A Quantitative Analysis of Unemployment Insurance in a Model with Fraud and Moral Hazard

David L. Fuller[*]

February 21, 2011

Abstract

In this paper I analyze the provision of unemployment insurance in an environment with unobservable job offers and unobservable employment status. Both rejection of suitable offers and unreported employment income represent unemployment insurance fraud. I analyze US data characterizing the prevalence of both types of fraud, and calibrate a model to match this data. I find novel implications from including fraud from unobservable employment status. For small increases in the benefit level (10%) the model is consistent with micro evidence on duration elasticities; however, larger increases in the benefit level *decrease* the unemployment rate and durations. Similarly, for a range of increases in the potential duration of benefits, the average duration of unemployment decreases. I calculate that actual occurrences of unemployment insurance fraud amount to 10% of total benefits paid and reduce welfare by around 1%. I also find the economy is better off relying on minimal welfare payments instead of the current U.S. system of unemployment benefits.

Keywords: unemployment insurance, fraud, moral hazard, hidden income

JEL classification: C61, D82, E61, J64, J65

---
[*]Department of Economics, Concordia University; dfuller@alcor.concordia.ca

1 Introduction

The provision of unemployment insurance is affected by many incentive problems. Moral hazard represents the most classic example of the inherent trade-off between insurance and incentives. Given unobservable search effort and job offers, the presence of insurance has perverse effects on an agent's incentives to search for employment, or to accept a job offer if one arrives. While moral hazard has dominated the literature on unemployment insurance, there exist other incentive problems. For example, when the unemployment insurance provider cannot perfectly observe an agent's employment status, an agent can continue collecting unemployment benefits after accepting a job offer. Both of the aforementioned incentive problems can result in *unemployment insurance fraud*, whereby an agent collects benefits she is not entitled to. In this paper, I quantitatively examine a model of unemployment insurance where both types of fraud occur. Using data on the U.S. economy and the unemployment insurance system, I first determine the prevalence of each type of fraud, and then examine the effects of changes in benefit levels and potential benefit durations.

According to data measuring the accuracy of paid unemployment benefits in the U.S., both informational frictions exist. The U.S. Department of Labor has a program called BAM, or Benefit Accuracy Measurement. It collects data through individual investigations and interviews on a sample of U.I. recipients, and determines whether the individual collected the appropriate amount of benefits. While moral hazard has received the most attention, the data indicate unobservable employment status may represent the more quantitatively relevant friction. In 2008, for example, approximately $8.4 million in benefits were paid to agents who received a suitable job offer, but did not accept it and continued collecting benefits. These are benefits overpaid due to the moral hazard friction, hereafter referred to as Type I fraud. The same year, $700 million in benefits were paid to agents who accepted jobs but continued to collect benefits; i.e. overpayment resulting from unobservable employment status, hereafter Type II fraud. Although these represent cases of detected fraud rather than actual occurrences, these data at least suggest that one cannot ignore the unobservable employment status friction.

To explain the occurrence of both types of fraud, I develop a parsimonious model similar to Hansen and Imrohoroglu (1992), extended to allow for unobservable employment status. Specifically, agents have preferences over consumption and leisure. Labor is supplied inelastically, so employed agents spend a fixed amount of time working. Agents face a stochastic employment pro-

cess, where they can accept or reject a job offer if one arrives. While employed, the job dissolves with some exogenous probability, so the model permits transitions into and out of employment. To partially self-insure against employment risk, agents have access to a simple storage technology that pays zero interest. If engaged in either type of fraud, I assume that the unemployment insurance provider has a verification technology, so that with some probability (unique to each type of fraud) an agent may be caught committing fraud.

Including the possibility of Type II fraud represents the key innovation in this paper, compared to the model in Hansen and Imrohoroglu (1992). In addition, there also exist key differences in the verification technologies used. Hansen and Imrohoroglu (1992) assume that when an agent receives a job offer and rejects, there exists a probability the agent does not collect benefits. After this initial "verification," an agent can continue to collect benefits fraudulently with no consequence. Moreover, outside of not being able to collect benefits, Hansen and Imrohoroglu (1992) assume there is no penalty for being caught committing Type I fraud. Thus, although not calculated explicitly in their analysis, the model necessarily overestimates the occurrences of Type I fraud. In contrast, I my model, for both Type I and Type II fraud, there exists a probability of being caught in *any* given period. If caught, the agent forfeits their rights to collect unemployment benefits and must pay a monetary fine. This assumption matches the basic structure of current U.S. unemployment fraud enforcement.

To calibrate the model, I interpret the data on fraud, which can be done in one of two ways. First, the data could be taken as indicating the *actual* occurrences of Type I and II fraud, whereby anyone committing fraud is detected in the data. Alternatively, one can interpret the data as only representing *detected* fraud, and the possibility that undetected cases remains. I adopt the second interpretation, and by doing so, the analysis determines the actual occurrences of both types of fraud implied by the data. Since the probability of being caught is not observed, I infer it from the model's predictions. I also calibrate the model to match the unemployment rate and average unemployment duration over the period from $2003 - 2006$. Finally, I also use the BAM data to calculate the average replacement rate and average potential duration of benefits for the agents in the sample.

Given these parameters, I find that Type I fraud occurs more frequently than Type II fraud. For Type I fraud, over the $2003 - 2006$ period, on average the detected fraud rate is 0.05%, where fraud

rates are reported as the fraction of total benefits collected fraudulently. The model predicts an actual Type I fraud rate of 9.16%. In dollar terms, from $2003-2006$, total benefits averaged \$34.225 billion a year; thus Type I fraud amounts to a yearly average of \$3.13 billion. For Type II fraud, the model predicts actual and detected fraud remain almost identical at 1.42%, or approximately \$486 million a year. Related, the results also present an intuitive finding: there exists a relatively effective verification technology for detecting Type II fraud, while Type I fraud remains difficult to detect. After calibrating the model, I then perform several policy experiments to determine the impact of each type of fraud.

