Aggregate Consumption in Times of Crisis: The Role of Financial Frictions

Edouard Challe, Xavier Ragot

Aggregate Consumption in Times of Crisis: The Role of Financial Frictions[1]

Edouard Challe* and Xavier Ragot[y]

*Economics Department, Ecole Polytechnique, 91128 Palaiseau, France; Banque de France, Paris, France, e-mail: edouard.challe@polytechnique.edu
[y]Banque de France, 39 rue des Petits Champs, 75001 Paris, France; PSE, Paris, France,

e-mail: xavier.ragot@banque-france.fr

Abstract

The joint behaviour of US aggregate consumption and saving over the period 2007–2009, and notably the pronounced U-shaped pattern of consumption together with the rise in saving, are difficult to reconcile with the view that financial markets are frictionless. We propose an alternative framework in which financial markets are incomplete and where households form a buffer stock of precautionary saving to self-insure against the (time-varying) risk of falling into unemployment, with the consequence of considerably amplifying and propagating crises. Our model can be solved in closed form because the wealth heterogeneity generated by uninsured income shocks remains minimal. We end the article by arguing that fully incorporating uninsured and time-varying individual risks into macroeconomic analysis may drastically alter our understanding of the business cycle, macroeconomic policy, and the role of financial intermediaries. (JEL code: E21)
Keywords: Incomplete markets, unemployment, consumption.

1 Introduction

Among the pillars of mainstream macroeconomic analysis shaken by the ongoing financial and economic crisis, perhaps that which has attracted most attention is the assumption that financial markets work perfectly. However, while few would now doubt that financial 'frictions' are important, it is still unclear which specific frictions have first-order implications and how to model them in a both realistic and tractable manner. In this article, we argue that a central such friction is that households face considerably more uninsurable income risk (basically, unemployment risk) than what is posited in the standard macroeconomic models. We then offer an example of a tractable model in which the implications of this friction can easily be analysed and use it to shed light on the recent evolution of US aggregate consumption and saving.

The basic motivation for our approach is the observation, reported in cross-sectional studies of household-level consumption and saving plans, that the latter depend substantially on idiosyncratic labour income risk,

---

and hence on the extent of the unemployment risk faced by individuals.[2] As we illustrate and discuss in Section 2, one potential explanation for the large and mean-reverting decline in US aggregate consumption that occurred from the mid-2008 onwards, and the contemporaneous rise in aggregate saving and the saving rate, is that individuals have cut spending in order to (re-)constitute a 'buffer stock' of precautionary wealth. This view is borne out by the rapid rise in the probability of involuntary layoffs that occurred at the same time, since the risk of falling into unemployment is the largest source of idiosyncratic income risk faced by most individuals. Inasmuch as it is only imperfectly insured, it is likely to lead to a substantial demand for 'self-insurance' via asset accumulation.

Although the effects of unemployment risk and the precautionary motive for holding assets may have first-order effects on aggregate fluctuations, they are particularly uneasy to model using current macroeconomic tools. For one thing, much of the theoretical/quantitative analysis of the business cycle is based on the complete-market, representative agent assumption. In this setup, a transitory increase in unemployment amounts to a transitory fall in aggregate labour income, and such changes have, by their very transitory nature, a limited impact on human wealth (i.e. the discounted value of current and future labour income flows). Then, 'permanent-income' behaviour implies that the representative agent should change consumption only by a small amount, with the cyclical variation in income being absorbed by large and procyclical variations in aggregate savings.

Sections 3 and 4 spell out a simple model of the precautionary motive, which is consistent with the observed U-shaped pattern of consumption and the persistent rise in savings. In our model, the risk of falling into unemployment is not socialized and entirely falls on every household's shoulders. Then, rising unemployment risk leads to a strengthening of the precautionary motive for holding assets, and the increase in aggregate savings that follows considerably amplifies the fall in current consumption, relative to that implied by the complete markets model. Although this amplifying mechanism differs from the textbook Keynesian explanation based on price rigidity, it is reminiscent of the Keynesian argument according to which the depth of recessions is due to large and persistent shortages of aggregate demand.[3]

---

[2] For example see Cochrane (1991), Carrol (1992), Guiso et al. (1996), and Gourinchas and Parker (2002). Also see Jappelli (1990) for direct evidence that a share of US households are liquidity-constrained, and hence cannot freely smooth out income fluctuations via borrowing.

[3] In our model (and as in much of the incomplete-market literature), the uninsured labour income risk that is associated with labour-market transitions is of purely transitory nature. In Krebs (2007), in contract, unemployment risk generates permanent earning losses (scars) for displaced workers. Introducing permanent income losses into our framework could only strengthen the precautionary motive for holding wealth.

