

Veracierto, Marcelo

Working Paper

Moral hazard, optimal unemployment insurance, and aggregate dynamics

Working Paper, No. WP 2022-07

Provided in Cooperation with:
Federal Reserve Bank of Chicago

Suggested Citation: Veracierto, Marcelo (2022) : Moral hazard, optimal unemployment insurance, and aggregate dynamics, Working Paper, No. WP 2022-07, Federal Reserve Bank of Chicago, Chicago, IL,
<https://doi.org/10.21033/wp-2022-07>

This Version is available at:
<https://hdl.handle.net/10419/264321>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Federal Reserve Bank of Chicago

Moral Hazard, Optimal Unemployment Insurance, and Aggregate Dynamics

Marcelo Veracierto

March 2022

WP 2022-07

<https://doi.org/10.21033/wp-2022-07>

**Working papers are not edited, and all opinions and errors are the responsibility of the author(s). The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.*

Moral Hazard, Optimal Unemployment Insurance, and Aggregate Dynamics

Marcelo Veracierto*

Federal Reserve Bank of Chicago, U.S.A

March, 2022

Abstract: In this paper, I explore how optimal aggregate dynamics can be shaped by the presence of moral hazard in unemployment insurance. I also analyze the optimal provision of unemployment insurance and the implications for the amount of cross-sectional heterogeneity. The economy that I consider embeds the Hopenhayn-Nicolini unemployment insurance model into a real business cycle model with search frictions. In a calibrated version I find that the presence of private information has large effects on optimal aggregate steady-state dynamics but not on aggregate fluctuations. In addition, I find that optimal consumption replacement ratios are approximately independent of the business cycle.

Keywords: Private information; Mechanism Design; Business Cycles; Moral hazard; Unemployment Insurance

JEL classification: D82; E32; J64; J65

*The views expressed here do not necessarily reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System. Address: Federal Reserve Bank of Chicago, Research Department, 230 South LaSalle Street, Chicago, IL 60604. E-mail: mveracie@frbchi.org. Phone: (312) 322-5695.

1 Introduction

The main purpose of this paper is to provide a theoretical exploration of how the presence of moral hazard in unemployment insurance (UI) can affect optimal aggregate dynamics, both at the steady state and over business cycle fluctuations. A novel feature of the paper is that instead of performing the analysis under exogenous forms of market incompleteness (such as borrowing constraints) and restricting policy instruments to belong to a certain class (e.g., constant replacement ratios or a fixed duration of benefits), it takes a more primitive approach. In particular, it solves the mechanism design problem of a real business cycle (RBC) economy subject to search frictions. Moral hazard problems arise in this economy because the search efforts of individuals are unobserved, creating a nontrivial trade-off for the social planner in terms of providing insurance and incentives. I am interested not only in studying how aggregate dynamics are affected by the presence of private information, but also in characterizing the optimal provision of unemployment insurance and exploring its implications for the optimal amount of cross-sectional inequality, both at the steady state and over the business cycle.

The model economy that I use embeds the optimal unemployment insurance model of Hopenhayn and Nicolini (1997) into a RBC model with search frictions. In this economy, all the production is done in a central island, where a representative firm produces output using capital and labor subject to an aggregate productivity shock. The aggregate shock follows a standard first-order autoregressive (AR(1)) process. Agents become exogenously separated from the production island; and in order to get back to it, they need to search. The probability of arriving at the production island depends on the search intensity of the agent, which is private information. While the search intensity of an agent is not observable, their location (either inside or outside the production island) is; and therefore, allocations can be made contingent on this information. Agents are risk-averse, value consumption, and dislike to search.

A social planner designs optimal dynamic contracts for the agents in this economy. In line with the literature, a dynamic contract is given a standard recursive formulation in which a promised value to the agent and the employment status of the agent describe its state. Given the current state of the contract and the aggregate state of the economy, the

contract specifies current consumption, current search intensity (if the agent is unemployed), and next-period state-contingent promised values. Since the model has a large number of agents and labor market shocks are idiosyncratic, the social planner needs to keep track of the whole distribution of individuals across promised values and employment statuses as a state variable. Given this distribution, the aggregate stock of capital, and the aggregate productivity level, the social planner seeks to maximize the present discounted utility of agents subject to incentive-compatibility, promise-keeping, and aggregate resource-feasibility constraints.

I calibrate this Hopenhayn-Nicolini RBC model to U.S. observations and compare its optimal aggregate dynamics under public and private information. I find that the presence of information frictions has important effects on its steady-state dynamics: The private information reduces aggregate employment, consumption, investment, and output by 2.4% and creates a significant amount of consumption inequality: The cross-sectional standard deviation of log consumption goes from zero to 0.083. However, at business cycle frequencies the private information has negligible aggregate effects: The impulse responses of aggregate employment, consumption, investment and output are approximately the same under public and private information.

The intuition for why the private information matters so much for steady-state dynamics is that the social planner faces a nontrivial trade-off between aggregate employment and unemployment insurance provision: Given that the separation rate is exogenous, the planner can only increase aggregate employment by inducing the unemployed agents to increase their search intensities. However, since their search intensities are private information, this can only be done by increasing the differences between the promised values of becoming employed and the promised values of continuing to be unemployed (thus reducing the amount of unemployment insurance provided). Given this trade-off, the social planner chooses to generate a lower aggregate employment level at the deterministic steady state. With business cycle dynamics things look different. It is still true that when there is a positive aggregate productivity shock, in order to induce unemployed agents to search more intensively, the social planner needs to punish them with the insurance that they receive. However, this is compensated by the higher insurance that they receive when there is a negative aggregate productivity shock. At a first-order approximation, the losses and gains in unemployment

insurance provision across expansions and contractions balance out and make the social planner respond to aggregate shocks approximately the same way under public and private information.

In terms of the design of optimal unemployment insurance, I find that there are no significant interactions with the state of the business cycle. Optimal consumption replacement ratios are determined at the deterministic steady-state of the economy and solely depend on idiosyncratic labor market histories. Over the business cycle, the consumption of agents at different stages of their idiosyncratic histories get roughly scaled up or down by a common factor, leaving all replacement ratios approximately unchanged. In fact, the variance of the distribution of log consumption levels remains roughly constant over time. What is significantly affected over the business cycle is the distribution of promised values of unemployed agents, since these agents are required to search more during booms and they must be given the incentives to do so.

This paper is related to a vast literature on optimal unemployment insurance. First, my paper is tightly related to Hopenhayn and Nicolini (2009), since it uses essentially the same model as theirs, but embedded in an RBC context.² However, Hopenhayn and Nicolini (2009) analyze optimal unemployment insurance abstracting from business cycle considerations, as do Hopenhayn and Nicolini (1997), Werning (2002), and Shimer and Werning (2008). Kroft and Notowidigdo (2011), Sanchez (2008) and Williams and Li (2015) all consider business cycle effects but perform their analysis within a principal-agent setup. In contrast, I consider an economy-wide social planning problem under aggregate uncertainty that takes into account not only the individual incentive constraints, but also the aggregate feasibility constraints. While there is a large general equilibrium literature analyzing optimal UI provision in a deterministic environment, the analysis in business cycle settings is more limited. Recent examples include Landais et al. (2018), Jung and Kuester (2015), and Mitman and Rabinovich (2015) in Mortensen-Pissarides frameworks and Boostani et al. (2017) in a directed search environment. However, all of these papers impose exogenous

²Actually, it is essentially the same model as the one in Section 2 of Hopenhayn and Nicolini (2009). However, instead of having indivisible search decision, search intensities in my model are continuous. Also, instead of characterizing optimal allocations only at low promised values, I characterize them over the whole support of the distribution (although my characterization is numerical instead of analytical).

restrictions to the financial markets that agents have available (e.g., exogenous borrowing limits or hand-to-mouth workers) and to the policy instruments that the government has available (e.g., homogeneous UI benefits, an infinite duration of UI benefits, or a constant probability of UI termination). In contrast, the Hopenhayn-Nicolini RBC model in this paper analyzes the true social optimum for the environment considered, in which the only restrictions to risk sharing arise because of the presence of private information. Thus, it provides a useful benchmark to which equilibria with exogenous restrictions to risk sharing and policy instruments may be compared.

Previous work analyzing optimal social insurance in aggregate models include Green (1987), Atkeson and Lucas (1992), Golosov et al. (2007), Farhi and Werning (2012), and Golosov et al. (2016). However, interactions with aggregate fluctuations have been mostly neglected in this literature. Notable exceptions are Phelan (1994), da Costa and Luz (2018), Werning (2007), Scheuer (2013), and Veracierto (2021). The most closely related paper is Veracierto (2021), since it also solves the mechanism design problem of an RBC economy subject to information frictions. The model that I considered in that paper is a standard RBC economy populated by agents who value consumption and leisure, and who receive idiosyncratic shocks to their value of leisure. These shocks, which are independent and identically distributed (i.i.d.) over time and across individuals, are private information. For that Mirrlees RBC economy, I showed that the presence of private information is completely irrelevant for both steady-state and business cycle dynamics: The stationary behavior of all macroeconomic variables (i.e., aggregate output, consumption, investment, hours worked, and capital) are the same under private information as under full information.³ On the contrary, in the Hopenhayn-Nicolini RBC model analyzed in this paper, the information frictions greatly affect steady-state aggregate dynamics. However, the information frictions are still irrelevant for business cycle dynamics.