In the first set of policy experiments, I consider the effects of increasing unemployment benefits on the unemployment rate and average unemployment duration. The existing literature on unemployment insurance, both theoretical and empirical, makes a clear prediction for these effects. Theoretically, search models along the lines of Pissarides (2000), or models of optimal unemployment insurance such as Hopenhayn and Nicolini (1997), all imply a monotonically increasing relationship between the average duration of unemployment (and/or the unemployment rate) and the level of benefits.[1] Empirically, this result is confirmed in general equilibrium studies such as Hansen and Imrohoroglu (1992) or Wang and Williamson (2002). There also exist micro studies that estimate duration elasticities with respect to the benefit level. For example Meyer (1990) finds a duration elasticity with respect to the benefit level of 0.8. Readers interested in a survey of this literature should consult Krueger and Meyer (2002) who provide a summary of similar studies. These empirical studies confirm the theoretical relationship between benefits and unemployment durations.

The policy experiments in this paper indicate, that in general, this monotonic relationship does not hold when both types of fraud are accounted for. For small increases in the benefit level, the model remains consistent with the elasticities calculated in Meyer (1990). Results from the calibration imply a 10% increase in the replacement rate raises the average unemployment duration by 1.5 weeks, an elasticity of 0.806. For a range of larger increases in the replacement rate however, the average duration of unemployment *decreases* from 4.65 weeks to 4.19 weeks, and the unemployment rate decreases from 5.34% to 4.94%. Specifically, the replacement rate can go from the current level of 45% to as high as 60% and the unemployment rate remains lower than

---

[1]The theoretical and empirical literature studying the provision of unemployment insurance is extensive. There exist many related articles I have not mentioned here.

the baseline level.

This result is due primarily to two factors. First, for low levels of accumulated assets, Type II fraud tends to dominate both Type I fraud, and accepting a job but not committing fraud. Second, there exists a distributional effect associated with the increasing level of benefits, whereby a larger fraction of agents have fewer accumulated assets. This occurs because the higher level of benefits reduces the precautionary savings of households. Thus, while increasing benefits increases the value of Type I fraud faster than the value of Type II fraud, the distributional effect moves a larger fraction of agents into the range of assets where Type II fraud still dominates the alternatives. Since agents substitute away from Type I fraud to Type II, more agents accept job offers and the unemployment rate and average unemployment duration both decrease. Eventually, Type I fraud dominates Type II for all levels of assets, and the unemployment rate begins to increase as benefits increase further.

I also explore the effects of changes in the potential duration of benefits. According to my estimates from the BAM data, the current U.S. system provides benefits for an average of 24 weeks, or given the model's time period of 1 month, approximately 6 months. The results indicate that increasing the potential duration of benefits from 6 months to as high as 8 months decreases the unemployment rate and durations. As in the case of the replacement rate, this result is driven by a precautionary savings effect. The effects of unemployment benefits on precautionary savings of households has been examined by the aforementioned studies of Hansen and Imrohoroglu (1992) and Wang and Williamson (2002). Micro studies such as Engen and Gruber (2001) have also found strong responses of precautionary savings to changes in unemployment benefits.

Finally, I also consider the welfare implications of unemployment insurance fraud. To make this comparison, I compute average welfare in the baseline economy, where both types of fraud occur, and in a hypothetical economy with perfect information. In the baseline parametrization, fraud has a welfare cost of around 1%. The policy experiments analyzing changes in the replacement rate and potential benefit duration all produce small changes in welfare (less than 0.1%). Along this dimension, eliminating unemployment benefits completely actually produces the largest increase in welfare of approximately 0.4%. While somewhat misleading, as agents in the model still receive minimum welfare payments (calibrated to match U.S. food stamp payments), the result suggests that given the incentive problems present, the current U.S. unemployment insurance system is too

generous.

Alvarez-Parra and Sanchez (2009) also examine a model with unobservable employment status. These authors' analysis differs from mine in several ways. First, Alvarez-Parra and Sanchez (2009) analyze the optimal contract using methodology similar to Hopenhayn and Nicolini (1997). In contrast, the focus of this paper is explaining observed outcomes in a general equilibrium framework. The source of unobservable employment status represents another important difference. Alvarez-Parra and Sanchez (2009), much like the model in Hopenhayn and Nicolini (2008) (although there employment status remains observable), assume there exists a hidden labor market, where agents *always* have the opportunity to be employed. Employment in the "informal" sector is associated with a lower productivity than employment in the "formal" sector. This friction has different implications from the Type II fraud studied in this paper. Specifically, in Alvarez-Parra and Sanchez (2009) the provision of incentives is designed to prevent agents from accepting the offer from the informal sector. This is distinct from incentives to prevent Type II fraud which require agents to accept an offer and to report doing so.

The remainder of the paper proceeds as follows. Section 2 describes the model and equilibrium. In Section 3 I describe the BAM data, and Section 4 calibrates the model. Section 5 presents the results of policy experiments, and Section 6 concludes.

2   Model

2.1   Preferences and Environment

There exists a unit mass of infinitely-lived agents. Time, t, is discrete. Agents have preferences over consumption and leisure given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t, h_t) \right]$$

where c_t denotes consumption in period t, and h_t denotes the number of hours worked. I assume agents are endowed with one unit of indivisible time with which to supply labor. Moreover, I assume labor is supplied inelastically, so that if an agent is employed $h_t = \bar{h}$, and if unemployed $h_t = 0$.

Employment status (alternatively income) remains stochastic and persistent. Denote an agent's

employment status by j, where $j = e$ refers to employment, and $j = u$ to unemployment. With probability π_j an agent in employment state $j \in \{e, u\}$ is employed at the end of the period, where $\pi_e > \pi_u$. For an unemployed agent, the interpretation is a job offer arrives each period with probability π_u. For an employed agent, π_e represents an exogenous probability the job dissolves.

Given their employment status, agents make consumption and savings decisions. I assume that there exists a zero interest storage technology where agents can store consumption goods from period t for use in period $t+1$. If employed, an agent produces y units of the consumption good and receives this output as a wage. When unemployed, agents collect benefits, b, which are modeled to capture features of the U.S. unemployment insurance system. Specifically, benefits are a constant fraction of the wage, $b = \theta y$, and have a potential duration of T periods. Benefits during this initial unemployment duration are considered "unemployment benefits." After T periods, benefits drop to a lower level, $b = dy$, where $d < \theta$. This lower level represents a more general welfare program (such as food stamps in the U.S.), and are not considered unemployment benefits.

