WORKING PAPERS
RESEARCH DEPARTMENT

WORKING PAPER NO. 06-18

ON THE AGGREGATE WELFARE COST OF GREAT DEPRESSION UNEMPLOYMENT

Satyajit Chatterjee
Federal Reserve Bank of Philadelphia
and
Dean Corbae
University of Texas, Austin

September 2006
This paper supersedes Working Paper 03-20.
Revision forthcoming in the *Journal of Monetary Economics*

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA
Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/econ/index.html

WORKING PAPER NO. 06-18

On the Aggregate Welfare Cost of Great Depression Unemployment

Satyajit Chatterjee
Federal Reserve Bank of Philadelphia
Dean Corbae
University of Texas, Austin

September 2006

Forthcoming: *Journal of Monetary Economics*. We are very grateful to Hide Ichimura, Bob King and Narayana Kocherlakota for very valuable comments on this project. 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, 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

The potential benefit of policies that eliminate a small likelihood of economic crises is calculated. An economic crisis is defined as an increase in unemployment of the magnitude observed during the Great Depression. For the U.S., the maximum-likelihood estimate of entering a depression is found to be about once every 83 years. The welfare gain from setting this small probability to zero can range between 1 and 7 percent of annual consumption in perpetuity. For most estimates, more than half of these large gains result from a reduction in individual consumption volatility.

Keywords: Depression, Unemployment, Welfare Cost
JEL: D52, E32, E63

1 Introduction

Lucas (1987) argued that the welfare gains from post-WWII stabilization policies were something on the order of one-tenth of 1 percent of annual U.S. consumption.[1] This paper takes a different approach to gauging the welfare gains from stabilization policies. Instead of focusing on the gains from a reduction in aggregate variability, it estimates the gains from the prevention of a Depression-style collapse of economic activity. By focusing on elimination of economic crises, the paper seeks to incorporate an important goal of real-world stabilization policies into welfare discussions. In the U.S. and elsewhere, institutions such as state insurance of bank deposits, unemployment insurance, and discretionary monetary and fiscal policies were a response to the Great Depression. An important goal of these institutions is to truncate the lower tail of the (probability) distribution of individual-level consumption processes as opposed to a mean-preserving truncation of both the upper and lower tails. Thus, the policies have both mean and higher-order effects. A more complete discussion of the welfare gains of stabilization policies therefore requires some quantification of the gains from an elimination of economic crises rather than simply the gains from a reduction in aggregate variability.[2]

Figure 1 further motivates this paper. The figure plots annual unemployment rates for the period 1900 to 1998. The striking aspect of these time series is the extraordinary rise in unemployment between the years 1930 and 1939, generally identified in history as the Depression years.[3] The rise is extraordinary not only because it has not been repeated but

[1] Most welfare measures constructed along the lines of Lucas suggest that the overall cost of aggregate variability is relatively small. On this point, see Otrok (2001).

[2] DeLong and Summers (1988) have previously argued for sizable welfare gains of stabilization policies on the grounds that such policies "fill in business cycle troughs without shaving off business cycle peaks" and thereby reduce the average unemployment rate.

[3] For the period 1900-1940, the Lebergott series for industrial unemployment was constructed by dividing the total number of unemployed workers reported in Lebergott (1964) Table A-3 by the sum of unemployed workers and nonfarm workers also reported in that table. 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

also because there is no corresponding 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. It is difficult to view Figure 1 and not think of the Depression as a potentially preventable lapse from the normal workings of the economy.

The question posed in this paper 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 the question, we study a model similar to Imrohoroglu's (1989). There is a unit measure of infinitely lived workers with identical preferences. Workers face idiosyncratic employment opportunities, and the probability of finding employment depends on the aggregate state of the economy. One difference is that the data generation process of our model includes an aggregate state that corresponds to an economic crisis in which the probability of finding employment in the business sector is much lower relative to the other aggregate states. We also consider alternative insurance arrangements concerning these adverse shocks to employment status. In our baseline model, workers cannot buy insurance against adverse shocks to their employment status in the depression state but can attempt to self-insure by holding stocks of an asset with a real rate of return lower than the rate of time preference of workers.[4] As a point of comparison we also consider other insurance arrangements such as complete markets.