The paper is organized as follows. In Section 2, I describe the economy. In Section 3, I describe the mechanism design problem. In Section 4, I parametrize the economy. Next, I analyze the optimal steady-state dynamics in Section 5. Then, I analyze the optimal

³For logarithmic preferences, I established the irrelevance result analytically. For other preferences, I showed it numerically.

fluctuations in Section 6. Finally, I provide conclusions in Section 7.

2 The economy

The economy is populated by a unit measure of agents subject to stochastic lifetimes. Whenever an agent dies they are immediately replaced by a newborn, leaving the aggregate population level constant over time. The preferences of an individual born at date T are given by

$$E_T \left\{ \sum_{t=T}^{\infty} \beta^{t-T} \sigma^{t-T} [r(c_t) - s_t] \right\},$$

where σ is the survival probability, $0 < \beta < 1$ is the discount factor, c_t is consumption, s_t is search intensity, r is given by

$$r(c) = \frac{c^{1-\alpha} - 1}{1 - \alpha},$$

and $\alpha > 0$.

All production in the economy is done in a central island.⁴ Agents get exogenously separated from the production island at a common rate ϕ , and in order to get back to it, they need to search. The probability of arriving at the production island at the beginning of the following period depends on the search intensity of the individual according to a strictly increasing and concave function $\eta(s_t)$. The search intensity s_t is assumed to be private information, creating a moral hazard problem similar to the one in Hopenhayn and Nicolini (2009). All agents are born outside the production island.

Output, which can be consumed or invested, is produced according to the following production function:

$$Y_t = e^{z_t} K_{t-1}^\gamma H_t^{1-\gamma},$$

where Y_t is output, z_t is aggregate productivity, K_{t-1} is capital, H_t is employment, and $0 < \gamma < 1$. The aggregate productivity level z_t follows a standard AR(1) process given by

$$z_{t+1} = \rho z_t + \varepsilon_{t+1},$$

where $0 < \rho < 1$ and ε_{t+1} is normally distributed with zero mean and standard deviation σ_ε .

⁴The fiction of a central island has been previously used by Alvarez and Veracierto (2001) and Ljungqvist and Sargent (2007).

Capital is accumulated using a standard linear technology given by

$$K_t = (1 - \delta) K_{t-1} + I_t, \tag{1}$$

where I_t is gross investment and $0 < \delta < 1$.

3 The mechanism design problem

In this section, I describe the problem of a social planner that maximizes utility subject to incentive-compatibility, promise-keeping, and aggregate resource-feasibility constraints. While the search intensities of agents are not observable, their locations (either inside or outside the production island) are; and therefore, recursive contracts can be made contingent on this information. For simplicity's sake, in what follows an agent will be called *employed* if they are located inside the production island, and *unemployed* if they are located outside of it.

In order to describe the mechanism design problem it will be important to distinguish between two types of agents: young and old. A young agent is one who has been born at the beginning of the current period. An old agent is one who has been born in some previous period. The social planner must decide recursive plans for both types of agents. The state of a recursive plan is given by the agent's employment status and the value (i.e., discounted expected utility) that the agent is entitled to at the beginning of the period. Given this state, the recursive plan specifies the current utility of consumption, the current search intensity, and the next-period state-contingent promised values. The social planner is fully committed to the recursive plans they choose and agents have no outside opportunities available.

A key difference between the young and the old is in terms of promised values. Because during the previous period the social planner has already decided on some recursive plan for a currently old agent, the planner is restricted to delivering the corresponding promised value during the current period. In contrast, the social planner is free to deliver any value to a currently young agent, since this is the first period they are alive. Reflecting this difference, I will specify the individual state of an old agent to be their promised value v and their current employment status n , where n is either e (employed) or u (unemployed). At date t , their current utility of consumption, search intensity, and next-period promised

value are denoted by $r_{ont}(v)$, $s_{ot}(v)$, and $w_{onn',t+1}(v)$, respectively, where $w_{onn',t+1}(v)$ is a random variable contingent on the realization of n' and z_{t+1} .⁵ Observe that there is no need to specify the search intensity of employed agents because it is always set to zero. In turn, since all young agents are unemployed and have no ongoing recursive contracts, their current utility of consumption, search intensity, and next-period promised value are simply denoted by r_{yut} , s_{yt} , and $w_{yun',t+1}$, respectively, where $w_{yun',t+1}$ is also contingent on the realization of n' and z_{t+1} .

The aggregate state of the economy is given by a vector $(z_t, K_{t-1}, \mu_{et}, \mu_{ut})$, where z_t is the aggregate productivity level, K_{t-1} is the stock of capital, μ_{et} is a measure describing the number of employed old agents across individual promised values v , and μ_{ut} is similar to μ_{et} but for unemployed old agents. The social planner seeks to maximize the weighted sum of the welfare levels of current and future generations of young agents (the welfare levels of old agents are predetermined by their promised values at the beginning of the period). That is, given the current state vector of the economy, the planner maximizes

$$V_t(z_t, K_{t-1}, \mu_{et}, \mu_{ut}) = \max \{ (1 - \sigma) (r_{yut} - s_{yt} + \beta \sigma E_t [\eta(s_{yt}) w_{yue,t+1} + (1 - \eta(s_{yt})) w_{yuu,t+1}]) + \theta E_t [V_{t+1}(z_{t+1}, K_t, \mu_{e,t+1}, \mu_{u,t+1})] \}, \quad (2)$$

subject to incentive-compatibility, promise-keeping, and aggregate resource-feasibility constraints, where E_t denotes the expectation conditional on z_t , V_t is the social planner's value function, and θ is the Pareto weight of the next-period generation relative to the current-period generation.⁶

Instead of providing the details for this economy-wide planning problem, I will describe it in terms of its component planning problems, as this is more convenient. In each period, there are two types of component planning problems—one concerned with providing insurance and incentives to individuals and the other concerned with making production and investment decisions. In these component planning problems, the joint stochastic process for the shadow price of labor (in terms of the consumption good), q_t , and the shadow price

⁵I follow the convention that a variable is dated t if it becomes known at date t .

⁶Observe that given the constant probability of dying $1 - \sigma$ and the immediate replacement with newborns, the number of young agents in the economy is always equal to $1 - \sigma$.

of the consumption good (in utiles), λ_t , are taken as given. The solutions to these component planning problems correspond to the solution of the economy-wide planning problem if certain side conditions are satisfied.

The component planning problem for individuals can be distinguished by the age and employment status of the agent. The date t component planning problem for employed old individuals is the following:

$$P_{oet}(v) = \max \left\{ \{q_t - c(r_{oet}(v))\} + \theta \sigma E_t \left[\frac{\lambda_{t+1}}{\lambda_t} (\phi P_{ou,t+1}(w_{oeu,t+1}(v)) + (1 - \phi) P_{oe,t+1}(w_{oee,t+1}(v))) \right] \right\}, \quad (3)$$

subject to

$$v = r_{oet}(v) + \beta \sigma E_t [\phi w_{oeu,t+1}(v) + (1 - \phi) w_{oee,t+1}(v)], \quad (4)$$

where $c(r) = [(1 - \alpha)r + 1]^{\frac{1}{1-\alpha}}$ is the consumption level corresponding to the utility of consumption r , $P_{oet}(v)$ is the current social value of an employed old individual with promised value v (in terms of the current consumption good), and $P_{out}(v)$ is similar to $P_{oet}(v)$ but for unemployed old agents. Observe that the flow social value of an employed old worker is the value of their labor input, net of the consumption goods transferred to them. The expected continuation social value of an employed old agent takes into account that, conditional on surviving, the agent will become unemployed with probability ϕ . This continuation social value is discounted according to the stochastic social discount factor $\frac{\lambda_{t+1}}{\lambda_t}$ and the relative Pareto weight θ . Equation (4) is the promise-keeping constraint.