An agent collecting unemployment benefits potentially has the opportunity to commit Type I and/or Type II fraud. If a job offer arrives, and the agent decides to reject the offer and remain unemployed, Type I fraud occurs and the agent enters state $j = uf$. In this case, the agent continues to collect unemployment benefits. There exists some probability, p_1, that the agent is found to be engaged in Type I fraud by the unemployment insurance agency. When caught, the agent forfeits her rights to continue collecting benefits, and must pay a monetary fine, f_1, from assets.

When an offer arrives and the agent decides to accept and become employed, she then has the opportunity to report unemployment and continue collecting benefits, which is referred to as Type II fraud, and denoted by the state $j = ef$. Analogous to Type I fraud, there exists some probability, p_2, the unemployment insurance agency detects the fraud. Again, if caught, the agent pays a monetary fine, f_2, from assets and forfeits rights to the remaining unemployment benefits. Notice, in the case of Type II fraud, this implies that if such an agent subsequently transitions to unemployment, she cannot collect unemployment benefits $b = \theta y$. This highlights an interesting incentive feature built into the existing U.S. unemployment insurance system. Specifically, an agent has additional incentives to accept a job offer simply to upgrade unemployment benefits.[2] In this

---

[2]This feature becomes potentially more interesting when agents can quit jobs. In this case, it may be possible for an agent to accept a job offer to upgrade unemployment benefits, and then quit to collect. Hopenhayn and Nicolini (2008) examine this issue in a model of optimal unemployment insurance.

model, this feature helps deter both types of fraud.

2.2 Value Functions

The agents' problem can be written recursively in the following manner. Let c_j denote current consumption, k accumulated "assets," and k'_j the savings decision for an agent in state $j = \{e, ef, u, uf\}$. For unemployed agents, and those agents deciding to commit fraud, their current employment status j, the current level of accumulated "assets" k, and the number of periods of benefits remaining, x, represent the state variables. For example, if an agent has been unemployed for 2 periods, $x = T - 2$. Finally, unemployment benefits are financed by a lump sum tax, denoted τ. When $x = 0$, $\tau = 0$, so only employed agents and unemployed agents collecting benefits are taxed.[3] Given this, the Bellman equations can be written as follows.

$$V_u(k, x) = \max_{c_u, k'_u} \mathrm{v}(c_u, 0) + \beta \left[\pi_u \max \left\{ V_e(k'_u), V_{ef}(k'_u, x - 1), V_{uf}(k'_u, x - 1) \right\} + (1 - \pi_u) V_u(k'_u, x - 1) \right]$$

(1)

$$\text{s.t.} \quad c_u + k'_u \leq \theta y + k - \tau \tag{2}$$

$$k'_u \geq 0 \tag{3}$$

where for $x = 0$, unemployment benefits have expired, so $b = dy$, and the problem becomes

$$V_u(k, 0) = \max_{c_u, k'_u} \mathrm{v}(c_u, 0) + \beta \left[\pi_u V_e(k'_u) + (1 - \pi_u) V_u(k'_u, 0) \right] \tag{4}$$

$$\text{s.t.} \quad c_u + k'_u \leq k + dy \tag{5}$$

$$k'_u \geq 0 \tag{6}$$

If an agent decides to reject a job offer and engage in Type I fraud, she enters state $j = uf$ and solves

---

[3] Many studies using similar models impose a lump-sum tax on only employed agents; however, since 1987 unemployment benefits have been considered taxable income. I have also calculated the baseline model when only employed agents pay taxes and the results are unaffected.

$$V_{uf}(k,x) = \max_{c_{uf},k'_{uf}} \mathrm{v}(c_{uf},0) + \beta\{p_1\left[\pi_u V_e(k'_{uf}-f_1)+(1-\pi_u)V_u(k'_{uf}-f_1,0)\right]+$$

$$(1-p_1)\left[\pi_u \max\left\{V_e(k'_{uf}),V_{ef}(k'_{uf},x-1),V_{uf}(k'_{uf},x-1)\right\}+(1-\pi_u)V_{uf}(k'_{uf},x-1)\right]\}$$

$$\tag{7}$$

$$s.t. \qquad c_{uf}+k'_{uf} \leq b+k-\tau \tag{8}$$

$$k'_{uf} \geq 0 \tag{9}$$

When an employed agent decides not to commit Type II fraud, or when benefits have expired, so the option to commit Type II fraud does not exist, her problem is

$$V_e(k) = \max_{c_e,k'_e} \mathrm{v}(c_e,\bar{h}) + \beta\left[\pi_e V_e(k'_e)+(1-\pi_e)V_u(k'_e,T)\right] \tag{10}$$

$$s.t. \qquad c_e+k'_e \leq y+k-\tau \tag{11}$$

$$k'_e \geq 0 \tag{12}$$

If an agent, with benefits remaining for $x \geq 1$ periods, transitions from unemployment to employment, she decides whether or not to commit Type II fraud. When engaged in Type II fraud the problem becomes,

$$V_{ef}(k,x) = \max_{c_{ef},k'_{ef}} \mathrm{v}(c_{ef},\bar{h}) + \beta\{p_2\left[\pi_e V_e(k'_{ef}-f_2)+(1-\pi_e)V_u(k'_{ef}-f_2,0)\right]$$

$$+ (1-p_2)\left[\pi_e V_{ef}(k'_{ef},x-1)+(1-\pi_e)V_u(k'_{ef},x-1)\right]\} \tag{13}$$

$$s.t. \qquad c_{ef}+k'_{ef} \leq y+b+k-\tau \tag{14}$$

$$k'_{ef} \geq 0 \tag{15}$$

where $V_{ef}(k,0) = V_e(k)$, since benefits have expired.

