and financial loss, the main intent of these policies was to prevent another occurrence of a depression. Formally speaking, the intent of these policies seems more consistent with an effort to eliminate the lower tail of the distribution of individual consumption rather than a mean-preserving shrinkage of both the upper and lower tails.³ Consequently, a calculation of the potential welfare gain from the pursuit of stabilization policies is incomplete if no attempt is made to quantify the gain from a reduction in the likelihood of economic crises.

To further motivate the project, Figure 1 plots the annual unemployment rate for the period 1900 to 1998. The striking aspect of this time series is the extraordinary rise in unemployment between the years 1930 and 1939, generally identified in history as the Depression years.⁴ The rise is extraordinary not only because it has not been repeated but also because there is no correspond-

³The focus of the welfare cost of business cycles literature is on this effect as is evident, for instance, in Alvarez and Jermann (2000).

⁴For the period 1900-1940, the Lebergott series for industrial unemployment was constructed by dividing the total number of unemployed workers reported in Lebergott's Table A-3 by the sum of unemployed workers and nonfarm workers also reported in that table. This construction assumes that most unemployed workers were in nonfarm occupations. The unemployment rates for 1941 and later are just those reported by the BLS. The Romer series was constructed by applying the corrections suggested by Romer (1986) to the industrial unemployment rate series.




Figure 1: Unemployment Rate 1900-1998

ing episode involving a steep decline in unemployment rates. Indeed, with an average unemployment rate of about 8 percent in non-Depression years, it is impossible for the unemployment rate to fall below this average as much as it rose above it during the Depression. In short, it is difficult to view Figure 1 and not think of the Depression as a huge, potentially preventable, lapse from the normal workings of the economy.

We take the defining characteristic of an economic crisis to be a very high unemployment rate of workers, similar to what was experienced in the U.S. during the Depression. The question we ask is: “What fraction of annual consumption would a worker be willing to pay to set the current probability of encountering a Depression-like event to zero?” To answer this question we construct an environment with the following features. There is a continuum of workers who encounter stochastic employment opportunities. The probability

of finding employment depends on the aggregate state of the economy. One of these aggregate states corresponds to an economic crisis where the probability of finding employment in the private sector is much lower relative to the other aggregate states. Workers cannot buy insurance against shocks to their employment status in the depression state, but they can self-insure by holding stocks of an asset whose return is lower than the (common) rate of time preference of individuals.⁵ This is perhaps the simplest environment that permits an analysis of the welfare consequences of eliminating the likelihood of economic crises.

Our calculations start with an estimate of the current likelihood of depressions, the likelihood that we set to zero in our welfare experiments. We obtain an estimate of this likelihood by fitting a three-state Markov chain to the observed monthly chronology of expansions, contractions, and depressions (in the U.S.) for the period 1900 to 1998.⁶ In fitting one Markov chain to the entire period we are ignoring any difference in the likelihood of depressions between pre- and post-Depression eras. We take the conservative position that the fact that no economic crisis has occurred since Depression-era stabilization policies went into place reflects luck rather than design. Under this assumption, we estimate the current likelihood of moving into a depression to be once every 1000 months (or once every 83 years).⁷ For the baseline calibration, the steady-state welfare gain from setting this small probability to zero is 1.87 percent of annual consumption, in perpetuity. We find that 58 percent of the total gain in welfare can be attributed to changes in second and higher-order moments of the (individual-level) consumption process, including a decline in variance. The remainder comes from an increase in mean consumption. Sensitivity analysis indicates that the gain could be as low as 1.3 percent or as high as 6.6 percent. Higher estimates are associated with larger contribution of changes in consump-

⁵This assumption is consistent with the general equilibrium implications of imperfect insurance, as shown, for instance, by Aiyagari (1994).

⁶An alternative approach to estimating the likelihood of a Depression-like event is to link it to the equity premium, as is done in Rietz (1988) and Danthine and Donaldson (1998).

⁷In fitting one Markov chain we are also ignoring any differences in the frequency of occurrences of non-depression states before and after the Depression.

tion volatility. For the experiment that generated the 6.6 percent gain, a full 80 percent of the total gain resulted from a reduction in consumption volatility (and changes in higher-order moments).

De Long and Summers (1988) made a criticism of Lucas-style welfare-cost calculations that may appear related to our work but is actually quite different. These authors took the view that successful stabilization policies “fill in business cycle troughs without shaving off business cycle peaks,” and so reduce the average unemployment rate and have a first-order effect on welfare (see also Cohen (2000)). In contrast, our welfare calculations assume that stabilization policies can eliminate crises but *not* ordinary recessions; in fact, we assume that stabilization policies work by turning what might have been a depression into an ordinary recession.⁸ Furthermore, our calculations allow for the effects of uninsured income risk. This element of realism is important because, as noted above, in most of our experiments the majority of the welfare gains come from changes in second and higher-order moments of the consumption process, including a decline in its variance. In fact, we document that a framework that ignores the risk of earnings loss from unemployment (as would, for instance, a representative agent model) predicts substantially lower welfare gains from an elimination of a Depression-like state.

Our study could also serve as a useful input into debates concerning the correct response of policy to the economic and financial crises that occurred around the world in the 1990s. At the risk of over-simplifying a complex issue, the types of policies discussed or implemented can be classified as belonging to one of two types: policies that seek to prevent a repeat occurrence of a crisis and policies that seek to implement a recovery in the midst of a crisis. The first

⁸Interestingly, in a commentary on DeLong and Summers’ paper, Martin Bailey pointed out that the most compelling case in favor of DeLong and Summers’ general point that stabilization policies may raise average incomes was the experience of the Great Depression, and he reasoned that “stabilization policies should be designed to avoid persistent downturns such as the Great Depression” (p. 494 of DeLong and Summers). Our paper makes the quantitative case for stabilization policies on precisely this ground.

group includes policies that seek to regulate financial trades thought to pose a risk of financial and economic crises. Since regulations generally interfere with the normal working of markets and impose efficiency costs, there is always a question as to how extensive they should be. By providing a methodology for assessing the gains from eliminating economic crises (and giving an estimate for a specific type of crisis), our paper provides an important ingredient currently missing in the cost-benefit analysis of crisis-related regulation.

2 Environment

The economy evolves through good (g), bad (b), and depression (d) times that have implications for employment prospects. The state of the economy $\eta \in \{g, b, d\}$ is assumed to follow a first-order Markov process. The transition matrix of η is given by:

$$\Lambda = \begin{bmatrix} \lambda_{gg} & \lambda_{bg} & \lambda_{dg} \\ \lambda_{gb} & \lambda_{bb} & \lambda_{db} \\ \lambda_{gd} & \lambda_{bd} & \lambda_{dd} \end{bmatrix}$$

where, for example, $\Pr\{\eta_{t+1} = g | \eta_t = b\} = \lambda_{gb}$.