The complete insurance assumption is still widely used in macroeconomics, even when labour market outcomes are involved. In particular, the joint assumption of complete insurance markets together with labour market (search) frictions were originally introduced by Merz (1995) and Andolfatto (1996), and subsequently used by Den Haan et al. (2000) and Walsh (2005), among many others. This class of models typically studies the dynamics of unemployment assuming either that households are risk-neutral, or that agents fully diversify unemployment risk within large 'families'. While these models yield important insights into the workings of the labour market and the rate of frictional unemployment, they have little to say about the effects of unemployment risk on the counter-cyclicality of savings and the strong procyclicality of consumption that we saw in the recent period.

An important class of models that dispenses with the complete markets assumption is that of 'heterogenous agents' models, in which uninsured idiosyncratic income risk coupled with borrowing constraints generate large-dimensional heterogeneities in individual asset holdings (Bewley 1983; Aiyagari 1994; Krusell and Smith 1998). While these models have proven extremely useful for sheding light on cross-sectional inequalities in income and wealth, their applicability to the study of the business cycle and macroeconomic policy has thus far remained limited. In these models indeed, the whole history of employment statuses contribute to the determination of any agent's wealth and infinitely many types of agents, each of whom having their own optimal consumption and saving plans, asymptotically co-exist in the economy. Because of the technical and computational difficulties involved when handling large-dimensional cross-sectional distributions, the first generation of heterogenous-agent models focused on stationary environments in which idiosyncratic labour income risk, while present, was assumed to be time-invariant. Variants of this approach include the prominent infinite-horizon models of Huggett (1993, 1997) and Aiyagari (1994), as well as overlapping generations models as in Huggett (1996) and Heer (2001). Krusell and Smith (1997, 1998) were the first to introduce aggregate shocks and time-varying labour income risk into this framework, but again computational limitations have narrowed the scope of issues that can be investigated.[4] In particular, the time-varying nature of labour income risk is in general limited to two or three states, and the continuous changes in unemployment risk that takes place along the business cycles as well as its gradual impact on precautionary saving cannot be studied explicitly.
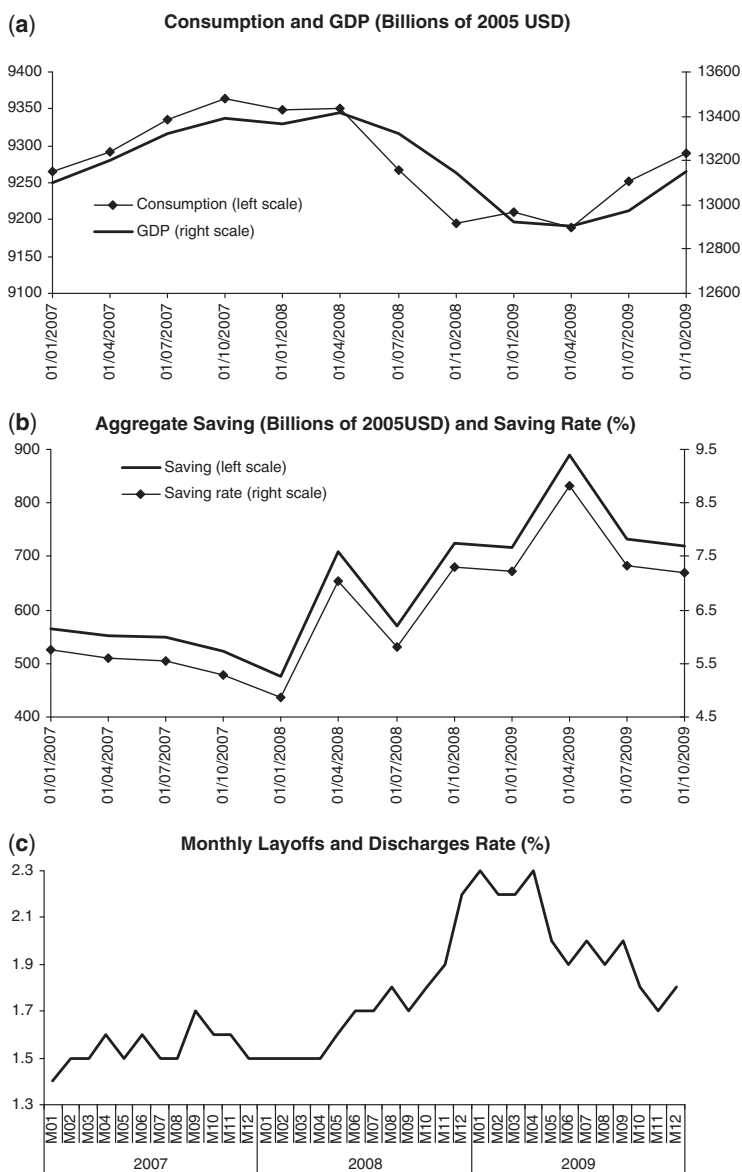
---