The decision to commit fraud (of either type) depends primarily on two factors. First, how quickly an agent finds employment. Given the limited duration of benefits, the gain from committing fraud evaporates as the unemployment duration increases. An agent who finds a job after only one period of unemployment has many weeks of benefits left to collect, making her more likely to continue reporting unemployment after finding a job. Similarly, when an offer arrives early in

the unemployment spell, rejecting it is more likely as the agent has more periods with benefits remaining to receive another offer. The decision to commit fraud also depends on an agent's asset level at the time she receives an offer.

In the case of Type I fraud, agents with higher levels of accumulated assets are more likely to reject a job offer if one arrives. This obtains because these agents can use their assets to smooth consumption while unemployed, so for large enough k, rejecting the offer may dominate accepting. For Type II fraud, the opposite is true; agents with low levels of accumulated assets are more likely to continue collecting benefits after accepting a job. The additional income is useful for these agents given their small assets holdings, and they have little to lose being caught engaged in Type II fraud, since they have relatively few assets to pay any penalties. For these agents, the primary penalty is not monetary, but rather the loss of eligibility for unemployment benefits during their next spell of unemployment. Of course these two effects, the remaining potential length of benefits and the level of accumulated assets, interact with each other, since the level of assets at the time of transition depends on the duration of the unemployment spell, as well as the entire history of employment.

Figures 1 and 2 display these features for Type I and Type II fraud, respectively. Denote by $k_i^\star, i \in \{uf, ef\}$, the cutoff level of assets where $V_i(k_i^\star, x) \geq V_j(k_i^\star, x), \forall j \in \{e, ef, u, uf\}$. For both types of fraud, as x decreases, the benefits of fraud decrease (value functions shift downwards), and for Type I fraud $k_i^\star$ increases, while for Type II it decreases. Moreover, for a given x, $k_{uf}^\star > k_{ef}^\star$; i.e. Type I fraud occurs for higher accumulated savings relative to Type II fraud.

2.3 Equilibrium

I now define an equilibrium in this economy. Let $\mathbf{S}$ denote the set of possible employment/benefit states. This includes an agent's employment status, $j \in \{e, ef, u, uf\}$, the current benefits collected, and the number of periods of eligibility remaining. Denote the time t employment/benefit state as $s \in \mathbf{S}$. Also denote by $g(k, s)$, the policy function for k' solving the agent's Bellman equations given above. Then, given a period t distribution of agents across asset and employment/benefit states, $\lambda_t(k, s)$, the policy functions $g(k, s)$ solving the above Bellman equations induce a mapping $\Gamma : \mathbf{S} \times \mathbb{R}^+ \to \mathbf{S} \times \mathbb{R}^+$. The distribution of agents across states evolves according to

$$\lambda_{t+1}(k, s) = \Gamma \lambda_t(k, s)$$

Figure 1: Type I Fraud Decision

Figure 2: Type II Fraud Decision

11

The focus here remains on steady states, so I solve for the stationary distribution $\lambda(k, s)$. Finally, let $B[\lambda(k, s)]$ denote the fraction of agents collecting unemployment benefits, and $E[\lambda(k, s)]$ the fraction of employed agents (including $j = e, ef$), given the distribution $\lambda(k, s)$. An equilibrium in this environment is defined as:

Definition 1 : *A **stationary equilibrium** is given by a policy function, $g(k, s)$, and a distribution, $\lambda(k, s)$, such that*

1. *$g(k, s)$ solves the agent's problem*

2. *$\lambda(k, s) = \Gamma\lambda(k, s)$*

3. *$B[\lambda(k, s)] b = (B[\lambda(k, s)] + E[\lambda(k, s)]) \tau$*

The first condition ensures that agents behave optimally, the second that we have a stationary distribution, and the last represents the balanced budget condition.

3   Data

This section describes the U.S. data on both Type I and Type II fraud. The U.S. Department of Labor publishes a yearly report detailing the accuracy of paid U.I. claims, a program referred to as BAM: Benefit Accuracy Measurement. To determine the accuracy of paid benefits, the BAM program chooses a weekly random sample of U.I. claims, and investigators audit these claims to determine their accuracy. According to the BAM State Operations Handbook ET No. 495, 4th Edition, the goal of the program, is in general, different from the goal of UI fraud investigators. While the fraud investigators look to recapture overpayments, BAM investigators are instead trying to calculate statistics on the UI program in general. Such investigations in the BAM data indicate cases of both over-payments (the fraud studied in this paper) as well as under-payments. Overpayment data includes both cases of Type I and Type II fraud, which is available for the years $1988 - 2009$. Table 1 displays data from the 2008 BAM report.

As Table 1 indicates, the "rates" of Type I and Type II fraud are calculated as a percentage of total benefits. Specifically, these rates represent the dollar value of benefits overpaid to individuals committing Type I (Type II) fraud, divided by the total benefits paid. According to the data, Type II fraud remains more prevalent (the other years display a similar pattern). In the next section, I

Table 1: Fraud Overpayments by Cause, 2008

Cause	Fraction of Total Benefits paid	$ Amount
Total Benefits Paid, U.S.	100 %	41,614,449,463
Unreported Earnings	1.68 %	698,608,274
Refused Suitable Offer	0.02 %	8,395,827

describe my interpretation of this data and how the model described in Section 2 captures these features of the data.

In the calibration, I use data from the time period $2003 - 2006$. Given the frequent changes in U.I. laws across states, using longer time periods can prove problematic as the data likely include significant changes to the U.I. system in certain states, but not in others. Data on Type I and Type II fraud are taken directly from the BAM yearly reports for these years, which are available at http://www.oui.doleta.gov/unemploy/bqc.asp; the reports also include a more detailed description of the BAM program scope and methodology.

I also use the individual level data provided by the BAM program to calculate other important parameters of interest. Each year includes around $24,000$ observations, and I am examining a total of $98,301$ observations for the $2003 - 2006$ time period. Specifically, I use information on earnings and weekly benefit amounts to calculate the replacement rate as well as the average potential duration of unemployment benefits. There exist many idiosyncracies among states regarding unemployment insurance laws, regulations, and benefit calculations; as a result, difficulties arise determining the actual replacement rate for a given state, let alone for the overall U.S. economy. In addition to the variability across states, a particular state often has complicated rules for calculating benefits. For example, a state may indicate a basic replacement rate of 50%. However, there typically exists a maximum benefit amount that binds for higher wage earners reducing the replacement rate, and there are deductions for dependents, etc. that increase the replacement rate.