The economy consists of a large number of infinitely lived individuals who differ at any point in time in their asset holdings and employment opportunities. Each individual maximizes

$$E \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\gamma}}{1-\gamma}$$

where c_t consumption in period t , $\beta \in (0, 1)$ is the discount factor and $\gamma > 0$ is the relative risk aversion parameter (for $\gamma = 1$, the function is taken to be \ln function).

Individuals are endowed with one indivisible unit of time each period. Each individual receives an employment opportunity that has one of two states, $i \in \{e, u\}$, where e stands for the employed state and u for the unemployed state. If

$i = e$, the individual produces y units of the consumption good in the business sector, and if $i = u$, the individual produces θy units of the consumption good in the non-business sector, where $0 < \theta < 1$. Without loss of generality, we set $y = 1$.

The individual-specific employment state is assumed to follow a first-order Markov process. The transition matrix is given by:

$$\Lambda^\eta = \begin{bmatrix} \lambda_{ee}^\eta & \lambda_{ue}^\eta \\ \lambda_{eu}^\eta & \lambda_{uu}^\eta \end{bmatrix}$$

where, for example, $\Pr\{i_{t+1} = e | i_t = u, \eta_{t+1} = g\} = \lambda_{eu}^g$ is the probability that an individual will be employed in good times at $t + 1$ given the individual was unemployed in period t .

The overall employment prospects faced by each individual depend on both the aggregate and individual states; that is, on the six pairs (η, i) , $\eta \in \{g, b, d\}$ and $i \in \{e, u\}$. These six pairs are denoted by $\omega^1, \dots, \omega^6$, where ω^1 stands for employed in a good state, ω^2 stands for unemployed in a good state, ω^3 stands for employed in a bad state, ω^4 stands for unemployed in a bad state, ω^5 stands for employed in a depression state, and ω^6 stands for unemployed in a depression state. The process governing ω is a first-order Markov process with transition matrix given by $\Phi = [\phi_{jk}]$, where $\Pr\{\omega_{t+1} = \omega^j | \omega_t = \omega^k\} = \phi_{jk}$. The transition probabilities are determined by Λ and Λ^η . For example, if $\omega_t = \omega^1$, then the probability of $\omega_{t+1} = \omega^2$, i.e., ϕ_{21} , is given by $\lambda_{gg}\lambda_{ue}^g$.

The market arrangement in the baseline model is as follows. Individuals can purchase unemployment insurance in the two non-depression states but not in the depression state. Given this (partial) incompleteness in insurance markets, individuals may have an incentive to self-insure and we assume they can do so by holding stocks of an asset with zero real return. We defer a discussion of these assumptions to later in the paper but note here that alternative asset market assumptions are explored in the study. Individuals enter period t with individual savings s_t held over from the previous period. An individual's budget

constraint can be written:

$$\begin{aligned} c(\omega_t) + s_{t+1} &= y(\omega_t) + s_t, \forall t, \omega \\ s_t &\geq 0 \end{aligned}$$

where $c(\omega_t)$ is the individual's consumption and $y(\omega_t)$ is the individual's post-insurance income.

The maximization problem faced by an individual in this economy can be represented as a discounted dynamic program. Let $s = s_t$, $\omega = \omega_t$, $s_{t+1} = s'$, and $\omega_{t+1} = \omega'$. Then, the Bellman equation for this program is:

$$V(s, \omega) = \max_{s' \geq 0} \frac{c_t^{1-\gamma}}{1-\gamma} + \beta \sum_{\omega'} \Phi(\omega', \omega) V(s', \omega') \quad (1)$$

subject to

$$c = y(\omega) + s - s' \geq 0. \quad (2)$$

Since individuals face idiosyncratic shocks in the depression state, they may hold different levels of savings. Let $\mu_t(s, \omega)$ be the probability that an individual attains the state (s, ω) . Then, the probability that state (s', ω') occurs is given by:

$$\mu_{t+1}(s', \omega') = \sum_{\omega} \sum_{s \in \Xi(s', \omega)} \Phi(\omega', \omega) \mu_t(s, \omega) \quad (3)$$

where $\Xi(s', \omega) = \{s : s' = s'(s, \omega)\}$. Under mild regularity conditions (ergodicity of the Markov process and the absence of cyclically moving subsets) the sequence of recursively defined distributions converges to a unique invariant distribution $\mu(s, \omega)$ from any initial distribution. The distribution $\mu(s, \omega)$ gives the fraction of time an individual is in state (s, ω) .

3 Estimates of the Aggregate State Transition Matrix

In order to estimate the aggregate state transition matrix we proceed by constructing a history of these aggregate states. We begin with the monthly NBER business-cycle chronology, which dates from December 1854. We associate NBER expansions with the good state and NBER contractions with the bad state. This *two-state* history is then augmented with a definition of what it means to be in a depression. If that definition is observed to be satisfied by some month, then that month's NBER classification is changed to the depression state.

As noted in the introduction, we take the defining characteristic of a depression to be a very high incidence of unemployment among industrial workers. But unemployment rate data are available only for the period beginning 1900, and for the pre-WWII portion of that period, it is available at an annual frequency only. Because of this data limitation, we confine our three-state history to the period 1900 to 1998.⁹

For our baseline calculation, we classified all months of any year in which the unemployment rate exceeded 17 percent as depression months. This definition simply picks out the 120 months corresponding to the 1930-1939 period generally known as the "Depression years."¹⁰ Accordingly, we changed the NBER classification of these months to the depression state. An alternative definition considered later in the paper classifies all months of any year in which the unemployment rate exceeded 20 percent as depression months.¹¹

⁹There is some fragmentary information on unemployment rates for the last decade of the nineteenth century (see, for instance, Lebergott (1964, Table A-15), Romer (1986, Table 9) and Keyssar (1977, Ch. 2)). It appears that for 5 out of those 10 years, industrial employment rates were very high. In an earlier version of this paper we showed that including this information raises the estimate of the welfare gains from eliminating depressions.

¹⁰Cole and Ohanian (1999) also identify the 10 years between 1930 and 1939 as the period during which output remained below trend.

¹¹A more sophisticated alternative would be to fit a 3-state regime to the unemployment

Given this three-state history, the maximum likelihood estimate of λ_{kj} , the (j, k) th element of the aggregate state transition matrix, is the ratio of the number of times the economy switched from state j to state k to the number of times the economy was observed to be in state j (Ross (1972) pp. 240-242).¹² Implementing this procedure for the whole sample yields the following estimate of Λ , with standard errors in parentheses below:

$$\hat{\Lambda} = \begin{bmatrix} 0.9766 & 0.0234 & 0 \\ (0.0053) & (0.0053) & (0) \\ 0.0745 & 0.9216 & 0.0039 \\ (0.0164) & (0.0168) & (0.0039) \\ 0.0083 & 0 & 0.9917 \\ (0.0083) & (0) & (0.0083) \end{bmatrix}$$

The estimated matrix has several noteworthy features. First, because there is only one depression episode in our sample, there is only one transition into and one transition out of the depression state. In the three-state history we construct, the depression follows contractionary months and is followed by expansionary months. Hence $\lambda_{dg} = \lambda_{bd} = 0$. Second, the estimated matrix implies that conditional on not being in a depression, the probability of falling into one is 0.0010. Third, the *unconditional* probability of a depression is 0.0975, which rate data using the procedure described in Hamilton (1989). We followed our simple procedure because for the pre-WWII period the NBER chronology is likely to be a better proxy for the frequency of good and bad times than any that can be inferred from the noisy unemployment series (see Romer (1986) for a discussion of the pitfalls of the pre-WWII unemployment series for cyclical analysis).