The date t component planning problem for unemployed old individuals is

$$P_{out}(v) = \max \left\{ -c(r_{out}(v)) + \theta \sigma E_t \left[\frac{\lambda_{t+1}}{\lambda_t} (\eta(s_{ot}(v)) P_{oe,t+1}(w_{oue,t+1}(v)) + (1 - \eta(s_{ot}(v))) P_{ou,t+1}(w_{ouu,t+1}(v))) \right] \right\}, \quad (5)$$

subject to

$$v = r_{out}(v) - s_{ot}(v) + \beta \sigma E_t [\eta(s_{ot}(v)) w_{oue,t+1}(v) + (1 - \eta(s_{ot}(v))) w_{ouu,t+1}(v)], \quad (6)$$

$$\eta'(s_{ot}(v)) \beta \sigma E_t [w_{oue,t+1}(v) - w_{ouu,t+1}(v)] = 1. \quad (7)$$

Observe that the flow social value of an old unemployed agent is simply the cost of the consumption goods transferred to them. The expected continuation social value of an unemployed old agent takes into account that, conditional on surviving, the agent will become employed with probability $\eta(s_{ot})$. Equation (6) is the promise-keeping constraint. Equation (7) is the incentive-compatibility constraint: It is the first-order condition for the individual's optimal choice of search intensity. Since the search intensity is not observable, the agent chooses it to maximize their private gains. In an interior solution, this is attained by equating the marginal private benefits (given by the marginal impact on the hazard rate multiplied by the expected discounted gains of becoming employed) to the marginal private cost (which is equal to one, given that s_{ot} enters linearly into the individual's utility function).⁷

Since all agents are born unemployed, there is only one type of component planning problem for young agents. It is given by

$$P_{yut} = \max \left\{ \frac{r_{yut} - s_{yt} + \beta\sigma E_t [\eta(s_{yt}) w_{yue,t+1} + [1 - \eta(s_{yt})] w_{yuu,t+1}]}{\lambda_t} - c(r_{yut}) + \theta\sigma E_t \left[\frac{\lambda_{t+1}}{\lambda_t} [\eta(s_{yt}) P_{oe,t+1}(w_{yue,t+1}) + [1 - \eta(s_{yt})] P_{ou,t+1}(w_{yuu,t+1})] \right] \right\}, \quad (8)$$

subject to

$$\eta'(s_{yt}) \beta\sigma E_t [w_{yue,t+1} - w_{yuu,t+1}] = 1. \quad (9)$$

Observe that the present value of the lifetime utility of the young agent directly enters the flow social value because the economy-wide social planner in equation (2) seeks to maximize the weighted sum of the welfare levels of current and future generations. The consumption goods transferred to the young agent are subtracted from the flow social value because the transfer tightens the aggregate consumption feasibility constraint of the economy-wide social planner. The incentive-compatibility constraint (9) is similar to equation (7).

The production component planning problem is simply

$$P_{pt}(K_{t-1}) = \max \left\{ e^{zt} K_{t-1}^\gamma H_t^{1-\gamma} - q_t H_t - I_t + \theta E_t \left[\frac{\lambda_{t+1}}{\lambda_t} P_{p,t+1}(K_t) \right] \right\}, \quad (10)$$

subject to

$$K_t = (1 - \delta) K_{t-1} + I_t. \quad (11)$$

⁷I describe the interior solution case for simplicity's sake. In general one needs to check for corner solutions (since $\eta(s_{ot}(v))$ cannot be negative or exceed one).

Observe that the flow social surplus in this planning problem is given by output net of the value of the labor input and the value of investment.

The economy-wide distribution of employed old agents across promised values is denoted by μ_{et} . The similar object for unemployed old agents is denoted by μ_{ut} . The law of motion for μ_{et} is given by

$$\begin{aligned} \mu_{e,t+1}(\mathcal{B}) &= (1 - \sigma) \sigma \eta(s_{yt}) \mathcal{I}[w_{yue,t+1} \in \mathcal{B}] + \sigma \int_{\{v:w_{oue,t+1}(v) \in \mathcal{B}\}} \eta(s_{ot}(v)) d\mu_{ut} \\ &\quad + \sigma \int_{\{v:w_{oee,t+1}(v) \in \mathcal{B}\}} (1 - \phi) d\mu_{et}, \end{aligned} \quad (12)$$

for every Borel set \mathcal{B} , where \mathcal{I} is an indicator function that takes a value of one if its argument is true and zero otherwise. This equation states that the number of employed old agents who have a promised value in the set \mathcal{B} at the beginning of the following period is given by the sum of three terms. The first term includes all young agents who do not die, find the production island, and get a promised value in the set \mathcal{B} . The second term includes all unemployed old agents who do not die, find the production island, and get a promised value in the set \mathcal{B} . The third term includes all employed old agents who do not die, do not get separated from the production island, and get a promised value in the set \mathcal{B} . The law of motion for μ_{ut} , which is given by

$$\begin{aligned} \mu_{u,t+1}(\mathcal{B}) &= (1 - \sigma) \sigma (1 - \eta(s_{yt})) \mathcal{I}[w_{yuu,t+1} \in \mathcal{B}] + \sigma \int_{\{v:w_{ouu,t+1}(v) \in \mathcal{B}\}} (1 - \eta(s_{ot}(v))) d\mu_{ut} \\ &\quad + \sigma \int_{\{v:w_{oeu,t+1}(v) \in \mathcal{B}\}} \phi d\mu_{et}, \end{aligned} \quad (13)$$

is similarly interpreted. Observe that because all next-period promised values are contingent on the realization of next-period aggregate productivity z_{t+1} , $\mu_{e,t+1}$ and $\mu_{u,t+1}$ are also state-contingent.

Given initial values for K_{-1} , μ_{e0} , and μ_{u0} , the side conditions that the stochastic process $\{q_t, \lambda_t\}_{t=0}^{\infty}$ needs to satisfy are the following:

$$H_t = \int d\mu_{et}, \quad (14)$$

and

$$(1 - \sigma) c(r_{yut}) + \int c(r_{out}(v)) d\mu_{ut} + \int c(r_{oet}(v)) d\mu_{et} + K_t - (1 - \delta) K_{t-1} = e^{z_t} K_{t-1}^{\gamma} H_t^{1-\gamma}. \quad (15)$$

Equation (14) is the aggregate labor feasibility constraint. It states that the input of hours into the production function equals the total number of employed old agents. Equation (15) describes the aggregate feasibility constraint for the consumption good. It states that the total consumption of the young, old unemployed, and old employed agents, plus aggregate investment, equals aggregate output.

A solution to the economy-wide mechanism design problem is a stochastic process for the measures $\{\mu_{et}, \mu_{ut}\}_{t=0}^{\infty}$, scalars $\{q_t, \lambda_t, K_{t-1}, I_t, H_t, P_{yut}, r_{yut}, s_{yt}, w_{yue,t+1}, w_{yuu,t+1}\}_{t=0}^{\infty}$, and functions $\{P_{out}, r_{out}, s_{ot}, w_{oue,t+1}, w_{ouu,t+1}, P_{oet}, r_{oet}, w_{oeu,t+1}, w_{oee,t+1}, P_{pt}\}_{t=0}^{\infty}$ such that starting from K_{-1} , μ_{e0} , and μ_{u0} , equations (3)-(15) are satisfied almost surely, for $t \geq 0$. In what follows I sketch the method that I use to compute a stationary solution.⁸

Computing a stationary solution is a difficult task not only because two of the state variables, μ_{et} and μ_{ut} , are infinitely dimensional but also because they are contingent on the realization of the aggregate shock z_t . Moreover, the distributions μ_{et} and μ_{ut} turn out to have odd shapes, which precludes the possibility of parametrizing them with polynomial forms and which are likely to make the laws of motion (12) and (13) highly non-linear. In order to avoid these difficulties I generalize the computational strategy introduced in Veracierto (2002) in a straightforward way. In particular, instead of keeping track of the distributions of promised values as state variables, what the computational method keeps track of is a long (but finite) history of realized decision rules:

$$\{w_{oee,t+1-m}, w_{oeu,t+1-m}, w_{ouu,t+1-m}, w_{oue,t+1-m}, w_{yuu,t+1-m}, w_{yue,t+1-m}, s_{o,t-m}, s_{y,t-m}\}_{m=1}^M. \quad (16)$$

Because all functions are parametrized as spline approximations, the computational method only needs to keep track of the values of the decision rules at the grid points $(v_j)_{j=1}^J$.⁹ Finite versions of the current distributions μ_{et} and μ_{ut} are then obtained by simulating the evolution of a large (but finite) panel of agents (and their descendants) using the realized history

⁸The general method is described in detail in Veracierto (2020).

⁹In Veracierto (2002) the decision rules were of the (S,s) variety. Thus, it was sufficient to keep track of the history of lower and upper thresholds that defined the decision rules. In this paper the decision rules are smooth functions that are approximated as splines.

of decision rules in (16) and the laws of motion (12) and (13).¹⁰ Given the current distributions μ_{et} and μ_{ut} (described by the updated panel of agents) and the current consumption allocation rule r_{out} , averages across agents can be then calculated to evaluate the aggregate feasibility constraints (14) and (15).

Because the distributions μ_{et} and μ_{ut} are completely described by the history of decision rules in equation (16), the laws of motion (12) and (13) can be replaced by the law of motion for the spline coefficients describing the functions in (16). Updating the history of spline coefficients with those chosen during the current period defines a (trivially) linear mapping that requires no further approximation.

The last piece of the computational strategy is to consider equations (3)-(11) (and their associated first-order and envelope conditions) only at the grid points $(v_j)_{j=1}^J$. Moreover, because all functions are described as spline approximations, terms like $P_{ou,t+1}(w_{oeu,t+1})$ in equation (3), which involves evaluating the function $P_{ou,t+1}$ outside a grid point, must be thought of as being determined by the vector of values of $P_{ou,t+1}$ at the grid points $(v_j)_{j=1}^J$ and by the scalar $w_{oeu,t+1}$.

Equations (3)-(11), their first-order and envelope conditions, and the aggregate feasibility conditions (14)-(15) can then be linearized (around the deterministic steady state) with respect to all scalar variables and function values at the grid points $(v_j)_{j=1}^J$.¹¹ The resulting system constitutes a linear rational expectations model that despite its high dimensionality, can be solved using standard methods. The only nonstandard feature of the system is that we seek a solution in which some state variables at date t are contingent on the realization of the productivity shock z_t and some decision variables at date t are contingent on the realization of z_{t+1} . However, a straightforward transformation of the deterministic solution delivers the state-contingent solution that we seek.¹²

¹⁰The initial panel of agents (M periods into the past) is randomly drawn from the deterministic steady-state distributions. However, because of the stochastic lifetimes, if M is large enough the initial panel of agents becomes immaterial because very few of them will survive M consecutive periods.