Given these difficulties, I calculate the replacement rate for each agent in the sample directly by dividing the Weekly Benefit Amount (WBA) by average weekly earnings. There exist several variables in the data set relevant for this task. They include "Base Period Earnings (BPE)," "High Quarter Earnings (HQE)," "Weeks Worked in Base Period (WWBP)," "Maximum Benefit Amount

(MBA)," and "WBA." Ideally, weekly earnings would be calculated as BPE divided by WWBP; unfortunately, many states do not record data on WWBP. Moreover, the Base Period differs state by state, so even calculating average weekly earnings over the base period (i.e. BPE divided by the number of total weeks in the base period) remains problematic. Thus, to calculate weekly earnings, I divide HQE by 13. The replacement rate is then calculated as the WBA divided by the estimate of weekly earnings.

Three states, MA, NJ, and OH do not record HQE.[4] NJ and OH do record WWBP, so weekly earnings for these two states is calculated as BPE divided by WWBP. MA does not record WWBP, so given the data there does not exist a reliable method for calculating weekly earnings given BPE. According to the states' website for the Executive Office for Labor and Workforce Development (http://www.mass.gov/?pageID=elwdhomepage&L=1&L0=Home&sid=Elwd), the state uses a replacement rate of 50%. Although not completely accurate given the aforementioned issues, I use this as the estimate for MA.

To calculate the average benefit duration in each state, I divide the MBA by the WBA. The MBA is the total dollar value of benefits an individual is entitled to. Table 11 in Appendix A displays the results for replacement rates and benefit durations for all states, the District of Columbia (DC) and Porta Rico (PR). For the U.S. overall, the average benefit duration is 24 weeks and the average replacement rate is 0.45. These are similar to the commonly used 26 week duration and 0.50 replacement rate.

4   Calibration

Given the model described above, I now calibrate it to U.S. data. The idea of the calibration is to choose a set of parameters, calculate optimal decision rules for agents, and then simulate the model. The simulation produces a steady state distribution of agents across asset level and employment states. Given these distributions, the model's predictions can be compared to U.S. data, and parameters chosen to reconcile these two.

---

[4]These states do not use high quarter earnings in their benefit calculations as many other states do; as a result, this data is not recorded. This also explains why WWBP is unavailable for many states.

4.1 Parameters

Table 2 lists the parameters to be determined in the calibration exercise. As in Hansen and Imrohoroglu (1992), I assume the agents' per-period utility function takes the form,

$$v(c_t, l_t) = \frac{[c_t^{1-\rho} l_t^{\rho}]^{1-\sigma} - 1}{1 - \sigma}$$

where $l_t = 1 - \bar{h}$ represents leisure, with $\bar{h}$ defining the fraction of time spent working. Following Hansen and Imrohoroglu (1992), I set $\rho = 0.67$, and $\bar{h} = 0.45$. To discipline σ, I compare the results for different values in the range $[0.5, 2.0]$. For the baseline case, I choose the value of σ that produces duration elasticities (with respect to the benefit level) closest to those reported in micro studies. Meyer (1990) finds a duration elasticity with respect to the benefit level of 0.8, which is the value I target. There exist many other studies calculating duration elasticities and Krueger and Meyer (2002) provide a summary of the relevant papers. They note that most micro studies find duration elasticities with respect to the benefit level between 0.2 and 1.0.

The discount factor is given by $\beta = \frac{1}{1+r}$, where r represents the risk-free interest rate. The time period in the model is one month, so $r = 0.04$ per-annum implies $\beta = 0.996$. The wage is set to $y = 0.5$, and the benefit $\theta = 0.45$, consistent with the average benefit in the U.S. calculated in Section 3. From the same calculations, the average benefit duration of 24 weeks implies $T = 6$. Once unemployment benefits expire, I set $d = 0.1$, consistent with welfare payments in the U.S.

Given the variability of state unemployment laws and enforcement, determining the penalties for being caught committing either Type I or Type II fraud remains difficult. In the baseline parametrization, I set $f_1 = f_2 = k'_j$. Thus, if caught committing either type of fraud, the agent forfeits her assets. While this is an ad hoc assumption, it does not differ significantly from a simple penalty where agents repay fraudulently collected benefits. Moreover, an agent's decision on whether or not to commit fraud depends on the relative sizes of p_1 and p_2 to f_1 and f_2, respectively. In this regard, I have simply normalized the penalties, and then choose p_1 and p_2 to match observed fraud rates.

With the aforementioned parameters determined, the parameters p_1, p_2, π_e, π_u must be chosen. In calibrating these parameters, I attempt to match the following moments in the data for the $2003 - 2006$ time period.

1. Unemployment rate

Table 2: Parameters

β	Discount factor
$\bar{h}$	Time spent working
p_1	Probability of being caught in Type I fraud
f_1	Penalty if caught committing Type I fraud
p_2	Probability of being caught in Type II fraud
f_2	Penalty if caught committing Type II fraud
π_u	Probability of receiving a job offer
π_e	Probability of separation from current employer
σ	Coefficient of relative risk aversion
ρ	Relative weight of leisure vs. consumption
y	Wage if employed
θ	Benefit if unemployed and eligible
T	Length of unemployment benefits
d	Benefit after T periods

2. Average unemployment duration

3. Type I fraud rate (detected)

4. Type II fraud rate (detected)

Table 3 displays the calibration across values of σ, along with the predicted duration elasticity with respect to the benefit level. This is calculated using the percent change in the average unemployment duration when benefits increase by 10% from 0.45. The model best matches duration elasticities for $\sigma = 1.5$; therefore, this is the baseline value used. Table 4 displays the parameter values for this baseline parametrization.