¹²The estimated transition probabilities are given by

$$\hat{\lambda}_{kj} = \frac{\sum_{t=1}^{T-1} 1\{\eta_{t+1} = k\}1\{\eta_t = j\}}{\sum_{t=1}^{T-1} 1\{\eta_t = j\}}$$

Given the Markov structure of our problem, the asymptotic standard errors of these estimates are given by:

$$\sqrt{\frac{\hat{\lambda}_{kj} (1 - \hat{\lambda}_{kj})}{\sum_{t=1}^T 1\{\eta_t = j\}}}$$

is orders of magnitude larger than the conditional probability. The large discrepancy between these two probabilities reflects the fact that the depression state is very persistent. This discrepancy is one reason why the welfare loss from the possibility of a Depression-like event is relatively large, even though the probability of encountering a Depression-like event, conditional on not being in one, is quite small.¹³

A word about the precision of these estimates. The fact that there is only one depression episode in our sample might be thought to imply that none of the parameters relating to the third state (the third column and row of the Λ matrix) can be reliably estimated. That’s not necessarily true. According to our history, the economy spent about 1,070 months in non-depression states. Thus, there were many instances in which the economy could have gone into a depression but didn’t. The fact that the depression state was encountered only once out of more than 1000 trials suggests we can be quite confident that the probability of moving into a depression state is very low. Similarly, the economy spent 120 months in the depression state before moving out of it. The fact that it took more than 100 trials for the economy to leave the depression state implies we can be reasonably confident that the probability of continuing in the depression state is quite high. As we shall see, these two features of a depression state, namely, the low probability of encountering one and its persistence once it is encountered, are the economically significant features.¹⁴

¹³The unconditional probability of a good state is 0.6951, and the unconditional probability of a bad state is 0.2074.

¹⁴We note, however, that the standard errors reported in parentheses are *asymptotic* standard errors and needn’t be good estimates of the sampling variance in “small” samples. To investigate the small sample properties of our maximum likelihood estimate of Λ , we ran Monte Carlo simulations where the data generation process is given by $\hat{\Lambda}$. As expected, the standard errors from the Monte Carlo simulations were larger than the asymptotic standard errors. Furthermore, we found an upward bias in the estimates of λ_{db} and λ_{gd} . Since correcting $\hat{\Lambda}$ for these biases only led to *higher* welfare gains of eliminating the depression-like state, we retained the more conservative estimates of $\hat{\Lambda}$ reported in the paper.

4 Calibration of Other Parameters

The calibration of the remaining parameters involves selecting parameter values for the elements of the individual-level transition matrices Λ^η , the preference parameters β and γ , the earnings-loss parameter θ , and the post-insurance income terms $y(\omega)$.

The Individual State Transition Matrix

The individual-level state transition matrix for each aggregate state is built up from two pieces of information pertaining to that state, namely the average unemployment rate in that state and the average duration of unemployment spells in that state.

The average unemployment rate in the good, bad, and depression states was fixed at the average unemployment rate for these states in the whole sample. These were 5.33 percent, 7.86 percent, and 23.48 percent, respectively. Since the unemployment rate data are available at only annual frequencies for the pre-WWII era, the average unemployment rate for each state was calculated for annual data. All non-Depression years in which there were at least nine expansionary months were classified as “good” years and all other non-Depression years as “bad” years.¹⁵

The duration of unemployment spells in good and bad times is based on the monthly average duration of unemployment spells reported by the BLS. These were determined to be 2.75 months during expansions and 3.75 months during

¹⁵Because the unemployment rate falls during expansions and rises during contractions, our procedure for calibrating U^g and U^b underestimates the true difference between these parameters. As a check, we estimated the average unemployment rate for the last six months of each expansion and the average unemployment rate for the last six months of each contraction in the postwar period (according to Romer (1986), the unemployment rate process for 1900-1928 period is not significantly different from that in the postwar era, once allowance is made for likely measurement errors in the prewar unemployment data). The estimates were 4.70 percent and 6.74 percent, respectively. Since this method of estimating U^g and U^b leads to uniformly lower values than what we estimate for our baseline calibration, we found that they led to higher welfare gains (of eliminating the likelihood of depressions) than those reported later in the paper.

contractions. The only data on the duration of unemployment spells that we could find for the Depression were for 1930 and 1931. By early 1930, 56 percent of male unemployed workers had been without work for at least nine weeks. The special census of unemployment undertaken in January 1931 reported that of the male workers unemployed in Boston, New York, Philadelphia, Chicago, and Los Angeles, 45.3 percent, 60.9 percent, 45.2 percent, 61.0 percent, and 33.2 percent, respectively, had been jobless for at least 18 weeks. In effect, the median unemployment duration had doubled in less than a year. The fact that the unemployment rate remained elevated for the next *seven* years suggests that the median duration of unemployment by the end of the Depression was probably a lot higher than 18 weeks. We fixed the average duration of unemployment spells in the depression state as 20 months, more than four times the median duration seen in 1931.¹⁶

The choice of average duration of unemployment spells for each aggregate state fixes λ_{uu}^η for $\eta \in \{g, b, d\}$ (and, also $\lambda_{eu}^\eta = 1 - \lambda_{uu}^\eta$). We chose the remaining elements to match the average unemployment rate in each aggregate state. Note that the evolution of the aggregate unemployment rate is given by:

$$U_t = U_{t-1}\lambda_{uu}^{\eta(t)} + (1 - U_{t-1})\lambda_{ue}^{\eta(t)}$$

where $\eta(t) \in \{g, b, d\}$. Since λ_{uu}^η etc. depend only on the current state, U_t converges to a constant if the state remains unchanged for some length of time. For each aggregate state, these limiting unemployment rates solve:

$$U^\eta = U^\eta\lambda_{uu}^\eta + (1 - U^\eta)\lambda_{ue}^\eta.$$

We chose the values of $\lambda_{ue}^\eta, \eta \in \{g, b, d\}$, so that U^g, U^b , and U^d matched 5.33 percent, 7.86 percent, and 23.48 percent, respectively.¹⁷

¹⁶We have experimented with setting the average duration of unemployment in the depression to 10 months. For reasons explained later in the paper, this change didn't affect the estimate of the welfare gain from elimination of depressions very much at all.