We start with an estimate of the current likelihood of depressions − the likelihood that is set to zero in the welfare experiments. This estimate is obtained 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. In fitting one Markov chain to the entire period, the paper assumes that the fact that no economic crisis has occurred since Depression-era stabilization policies went into place reflects luck rather than design. Under

Romer (1986) to the industrial unemployment rate series.

[4]This assumption on the real return is consistent with the general equilibrium implications of imperfect insurance, as shown, for instance, by Aiyagari (1994).

this assumption, the current likelihood of moving into a depression is estimated to be once every 1000 months (or once every 83 years). With complete markets, the welfare gains from setting this small probability to zero is 0.97 percent of annual consumption in perpetuity. Of this total gain, 0.78 percentage points results from an increase in mean consumption and 0.19 percentage points from a reduction in uncertainty (or volatility) in consumption. For our baseline model with incomplete markets, though, the contribution of lower uncertainty is much higher – 1.09 – leading to a total gain of 1.87 percent. The contribution of lower uncertainty is about 5.5 times larger ($= 1.09 \div 0.19$) under incomplete markets because households engage in a lot of precautionary savings (self-insurance) during a Depression episode – savings that are dissipated when the economy emerges from a Depression. The resulting large swing in individual consumption is the main reason why consumption is considerably more volatile with even a small probability of Depression-like events.

Sensitivity analysis indicates that the gains from elimination of economic crises could be as low as 1.3 percent or as high as 6.6 percent. Higher estimates are associated with a larger contribution of reduction in uncertainty. For the experiment that generated the 6.6 percent welfare gain, 80 percent of that gain can be attributed to a reduction in consumption volatility.

2 Environment

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

$$\Lambda = \begin{bmatrix} \lambda_{gg} & \lambda_{bg} & \lambda_{dg} \\ \lambda_{gb} & \lambda_{bb} & \lambda_{db} \\ \lambda_{gd} & \lambda_{bd} & \lambda_{dd} \end{bmatrix} \tag{1}$$

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

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

$$E\sum_{t=0}^{\infty}\beta^{t}\frac{c_t^{1-\gamma}}{1-\gamma} \tag{2}$$

where c_t consumption in period t, $\beta \in (0,1)$ is the discount factor, and $\gamma > 0$ is the relative risk aversion parameter.

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 also follows a Markov process whose transition matrix is given by:

$$\Lambda^{\eta} = \begin{bmatrix} \lambda_{ee}^{\eta} & \lambda_{ue}^{\eta} \\ \lambda_{eu}^{\eta} & \lambda_{uu}^{\eta} \end{bmatrix} \tag{3}$$

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 that the individual was unemployed in period t.

Thus 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 $\omega^1, \ldots, \omega^6$, where ω^1 stands for "employed in a good state," ω^2 for "unemployed in a good state," ω^3 for "employed in a bad state," ω^4 for "unemployed in a bad state," ω^5 for "employed in a depression state," and ω^6 for "unemployed in a depression state." The process governing ω is also a first-order Markov process with transition matrix given by $\Phi = \left[\phi_{jk}\right]$, where $\Pr\{\omega_{t+1} = \omega^j \mid \omega_t = \omega^k\} = \phi_{jk}$. These 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 $\lambda_{gg}\lambda^g_{ue}$.

Environments with varying levels of insurance against adverse shocks to employment status are studied in this paper. Denote the post-insurance income of a worker in state ω in period t by $x(\omega_t)$. Given the potential incompleteness of insurance markets, workers have an incentive to self-insure, and it is assumed they can do so by holding stocks of an asset with zero real return (alternative asset market assumptions are investigated later in the paper). Individuals enter period t with individual savings s_t held over from the previous period. An individual's budget constraint can be written:

$$c_t + s_{t+1} = x(\omega_t) + s_t, \; s_t \geq 0 \tag{4}$$

where c_t is the individual's consumption in period t.

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:

$$V(s, \omega) = \max_{s' \geq 0} \frac{c_t^{1-\gamma}}{1-\gamma} + \beta \sum_{\omega'} \Phi(\omega', \omega) V(s', \omega') \tag{5}$$

subject to

$$c = x(\omega) + s - s' \geq 0. \tag{6}$$

The solution to this problem yields decision rules $c(s, \omega)$ and $g(s, w)$ for c and s', respectively.