[4] Also see Heathcote (2005) on the effect of aggregate uncertainty about taxes in the infinite-horizon framework, and Storesletten et al. (2007) for a stochastic overlapping-generations asset-pricing model.

Our model differs from this latter approach by focusing on a class of equilibria that endogenously generates a limited (in fact, minimal) degree of cross-sectional heterogeneity, and in which the interactions between idiosyncratic and aggregate shocks can straightforwardly be analysed. The particularity of our model is that it can be solved with paper and pencil and is thus liable to a variety of theoretical and quantitative explorations. Formally, this model belongs to broader class that we have been developing and which relies on reducing the cross-sectional distribution of heterogenous households from an infinite to a finite number of types (see Algan et al. 2009, forthcoming; Challe and Ragot 2010; Challe et al. 2010 for applications of this framework pertaining to asset pricing and macroeconomic policy issues). More specifically, the class of equilibria that we focus on is based on two assumptions. The first is that unemployed households always end up facing a binding borrowing constraint if they stay unemployed for a sufficiently large number of periods. The second assumption is that the marginal utility of consumption is constant above a certain threshold consumption level, which can be seen as an extreme (but particularly simple) form of decreasing relative risk aversion. While we focus of the simplest equilibrium having this joint property here (i.e. one in which the constraint is binding from the very first period of unemployment while marginal utility is constant for all high-income agents), it is important to keep in mind that the framework can handle larger (but, crucially, finite) cross-sectional wealth distributions.

Section 5 discusses some alternative potential applications of our framework and, more generally, how incorporating the large and time-varying idiosyncratic income uncertainty that agents face may alter our understanding of the business cycle, macroeconomic policy, and the role of financial intermediaries.

2 Consumption, saving, and unemployment risk over the period 2007–2009

The present Section summarizes the evidence that motivates the theoretical framework developed below. Figure 1a displays the joint behaviour of personal consumption expenditures and gross domestic product (GDP) in the US over the years 2007–2009. Figure 1b shows the behaviour of aggregate saving and the saving rate over the same period (all four series are from the National Income and Product Accounts). Finally, Figure 1c is the Layoffs and Discharges Rate, extracted from the Job Openings and Labor Turnover Survey of the Bureau of Labor Statistics, which we take as a proxy for the unemployment risk faced by US households.

(a) **Consumption and GDP (Billions of 2005 USD)**

Consumption (left scale)
GDP (right scale)

(b) **Aggregate Saving (Billions of 2005USD) and Saving Rate (%)**

Saving (left scale)
Saving rate (right scale)

(c) **Monthly Layoffs and Discharges Rate (%)**

2007   2008   2009

Figure 1 Consumption, saving, GDP, and unemployment risk. Source: BEA/ NIPA (2007Q1-2009Q4) and BLS/JOLTS (2007M1-2009M12). Note: 'Consumption' denotes personal consumption expenditures, and 'Saving' is defined as disposable personal income minus personal consumption expenditures (and divided by disposible personal income to obtain the 'Saving Rate'). The 'Layoffs and Discharges Rate' is that of the private sector.

The behaviour of aggregate consumption displays three striking features. First, the drop in consumption that occurred from the mid-2008 onwards is large, both as a share of consumption before the shock and proportionally to aggregate output. Second, it is U-shaped, with consumption starting recovering about a year after its initial fall. Third, consumption closely tracks GDP over this period. Figure 1b illustrates that the fall in aggregate consumption has been associated with a substantial rise in aggregate savings.

This pattern of consumption and saving appears puzzling from the point of view of the frictionless, representative agent model. First, the strong motive for intertemporal consumption smoothing predicted by that model implies that aggregate consumption should react strongly to permanent changes in aggregate (labour and capital) income but relatively little to transitory income changes—i.e. the representative agent behaves as a 'permanent income' consumer. Hence, the baseline representative agent model can be made consistent with the large fall in aggregate consumption that took place in the second half of 2008 only to the extent that permanent income has itself fallen sharply over this period. Second, and for the very same reason, the U-shaped pattern of aggregate consumption over the period can be reconciled with permanent income behaviour only if perceived permanent income has itself been subject to very large revisions. Only equally large news—first downwards in the mid-2008, then upwards from the mid-2009 onwards—could have generated such shifts.