¹¹The laws of motion (12) and (13) have already been given a linear representation.

¹²See Proposition 1 in Veracierto (2020).

4 Parametrization

In order to parametrize the model, I consider its full information version, which is obtained by dropping the incentive compatibility constraints (7) and (9) from the mechanism design problem. The reason for doing this is that the full information economy is far simpler to compute than the private information case.¹³ In fact, it turns out that when the relative Pareto weight θ equals the discount factor β , the full information mechanism design problem reduces to the following representative agent planning problem:

$$V(z_t, K_{t-1}, H_t) = \max \{r_t - s_t(1 - H_t) + \beta E_t [V(z_{t+1}, K_t, H_{t+1})]\}, \quad (17)$$

subject to

$$c(r_t) + K_t - (1 - \delta) K_{t-1} \leq e^{z_t} K_{t-1}^\gamma H_t^{1-\gamma}, \quad (18)$$

$$H_{t+1} \leq \sigma [\eta(s_t)(1 - H_t) + (1 - \phi) H_t], \quad (19)$$

$$z_{t+1} = \rho z_t + \varepsilon_{t+1}, \quad (20)$$

where all variables are exactly the same as in Section 2. Observe that in equation (19) only a fraction σ of the total number of agents who would otherwise find a job $\eta(s_t)(1 - H_t)$ or remain employed $(1 - \phi) H_t$ end up actually employed at the beginning of the following period. The reason is that in the original heterogeneous agents economy a fraction $1 - \sigma$ of the population dies and is replaced by young agents who start their lives unemployed.

In parametrizing this economy I set the time period to one year, assume that the search technology is given by

$$\eta(s_t) = D(1 - e^{-\tau s_t}),$$

and choose parameter values as described next.¹⁴ Following the RBC literature, I select a labor share $1 - \gamma$ of 0.64, a depreciation rate δ of 0.10, a private discount factor β of 0.96, a

¹³The parametrization strategy below requires computing the business cycle fluctuations of the model under each possible set of parameters considered. Applying this strategy to the economy with private information would be too costly.

¹⁴I experimented using other functional forms for the search technology, but in all cases the model was able to generate much lower employment fluctuations.

persistence of aggregate productivity ρ of 0.95, and a variance of the innovations to aggregate productivity σ_ε^2 equal to 4×0.007^2 . The social discount factor θ is chosen to be the same as the private discount factor β (an assumption already embodied in the representative agent problem (17)-(20)). Consumption preferences are assumed to be logarithmic (i.e., α is set to 1). In terms of the life-cycle structure, I choose $\sigma = 0.975$ in order to generate an expected life span of 40 years. There are three parameters closely related to the labor dynamics of the model: the employment separation rate ϕ and the search technology parameters D and τ . To determine ϕ I turn to Krusell et al. (2017), who measured a quarterly employment-to-employment transition rate using Current Population Survey (CPS) data. This transition rate implies an average duration of an employment spell equal to 8.9 years.¹⁵ Reproducing this observation requires setting $\phi = 0.112$. In turn, the parameters D and τ are chosen to satisfy two criteria: that aggregate employment under full information be equal to 0.60 (the ratio of employment to the working-age population in U.S. data) and that the standard deviation of aggregate employment be as large as possible (since models with search frictions tend to have difficulties generating large employment fluctuations). The selected values of D and τ turn out to be 8.0 and 0.016, respectively.¹⁶

5 Steady-state dynamics

I now turn to describing the deterministic steady-state properties of the economy with private information under the parameters determined in the previous section. I start with the blue line in Figure 1.A, which depicts the job-finding probability for old unemployed agents $\eta(s_o)$ across promised values v . Not surprisingly, this function is generally decreasing in v . That is, agents with higher promised values are required to search less. What is interesting is that there exists a threshold promised value above which the job-finding probability becomes zero (marked as a vertical red line) and another threshold promised value below which the job-

¹⁵Observe that average employment spells are much longer than average job spells, since many workers experience job-to-job transitions without going through nonemployment.

¹⁶Under this parametrization aggregate employment fluctuates 32% as much as output (not as much as in the data, but a much better performance than the standard Mortensen-Pissarides model).

finding probability becomes one (marked as a vertical green line). The gold line in Figure 1.A shows the job-finding probability $\eta(s_y)$ for young agents. We see that agents start their lives with a job-finding probability of 0.23.

Figure 1.B depicts next-period promised values for unemployed agents as a function of v . The purple line is the 45-degree line, while the vertical lines mark the threshold promised values. In turn, the gold line describes w_{ouu} , while the dark blue line describes w_{oue} .¹⁷ We see that w_{ouu} coincides with the 45-degree line above the zero-search threshold. Thus, if some agent enters unemployment with a promised value larger than that threshold, they remain unemployed forever and their promised value never changes (effectively retiring at that promised value). However, the w_{ouu} function remains uniformly below the 45-degree line for promised values below the zero-search threshold. Thus, an unemployed agent's promised value decreases during their unemployment spell. Moreover, the punishment for remaining unemployed increases during the unemployment spell: The vertical difference between the 45-degree line and w_{ouu} increases with lower values of v . However, when an unemployed agent finds employment, the agent gets a significant reward in terms of next-period promised value (the vertical difference between w_{oue} and the 45-degree line). Since w_{oue} is parallel to the 45-degree line, this reward is the same at all promised values below the zero-search threshold. Observe that this reward is needed even when the agents put a tiny amount of search (i.e., when their promised value v is an epsilon below the zero-search threshold). The reason for this is that the search technology chosen has a finite slope at $s = 0$; and therefore, the incentive compatibility constraint (7) requires a positive reward for inducing the agent to do even an infinitesimal amount of search. Figure 1.B also displays the next-period promised values of young agents in the event that they become employed w_{ye} or continue to be unemployed w_{yu} . We see that they roughly correspond to the values w_{oue} and w_{ouu} received by an old agent with a promised value v in the middle of the distribution μ_u .

¹⁷Both functions are depicted over the support of the invariant distribution μ_u . However, the function w_{ouu} is not shown for promised values smaller than the low threshold because no unemployed agent remains unemployed in that range. Similarly, the function w_{oue} is not depicted for promised values larger than the high threshold because no unemployed agent becomes employed in that range.

Figure 1.C depicts next-period promised values for employed agents as a function of v .¹⁸ Again the purple line represents the 45-degree line and the vertical lines mark the threshold promised values. The blue line describes w_{oee} , while the gold line describes w_{oeu} . We see that w_{oee} coincides with the 45-degree line at all values of v . That is, the promised values of employed agents do not change while they remain employed, which is quite intuitive because employed agents face no incentive problems. Observe that w_{oeu} also coincides with the 45-degree line but to the right of the zero-search threshold. That is, employed agents with promised values higher than the threshold do not see their promised values change when they become unemployed. This is also quite intuitive because these agents will never search again (effectively retiring from the labor force). However, we see that w_{oeu} remains below the 45-degree line to the left of the zero-search threshold, so in this range employed agents get punished when they become unemployed. In fact we see that the punishment, given by the vertical difference between the 45-degree line and w_{oeu} , increases with lower values of v . This is also intuitive because, besides what may happen with their consumption levels at the time that they become unemployed (an issue that will be addressed below), their search intensities (which reduce their welfare levels) increase with lower values of v .

Figure 2.A shows the log of consumption as a function of v . The vertical lines once again represent the search thresholds, and the blue line shows u_{oe} and the gold line u_{ou} . The interesting feature of this figure is that to the right of the zero-search threshold, u_{oe} and u_{ou} coincide: The consumption level in that region is identical for employed and unemployed agents. Thus, when an employed agent retires, their consumption level remains exactly the same as when they were last employed. Not surprisingly, to the left of the zero-search threshold, u_{ou} must be higher than u_{oe} because in order to obtain the same promised value v , the unemployed agents must be compensated with higher consumption for their positive search effort.