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:

$$\mu_{t+1}(s',\omega') = \sum_{\omega} \sum_{s \in g^{-1}(s',\omega)} \Phi(\omega',\omega)\mu_t(s,\omega) \tag{7}$$

where $g^{-1}(s',\omega) = \{s : s' = g(s,\omega)\}$. 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 Mapping the Model to the Data

Estimation of the aggregate state-transition matrix Λ is presented first, followed by the calibration of other parameters of the model.

3.1 Estimation of the Aggregate State Transition Matrix

The starting point of the estimation is a construction of the history of the aggregate states. This is accomplished by taking the monthly NBER business-cycle chronology since January 1900 and associating NBER expansions with the good state and NBER contractions with the bad state. This two-state history is then augmented by a third crisis state with a very high incidence of unemployment. Specifically, all months of any year in which the unemployment rate exceeded 17 percent were re-classified as depression months. This definition picks out the 120 months corresponding to the 1930-1939 period – generally known as the Depression years.[5] 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.

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

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

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).[6] Implementing this procedure for the whole sample yields the estimate, $\widehat{\Lambda}$, of the aggregate state transistion matrix reported in Table 1 (standard errors appear in parentheses in the table).[7]

The estimated matrix has several noteworthy features. First, because there is only one depression episode in the sample, there is only one transition into and one transition out of the depression state. Also, 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 about 0.001.[8] But the *unconditional* probability of a depression is 0.0975, which 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

[6] The 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:

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

[7] The reported standard errors 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 $\widehat{\Lambda}$. 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 $\widehat{\Lambda}$ for these biases only led to *higher* welfare gains of eliminating the depression-like state, we retained the more conservative estimates of $\widehat{\Lambda}$ reported in the paper.

[8] This number was obtained by multiplying the unconditional probability of being in the bad state, which is 0.2074, by 0.0039.

though the probability of encountering a Depression-like event, conditional on not being in one, is quite small.

3.2 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, and the earnings-loss parameter θ.

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 assembled by Lebergott (1964). 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.

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 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. The average duration of unemployment spells in the depression state was fixed at 20 months, more than four times the median duration seen in 1931.

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

$$U_t = U_{t-1}\lambda^{\eta(t)}_{uu} + (1 - U_{t-1})\lambda^{\eta(t)}_{ue} \tag{8}$$

where $\eta(t) \in \{g, b, d\}$. Since λ^{η}_{uu} 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:

$$U^{\eta} = U^{\eta}\lambda^{\eta}_{uu} + (1 - U^{\eta})\lambda^{\eta}_{ue}. \tag{9}$$

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

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 $1 - 1/(40 * 12)$, 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.[9]

4 Welfare Measures

We wish to estimate the aggregate welfare gain from moving to an environment in which the $\widehat{\Lambda}$ matrix is replaced by

$$\Lambda^* = \begin{bmatrix} 0.9766 & 0.0234 \\ 0.0745 & 0.9216 + 0.0039 \end{bmatrix}. \tag{10}$$

The off-diagonal elements of this matrix are identical to the corresponding elements of $\widehat{\Lambda}$, 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. In the first type of welfare calculation, which is our preferred type, it is assumed 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 begins in that state. But if it is in the depression state, then

[9] Darby (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.