¹⁷These choices imply that average unemployment rate in the good state is somewhat larger than U^g , and the average unemployment rate in the bad and depression states is somewhat less than U^b and U^d , respectively. However, since all three states are highly persistent, these discrepancies are minor.

Preference and Earning-Loss Parameters

We set $\beta = 0.9946$, which is equivalent to an annual discount rate of 6 percent. We arrived at this number by assuming a rate of time preference equal to 4 percent at an annual rate as well as assuming that the constant monthly survival probability is equal to $1 - 1/(40 \cdot 12)$, so that individuals have a working life of 40 years. We set the risk aversion parameter, γ , to 3.

The value of θy is given by “home production.” According to Greenwood, Rogerson, and Wright (1995), “attempts to measure the value of the output of home-production come up with numbers between 20 and 50 percent of the value of measured market GNP.” To be conservative, we set the earning loss parameter θ to 0.5 in the baseline calibration.¹⁸

Insurance Payments

While we take a high unemployment rate of workers as a defining characteristic of a depression, such an event is likely to have consequences for the functioning of asset markets as well. In particular, suppose there are two kinds of assets, one of which is issued by the business sector and another by the government. In normal times, the return on the business-sector asset is close to the rate of discount while the return on the government asset is zero. In a depression, the return on the government asset is still zero, but the business-sector asset becomes worthless. In this world, workers will accumulate stocks of the business-sector asset to self-insure against the risk of unemployment in normal times but use the government asset to insure against unemployment during depressions. If the rate of return on the business-sector asset is close to the rate of discount, we know from Bewley (1977) that the worker will accumulate enough of the asset to almost perfectly insure against unemployment risk during normal times. Therefore, a rough way to capture this situation is to assume that both employed and unemployed workers receive the per capita endowment in the

¹⁸Darby (1976) pointed out that workers engaged in government relief programs during the Depression were counted as unemployed. Darby also reports that the average wage earned by these “unemployed” workers during the years 1930-1939 was about 41 percent of the average wage during those years, which is lower than our baseline calibration of 50 percent.

good and bad states (so there is no risk of loss of income due to unemployment in these times) but confront the risk of unemployment in the depression state. Specifically, for $\eta = g$ and b we set $y(\eta, i) = (1 - U^\eta) + \theta U^\eta$ for all i , while we set $y(d, e) = 1$ and $y(d, u) = \theta$.¹⁹ Later, we examine what happens if insurance is not available in any of the states and if it's available in all states. This last case corresponds to doing welfare calculation in the context of representative agent.

5 The Response of Per Capita Consumption in a Depression

Given our objective, it's important for the calibration to deliver the decline in per capita consumption observed during the Depression years. To see how well the model captures the decline in per capita consumption during the Great Depression, we simulated our model with the observed history of aggregate states, starting with an initial distribution of asset holdings corresponding to the average over good states (since 1900 was an expansionary year).²⁰ Figure 2 plots the computed percentage deviations of the simulated per capita consumption against the percentage deviation in actual per capita consumption.²¹ This is done for two polar opposite assumptions concerning the extent to which non-business-sector income is measured in real GDP. The simulated path shown by

¹⁹The aggregate unemployment rate is not equal to U^η in the first few periods following the economy's arrival into state η . However, switches between aggregate states are relatively rare and convergence to U^η is always very rapid. Living with this minor discrepancy saves us from putting U as another state variable in the dynamic program.

²⁰We obtain decision rules for optimal asset holdings by successive approximations on the value function $V(s, \omega)$. We discretize the state space of asset holdings to lie between 0 and 10.8 in increments of 0.027 for a total of 401 grid points. The upper bound is roughly equal to 11 months of income if the employed state continues for that long. In equilibrium, this constraint is never binding.

²¹The consumption series is based, in part, on the annual Kendrick real consumption series for 1889-1953 reported in Appendix B of Gordon (1986), deflated by population. The percentage deviations shown in the figure are taken from a quadratic trend.




Figure 2: Simulated & Actual Per Capita Consumption Paths with Home Production Measured or Not Measured

the dashed line assumes that all of the non-business-sector income received by an unemployed individual is measured in real GDP. Under this assumption, per capita consumption drops about 16 percent but gradually recovers to a decline of about 10-12 percent by the end of the Depression. When the economy emerges from the Depression, per capita consumption rises sharply to around 10 percent above trend and then declines to its normal level value by around 1945. The dotted line shows what happens if non-business-sector income is not measured in real GDP. In this case, consumption declines about 24 percent in 1930, about 26 percent in 1931, and then recovers to a decline of about 24 percent toward the end of the Depression. The drop in per capita consumption is now much steeper because for unemployed individuals only consumption in excess of non-business-sector income is included in aggregate consumption and real GDP.

As is evident, these two cases “bracket” the actual decline in per capita consumption during the Depression. In the data, aggregate per capita consumption doesn’t fall below trend until 1931 and reaches its trough of around 19 percent in 1933. Then there is a recovery, with the path of consumption ending up in the neighborhood of the simulated consumption path by around 1945. The fact that, in comparison to the simulations, actual Depression-era per capita consumption is more volatile and reaches its trough later should not be too surprising. In the data, the Depression-era unemployment rate is more volatile, peaking at 30 percent. Furthermore, in the simulations individuals know right away they are in the Depression whereas the realization that something had gone very wrong was gradual in reality. These differences suggest that a better metric for judging how well the Depression is captured is to compare the cumulative consumption loss between 1930 and 1945. If non-business-sector income is measured, the cumulative consumption loss is 76 percent of mean aggregate consumption, and if it’s not measured it’s 227 percent. In the data, the cumulative consumption loss over the same period is 107 percent. If we assume that 23 percent of non-business-sector income goes unmeasured, the cumulative loss in consumption in the model matches that in the data.

In summary, the predictions of the baseline model for the path of per capita consumption during a depression does not appear to be grossly inconsistent with observations. We now turn to our welfare comparisons.

6 Welfare Estimates

We wish to estimate the utility gain from moving to an environment for which the $\hat{\Lambda}$ matrix is replaced by

$$\Lambda^* = \begin{bmatrix} 0.9766 & 0.0234 \\ 0.0745 & 0.9216 + 0.0039 \end{bmatrix}$$

The off-diagonal elements of this matrix are identical to the corresponding elements of $\hat{\Lambda}$, as is λ_{gg} . But the probability of remaining in the bad state is now

higher by 0.0039, the probability of moving into a depression from a bad state in the $\widehat{\Lambda}$ matrix. The assumption here is that stabilization policies prevent ordinary recessions from turning into depressions. The individual level transition matrices for the good and bad state remain the same, and the parameters γ , β and θ are assumed to be the same as well. Let $V^*(s, \omega)$ be the value function for this new, depression-proof, economy.