More interesting is to analyze the consumption changes that take place as agents transition between the different employment states, since this informs us about the amount of insurance provided. This is shown in Figure 2.B. The light blue line shows the consumption changes that take place when an employed agent continues to be employed in the next pe-

¹⁸Both functions are depicted over the support of the invariant distribution μ_e .

riod.¹⁹ We see that their consumption levels do not change. This is quite obvious: Since these agents experience no idiosyncratic shock to their employment state and since being employed entails no incentive problems, there is no reason to change their consumption levels. Much more interesting is the dark blue line, which shows the consumption changes that take place when an employed agent becomes unemployed in the following period. Again, we see that their consumption levels are not affected. This is actually quite intuitive: Since the idiosyncratic shocks that determine an employment-to-unemployment transition are completely exogenous and current search decisions are not affected by current consumption levels, there is no reason not to fully insure agents against those shocks. The gray line shows the consumption changes that take place when an unemployed agent continues to be unemployed the following period. For promised values larger than the zero-search threshold, we see that there are no consumption changes: Agents receive a constant consumption stream during their retirement. However, to the left of the zero-search threshold, we see that consumption drops if an unemployed agent continues to be unemployed. The reason is that in order to induce the agent to search, the planner needs to punish them in case that they continue to be unemployed (and reward them in case they become employed) and partially does this by reducing the agent's consumption level (recall that the agents are also induced to increase their search intensity if they continue to be unemployed). Observe that for promised values close to the zero-search threshold, the consumption change is small but it becomes significant at lower promised values. Thus, if an agent remains unemployed for a long period, the accumulated consumption loss can become quite substantial. The gold line shows the consumption changes that take place when unemployed agents find employment. We see that for promised values smaller than the low threshold there are no changes to their consumption levels. The reason is that because these agents find employment with a probability of one, there is no reason not to smooth their consumption levels perfectly across the two employment states (the agent is induced to find employment with a probability of one by being promised a sufficiently low continuation value w_{ouu}). To the right of the low threshold we see that the gain in promised values shown in Figure 1.B (the vertical differences between the w_{oue} and the 45-degree lines) are achieved not only by the agents not having to search

¹⁹The light blue line happens to be completely hidden under the dark blue line.

once they become employed, but also by their consumption levels increasing. Once promised values are sufficiently close to the high threshold, the search intensities are so close to zero that the gains in promised values must be achieved by large increases in consumption levels (although these gains actually take place with very low probabilities).

Figures 3.A and 3.B report the implied invariant distributions μ_e and μ_u , respectively. We have already encountered the least upper bound for the support of μ_e : In Figure 1.B the value of the w_{oue} function at the zero-search threshold gives the highest promised value that an unemployed agent can possibly get by becoming employed. Since we know from Figure 1.C that an employed agent with that promised value does not get punished when they become unemployed and retire, this value is also the least upper bound for the support of μ_u . Observe that the greatest lower bound for the promised value of an unemployed agent who could continue to be unemployed during the following period (with a negligible probability) is the low threshold v_{low} . In the following period, conditional on continuing to be unemployed, this agent will transit to a promised value given by $w_{ouu}(v_{low})$ and exert the maximum search effort s_{max} .²⁰ Subsequently, this agent will forever receive a constant consumption stream (see Figure 2.B) and exert search effort s_{max} whenever they become unemployed again in the future. Since this agent will transit between $w_{ouu}(v_{low})$ and $w_{oue}(w_{ouu}(v_{low}))$ for the rest of their lifetime, these values constitute the greatest lower bounds for the distributions μ_u and μ_e , respectively.²¹

While the greatest lower bound for the support of the distribution μ_u is $w_{ouu}(v_{low})$, Figure 3.B shows that very few unemployed agents reach the area to the left of the low threshold v_{low} . This is because promised values drift down slowly while unemployment lasts (as shown by the vertical distance between the 45-degree line and w_{ouu} in Figure 1.B) and job-finding probabilities increase sharply as promised values approach the lower threshold v_{low} . Observe the large mass that μ_e has at $w_{ye} = -9.7$. This is because 23% of young agents find employment right away, every new generation gets the same w_{ye} , and they accumulate at that value given the long average duration of employment. In contrast, the distribution

²⁰The maximum search effort is defined by $\eta(s_{max}) = 1$.

²¹Since this agent always enjoys the same amount of consumption, the difference between $w_{ouu}(v_{low})$ and $w_{oue}(w_{ouu}(v_{low}))$ merely reflects the search effort s_{max} that is exerted while unemployed.

μ_u gets a much lower mass at $w_{yu} = -18.1$ because if young agents don't find employment right away their promised values drift down. Also observe the rather odd shapes for μ_e and μ_u . These odd shapes arise because of the following: promised values do not change while employed; the average duration of employment is high; promised values drift down slowly during the relatively few unemployment episodes that agents experience during their lifetimes; when they get reemployed their promised values jump by the same large amount; agents get absorbed into retirement or death; and newborns always start their second period of their lives at either w_{ye} or w_{yu} . Thus, there is not enough mixing in the distributions μ_e and μ_u .

Having described the steady-state dynamics of the economy with private information, I turn to evaluate the role of the private information in shaping these dynamics. I do this by comparing this steady-state with the steady-state of the economy with full information (described by equations (17)-(19)) under identical parameter values. Table 1 shows the results. We see that the information frictions have large steady-state effects: They reduce aggregate employment, capital, investment, consumption, and output by 2.4%.²² The intuition for this result is straightforward. Given the constant separation rate, the only way that the planner can generate a high aggregate employment level is by inducing agents to search more intensively. However, since the planner does not observe the search intensity of the agents, this can only be done by increasing the difference between the promised values of becoming employed and the promised values of continuing to be unemployed. That is, the social planner needs to punish the agents by reducing the amount of insurance that they receive. Since agents are risk-averse and the social planner cares about the welfare of its agents, the planner decides to generate a lower aggregate employment level (and with it, lower capital, investment, and output) than in the economy with full information. Despite this, the planner needs to tolerate significant consequences for the amount of consumption heterogeneity: While everybody consumes the same amount in the economy with full information (since they are fully insured), the standard deviation of log consumption becomes 8.3% in the economy with private information.

²²Because the steady-state Euler equation holds in both economies, the effects on employment, capital, investment, consumption, and output must be the same.

6 Business cycle dynamics

In this section I reintroduce the aggregate productivity shocks and analyze the resulting business cycle dynamics. In order to evaluate the role of the information frictions in determining this type of dynamics, I compare the business cycles of the economies with private information and public information. This is done in Figures 4.A through 4.D, which show impulse responses of aggregate employment, output, consumption, and investment, respectively, to a one standard deviation increase in aggregate productivity. The red lines report impulse responses for the economy with private information, while the blue lines report them for the economy with full information. We see that the information frictions play no important role in business cycle dynamics: The impulse responses of all variables are roughly the same in both economies. This is a surprising result, especially after having established that the information frictions play a significant role in steady-state dynamics. However, the intuition for it is quite simple. When a positive productivity shock hits the economy, the planner wants to get people out of unemployment quickly; however, because the search intensity of agents is private information, the planner can only do this by increasing the difference between the promised values of becoming employed and continuing to be unemployed, reducing the amount of insurance provided. Since agents are risk averse, this would be a reason for the social planner not to increase aggregate employment as much as in the economy with full information. However, the welfare losses from providing lower insurance during a boom are compensated by the welfare gains from providing more insurance during a recession. At a first-order approximation, these welfare losses and gains cancel out perfectly and the planner decides to adjust aggregate employment as much as in the economy with full information.

I now investigate the implications of the private information for the optimal amount of inequality over the business cycle. I focus on the cross-sectional variances of two variables: log-consumption and promised values. In order to gain a better picture of their behavior I consider two groups of agents—the employed and the unemployed—and decompose the total variance of each variable into a within-groups variance and a between-groups variance. I further decompose the within-groups variance into the variance of employed agents and the variance of unemployed agents. More specifically, for each variable x considered, I decompose

the total cross-sectional variance σ_t^2 into the following terms:

$$\sigma_t^2 = H_t \sigma_{et}^2 + (1 - H_t) \sigma_{ut}^2 + \sigma_{between,t}^2,$$

where H_t is the fraction of the population that is employed, σ_{et}^2 is the variance of x within employed agents, σ_{ut}^2 is the variance of x within unemployed agents,

$$\sigma_{between,t}^2 = H_t (\bar{x}_{et} - \bar{x}_t)^2 + (1 - H_t) (\bar{x}_{ut} - \bar{x}_t)^2,$$

and \bar{x}_{et} , \bar{x}_{ut} , and \bar{x}_t are the mean value of x within employed agents, within unemployed agents, and in the total population, respectively.

Figure 5.A shows the impulse responses of σ_t^2 , σ_{et}^2 , σ_{ut}^2 , and $\sigma_{between,t}^2$ for the logarithm of consumption. These variances are depicted in levels to show their relative importance. We see that the variance between groups is always small and that the variance of employed agents is larger than the variance of unemployed agents. The reason why the variance of log-consumption levels is lower for unemployed agents than for employed agents is that unemployed agents get compensated for having a significant variance in their search intensities. Figure 5.B is similar to Figure 5.A, but for promised values. We see that the variance of promised values between groups is again small, but now the variance of promised values for unemployed agents exceeds the one for employed agents (reflecting that the variance in promised values includes variability in search intensities).

The changes to the variances in Figures 5.A and 5.B turn out to be so small that they are hardly noticeable in levels. In order to make these changes more transparent, Figures 6.A and 6.B show them as percentage differences from their steady-state values (with a common y-axis to facilitate comparisons). Figure 6.A shows that the variance of log-consumption between groups decreases quite significantly during the first three years after the aggregate shock hits the economy, but then stabilizes and starts to increase, surpassing its initial level by the ninth year. The reason for the initial decrease is that in order to induce unemployed agents to search more intensively, the planner needs to increase the differences between the promised values of becoming employed and continuing to be unemployed. As a result, while the average log-consumption levels of both groups increase with the aggregate shock, the planner increases the average log-consumption of employed agents more than the average log-consumption of unemployed agents. Since the steady-state average log-consumption level

of employed agents is lower than for unemployed agents, this decreases the between-groups variance.²³ After aggregate employment reaches its peak the unemployed agents start to decrease their search intensities, allowing the planner to start increasing the average log-consumption level of unemployed agents relative to that of employed agents. Contrary to the variance between groups, the variances of log-consumption within the groups of employed and unemployed agents are hardly affected. After a drop of 0.1% on impact, both of these variances start to increase very slowly, reaching a total increase of 0.4% by the tenth year after the aggregate productivity shock. The reason why the within group log-consumption variances increase slowly is that the social planner chooses decision rules that induce unemployed agents to search more early on and less later on, but this translates into significantly more heterogeneity only when agents experience different idiosyncratic histories (which, given the persistence of the labor market shocks, takes considerable time to realize).