it begins in the bad state as well. Under these assumptions, the fraction of consumption an individual would be willing to give up if he is currently in state (s, ω) is $1 - \alpha(s, \omega)$, where $\alpha(s, \omega)$, $\omega \leq 4$, solves $V(s, \omega) = \alpha(s, \omega)^{1-\gamma} V^*(s, \omega)$, and $\alpha(s, 5)$ and $\alpha(s, 6)$ solve $V(5, \omega) = \alpha(5, \omega)^{1-\gamma} V^*(s, 3)$ and $V(6, \omega) = \alpha(6, \omega)^{1-\gamma} V^*(s, 4)$, respectively. The aggregate welfare gain is given by $1 - \bar{\alpha} = 1 - \sum_s \sum_\omega \widehat{\mu}(s, \omega)\alpha(s, \omega)$, where $\widehat{\mu}(s, \omega)$ is the invariant measure for the three-state depression-prone environment. In the second type of calculation it is assumed that each individual is offered the *average* lifetime utility of the depression-proof environment. In this case $\alpha^{SS}(s, \omega)$ solves $V(s, \omega) = \alpha^{SS}(s, \omega)^{1-\gamma} \bar{V}^*$ for all s, ω where $\bar{V}^*$ is $\sum_s \sum_\omega \mu^*(s, \omega) V^*(s, \omega)$ with $\mu^*(s, \omega)$ being the invariant distribution of the two-state depression-proof economy. Then the aggregate welfare gain is given by $1 - \bar{\alpha}^{SS} = 1 - \sum_s \sum_\omega \widehat{\mu}(s, \omega)\alpha^{SS}(s, \omega)$. This latter measure is referred to as the *steady-state* gain in welfare and has been used before by Imrohoroglu. The difference between the preferred measure and the steady-state measure is that the former takes account of the transition path to the new depression-proof steady state.

5 Welfare Gains Under Alternative Insurance Arrangements

In what follows, the two welfare measures are calculated for environments with alternative market arrangements. The results for all these experiments are collected in Table 2.

5.1 Welfare Gain With Complete Markets

As an important point of comparison, the welfare gain is calculated first for the case where 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 $\eta = g, b, d$. Since this leaves only aggregate uncertainty to play a role, this effectively corresponds to the representative agent case.[10]

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

The top entry in Table 2 reports the operating characteristics of the full insurance economy in the depression-prone and depression-free regimes. In both regimes, eliminating uninsured employment risk eliminates the need to save. The uncertainty in post-transfer earnings in both regimes is too low relative to the cost of saving (the difference in the rate of return on savings (zero) and the rate of discount (6 percent, annualized)) to motivate people to save. Eliminating the depression state raises mean consumption from 0.9628 to 0.9704 and reduces the standard deviation of consumption from 0.027 to 0.005. The gain in aggregate welfare according to our preferred measure (labeled TG for total gain) is 0.97 percent and according to the steady state measure (labeled SS) it is 0.98. The steady state measure is higher because for currently unemployed people, the unconditional probability of being unemployed is significantly lower than the conditional probability of being unemployed. Only the former matters in the calculation of the steady state measure.

The final column (labeled RU) gives an estimate of the welfare gain that is due to the reduction in uncertainty. Observe that the increase in mean consumption (0.0076 units) is 0.78 percent ($= 0.0076 \div 0.9704 \times 100$) of mean consumption in the depression-free regime. Since people will pay exactly 0.78 percent of mean consumption to obtain this increase, we may infer that approximately 19.6 percent ($= (0.97 - 0.78) \div 0.97 \times 100$) of total gain in welfare is due to reduction in volatility.[11]

5.2 Welfare Gain With Partial Insurance (Baseline)

The middle entry in Table 2 reports the results for the economy in which there is full insurance against unemployment in the good and bad states but there is no insurance against unemployment in the depression state; i.e., the post-transfer earnings of workers are $(1-U^{\eta})+\theta U^{\eta}$

[11] It's worth noting that our representative agent estimate of 0.2 ($= (0.98 - 0.78)$) percent for the steady state welfare gain resulting purely from a reduction in volatility is very close to Lucas's (representative agent) estimate of the welfare gain from elimination of pre-WWII cyclical volatility. As noted in Lucas (1987) (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.

for $\eta = g, b$ but 1 and θ for employed and unemployed workers, respectively, in the depression. This baseline environment is motivated by the thought that a depression will adversely affect asset markets as well. In particular, consider a world with two kinds of assets, one of which is issued by the business sector and the other 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 situation, 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 workers will accumulate enough of the business-sector asset to almost perfectly self-insure against unemployment risk during normal times. A rough way of capturing this situation is to simply assume that both employed and unemployed workers receive the per capita endowment in the good and bad state (so there is no risk of earnings loss due to unemployment in these times) but confront workers with the risk of earnings loss due to unemployment in depressions.