The explanation for this pattern that we pursue here is that households have engaged in strong precautionary saving behaviour. In short, the unemployment risk faced by US households has risen considerably over the period, and they have responded to this greater uncertainty by cutting their consumption spending in order to (re-)constitute a buffer stock of precautionary wealth. By definition, an increase in the precautionary saving motive implies that, for a given level of income, agents substitute current consumption for current saving, so the former may fall while the latter rises. Moreover, the precautionary motive is stronger when labour income risk increases, so this explanation is consistent with the observed joint behaviour of aggregate consumption and the layoffs and discharges rate—see Figure 1c.[5] Third, since the precautionary saving motive is

[5] We show the layoffs and discharges rate rather than the overall separations rate because the former is arguably a more accurate indicator of the risk of unemployment perceived by individuals. Since the beginning of 2007, the separation rate has decreased mildly, but this resulted from the composition of layoffs and discharges (which have risen sharply, as we document) and voluntary quits (which have fallen). See Shimer (2005) for historical evidence on the separation and job-finding rates.

greater during recessions (because so is individual unemployment risk), it amplifies the consumption fall that would prevail without it, possibly to the point of rendering aggregate saving counter-cyclical. Finally, the counter-cyclicality of individual unemployment risk and its impact on precautionary saving imply that changes in aggregate income that are purely transitory (and fully anticipated as such) may bring about an equally cyclical response of consumption; this provides a natural explanation for the U-shaped pattern of consumption over the period and its close tracking of GDP. Section 3 develops a simple model of precautionary saving behaviour that is consistent with this view.

3 A simple model of precautionary saving behaviour

We consider the behaviour of a perfectly competitive, open economy facing the world gross interest rate $R \geq 1$.[6] Firms produce the unique output good thanks to the constant-return to scale technology $Y_t = F(K_t, N_t) = N_t f(k_t)$, where F is a neoclassical production function and $k_t = K_t/N_t$. Capital depreciates at rate δ. The optimal demand for inputs under perfect competition and constant returns to scale implies that k_t is time-invariant and equal to

$$k = f'^{-1}(R - 1 + \delta). \tag{1}$$

Since the capital/labour ratio is constant, so is the equilibrium real wage w, which is

$$w = f(k) - kf'(k). \tag{2}$$

Our assumption that the interest rate (and hence the wage rate) is constant is clearly unrealistic. However, it greatly simplify the analysis, and is hopefully relatively innocuous provided that unemployment risk, rather than changes in the interest rate, is the primary driver of the precautionary motive.

Total population is normalized to one. Every household is endowed with one unit of labour, which is supplied inelastically to firms provided that the household is employed; the labour income of the employed is thus w. From the second period of unemployment onwards, households become 'home producers' and earn the fixed income $\lambda \in (0,w)$. In every

---

period, employed household have a (time-varying) probability $\alpha_t \in (0,1)$ of staying employed in the next period, and unemployed households have a (constant) probability $\rho \in (0,1)$ of staying unemployed in the next period. While it would be straightforward to introduce time-variations in the job-finding rate $1-\rho$ (with little changes in our results), we focus on changes in $1-\alpha_t$ here, which we interpret as the involuntary layoffs rate and which, as shown above, has undergone large and persistent variations since the beginning of the current crisis. Note that while we take these transitions as exogenous and *ad hoc*, they are obviously not so in the real world. These probabilities are, thus, best interpreted as reflecting unmodelled frictions in job creation and destruction as well as in the uncoordinated process of search that takes place between workers and firms.

The crucial feature of our model is that agents have limited ability to insure against unemployment risk. For simplicity, we make the extreme assumption that unemployment insurance, either private or public, is infeasible. Here again, the qualitative properties of the model would be similar if we assumed that some (but imperfect) insurance was available. Second, we assume that unemployed agents cannot borrow against future income. The joint assumption of market incompleteness and borrowing constraints implies that agents falling into unemployment cannot limit their current income fall by collecting insurance payments or finance current consumption against their (expectedly higher) future income. This creates a precautionary saving motive for the employed whose intensity depends on the extent of unemployment risk.