Figure 6.B shows that the component of the variance of promised values most affected by the aggregate productivity shock is the variance between groups. This variance increases on impact because, while the average promised values of both groups increase with the positive aggregate productivity shock, initially the average promised value to unemployed agents increases much less because they are required to search more. Given that at the steady state unemployed agents have a lower average promised value than employed agents, this difference increases the between-groups variance.²⁴ Subsequently, unemployed agents are allowed to reduce their search intensities and, consequently, start to catch up to the increase in the average promised value of employed agents (eventually surpassing this increase and reducing the between-groups variance). The second largest proportionate effects in Figure

²³The reason why the average log-consumption level is lower for employed agents than for unemployed agents is straightforward: To the right of the zero-search threshold employed agents escape their state at the exogenous separation rate, but unemployed agents get absorbed into it (effectively retiring from the labor force). Because the resulting mass of agents to the right of the zero-search threshold is larger for the unemployed than for the employed, because both groups of agents consume the same in that region, and because consumption levels are the largest in that region, the average log-consumption level ends up being higher for unemployed agents than for employed agents.

²⁴The steady-state differences in the average promised values of both groups is dominated by the fact that unemployed agents have to search while employed agents don't.

6.B correspond to the variance of promised values of unemployed agents. We see that this variance decreases 1% on impact. The reason is that, because promised values can be made contingent on the realization of the aggregate shock, the social planner decides to compress the unemployed agents' distribution before providing them with more incentives to search.²⁵ After the initial period, the variance of promised values of unemployed agents increases as the promised values of those who do not find employment start to drift down quickly early on (when the search intensities increase) and as the promised value of agents who become unemployed later on (when the search intensities decrease) become relatively larger. Ten years after the aggregate productivity shock has hit the economy, the variance of promised values to unemployed workers is 4.5% above its initial level. The smallest proportionate effects in Figure 6.B correspond to the variance of promised values to employed agents. Given that these promised values are state-contingent, the planner also decides to decrease their dispersion on impact in anticipation that some of these agents will be transiting to unemployment in the near future (when they will be receiving less insurance). However, the impact effect is small: It decreases only by 0.3%. After the initial period, the variance of promised values to employed agents starts to increase slowly as the variance of their log-consumption levels goes up. By the tenth year after the initial shock, the variance of promised values to employed agents has become slightly higher than its initial steady-state level.

In order to complete the analysis of the optimal provision of unemployment insurance over the business cycle I evaluate the behavior of consumption replacement ratios. To this end I perform the following experiment. I consider two agents, named E and U, who were both employed the period before the aggregate shock hit the economy (i.e., period 0 in Figures 4-6). Moreover, I assume that in that period both agents had identical promised values given by the steady-state average promised value of employed agents. Once the aggregate productivity shock hits the economy, the employment histories of both agents start to diverge. In particular, agent E remains continuously employed, while agent U becomes unemployed in period 1. In what follows I will compare the consumption levels

²⁵This is also the reason why the variance of log-consumption within unemployed agents slightly decreases on impact in Figure 6.A.

of both agents under different possible reemployment scenarios for agent U. The gray line in Figure 7 reports the consumption level of agent U relative to agent E assuming that U remains continuously unemployed during the following ten years. We see that the first period that agent U becomes unemployed they receive a consumption replacement ratio of 100%. However, this replacement ratio drops to 95% after five years of unemployment and to 76% after ten years. While 76% is a low replacement ratio, it is important to keep in mind that the ex-ante probability of agent U reaching ten years of continuous unemployment is rather small: only 1.1%. The gold line in Figure 7 shows the consumption replacement ratio of agent U during the initial period of time when they become reemployed. This replacement ratio is shown for each possible period of reemployment after becoming unemployed in period 1. We see that if agent U becomes reemployed after only one period, they receive an initial consumption level that is 4% higher than agent E's. If agent U's reemployment occurs five years after becoming unemployed they receive roughly the same consumption level as agent E. However, if agent U's reemployment happens ten years after becoming unemployed they consume only 85% as much.

While the previous consumption replacement ratios were calculated after the economy received a positive aggregate productivity shock, it is important to compare them with the replacement ratios that are obtained in the absence of an aggregate shock (i.e., at the deterministic steady state). The reason is that this will provide information of how the optimal replacement ratios should change over the business cycle. This is done in Figure 7, where the blue line reports consumption replacement ratios while continuing to be unemployed and the orange line does the same after reemployment, both under the assumption that the aggregate productivity level z_t stays at zero. We see that both lines are almost completely hidden by the corresponding lines obtained under the aggregate productivity shock (only after seven years do the lines become somewhat distinct). We conclude that optimal consumption replacement ratios are approximately independent of the state of the business cycle.

7 Conclusions

In this paper I evaluated how the presence of moral hazard in unemployment insurance may affect optimal aggregate dynamics. To this end I embedded the Hopenhayn-Nicolini optimal unemployment insurance model into an RBC model with search frictions and solved the associated mechanism design problem. I found that private information in the search intensities of agents has large effects on the optimal steady-state dynamics of the model, decreasing employment, consumption, investment, and output by 2.4%. However, the effects on the optimal aggregate business cycle fluctuations are negligible. In terms of the design of the optimal unemployment insurance system, I found that optimal consumption replacement ratios are roughly independent of the state of the business cycle. Moreover, the cross-sectional variability of log consumption levels is roughly constant over time.

For simplicity's sake, I considered a version of the Hopenhayn-Nicolini unemployment insurance model in which all separations from the production island were due to an exogenous separation shock. As a consequence, agents were fully insured when they became unemployed and their subsequent consumption levels did not depend on the length of their last employment spell. However, if agents suffered disutility from work and the social planner could not distinguish between layoffs and quits, agents may want to separate opportunistically in order to collect generous unemployment insurance. In this scenario the optimal unemployment insurance system would not fully insure workers when they become unemployed and the benefits received would depend on the length of their last employment spell. Hopenhayn and Nicolini (2009) have already considered this realistic case within a principal agent setting, and it would be extremely interesting and perfectly feasible to incorporate it to an RBC model following what's been described in this paper. Such an extension is left for future research.

In this paper I compared the optimal allocation of an economy with private information with its optimal allocation under public information. While this allowed me to isolate the effects of the private information, it would be useful to compare the economy with private information with other benchmarks. In particular, comparing it with an economy with realistic financial markets and labor market policies would help determine how far actual economies may be from their constrained social optima. This is also left for future research.

References

- ALVAREZ, F. AND M. VERACIERTO, “Severance payments in an economy with frictions,” *Journal of Monetary Economics* 47 (2001), 477–498.
- ATKESON, A. AND R. E. LUCAS, “On Efficient Distribution With Private Information,” *Review of Economic Studies* 59 (1992), 427–453.
- BOOSTANI, R., M. GERVAIS AND L. WARREN, “Optimal Unemployment Insurance in a Directed Search Model,” mimeo, University of Iowa, 2017.
- DA COSTA, C. AND V. F. LUZ, “The Private Memory of Aggregate Uncertainty,” *Review of Economic Dynamics* 27 (January 2018), 169–183.
- FARHI, E. AND I. WERNING, “Capital Taxation: Quantitative Explorations of the Inverse Euler Equation,” *Journal of Political Economy* 120 (2012), 398 – 445.
- GOLOSOV, M., M. TROSHKIN AND A. TSYVINSKI, “Redistribution and Social Insurance,” *American Economic Review* 106 (February 2016), 359–386.
- GOLOSOV, M., A. TSYVINSKI AND I. WERNING, “New Dynamic Public Finance: A User’s Guide,” in *NBER Macroeconomics Annual 2006, Volume 21* NBER Chapters (National Bureau of Economic Research, Inc, 2007), 317–388.
- GREEN, E. J., “Lending and the Smoothing of Uninsurable Income,” in E. C. Prescott and N. Wallace, eds., *Contractual arrangements for intertemporal trade* (Minneapolis: University of Minnesota Press, 1987), 3–25.
- HOPENHAYN, H. A. AND J. P. NICOLINI, “Optimal Unemployment Insurance,” *Journal of Political Economy* 105 (April 1997), 412–38.
- , “Optimal Unemployment Insurance and Employment History,” *Review of Economic Studies* 76 (2009), 1049–1070.
- JUNG, P. AND K. KUESTER, “Optimal Labor-Market Policy in Recessions,” *American Economic Journal: Macroeconomics* 7 (April 2015), 124–156.