The welfare calculations are done in two ways. Imagine that the three-state economy has attained its stochastic steady state. At some random date, individuals are given the choice of living in an environment with Λ^* . At that instant, the economy will be in one of three possible states, and there will be a joint distribution of individuals across asset holdings and employment status. We can imagine asking each individual in this distribution the maximum he is willing to pay each period in the two-state depression-proof environment for the privilege of living in that environment.

In the first type of welfare calculation, which is our preferred type, we assume that each individual begins the new regime with his current asset-holding and employment status. In addition, we assume that if the economy is in the good or bad state, then the new regime will begin in that state as well, and if the economy is in the depression state, then the new regime will begin in the bad state. Thus, the fraction of consumption the individual is willing to give up if he is currently in state $(s, \omega), \omega \leq 4$ is found by computing $1 - \alpha(s, \omega)$, where $\alpha(s, \omega)$ solves:

$$V(s, \omega) = \alpha(s, \omega)^{1-\gamma} V^*(s, \omega).$$

If the economy is in a depression, then $\alpha(s, 5)$ and $\alpha(s, 6)$ are computed as follows:

$$\begin{aligned} V(s, 5) &= \alpha(s, 5)^{1-\gamma} V^*(s, 3) \\ V(s, 6) &= \alpha(s, 6)^{1-\gamma} V^*(s, 4) \end{aligned}$$

Denoting the invariant measure for the (three aggregate state) depression-prone environment by $\widehat{\mu}(s, \omega)$ (this probability distribution is the unconditional prob-

ability of an individual having assets s in state ω) the average gain in welfare across individuals is given by $1 - \bar{\alpha} = 1 - \sum_s \sum_\omega \hat{\mu}(s, \omega) \alpha(s, \omega)$.

In the second type of calculation we assume that each individual is offered the *average* lifetime utility in the depression-proof environment. In this case $\alpha^{SS}(s, \omega)$ is given by:

$$V(s, \omega) = \alpha^{SS}(s, \omega)^{1-\gamma} \bar{V}^*$$

where \bar{V}^* is $\sum_s \sum_\omega \mu^*(s, \omega) V^*(s, \omega)$ with $\mu^*(s, \omega)$ being the invariant distribution in the depression-proof economy. Then, $1 - \bar{\alpha}^{SS} = 1 - \sum_s \sum_\omega \mu(s, \omega) \alpha^{SS}(s, \omega)$.

We refer to this measure as the *steady-state* gain in welfare. The difference between our preferred measure and the steady-state measure is that the former takes account of the fact that after the regime change individuals will want to decumulate assets, since there is less uncertainty in the new depression-free regime. The additional consumption permitted by this decumulation along the transition path to the new steady state is taken into account in our preferred measure but ignored by the steady-state one.

Table 1
Welfare Gains From Eliminating Depressions and Cycles
(As a Percentage of Per Period Consumption)

Eliminating Depressions		Gains From Eliminating Cycles	
Total Gain	SS Gain	Lucas's Estimate	Imrohoroglu's Estimate
1.87	1.70	0.01	0.3

Both calculations are reported in Table 1. The total welfare gain (including the gain from the decumulation of assets along the transition path) is 1.87 percent of consumption per month (or per year) and the steady-state gain is 1.70 percent.²² To put these numbers in perspective, note that Lucas estimated the welfare gain from eliminating all cyclical volatility in the postwar era to be 0.01 percent of consumption and Imrohoroglu estimated it to be 0.3 percent. These authors computed steady-state gains so the relevant comparison is with our steady-state gain measure. We find that the gain from getting rid of a Depression-like state is 170 times Lucas's (1987) estimate of the gains from eliminating cycles and about 6.5 times Imrohoroglu's (1989) estimate.²³

Experimentation reveals that the welfare gains from eliminating depressions vary approximately linearly with the likelihood of encountering a depression. If we thought that the true likelihood of encountering a Depression-like event was actually once every 1600 years (rather than once every 83 years), that would cut our estimated welfare gains by a factor of about 20. Note that if the true likelihood of encountering a depression was really once in 1600 years, the chance of encountering a depression episode in a 83-year sample would be around 5 percent. Thus, a welfare gain of around 0.094 percent ($= 1.87 \div 20$) corresponds to the lower bound of a 95 percent confidence interval of our point estimate.

²²It may be interesting to note that *conditional on being in a depression*, individuals are willing to pay, on average, 6.45 percent of annual consumption to receive the steady-state utility of the depression-free economy.

²³The total gain from elimination of the depression state depends on the value of θ . When this number is set closer to the value assumed in Imrohoroglu (0.25), the total welfare gain is around 20 times her estimate of the cost of business cycles.

Even under a most conservative estimate of the likelihood of depressions, the welfare gain is more than nine times larger than Lucas’s estimate of the welfare gain from eliminating cycles.

Where do these gains come from? Table 2, which lists the key operating characteristics of the two economies, indicates three relevant differences. First, average asset holdings go from being 0.41 of monthly earnings in the 3-state environment to 0 in the 2-state economy. Because the difference between the total gain in welfare and the steady-state gain is 0.17 percent, we can infer that 9.1 ($= (1.87 - 1.7) \div 1.87$) percent of the total welfare gain is due to the fact that individuals need to hold fewer assets in the new regime.

Table 2
Steady-State Properties of the 2-State and 3-State Environments

Models	\bar{s}	$\sigma(s)$	\bar{c}	$\sigma(c)$
3 – S	0.4077	1.3057	0.9628	0.0563
2 – S	0	0	0.9704	0.0054

Second, average consumption rises by 0.0076 units in the 2-state economy, an increase of 0.78 percent ($= 0.0076 \div 0.9704$) relative to mean consumption in the 2-state economy. Since an individual would be willing to forgo exactly 0.0076 units of consumption for a gain of 0.0076 units, we can infer that 0.78 percentage point of the total gain welfare gain results from this source. Thus, increase in mean income accounts for 41.7 ($= (0.78 \div 1.87) \times 100$) percent of the total welfare gain.

Third, the volatility of individual consumption is lower by a factor of 10 in the 2-state economy. Since the first two effects account for 50.8 percent of the total gain, we can infer that the remaining 49.2 percent of the gain must be due to reduction in the variance of consumption and changes in other higher-order moments of the consumption process. Thus the most important contributor to the total welfare gain is the reduction in the volatility of the consumption process. However, since the only reason individuals accumulate a

buffer stock of assets is to dampen fluctuations in consumption, the reduction in uncertainty associated with the elimination of the depression state accounts for 58.3 ($= 9.1 + 49.2$) percent of the total welfare gain.