The budget constraint of household i is

$$c_t^i + a_t^i = y_t^i + a_{t-1}^i R,$$

where c_t^i denotes consumption, a_t^i asset holdings, and y_t^i current income, with $y_t^i = w$ if the household is employed, $y_t^i = 0$ if the household has just fallen into unemployment, and $y_t^i = \lambda$ otherwise. Household i is assumed to maximize

$$E_t \sum_{t=0}^{\infty} \beta^t u(c_t^i), \text{ with } u(c) = \begin{cases} \gamma c^* \ln(c) & \text{if } c \leq c^*, \\ \Omega + \gamma c & \text{if } c > c^*, \end{cases}$$

and $\Omega \equiv \gamma c^* (\ln(c^*) - 1)$, $c^* > 0$, $\beta \in (0, 1)$. The instant utility function $u(.)$ is parameterized by (c^*, γ). It is increasing and continuously differentiable over $(0, +\infty)$, strictly concave up c^* and linear with slope γ at and above c^*. The linearity of $u(.)$ above c^* is crucial for the construction of our equilibrium with limited households heterogeneity, for it implies that above a certain level of consumption all households share the same marginal utility and hence the same demand for assets. The concavity of $u(c)$ below c^* is also essential because employed households must dislike large

consumption falls if they are to form any precautionary saving in the first place. Finally, we assume that β and R are jointly sufficiently low so that

$$\beta R < 1 \qquad (3)$$

Assumption (3) ensures that some agents may wish to borrow and hence may face a binding borrowing constraint, and is necessary for the equilibrium that we study to exist. Note that condition (3) is robust to the introduction of an endogenous interest rate. In particular, this inequality arises naturally in general equilibrium when agents face incomplete markets and borrowing constraints, because the precautionary demand for wealth pushes the interest rate down relative to the complete market case (see Huggett 1997, for a full discussion of this point in the context of the neoclassical growth model).[7]

We focus on the simplest possible equilibrium generated by our model given our assumptions about preferences and technology, and construct this equilibrium via a simple 'guess and verify' method. More specifically, our equilibrium has the property that all employed households end the current period with the same level of precautionary wealth a_t, while all unemployed households end the current period with zero wealth (consequently, the wealth distribution has two states). The first part of the conjecture will be satisfied if all employed households reach a consumption level above c^*. The second part will also be satisfied if all unemployed households face a binding borrowing constraint, and hence optimally choose to liquidate their entire asset wealth instantaneously.

Since the type of a particular household depends on both its beginning- and end-of-period wealth, it follows that under our conjecture households can be of four different types only, depending on their employment status in the current period (which determines their end-of-period wealth) and that in the previous period (on which their wealth at the end of the previous period, and hence that at the beginning of the current period, depends). We denote by ij, with $i, j = e, u$, the type of a household, where i and j refer to their employment statuses in the previous and the current period, respectively, and where e and u stand for 'employed' and 'unemployed', respectively (e.g. a eu 'household' is currently unemployed but was employed in the previous period). In short, our conjecture about the cross-sectional distribution of wealth, which will turn out to

---

[7] Whether condition (3) is satisfied or not in general equilibrium ultimately depends on the quantity of both inside (i.e. private debt) and outside (i.e. capital, public debt, money, etc.) assets that agents may use for self-insurance against idiosyncratic shocks. A large quantity of assets may allow them to achieve full insurance via asset accumulation, in which case no agent faces a binding borrowing constraint and the complete markets relation $\beta R = 1$ will prevail in the steady state. Otherwise, insurance will be imperfect and $\beta R < 1$ (Woodford 1990 and Huggett 1993, 1997).

be true in equilibrium under appropriate parameter restrictions, drastically limits the length of individual history of unemployment shocks relevant for the determination of the type of a household; this is in contrast with much of the heterogenous agents literature, where the entire individual history matters, and hence the economy is asymptotically populated by infinitely many types of agents (e.g. Aiyagari 1994).

Under our joint conjecture, the budget constraints of the different types of households are given by

$$ee : c_t^{ee} + a_t = w + a_{t-1}R \tag{4}$$

$$ue : c_t^{ue} + a_t = w \tag{5}$$

$$eu : c_t^{eu} = a_{t-1}R \tag{6}$$

$$uu : c_t^{uu} = \lambda \tag{7}$$

This can be explained as follows. uu households, who were unemployed in the previous period, ended that period with no wealth and thus enjoy neither asset nor labour market income; however, they have turned into home producers and thus get the income λ. eu households were employed in the previous period and hence left it with wealth a_{t-1}; in the current period, they get no labour income but the asset income $a_{t-1}R$, which is entirely consumed (since these households do not replete their asset wealth, by conjecture). ue households, who were unemployed in the previous period and hence start the current one with no wealth, must both build-up their asset wealth, a_t, and consume, c_t^{ue}, out of their labour income w. Finally, ee households ended the previous period with wealth a_{t-1}. In the current period, they get the asset income $a_{t-1}R$, the labour income w, and allocate this to the repletion of assets, a_t, and the consumption of goods, c_t^{ee}. In equations (4–7), w, λ and R are given and constant; hence, once a_t is known, the consumption levels of the four types of agents (the c_t^{ij}s) can be determined as residuals.