- KROFT, K. AND M. J. NOTOWIDIGDO, “Should Unemployment Insurance Vary With the Unemployment Rate? Theory and Evidence,” NBER Working Papers 17173, National Bureau of Economic Research, Inc, June 2011.
- KRUSELL, P., T. MUKOYAMA, R. ROGERSON AND A. AHIN, “Gross Worker Flows over the Business Cycle,” *American Economic Review* 107 (November 2017), 3447–3476.
- LANDAIS, C., P. MICHAILLAT AND E. SAEZ, “A Macroeconomic Approach to Optimal Unemployment Insurance: Theory,” *American Economic Journal: Economic Policy* 10 (May 2018), 152–181.
- LJUNGQVIST, L. AND T. J. SARGENT, “Understanding European unemployment with matching and search-island models,” *Journal of Monetary Economics* 54 (2007), 2139–2179.
- MITMAN, K. AND S. RABINOVICH, “Optimal unemployment insurance in an equilibrium business-cycle model,” *Journal of Monetary Economics* 71 (2015), 99–118.
- PHELAN, C., “Incentives and Aggregate Shocks,” *Review of Economic Studies* 61 (1994), 681–700.
- SANCHEZ, J. M., “Optimal state-contingent unemployment insurance,” *Economics Letters* 98 (March 2008), 348–357.
- SCHEUER, F., “Optimal Asset Taxes in Financial Markets with Aggregate Uncertainty,” *Review of Economic Dynamics* 16 (July 2013), 405–420.
- SHIMER, R. AND I. WERNING, “Liquidity and Insurance for the Unemployed,” *American Economic Review* 98 (December 2008), 1922–1942.
- VERACIERTO, M., “Computing Equilibria of Stochastic Heterogeneous Agent Models Using Decision Rule Histories,” Working Paper Series WP 2020-05, Federal Reserve Bank of Chicago, February 2020.
- , “Business cycle fluctuations in Mirrlees economies: The case of i.i.d. shocks,” *Journal of Economic Theory* 196 (2021).

VERACIERTO, M. L., “Plant-Level Irreversible Investment and Equilibrium Business Cycles,” *American Economic Review* 92 (2002), 181–197.

WERNING, I., “Optimal Unemployment Insurance with Unobservable Savings,” mimeo, MIT, 2002.

———, “Optimal Fiscal Policy with Redistribution,” *The Quarterly Journal of Economics* 122 (2007), 925–967.

WILLIAMS, N. AND R. LI, “Optimal unemployment insurance and cyclical fluctuations,” FRB Atlanta CQER Working Paper 2015-2, Federal Reserve Bank of Atlanta, August 2015.

Table 1

Steady state effects

Information	Y	C	I	H	K	$\sigma(\ln c)$
Public	100.0	100.0	100.0	100.0	100.0	0.0
Private	97.6	97.6	97.6	97.6	97.6	0.083