The total welfare gain is 1.87 percent of consumption per month (or per year) and the steady state gain is 1.70 percent. 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 the steady state gain measure. Thus 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.[12]

We find that the welfare gains from eliminating depressions vary approximately linearly with the likelihood of encountering a depression. If the true likelihood of encountering a

[12]The 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.

Depression-like event was actually once every 1600 years (rather than once every 83 years), that would cut the 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 an 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. 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.

These large gains stem from three sources. First, average asset holdings go from being 0.41 of monthly earnings in the depression-prone regime to 0 in the depression-free regime. Because the difference between the total gain in welfare and the steady state gain is 0.17 percent, it follows that at least 9.1 ($= (1.87 - 1.7) \div 1.87 \times 100$) percent of the total welfare gain is due to the fact that individuals need to hold fewer assets in the new regime.[13] Second, average consumption rises by 0.0076 units (same as in the complete markets case) in the depression-free regime. Thus, the increase in mean consumption 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 depression-free regime. 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 single 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

[13] As noted earlier, the difference between the total and steady state gains also reflects the fact that the unconditional probability of unemployment is lower than the conditional probability of unemployment for someone who is currently unemployed. This fact tends to *raise* the steady-state measure above the total measure. Therefore, the welfare gain attributable to asset decumulation is actually somewhat greater than noted in the text.

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 regimes is in the volatility of individual-level consumption. In the depression-free regime, 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. In the depression-prone regime, 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 change dramatically. 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 of both unemployed *and* employed people contribute to the relatively high volatility of individual consumption in the depression-prone economy.[14]

Figures 1 and 2 show the simulated paths of unemployment and per capita consumption implied by our model.[15] In Figure 2, the simulated path shown by the dashed line assumes that all of the non-business-sector income received by an unemployed individual is included in measured GDP. Under this assumption, per capita consumption drops about 16 percent

[14] These changes in individual and aggregate consumption occur even though a depression is assumed not to affect the earnings of employed and unemployed agents.

[15] The 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. In both figures, the lightly shaded bars correspond to the bad state and the darker bars correspond to the depression state.

but gradually recovers to a decline of about 10-12 percent by the end of the Depression. The recovery occurs because the rate of asset accumulation by workers begins to decline as workers get closer to their target buffer stock of assets (of about 8 months of employed income). Another factor that contributes to the recovery is that workers who become unemployed later in the Depression experience less of a decline in consumption because they have higher levels of precautionary savings. Figure 2 also shows that when the model economy emerges from the Depression, per capita consumption rises sharply above trend as workers reduce their precautionary savings. The dotted line displays a similar pattern but the pattern is more pronounced because, in this case, 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.[16]

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 (2003) 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 three-state economy.

5.3 Welfare Gain With No Insurance

The next to last row in Table 2 presents the results for the case where there is no unemployment insurance in any aggregate state. 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

[16] If it is assumed that 23 percent of non-business-sector income goes unmeasured, the cumulative loss in consumption between 1930-1945 in the model matches that in the data.

the contribution of the three different channels, the reduction in precautionary savings 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 depression-free 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.

The reason gains are somewhat lower than in the partial insurance case 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 partial insurance (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. 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.

5.4 Welfare Gain Under Autarky

The last row in Table 2 presents the results for the case where individuals have no insurance against unemployment and cannot save. The fact that workers cannot save reduces their life-time utility in both the depression-prone and depresssion-free regimes – witness the large increase in consumption volatility in the two regimes. It is interesting to compare this case to the previous case where there is savings but no insurance. The gains are larger in comparison – indicating that the opportunity to save lowers the benefit from elimination of economic crises. In a sense, self-insurance is a substitute for stabilization policies.

6 Sensitivity Analysis

In this section we report the sensitivity of our results to changes in key parameter values. The results are collected in Table 3. 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 (TG) from eliminating the depression state is now 6.59 percent while the steady state (SS) gain is 6 percent. With regard to the sources 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 – labeled RU), which is 80.4 percent of the total gain in welfare.

The next experiment reduces the relative risk aversion parameter to 1.5. 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 40 percent.