The most dramatic difference in the operating characteristics of the two economies is in the volatility of individual consumption. In the 2-state economy, the volatility of consumption is low because unemployment insurance makes the volatility in an individual's post-transfer income equal to the cyclical volatility in per capita earnings. The cyclical variability in per capita earnings is low enough that individuals do not find it in their interests to accumulate the zero-return asset to buffer their consumption against these fluctuations. Hence, in every period individuals set their consumption equal to their post-transfer income (which leads to the mean level of asset holdings being zero). In the 3-state economy individuals do not find it optimal to accumulate assets in the two non-depression states even though they are aware that if the depression materializes unemployment insurance will cease. Consequently, when a depression does materialize the consumption paths of all individuals changes dramatically. Evidently, the probability of the depression state is low enough that individuals do not find it worth their while to save for it in advance via a low-return asset. Those who become unemployed at the start of the depression are the worst affected: they have no buffer stock of assets and no insurance and their consumption moves down with their earnings one-for-one. Those who continue to be employed recognize the possibility of earnings loss due to unemployment and also reduce their consumption in an effort to accumulate a buffer stock of assets. These big drops in consumption at the start of the depression contribute to the relatively high volatility of individual consumption in the 3-state economy.²⁴

²⁴These changes in individual and aggregate consumption occur even though a depression is assumed not to affect the earnings of employed and unemployed agents. Thus, we ignore any decline in productivity that may have occurred during depressions. Taking such effects into account would only raise our welfare gain estimates. Also, our study ignores any interaction between an individual's unemployment history and his or her business-sector earnings. If unemployment spells reduce future employment earnings, the welfare gain from elimination of depressions will be much larger.

The drop in *everyone's* consumption explains the drop in per capita consumption at the start of a depression in Figure 2. As the depression proceeds, the rate of asset accumulation of employed individuals begins to decline as they get closer to their target buffer stock of assets (of about 8 months of employed income) and their consumption begins to recover. This is the main reason aggregate consumption recovers after the initial drop. Another factor that contributes to the recovery is that individuals who become unemployed later in the depression experience less of a decline in consumption because they get time to accumulate assets.

The fact that a significant welfare gain from elimination of the depression state comes from a reduction in consumption volatility gives our findings a flavor similar to more recent studies of the welfare cost of business cycles. Storesletten, Telmer and Yaron (2001) and Krebs (forthcoming) show that the welfare gain from the elimination of cyclical variation in uninsured idiosyncratic risk can be quite large if permanent income shocks are an important component of this risk. While we don't model permanent idiosyncratic income shocks, unemployment during a depression, and the depression itself, are quite persistent states. People who become unemployed during a depression can expect their income to be low for a relatively long period of time, a fact that contributes to the volatility of consumption in the 3-state economy.

7 The Role of Uninsured Income Risk

The size of the welfare gain depends on the assumption that earnings loss from unemployment cannot be insured in the depression state. This becomes evident if we re-run our experiment under the assumption that unemployment insurance is available in all states, i.e., an individual's post-transfer income in state η is $(1 - U^\eta) + \theta U^\eta$, for $\eta = g, b, d$. Table 3 displays the operating characteristics of the two economies.

Table 3

The Effect of Full Unemployment Insurance

Total Gain = 0.97%, Steady State Gain = 0.98%

Models	\bar{s}	$\sigma(s)$	\bar{c}	$\sigma(c)$
3 - S	0	0	0.9628	0.0266
2 - S	0	0	0.9704	0.0054

Eliminating uninsured employment risk eliminates the need to save in all aggregate states. The uncertainty in post-transfer earnings is too low to overcome the difference in the rate of return on savings (zero) and the rate of discount (6 percent, annualized). Since decumulation of assets is no longer an effect, there is very little difference between the total and steady-state gain in welfare. There is still a substantial reduction in consumption volatility in the 2-state economy arising purely from the fact that variation in post-transfer earnings (or, equivalently in per capita earnings) is greater in the 3-state economy compared with the 2-state economy. However, the increase in average consumption in the 2-state economy (which is still 0.0076 units or 0.78 percent of mean consumption in 2-state economy) now accounts for 80.4 percent of the total welfare gain. Since welfare calculations for the economy with full unemployment insurance is equivalent to welfare calculations with a representative agent, these findings establish that a representative agent approach to our problem will lower the estimate of the total gain in welfare almost 50 percent and attribute most of the gain to an increase in average consumption resulting from the elimination of the depression state. In this sense, uninsured unemployment risk during depression matters.²⁵

²⁵It's worth noting that our representative agent estimate of 0.2 (= (0.98 - 0.78)) percent for the welfare gain resulting purely from a reduction in volatility is very close to Lucas' (representative agent) estimate of the welfare gain from elimination of pre-WWII cyclical volatility. As noted in his 1987 book (pp. 26-28), variance in de-trended (log) aggregate consumption in the pre-WWII era is 0.0015, which, when multiplied by $\frac{1}{2}\gamma$ for $\gamma = 3$, gives a welfare gain estimate of 0.0023, or 0.23 percent.

How do the results change if we assume that insurance against earnings loss is not available in *any* state? Table 4 displays the operating characteristics of such an economy. The total gain in welfare from the elimination of the depression state is now 1.56 percent and the steady-state gain is 1.4 percent. With regard to the contribution of the three different channels, the reduction in buffer stock assets contributes about 10.3 ($= (1.56 - 1.4) \div 1.56$) percent, the increase in mean consumption (which is still 0.78 percent of average consumption in the 2-state economy) contributes 50 ($= 0.78 \div 1.56$) percent, and the remaining 39.7 percent results from a reduction in variance and other changes in the higher-order moments of the consumption process.

Table 4

The Effect of the Absence of Unemployment Insurance

Total Gain = 1.56%, Steady State Gain = 1.40%

Models	\bar{s}	$\sigma(s)$	\bar{c}	$\sigma(c)$
3 - S	2.23	1.16	0.9628	0.0716
2 - S	1.86	0.53	0.9704	0.0543

The reason gains are somewhat lower than in our baseline model is evident from the behavior of asset holdings. Because individuals now face the risk of earnings loss from unemployment in all aggregate states, they find it in their interests to accumulate a buffer stock of assets not only in the depression state but also in good and bad states. One consequence of this behavior is that when the depression materializes, all individuals are somewhat better prepared than in the baseline model; individuals who lose their jobs at the start of the depression now have some savings to cushion the blow, and individuals who continue to remain employed have only to add to their existing buffer stock of assets rather than start from scratch. Consequently, the variability of individual consumption in the depression-prone economy is somewhat less now than in the baseline model. This works to reduce the welfare gain from an elimination of the depression state. In this case, the reduction in uncertainty accounts for 50 percent ($10.3 + 39.7$) of the welfare gain. The key lesson here is that

improvements in risk-sharing that are unlikely to survive a depression-like event make it more important to eliminate the possibility of such events through stabilization policies.²⁶

8 Sensitivity Analysis

In this section we report the sensitivity of our results to changes in key parameter values. The results are collected in Table 5.