According to Table 4, there exists a relatively large probability of being caught committing Type II fraud (0.676), and a relatively small probability of being caught committing Type I fraud (0.048). This represents an intuitive finding, which could be interpreted several ways. First, p_j could represent the probability an agent is verified or audited by the unemployment insurance

Table 3: Parameter Values

Parameters	$\sigma = 0.5$	$\sigma = 1.0$	$\sigma = 1.5$	$\sigma = 2.0$
p_1	0.0872	0.0635	0.048	0.03699
p_2	0.81	0.739	0.676	0.594
π_u	0.186	0.187	0.19	0.19
π_e	0.99	0.99	0.99	0.99002
Elasticity	1.17	1.03	0.806	-0.408

agency, and if verified while engaged in fraud, the agent is caught for certain. Alteratively, it could be that agents are verified every period, but may not always be caught. Finally, it could be a combination of these two. It remains impossible to determine which of these scenarios describes reality, since at best, we only observe the probability of verification and the number of agents caught. The true number committing fraud remains unknown, however, and the true probability of apprehension undetermined. Regardless of which case does obtain, the calibration indicates a good verification technology for Type II fraud, but not for Type I. In practice, for Type II fraud, verification is often accomplished by a cross-referencing of employment records filed by employers with unemployment benefit records, as well as tips followed by fraud investigators. In the case of Type I fraud, there is no technology equivalent to cross-referencing records, so detection relies primarily on fraud investigators.

With this particular interpretation of the data and calibration strategy, however, the calibration does not necessarily remain identified. For any given observed fraud rate, there may exist an equilibrium with a high probability of being caught, and another with a low probability of being caught. In the baseline calibration, I always choose the highest detection probability that matches the data. For Type II fraud, the alternative is a very low detection probability which given modern technology for cross-referencing employer information with unemployment benefit registrations, seems implausible. For Type I fraud, the probabilities of being caught in the "high probability" case are already quite low. The case with very low probabilities of being caught implies almost all agents who receive a job offer reject it, and the majority of benefits paid are fraudulent. Again this possibility appears implausible.

Table 4: Parameter Values

β	0.9967
$\bar{h}$	0.45
p_1	0.048
p_2	0.676
π_u	0.19
π_e	0.99
σ	1.5
ρ	0.67
y	0.5
θ	0.5
T	6
d	0.1

4.2 Results

The first three columns Table 5 display the moments from the data, along with those predicted by the baseline model. It shows both the detected cases of fraud, and the actual occurrences. Thus, the *actual* occurrences of Type I fraud (9.16% of total benefits) remain more common than Type II fraud (1.42%), opposite what the data indicate. In dollar terms, on average over the period $2003-2006$, \$34.225 billion of benefits, per year, were paid; therefore, each year actual occurrences of Type I fraud averaged \$3.14 billion. Similarly, actual occurrences of Type II fraud accounted for \$486 million. The last three columns of Table 5 show the model's predictions for different values of σ.

Notice, detected fraud is not simply the probability of being caught multiplied by the actual occurrences of fraud. This occurs because the fraud rates reported in Table 5 are calculated as the percentage of total benefits collected fraudulently. Given this, several factors determine how detected fraud and actual fraud compare relative to each other. First, it depends when the fraud opportunities arrive (i.e. π_u), and then on how long the agent collects benefits fraudulently for. The latter is determined by the length of benefits (T) and the probability of apprehension. Thus,

Table 5: Calibration Results

Moment	Data	Model	$\sigma = 0.5$	$\sigma = 1$	$\sigma = 2$
Unemployment Rate (%)	5.3	5.3	5.3	5.3	5.3
Average Unemployment Duration	4.63	4.65	4.64	4.65	4.63
Type I Fraud (%)	0.06	0.05	0.086	0.017	0.017
Type I Fraud Actual (%)	-	9.16	5.54	6.61	8.55
Type II Fraud (%)	1.44	1.42	1.46	1.46	1.45
Type II Fraud Actual (%)	-	1.42	1.46	1.46	1.45
Duration Elasticity	0.80	0.806	1.17	1.034	-0.408

there is no simple calculation for determining detected fraud as a fraction of actual fraud. Similarly, this explains why actual and detected Type II fraud remain the same. As the fraction of benefits due to Type II fraud increase (in policy experiments for example), the difference between detected and actual Type II fraud increases.

5 Policy Experiments

This section considers two counterfactuals: what happens to the outcomes predicted by the model when (i) the replacement rate (θ) changes and (ii) the potential benefit duration (T) changes. I focus in particular on the effects of these changes on the unemployment rate and average unemployment duration.

5.1 Replacement Rate

In the first set of policy experiments I consider the effects of increasing the replacement rate. In a standard moral hazard model of unemployment insurance, the unemployment rate and average duration of unemployment increase monotonically with the benefit level. The results of the first policy experiment, displayed in Table 6, indicate this relationship no longer holds in a model with both types of fraud (the baseline case, $\theta = 0.45$, is in bold).

Initially the unemployment rate and duration increase as the benefit increases, then decrease to

Table 6: Changing the Replacement Rate

Benefit (θ)	U.R.	Duration	Type I Actual	Type II Actual
0.11	4.94	4.19	0.00	0.07
0.2	4.94	4.19	0.00	0.20
0.3	4.94	4.19	0.00	0.33
0.405	4.94	4.19	0.00	0.87
0.45	**5.34**	**4.65**	**9.16**	**1.42**
0.495	5.66	5.03	15.12	3.05
0.55	4.94	4.19	0.00	20.13
0.6	4.97	4.22	0.55	19.90
0.65	5.6	4.95	12.14	14.76
0.75	6.79	6.31	28.82	9.14
0.85	8.20	8.07	45.54	1.59
0.95	8.40	8.35	48.09	0.49

the lowest level permitted by π_u and π_e, and finally begin increasing again. Increases in θ between 0.495 and 0.6 actually decrease the unemployment rate and average unemployment duration from the baseline level. These results represent a novel implication of accounting for Type II fraud, as standard models with only Type I fraud *always* imply a monotonic relationship between θ and these two moments.