In the third experiment, depression months were redefined to be all months of any year in which the unemployment rate exceeded 20 percent. Now, the period 1930-1939 is broken up into two depression episodes: one between 1930 and 1935 and another between 1937 and 1938. This has the effect of lowering the persistence of a depression state but raises the conditional probability of encountering a depression. 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.

In the fourth experiment, income is permitted 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.

In the fifth experiment, the average duration of unemployment in the depression state was lowered to 10 months. Surprisingly, this had a relatively small effect on the welfare gain. The reason is that any shortening of the average duration needs to be offset by an increase in the probability of entering unemployment in the depression so as to keep the average depression unemployment rate at 23.48 percent. The offset in the incidence of unemployment roughly cancelled the welfare effects of changes in unemployment duration.

Finally, two other experiments (not reported in the table) were conducted. In one experiment the return on the asset was allowed to vary with the occurrence of a depression. In particular, the asset was viewed as money, and it was 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 end of the depression it fell (because the decline in the price level is reversed). Again, the overall change in welfare gains was small. In the final experiment, the asset was assumed to have a 2 percent rate of return instead of zero, but this too had a very small effect on welfare.

7 Conclusion

Many macroeconomists believe that unemployment during the Great Depression was socially very costly. Yet, welfare cost measures of unemployment constructed along the lines of Lucas (1987) suggest that the overall cost of business cycles, the Great Depression included, is relatively small. This paper approached the welfare-cost debate from a different angle and sought to obtain an estimate of the cost of low-probability Depression-like episodes. It was found that these costs are quite large. For the U.S., the probability of moving into a Depression-like state was estimated to be about once in every 83 years and the welfare gain from setting this small probability to zero ranged between 1.3 to 6.6 percent of annual consumption, in perpetuity. For the baseline calibration, the welfare gain was estimated to be 1.87 percent, with 58 percent of this 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 this paper quantifies the potential gain from pursuing policies that reduce the likelihood of economic crises, it does not say 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. 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, 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. A further example where coordination problems are associated with a financial collapse that spills over to the real side of the economy is present in Cooper and Corbae (2002). In their environment, adding liquidity to the banking system is a form of stabilization policy that can overcome strategic uncertainty and avoid financial collapse. In this case, the cost of eliminating crises is negligible.

References

Aiyagari, S. R., 1994. Uninsured idiosyncratic risk and aggregate saving. Quarterly Journal of Economics 109, 659-684.

Bewley, T., 1977. Permanent income hypothesis: A theoretical formulation. Journal of Economic Theory 16, 252-292.

Cole, H. L., Ohanian, L., 1999. The Great Depression in the United States from a neoclassical perspective. Quarterly Review, Federal Reserve Bank of Minneapolis Winter, 2-21.

Cooper, R.W., Corbae, D., 2002. Financial collapse: A lesson from the Great Depression. Journal of Economic Theory 107, 159-190.

Darby, M. R., 1976. Three-and-a half million U.S. employees have been mislaid: Or, an explanation of unemployment, 1934-1941. Journal of Political Economy 84, 1-16.

DeLong, J. B., Summers, L.H., 1988. How does macroeconomic policy affect output? Brookings Papers on Economic Activity 2*, 433-480.*

Diamond, P., 1982. Aggregate demand management in search equilibrium. Journal of Political Economy 90, 881-894.

Diamond, D., Dybvig, P., 1983. Bank runs, deposit insurance, and liquidity. Journal of Political Economy 91, 401-419.

Gordon, R. J., 1986. The American business cycle: Continuity and change. The University of Chicago Press, Chicago.

Greenwood, J., Rogerson, R., Wright, R. 1995. Household production in real business cycle theory. In: Cooley, F., (Ed.), Frontiers of Business Cycle Research. Princeton University Press, Princeton, pp. 157-174.

Imrohoroglu, A., 1989. Cost of business cycles with indivisibilities and liquidity constraints. Journal of Political Economy 97, 1364-1383.

Keyssar, A., 1977. Men out of work: A social history of unemployment in Massachusetts. Doctoral Thesis, Harvard University.

Krebs, T., 2003. Growth and welfare effects of business cycles in economies with idiosyncratic human capital risk. Review of Economic Dynamics 6, 846-868.