The optimal asset holdings of employed households, a_t, can be inferred directly from their (common) Euler equation. If they fall into unemployment in the next period, which occurs with probability $1-\alpha_t$, they will become eu households in the next period, so that their marginal utility of consumption will be $u'(c_{t+1}^{eu}) = u'(a_tR)$ (see (6)). If they stay employed, which occurs with complementary probability, they will be ee households, and thus enjoy marginal utility $u'(c_t^{ee})$, in the next period. A key condition in the construction of our equilibrium (and one that will have to be verified once the whole equilibrium is worked out) is that the consumption of the employed is sufficiently high so that $c_t^{ee}, c_t^{ue} > c^*$, while the

consumption of the unemployed is so low that $c_t^{eu}, c_t^{uu} < c^*$. In this case, $u'(c_t^{ee}) = u'(c_t^{ue}) = \gamma$ while $u'(a_t R) = \gamma c^*/a_t R$. Hence, the Euler equation characterizing the optimal asset demand of both *ue* and *ee* households is

$$\gamma = \beta R(\alpha_t \gamma + (1 - \alpha_t)(\gamma c^*/a_t R)). \tag{8}$$

Note that there is no expectations operator in (8). This is because the only source of uncertainty that households face when choosing their asset holdings is about whether they will keep their job or not in the next period, and that the probabilities associated with either event is known in the current period. Solving (8) for a_t, we get the following expression for the optimal asset demand of employed agents,

$$a_t = \beta c^* \left(\frac{1 - \alpha_t}{1 - \beta R \alpha_t} \right). \tag{9}$$

This expression is the central equation of the model. Importantly, it says that individual asset holdings, a_t, are increasing in $1-\alpha_t$: a rise in the risk of falling into unemployment causes households to raise precautionary savings. This mechanism will provide an amplification mechanism for the fall in aggregate consumption originally caused by a lower aggregate labour endowment.

The description of the model is complete once we have characterized the evolution of the shares of each agent type implied by the dynamics of α_t. Calling ω_t^{ij} the share of *ij* households at date *t*, simple flow accounting indicates that these shares evolves as follows

$$\omega_t^{ee} = \alpha_{t-1}(\omega_{t-1}^{ee} + \omega_{t-1}^{ue}), \tag{10}$$

$$\omega_t^{ue} = (1 - \rho)(\omega_{t-1}^{eu} + \omega_{t-1}^{uu}), \tag{11}$$

$$\omega_t^{eu} = (1 - \alpha_{t-1})(\omega_{t-1}^{ee} + \omega_{t-1}^{ue}), \tag{12}$$

$$\omega_t^{uu} = \rho(\omega_{t-1}^{eu} + \omega_{t-1}^{uu}). \tag{13}$$

We show in Appendix A that our conjectured four-agent type equilibrium exists for large (and plausible) parameter ranges. More specifically, we check that under such values the steady-state counterpart of our model (i.e. one in which $\alpha_t = \alpha$) satisfies our conditions about the ranking of consumption levels and the bindingness of the borrowing constraint for unemployed agents. Then, if deviations of α_t from the steady state are of sufficiently small magnitude, these conditions will also hold along the stochastic model, and will thus be consistent with our conjectured equilibrium.

Having derived the individual consumption and asset accumulation rules, we may now turn to the aggregate dynamics of the model.

4 Dynamic effects of a shock to unemployment risk

Section 3 showed how the precautionary motive for holding wealth affects the way households react to changes in unemployment risk. We now turn to the aggregate implications of the precautionary motive, in particular with respect to aggregate consumption, wealth and saving.

We take the path of α_t to be the exogenous forcing variable here. At any date t, a share α_t of currently employed households will stay so in the next period, while a share $1-\rho$ of currently unemployed households, who are in number $1-N_t$, will leave unemployment in the next period. Hence, the dynamics of total employment is given by

$$N_t = \alpha_{t-1}N_{t-1} + (1 - \rho)(1 - N_{t-1}). \tag{14}$$

The most direct interpretation of unemployment risk as modelled here is to think of N_t as the number of firms that are operative in the economy at date t (or, equivalently, the number of divisions of the 'representative firm'), each of which being able to hire one household. Then, a shock to α_t is one about the number of such firms that will remain operative in the next period, while ρ is the rate at which potential firms restart being operative.