These decreases occur because over a certain range of increases in θ, agents substitute Type II fraud for Type I fraud. The last two columns of Table 6 confirm this; for example, when benefits increase from 0.495 to 0.55, Type I fraud decreases while Type II fraud increases. This happens for two reasons. First, for low levels of accumulated savings, $V_{ef}(k, x) > V_{ur}(k, x)$; the value of Type II fraud exceeds Type I fraud. Second, there exists a distributional effect moving the majority of agents into this asset range. As θ increases, precautionary savings decrease, and the fraction of agents in the range of savings where Type II dominates Type I fraud increases; as a result, they prefer accepting the job but reporting unemployment. Eventually this effect dies out and the unemployment rate begins increasing with θ. This results from the strictly concave per-period

Table 7: Economy with no Type II Fraud

Benefit (θ)	U.R.	Duration	Type I Actual
0.11	4.84	4.10	0.00
0.405	4.84	4.10	0.00
0.45	**5.29**	**4.62**	**8.20**
0.495	5.88	5.31	22.82
0.55	6.31	5.79	29.76
0.6	6.82	6.39	36.62
0.65	7.13	6.79	39.85
0.75	7.44	7.17	43.52
0.85	8.02	7.83	48.62
0.95	8.22	8.13	50.24

utility function, which implies that for any given x, $V_{ur}(k,x)$ increases faster than $V_{ef}(k,x)$ as θ increases, implying it eventually dominates at all levels of savings.

Figure 3 displays how the value functions $V_e(k), V_{ef}(k,5)$, and $V_{ur}(k,5)$ change as benefits increase from $\theta = 0.45$ to $\theta = 0.55$. As expected, the shift in V_{ur} is larger than the shift in V_{ef}, although the latter dominates for lower levels of k. Then, Figure 4 plots the distibution of assets for $\theta = 0.45$ and $\theta = 0.55$. Notice how the distribution shifts to the left, and that most agents now lie in the range where V_{ef} dominates V_{ur}.

For comparison, Table 7 displays the results from the same policy experiment in an economy where only Type I fraud exists (i.e. employment status is observable). In this economy, the calibration implies the following parameters (different from the baseline calibration above): $\pi_e = 0.99002, \pi_u = 0.192$, and $p_1 = 0.0096$. Notice that the unemployment rate and duration increase monotonically with the replacement rate. While increases in θ from 0.495 to 0.6 decrease the unemployment rate and duration in the economy with Type II fraud, they both increase in this range for the economy with only Type I fraud. For large increases in θ (0.85 and above), the unemployment rate and duration is actually lower in this economy compared to the economy with both types of fraud. This obtains because the calibrated economy with only Type I fraud has a

Figure 3: Value functions

Figure 4: Asset Distributions, $\sigma = 1.5$

Table 8: Effects of changing potential benefit duration

Benefit Duration	U.R.	Avg. Dur.	Type I	Type II
1	4.94	4.19	0.00	0.00
2	4.94	4.19	0.00	0.00
3	4.94	4.19	0.00	0.00
4	4.94	4.19	0.00	0.09
5	4.94	4.19	0.00	0.46
6	**5.34**	**4.65**	**9.16**	**1.42**
7	4.94	4.19	0.00	20.59
8	4.94	4.19	0.00	21.01
9	5.49	4.83	9.69	19.26
10	5.67	5.04	12.07	18.57
11	6.59	6.11	23.74	15.92

higher probability of a job offer; as a result, the "maximum" unemployment rate in this economy is lower.

5.2 Potential Benefit Duration

Increasing the duration of benefits (T) represents the other policy experiment I perform. In this case, similar results obtain, and increasing benefit durations actually decreases the unemployment rate and average unemployment duration. Table 8 displays the results for the baseline case of changing the benefit duration (the baseline case, $T = 6$, is in bold). Similarly to the case of increasing θ, the decrease in average unemployment durations is explained by a substitution of Type II fraud for Type I fraud, as the average level of precautionary savings decreases. The last two columns of Table 8 display this feature. For comparison, Table 9 shows the effects of changing the potential benefit duration in the economy with only Type I fraud. Here, as expected, there exists a monotonically increasing relationship with the unemployment rate and duration.

Several micro studies have also examined the effects of increased benefit durations on the average unemployment duration. For example, Katz and Meyer (1990) find a duration elasticity with respect

Table 9: Economy with no Type II Fraud

Benefit Duration	U.R.	Avg. Dur.	Type I
1	4.84	4.10	0.00
5	4.94	4.10	0.00
6	**5.29**	**4.62**	**8.20**
7	5.52	4.90	15.83
8	6.17	5.63	27.92
9	6.67	6.22	33.30
10	6.96	6.54	36.83
11	7.52	7.25	43.27

to the potential duration of benefits of about 0.5. In this sense, the baseline model is inconsistent with these micro studies. This may be interpreted several ways. First, an increase in T by one period corresponds to an increase in the potential benefit duration of 4 weeks. It remains possible that the model does predict an increase in the unemployment rate and durations for a 1 week increase in the potential benefit duration, but we miss this effect given the 1 month time period. An alternative possibility is that the micro estimates are biased and therefore incorrect. Since these estimates do not control for or account for Type II fraud, or allow for changes in agent behavior (such as precautionary savings), it could be the case that the estimates are incorrect. Finally, the inconsistency could stem from the baseline model in this paper overestimating the effect of potential benefit durations on precautionary savings.

Comparing Tables 6 and 8 offers insights into the most effective policies for deterring both types of fraud. In Table 6, although Type I fraud abates quickly as the replacement rate decreases, Type II fraud remains more persistent. For even low levels of the replacement rate ($\theta = 0.11$) Type II fraud still occurs. Table 8 displays that decreases in the potential benefit duration have a stronger effect on the two fraud rates. Again Type I fraud disappears faster than Type II, however in this case Type II fraud also goes to zero. These results imply policy makers interested solely in reducing fraud should reduce potential benefit durations instead of decreasing benefits.

Table 10: Welfare Comparisons

Economy	Unemployment Rate	Avg. Duration	Welfare Change (%)
Baseline	5.34	4.65	–
Perfect Info	4.94	4.19	0.877
No U.I.	4.94	4.19	0.337

5.3 Welfare

The previous policy experiments focused on how changes in the replacement rate and potential duration of benefits affect the unemployment rate, unemployment durations, and Type I and II fraud rates. In this section I analyze the effects of Type I and II fraud on average welfare, and I also consider the welfare implications of changes to the current U.S. unemployment insurance system.