Lebergott, S. 1964. Manpower in economic growth: The American record since 1800. McGraw Hill, New York.

Lucas, R. E. Jr., 1987. Models of Business Cycles. Basil Blackwell, Oxford.

Otrok, C., 2001. On measuring the welfare cost of business cycles. Journal of Monetary Economics 47, 61-92.

Romer, C. D., 1986. Spurious volatility in historical unemployment data. Journal of Political Economy 94, 1- 37.

Ross, S.M., 1972. Introduction to Probability Models. Academic Press, New York.

Storesletten, K., Telmer,C., Yaron, A., 2001. The welfare cost of business cycles revisted: Finite lives and cyclical variations in idiosyncratic risk. European Economic Review 45, 1311-1339.

Table 1

	Expansion	Recession	Depression
Expansion	0.9766	0.0234	0
	(0.0053)	(0.0053)	(0)
Recession	0.0745	0.9216	0.0039
	(0.0164)	(0.0168)	(0.0039)
Depression	0.0083	0	0.9917
	(0.0083)	(0)	(0.0083)

Table 2

	3-State Regime			2-State Regime			Welfare Gain in %		
	$\bar{s}$	$\bar{c}$	$\sigma(c)$	$\bar{s}$	$\bar{c}$	$\sigma(c)$	TG	SS	RU
Full Ins.	0.00	0.9628	0.027	0.00	0.9704	0.005	0.97	0.98	19.4
Partial Ins.	0.41	0.9628	0.056	0.00	0.9704	0.005	1.87	1.70	58.3
No Ins.	2.23	0.9628	0.072	1.86	0.9704	0.054	1.56	1.40	50.0
Autarky	0.00	0.9628	0.017	0.00	0.9704	0.014	1.79	1.82	58.1

Table 3

	3-State Regime			2-State Regime			Welfare Gain in %		
	$\bar{s}$	$\bar{c}$	$\sigma(c)$	$\bar{s}$	$\bar{c}$	$\sigma(c)$	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
$\gamma = 1.5$	0.17	0.96	0.06	0	0.97	0.01	1.30	1.22	40.0
$\overline{U} \geq 0.20$	0.23	0.96	0.05	0	0.97	0.01	1.58	1.48	60.8
2 % Growth	0.25	0.96	0.06	0	0.97	0.01	1.96	1.80	59.9
Duration=10	0.37	0.96	0.05	0	0.97	0.01	1.69	1.57	53.8

Figure 1: Actual and Simulated Unemployment

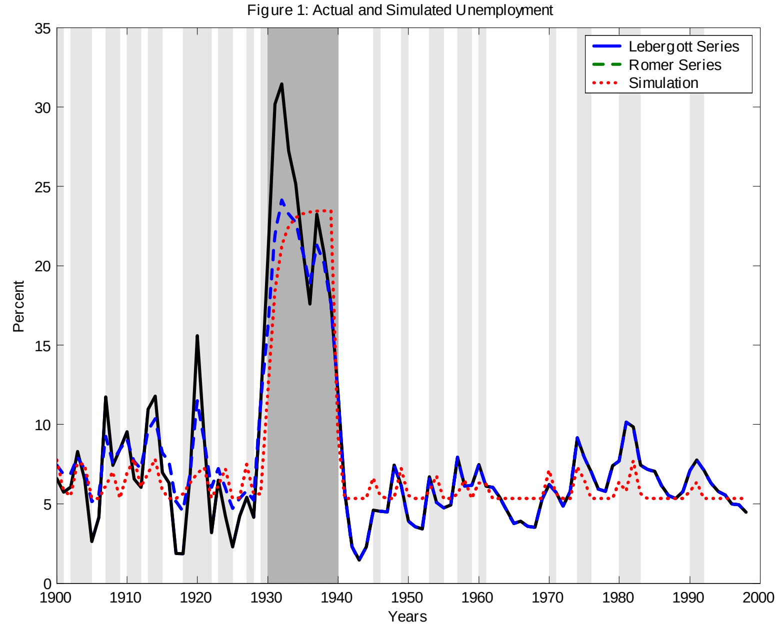

Figure 2: Actual and Simulated Per-Capita Consumption