Figure 1: Job finding probabilities and promised values

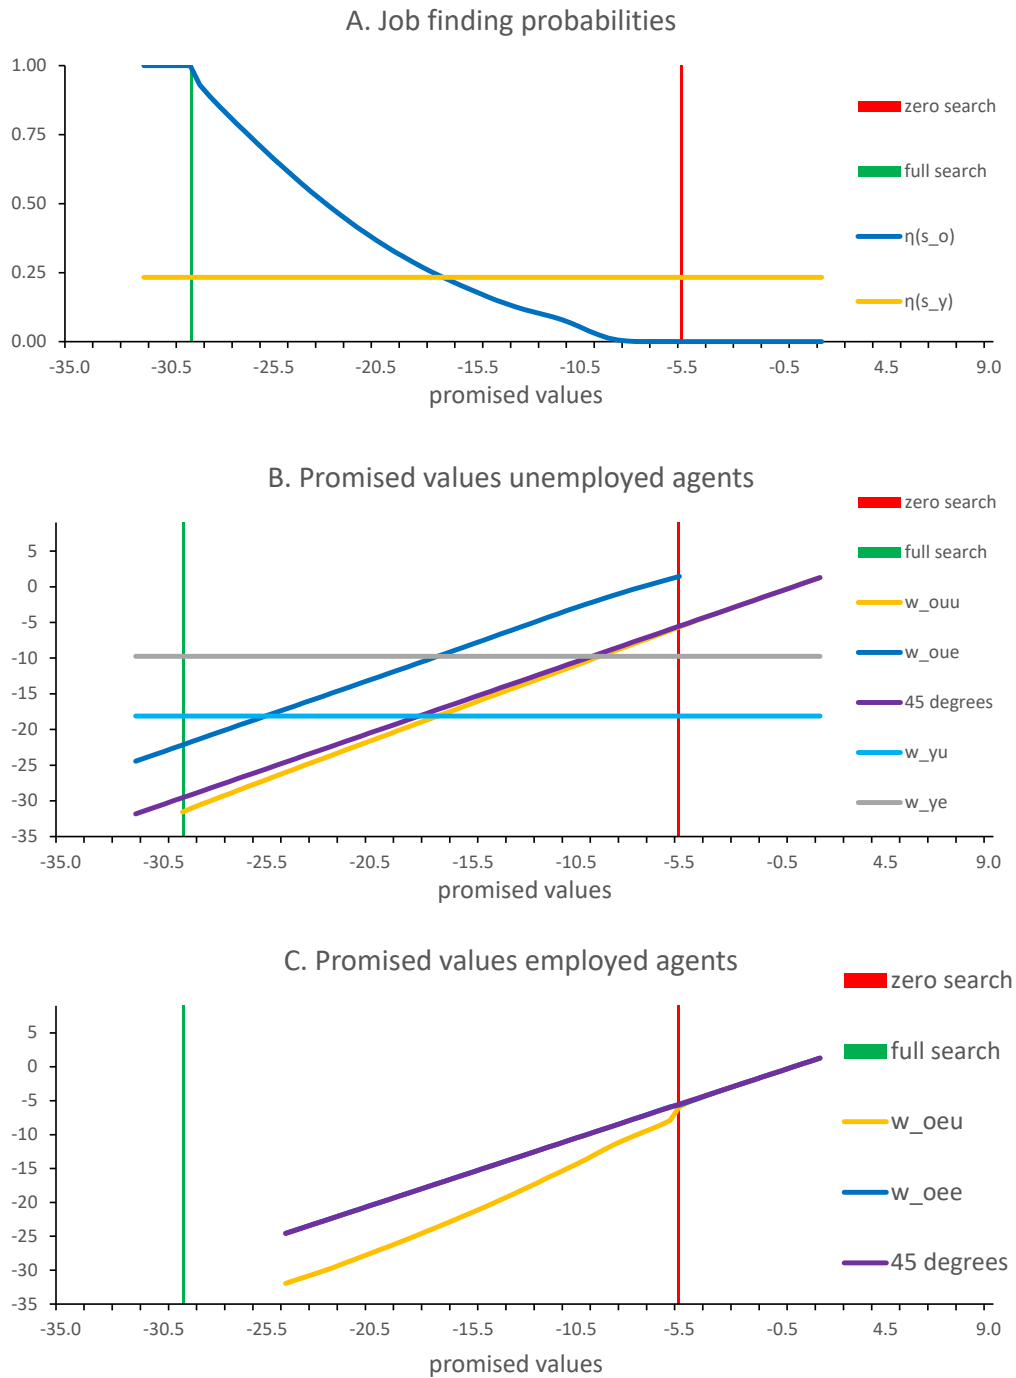


Figure 2: Log consumption levels and changes

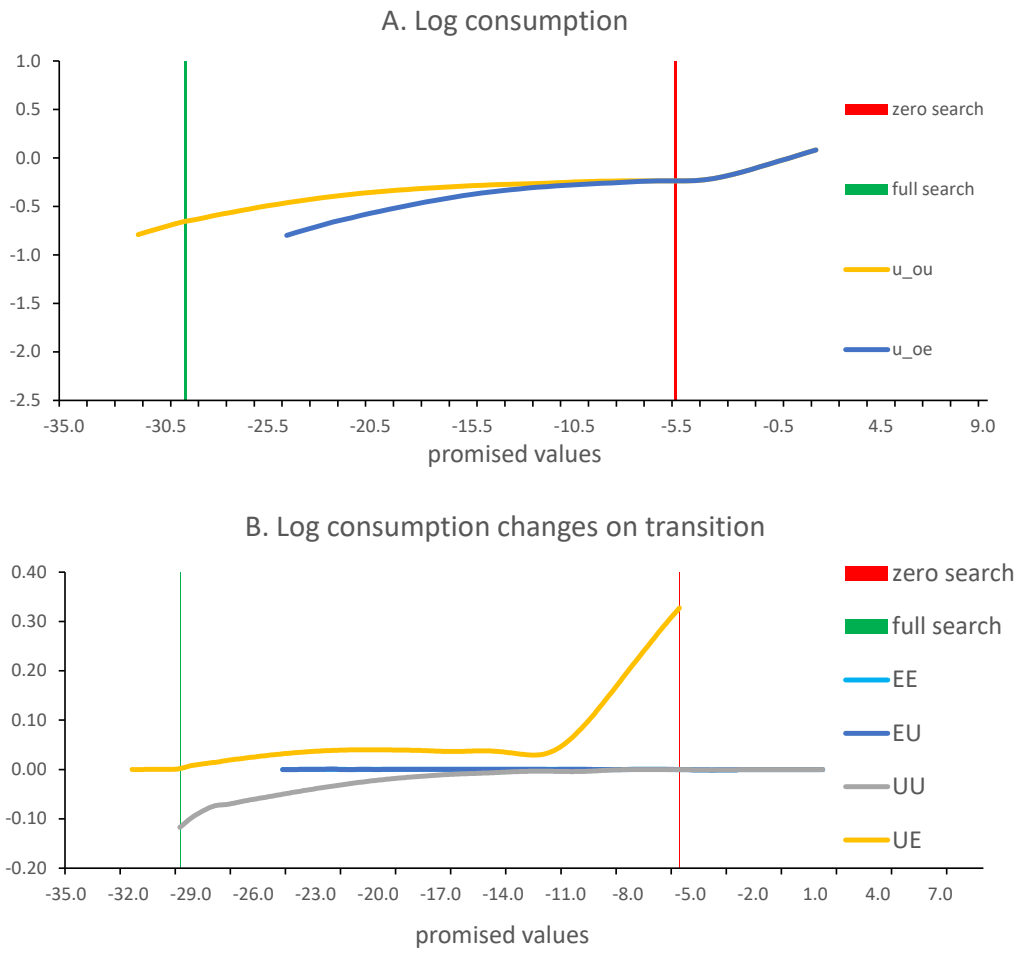



Figure 3: Invariant distributions

A. Distribution employed agents



B. Distribution unemployed




Figure 4: Impulse responses

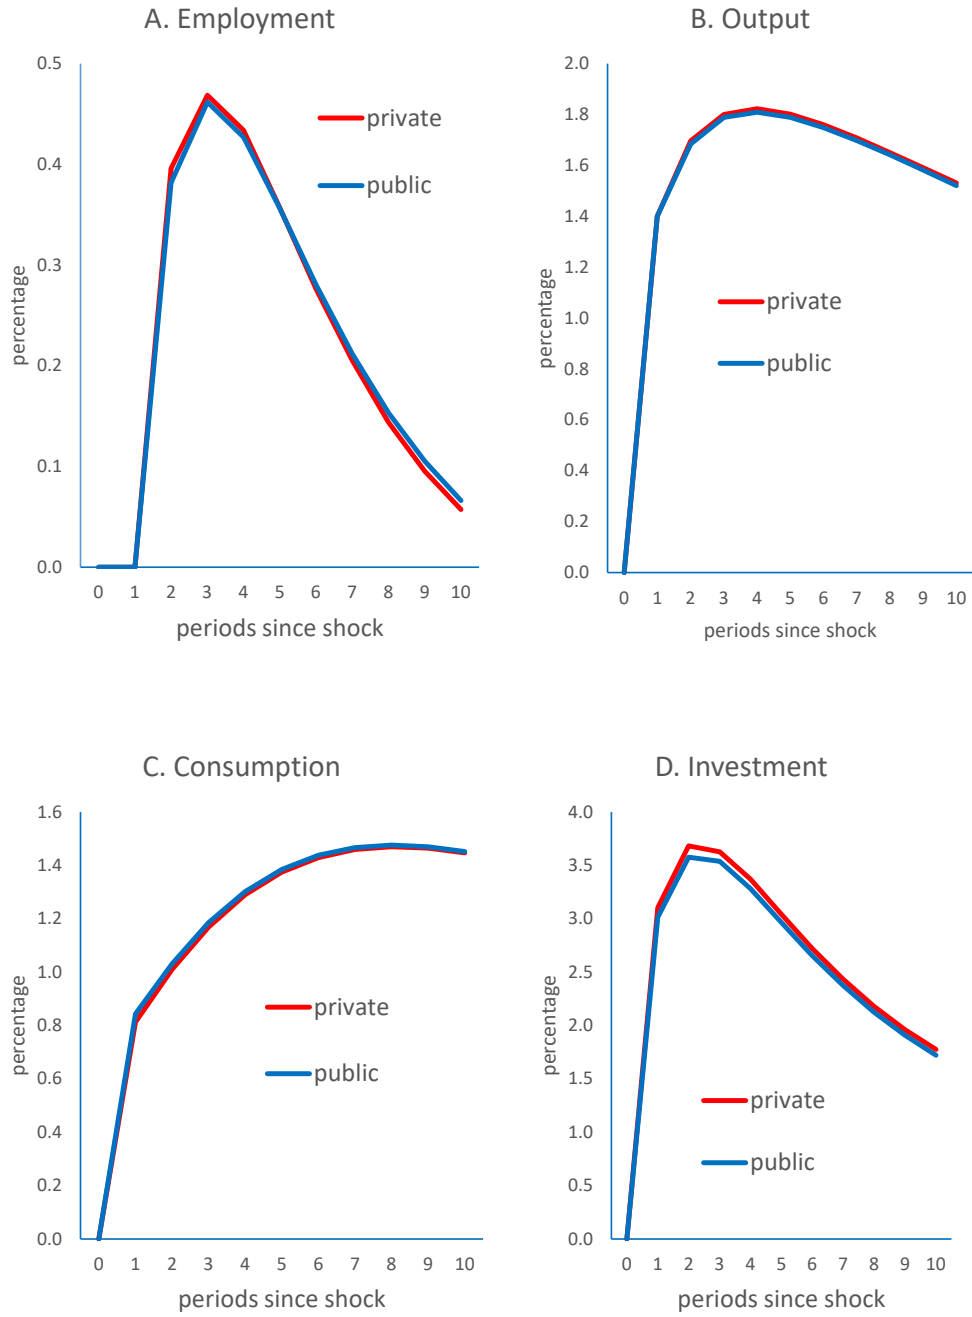


Figure 5: Variances of log consumption and promised values (levels)

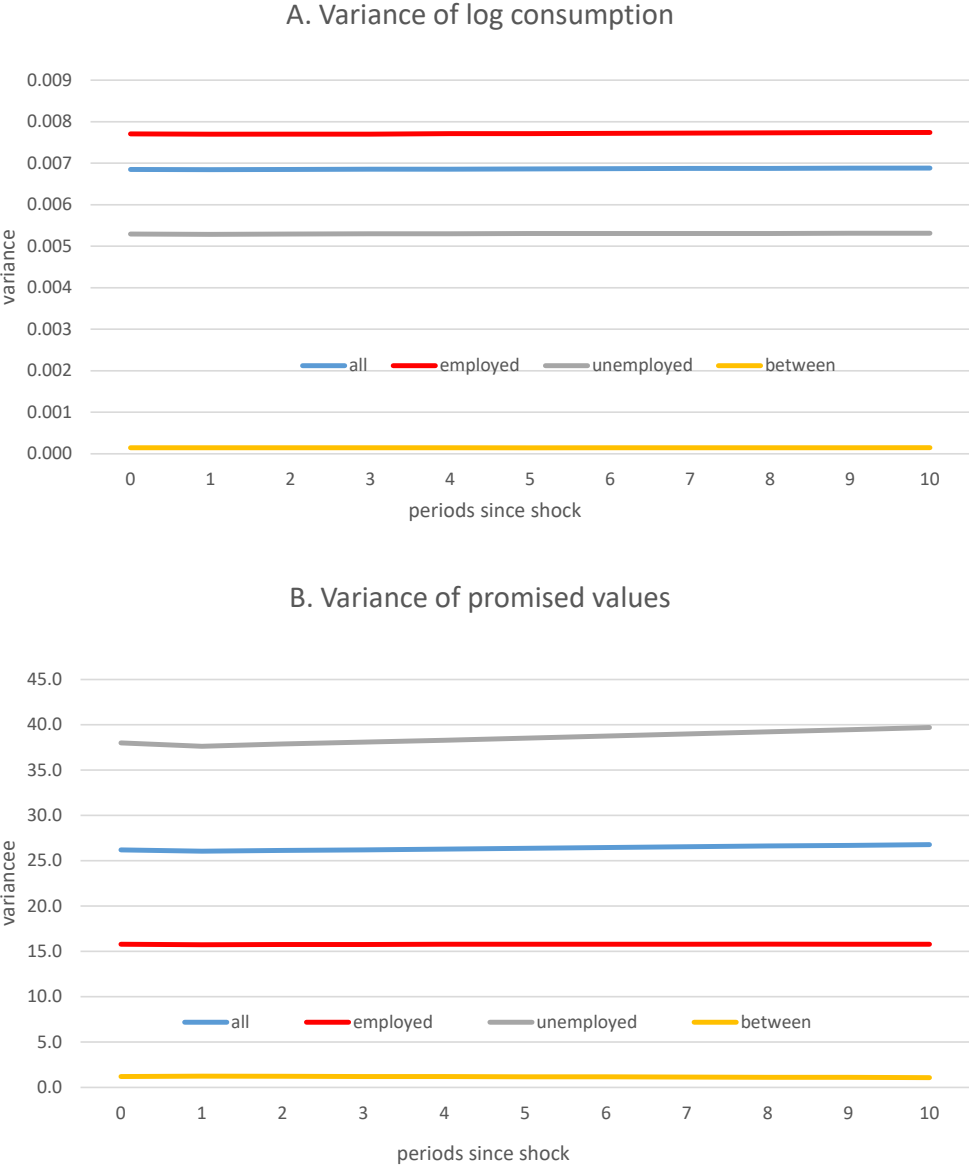


Figure 6: Variances of log consumption and promised values (% deviations from trend)



Figure 7: Consumption replacement ratios