All other variables of interest in the model depend on α_t and N_t. For example, since only employed households hold assets, end-of-period total assets are $A_t = a_t N_t$, where a_t is given by (9). Similarly, from the constant-returns-to-scale assumption, aggregate output is $Y_t = f(k)N_t$, so the asset/output ratio is simply

$$A_t/Y_t = f(k)^{-1}a_t.$$

This ratio is countercyclical, due to the precautionary nature of the demand for wealth: when α_t decreases (i.e. unemployment risk is higher), individual wealth a_t rises while total employment N_t falls. However, Y_t is also scaled by N_t and hence the asset/output ratio exactly tracks individual wealth.

What we are primarily concerned with here is the behaviour of aggregate consumption, C_t, and the saving rate, S_t/Y_t, induced by this precautionary demand for wealth. Aggregating the budget constraints (4–7) under the population sizes given by equations (10–13), we find that aggregate consumption is

$$C_t = -(A_t - RA_{t-1}) + \left(wN_t + \lambda\omega_t^{uu}\right), \tag{15}$$
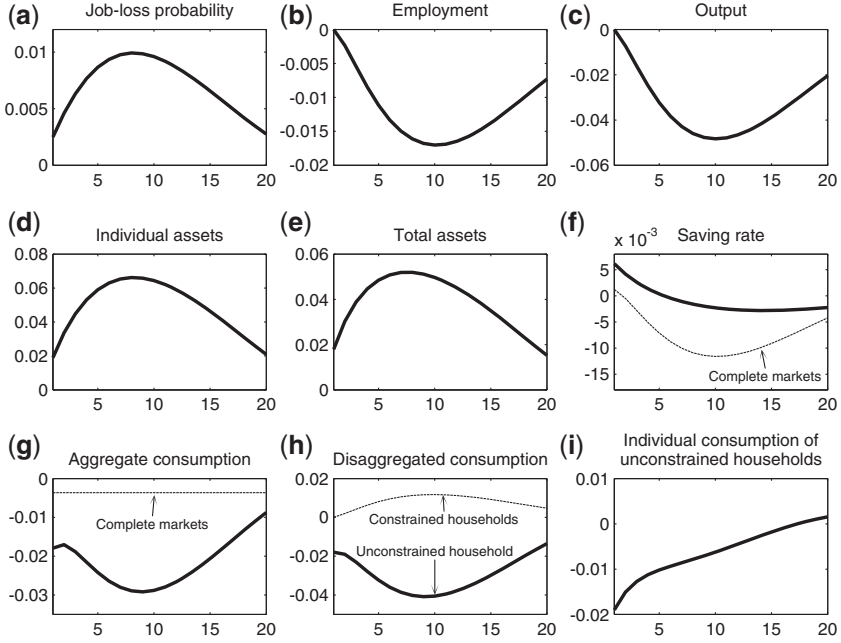
Figure 2 Impulse-response functions. The panels display the level-deviations from the steady states of all variables of interest following an exogenous change in the path of unemployment risk, as represented in Figure 2a.

where $A_t - RA_{t-1}$ is the change in total assets held by the households and $wN_t + \lambda \omega_t^{uu}$ the total labour income. Finally, aggregate saving, $S_t = Y_t - C_t$, and the saving rate, S_t/Y_t, can be computed directly from (15) and the fact that $Y_t = f(k)N_t$.

Figure 2 displays the dynamics of the variables of interest following a change in the path for the job-loss probability, $1 - \alpha_t$. More specifically, we assume, clearly for illustrative rather than quantitative purpose, that $\alpha_{T-1} = \alpha$ (i.e. the job-loss rate is at its steady-state value at date $T-1$), $\alpha_{T-\alpha} = -0.0025$ (i.e. a shock of size 0.25% occurs at date T), while $\alpha_t - \alpha$ obeys the following AR(2) process from date $t = T+1$ onwards

$$\alpha_t - \alpha = 1.84(\alpha_{t-1} - \alpha) - 0.86(\alpha_{t-2} - \alpha)$$

This path generates a hump-shaped pattern followed by a gradual (and mildly oscillatory) mean-reversion of the job-loss rate, which peaks at +1%—as is roughly consistent with the evidence reported in Section 2. The production function is assumed to be of the Cobb–Douglas form, i.e. $Y_t = K_t^\phi L_t^{1-\phi}$, with $\varphi = 1/3$. We set the gross interest rate R to 1.01 and the

discount rate β to 0.97, so that $R\beta = 0.9797 < 1$ (as required by inequality (3)). The threshold level of consumption c^* and the slope of the instant utility function at c^* are both set to 1. Finally, the job-finding rate $1-\rho$, the mean job-loss probability $1-\alpha$, and the depreciation rate δ are assumed to be 0.5, 0.03, and 0.025, respectively. Given these parameters, which satisfy conditions (A1–A3) in Appendix A, the dynamics of the variables of interest directly follow from the path for α_t.