In Section 4.2 I find that actual Type I and II fraud amount to around 10% of total benefits paid, or approximately $4 billion. To determine the welfare implications of fraud, I perform the following counterfactual experiment. I assume there is perfect information, so when an offer arrives, an agent must accept and cannot continue collecting benefits; therefore, neither Type I or Type II fraud occurs. This exercise takes as given the structural parameters from the calibrated economy. I then compare the average welfare of an agent in each economy. Table 10 displays the results from this counterfactual.

In the perfect information economy, since neither type of fraud occurs, the unemployment rate and average unemployment duration are at the lowest levels permitted by π_e and π_u. The results imply that Type I and II fraud have a welfare cost of nearly 1%. The final row of Table 10 displays another counterfactual experiment. Here I consider the economy with no unemployment benefits, $\theta = 0$. Again, since there are no benefits to collect, both types of fraud remain nonexistent. Surprisingly, the economy with no unemployment benefits represents a welfare improvement over the baseline economy. Comparatively, the policy experiments in Section 5 all produce changes in welfare of less than 0.1%, and the economy with no unemployment benefits dominates all of them. This result is somewhat misleading, however, since agents still collect minimum welfare payments, d, and the model does not consider the financing of these benefits. It remains possible that there exists an optimal combination of replacement rate and potential benefit duration that dominates

the economy without unemployment benefits. The results do suggest, however, that current U.S. unemployment benefits are too generous.

6   Conclusion

In this paper I develop a parsimonious model to explore unemployment insurance fraud resulting from rejection of suitable job offers (Type I) and unreported employment income (Type II). I find that although the data show detected Type II fraud is more common than Type I, the detection probability for Type I fraud remains low enough that actual occurrences of Type I fraud dominate Type II fraud. Both types of fraud amount to over 10% of total benefits paid.

The analysis finds novel effects from accounting for Type II fraud. While the model remains consistent with micro estimates of duration elasticities for small changes in the benefit level, larger changes actually decrease the unemployment rate and average duration of unemployment. This effect occurs through a change in precautionary savings that causes agents to substitute away from Type I fraud into Type II fraud. Similarly, I find that a range of increases in the potential duration of benefits also reduces the unemployment rate and durations. In terms of welfare, I find that current levels of fraud have a welfare cost of just under 1%, and an economy without unemployment benefits dominates the current system.

There exist several interesting directions for future research. First, while I perform some basic welfare experiments, I do not consider the fully optimal scheme. One direction for future research is to methodology from the dynamic mechanism design literature to determine the optimal scheme and the welfare benefits of adopting it. Similarly, the calibration implies there exists a relatively effective verification technology for detecting Type II fraud, but not for Type I fraud. Exploring the optimal scheme in an environment with verification for Type II and unobservable job offers for Type I fraud represents another interesting direction for future research.

References

Alvarez-Parra, F., Sanchez, J. M., 2009. Unemployment insurance with a hidden labor market. Journal of Monetary Economics 56, 954–967.

Engen, E. M., Gruber, J., 2001. Unemployment insurance and precautionary saving. Journal of Monetary Economics 47, 545–579.

Hansen, G., Imrohoroglu, A., 1992. The role of unemployment insurance in an economy with liquidity constraints and moral hazard. Journal of Political Economy 100, 118–142.

Hopenhayn, H., Nicolini, J., 1997. Optimal unemployment insurance. Journal of Political Economy 105, 412–438.

Hopenhayn, H., Nicolini, J., 2008. Optimal unemployment insurance and employment history. Forthcoming, Review of Economic Studies.

Katz, L. F., Meyer, B. D., 1990. The impact of the potential duration of unemployment benefits on the duration of unemployment. Journal of Public Economics 41, 45–72.

Krueger, A. B., Meyer, B. D., 2002. Labor supply effects of social insurance. Handbook of Public Economics 4, 2327–2392.

Meyer, B., 1990. Unemployment insurance and unemployment spells. Econometrica 58, 757–782.

Pissarides, C. A., 2000. Equilibrium Unemployment Theory. MIT Press, Cambridge, MA.

Wang, C., Williamson, S., 2002. Moral hazard, optimal unemployment insurance, and experience rating. Journal of Monetary Economics 49, 1337–1371.

A   Appendix

Table 11 describes the BAM data on the average replacement rate and benefit duration for each U.S. state, the District of Columbia (DC), and Porta Rico (PR). The U.S. average for the replacement rate is 0.45 and for the potential benefit duration is 24 weeks

Table 11: Replacement Rate and Benefit Duration, by State

State	Benefit Duration	Replacement Rate	State	Benefit Duration	Replacement Rate
AK	21.09	0.32	MT	21.95	0.43
AL	24.71	0.41	NC	23.15	0.48
AR	23.63	0.48	ND	20.67	0.40
AZ	23.50	0.40	NE	23.90	0.44
CA	24.57	0.45	NH	26.00	0.38
CO	24.20	0.45	NJ	24.84	0.54
CT	26.07	0.47	NM	25.39	0.62
DC	25.65	0.42	NV	23.73	0.44
DE	25.68	0.43	NY	26.00	0.44
FL	22.17	0.41	OH	25.65	0.44
GA	21.23	0.44	OK	22.99	0.52
HI	25.99	0.54	OR	25.11	0.47
IA	24.18	0.48	PA	25.92	0.50
ID	21.62	0.46	PR	26.00	0.43
IL	25.98	0.39	RI	22.69	0.55
IN	20.82	0.52	SC	24.10	0.46
KS	23.37	0.50	SD	24.63	0.43
KY	25.66	0.46	TN	22.29	0.41
LA	23.25	0.40	TX	22.30	0.44
MA	27.35	0.50	UT	21.61	0.45
MD	25.99	0.46	VA	21.82	0.43
ME	23.21	0.47	VT	26.15	0.49
MI	24.77	0.45	WA	25.80	0.43
MN	24.03	0.44	WI	24.36	0.46
MO	23.71	0.43	WV	26.00	0.38
MS	24.18	0.44	WY	22.43	0.46