Table 5

Experiments	3-State Economy			2-State Economy			Welfare		
	\bar{s}	\bar{c}	$\sigma(c)$	\bar{s}	\bar{c}	$\sigma(c)$	TG	SS	RU
Baseline	0.41	0.96	0.07	0	0.97	0.01	1.87	1.70	58.3
$\theta = 0.2$	1.47	0.94	0.09	0	0.95	0.01	6.59	6.00	80.4
$\gamma = 1.0$	0.17	0.96	0.06	0	0.97	0.01	1.30	1.22	40.0
$\bar{U} \geq 0.20$	0.23	0.96	0.05	0	0.97	0.01	1.58	1.48	60.8
Growth = 2 %	0.25	0.96	0.06	0	0.97	0.01	1.96	1.80	59.9

For ease of comparison, the first set of results are those from the baseline model. The next line reports results if the income in the unemployed state is set at 20 percent of income in the employed state (this is the lower bound on the income from home production reported in the Greenwood, Rogerson and Wright study mentioned earlier). As one would expect, average consumption is now lower, average savings higher, and volatility of consumption higher than in the baseline model. The total gain from eliminating the depression state is now 6.59 percent while the steady state gain is 6 percent. With regard to the sources

²⁶A second lesson is that real-world features that impinge on a household's level of precautionary savings, but which are ignored in this paper, could have a bearing on the calculation of welfare gains. One such feature is habit formation, which is known to increase the level of precautionary savings by hefty amounts (Diaz, Pijoan-Mas, and Rios-Rull (forthcoming)). However, because habit formation also increases the utility loss from steep declines in consumption, introducing this feature may increase both precautionary savings *and* our estimate of the welfare gain.

of the gain, the final column reports the combined percentage contribution due to reduction in uncertainty (the contribution from the decumulation of assets plus the contribution from the reduction in volatility of consumption), which is 80.4 percent of the total gain in welfare.

The next experiment reduces the relative risk aversion parameter to 1. Relative to the baseline model, average asset holdings fall, and there is a modest decline in the volatility of consumption. Predictably, the welfare gain estimates are now lower, with the total gain being 1.3 percent and the steady state gain 1.22 percent. The contribution due to reduction in uncertainty is now only 40 percent.

In the third experiment, we defined depression months to be all months of any year in which the unemployment rate exceeded 20 percent. We re-estimated the aggregate state transition matrix based on this new history. Now, the period 1930-1939 is broken up into two depression episodes, one between 1930 and 1935 and another between 1937 and 1938. This alters the estimated aggregate transition matrix to

$$\hat{\Lambda} = \begin{bmatrix} 0.9766 & 0.0234 & 0 \\ (0.0053) & (0.0053) & (0) \\ 0.0715 & 0.9214 & 0.0071 \\ (0.0154) & (0.0161) & (0.0050) \\ 0 & 0.0208 & 0.9792 \\ (0) & (0.0146) & (0.0146) \end{bmatrix}.$$

Notice that since the duration of the depression state has fallen, λ_{dd} has fallen relative to the baseline. On the other hand, the two instances of transition to the depression state raise the conditional probability of entering a depression λ_{db} relative to the baseline, and the average unemployment rate in the depression-like state is now slightly higher as well (24.98 percent versus 23.48 percent in the baseline model). The welfare calculations reveal that these changes roughly offset each other, leading to a modest decline in the welfare gains from elimination of the depression-like state. The contribution of reduction in uncertainty

to this gain is about 61 percent.

This experiment also sheds light on the benefit of stabilization policies that cannot eliminate depressions but can accelerate recovery from them. In the above experiment, the average duration of the depression state is close to 4 years, but the depression state is about twice as likely to occur as in the baseline case. Given the linearity of the welfare estimates in the probability of encountering a depression, we may infer that if we reduce the probability of encountering a depression in the above experiment so that it occurs half as often, the welfare estimate from eliminating depressions would be cut in half. Hence, the welfare gain from eliminating a depression state that occurs once every 83 years but has an average duration of 4 years would be around $1.58/2 = 0.79$ percent. This implies that simply reducing the average duration of the depression state to 4 years without changing its frequency of occurrence (relative to the baseline model) would result in a total welfare gains of around 1.1 ($= 1.87 - 0.79$) percent. This is still much larger than the gains reported in Lucas (1987).

In the final experiment, we allow income to grow at a 2 percent annual rate. There is a modest increase in the welfare gains from elimination of the depression state relative to the baseline model and also a modest increase in the percentage contribution from the reduction in uncertainty. If we denote the monthly growth rate by g , the individual acts as if his purchases of assets are taxed at the rate g and he faces a discount factor $\beta(1+g)^{1-\gamma}$. Since $g > 0$ and $\gamma > 1$, both forces act to restrain purchases of the asset. Consequently, people end up being less protected from loss in earnings during a depression episode. Consistent with this intuition, we find that average asset holdings are about 27 percent lower in the economy with growth. On the other hand, the standard deviation of consumption is lower in the economy with growth, which seems to run counter to this intuition. However, the standard deviation of consumption is importantly affected by the consumption path of individuals who never lose their jobs during the depression (and these are the majority). Since these agents don't accumulate as many assets as in the baseline model their consumption

doesn't drop as much as in the baseline model, and this shows up in less overall volatility in aggregate consumption.

We also ran other experiments whose results are not reported because they made little difference to the estimate of the welfare gain. In one experiment, we varied the average duration of unemployment in the depression state. This had a very small effect on welfare because any lengthening or shortening of the average duration needed to be offset by a decrease or increase in the probability of entering unemployment (in the depression) in order to keep the average depression unemployment rate at 23.48 percent. These offsets in the incidence of unemployment cancelled the welfare effects of changes in unemployment duration. In another experiment, we let the return on the asset vary with the depression. In particular, we viewed the asset as money and assumed that at the onset of the depression the real value of money rose (because of a fall in the price level) while at the the end of the depression it fell (because the decline in the price level is reversed). This lowered the benefit accruing from decumulation of assets, since there is now a decline in the real value of assets as the economy emerges from a depression. However, the overall change in welfare was quite small. In the final experiment, we assumed that the asset had a 2 percent rate of return instead of zero, and this had a very small effect on welfare as well.

9 Conclusion

Our aim in this paper was to obtain an estimate of the benefit of stabilization policies that reduce the likelihood of a Depression-style collapse of economic activity. For the U.S., we estimate the probability of moving into a Depression-like state to be about once in every 83 years and the welfare gain from setting this small probability to zero ranges between 1.3 percent and 6.6 percent of annual consumption, in perpetuity. For our baseline calibration, the welfare gain is about 1.87 percent, with 58 percent of the gain coming from changes in second and higher-order moments of the consumption process, including a substantial decline in its variance. Higher estimates of the gain imply larger

contributions from the induced reduction in consumption volatility.

While we have quantified the potential gain from pursuing policies that reduce the likelihood of economic crises, we have not said anything about the potential costs of doing so. To take that step would require a theory of economic instability. This is a controversial issue, but one plausible theory locates the source of instability in the difficulties of coordinating trade.²⁷ One influential example of such a theory is Diamond's (1982) model of uncoordinated trade in which he showed that pessimism about the possibility of meeting trading partners can lead to self-fulfilling trade collapse. Another influential example is Diamond and Dybvig's (1983) theory of bank runs, in which pessimism about the likelihood of getting one's money back can lead to self-fulfilling banking panics. Both models suggest microeconomic interventions that can eliminate these undesirable outcomes, with deposit insurance in the Diamond-Dybvig model being a clear example. If these models are relevant for thinking about real-world economic crises, then it is the cost of microeconomic interventions such as deposit insurance that would have to be weighed against the benefits of eliminating the likelihood of economic crises.

References

- [1] Aiyagari, S. Rao. "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 1994 (pp. 659-684).
- [2] Alvarez, Fernando and Urban Jermann. "Using Asset Prices to Measure the Cost of Business Cycles," National Bureau of Economic Research, Working Paper No. 7978, October 2000.

²⁷Cole and Ohanian (1999) argue that a neoclassical model (i.e., one without coordination of trade problems) cannot account quantitatively for the severity and persistence of the Depression.

- [3] Atkeson, Andrew and Chris Phelan. "Reconsidering the Costs of Business Cycles with Incomplete Markets," in S. Fischer and J. Rotemberg eds., *NBER Macroeconomics Annual 1994*, (pp. 187-207).
- [4] Barlevy, Gadi. "Evaluating the Costs of Business Cycles in Models of Endogenous Growth," mimeo, Department of Economics, Northwestern University, April 2000.
- [5] Beaudry, Paul and Carmen Pages. "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance," manuscript, May 1999.
- [6] Bewley, Truman. "Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16, 1977 (pp. 252-292).
- [7] Cohen, Darrel. "Quantitative Defense of Stabilization Policy," Finance and Economics Discussion Series No. 2000-34, Federal Reserve Board, Washington D.C.
- [8] Cole, Harold L., and Lee Ohanian. "The Great Depression in the United States From a Neoclassical Perspective," *Quarterly Review*, Federal Reserve Bank of Minneapolis, Winter 1999, (pp. 2-21).
- [9] Danthine, Jean-Pierre and John B. Donaldson. "Non-Falsified Expectations and General Equilibrium Asset Pricing: The Power of the Peso," Working Paper No. PW-97-19, Columbia Business School, August 1998.
- [10] Darby, Michael R. "Three-and -a Half Million U.S. Employees Have Been Mislaid: Or, an Explanation of Unemployment, 1934-1941," *Journal of Political Economy*, 84, 1976 (pp. 1-16).
- [11] De Long, J. Bradford and Lawrence H. Summers. "How Does Macroeconomic Policy Affect Output," *Brookings Papers on Economic Activity*, 2, 1988 (pp. 433-480).
- [12] Diamond, Peter. "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90, 1982 (pp. 881-894).

- [13] Diamond, Douglas. and Philip Dybvig. “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy*, 91, 1983 (pp. 401-419).
- [14] Diaz, Antonia, Josep Pijoan-Mas and Jose-Victor Rios-Rull. “Habit Formation: Implications for the Wealth Distribution,” forthcoming, *Journal of Monetary Economics*.
- [15] Dolmas, Jim. “Risk Preferences and the Welfare Costs of Business Cycles,” *Review of Economic Dynamics*, 1, 1998 (pp. 646-676).
- [16] Gordon, Robert J. *The American Business Cycle: Continuity and Change*. Chicago: The University of Chicago Press (for NBER), 1986.
- [17] Greenwood, Jeremy, Richard Rogerson, and Randall Wright. “Household Production in Real Business Cycle Theory,” in Thomas F. Cooley ed., *Frontiers of Business Cycle Research*, Princeton: Princeton University Press, 1995, (pp. 157-174).
- [18] Hamilton, James D. “A New Approach to the Economic Analysis of Non-stationary Time Series and Business Cycles,” *Econometrica*, 57, 1989 (pp. 357-384).
- [19] Imrohorglu, Ayse. “Cost of Business Cycles with Indivisibilities and Liquidity Constraints,” *Journal of Political Economy*, 97, 1989 (pp. 1364-1383).
- [20] Jones, Larry E., Rodolfo Manuelli and Ennio Stacchetti. “Technology (and Policy) Shocks in Models of Endogenous Growth,” NBER Working Paper No. 7063, April 1999.
- [21] Keyssar, Alexander. *Men Out of Work: A Social History of Unemployment in Massachusetts*. Doctoral Thesis, Harvard University, 1977.
- [22] Krebs, Tom. “Growth and Welfare Effects of Business Cycles in Economies with Idiosyncratic Human Capital Risk,” forthcoming, *Review of Economic Dynamics*.

- [23] Krusell, Per and Tony A. Smith Jr. "On the Welfare Effects of Eliminating Business Cycles," *Review of Economic Dynamics*, 2, 1999 (pp. 245-272).
- [24] Lebergott, Stanley. *Manpower in Economic Growth: The American Record since 1800*. Economics Handbook Series. New York: McGraw Hill, 1964.
- [25] Lucas, Robert E. Jr. *Models of Business Cycles*. Oxford, UK: Basil Blackwell, 1987.
- [26] Matheron, Julien and Tristan-Pierre Maury. "The Welfare Cost of Fluctuations in AK Growth Models," mimeo, MODEM, Universite Paris X, June 2000.
- [27] Obstfeld, Maurice. "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability," *European Economic Review*, 38, 1994 (pp. 1471-1486).
- [28] Otrok, Christopher. "On Measuring the Welfare Cost of Business Cycles," *Journal of Monetary Economics*, 47, 2001 (pp. 61-92).
- [29] Pemberton, James. "Growth Trends, Cyclical Fluctuations, and Welfare with Non-expected Utility Preferences," *Economic Letters*, 50, 1996 (pp. 387-392).
- [30] Portier, Franck and Luis A. Puch. "The Welfare Cost of Fluctuations in Representative Agent Economies," mimeo, Universite de Toulouse, July 1999.
- [31] Ramey, Valerie and Gary Ramey. "Cross-Country Evidence on the Link Between Volatility and Growth," *American Economic Review*, 85, 1995 (pp. 1138-1151).
- [32] Rietz, Thomas A. "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22, 1988 (pp. 117-131).
- [33] Romer, Christina D. "Spurious Volatility in Historical Unemployment Data," *Journal of Political Economy*, 94, 1986 (pp. 1- 37).

- [34] Ross, Sheldon M. *Introduction to Probability Models*, New York: Academic Press, 1972.
- [35] Storesletten, Kjetil, Chris Telmer and Amir Yaron, "The Welfare Cost of Business Cycles Revisted: Finite Lives and Cyclical Variations in Idiosyncratic Risk," *European Economic Review*, 45, 2001, (pp. 1311-1339).
- [36] Tallarini, Thomas D., Jr. "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics*, 45, 2000 (pp.507-532).