Employment, N_t, and (market) output, Y_t, are given by equation (14) and $f(k)N_t$, respectively. Individual assets, given by (9), reflect the precautionary saving behaviour of employed (i.e. currently unconstrained) households, and thus track the job-loss probability. In our model, only employed households are unconstrained and hold assets; then, total assets, $A_t = a_t N_t$, rise because individual assets change more than the number of savers (i.e. the 'intensive' asset holding margin dominates the 'extensive' margin). This is in turn reflected in a persistently high saving rate.

The amplifying effect of precautionary savings on aggregate consumption and its components is depicted in Figure 2g, h, and i. In Figure 2g, the consumption response implied by our model is compared with that implied by the complete-markets analogue (see Appendix C for details). Essentially, in the complete-markets model unemployment risk is fully diversified between individuals and hence no precautionary saving behaviour takes place; formally, the aggregate labour income $wN_t + \lambda\omega_t^{uu}$ is given to the representative agent, who aligns current consumption to total (i.e. financial and human) wealth. The partial equilibrium nature of our model implies that complete-markets consumption follows a random walk (as in Hall 1978), while the transitory nature of the shock means that it moves by a relatively small amount (because the impact of the shock on human wealth is limited). In short, the fall in aggregate consumption that would be incurred by the representative agent is both *mild* and *acyclical*. The situation is very different under incomplete markets; in this situation, the precautionary saving motive leading to the gradual increase in total assets, A_t, is reflected in a large, gradual, and persistent fall in aggregate consumption. Moreover, inasmuch as the strength of the precautionary saving motive tracks unemployment risk and that the latter is hump-shaped, aggregate consumption is mean-reverting and U-shaped, as is consistent with the evidence shown in Section 2.

Figure 2h and i look at disaggregated consumption levels to illustrate that precautionary saving, rather than other factors, is indeed chiefly responsible for the prolonged fall in current consumption. Indeed, one legitimate concern one may have by looking at the path aggregate consumption is that the latter could be driven by the consumption of borrowing-constrained agents (i.e. those who do not form precautionary
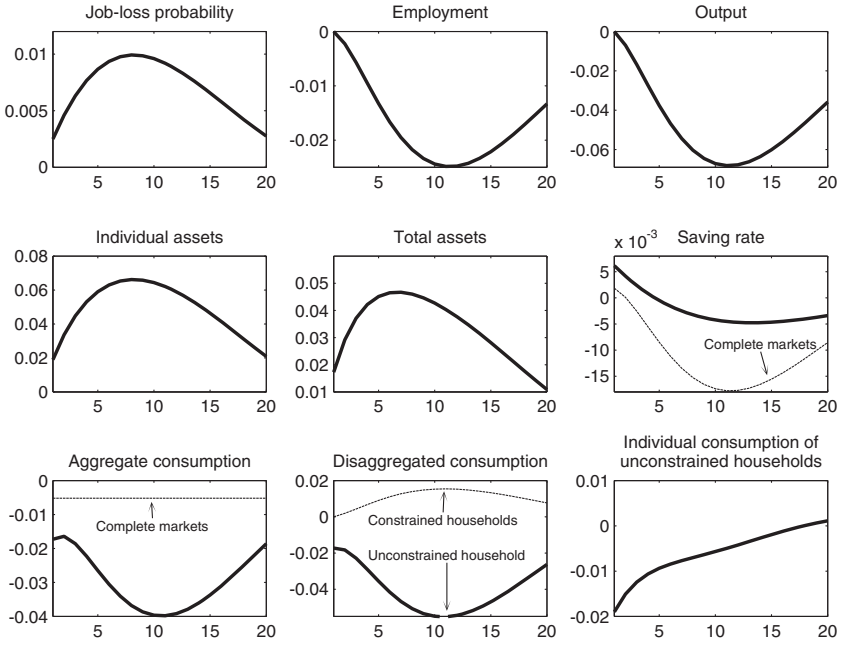
Figure 3 Sensitivity. The panels are as in Figure 2. All parameters are at their baseline value except for the job-finding rate, which is set at 0.3 (instead of 0.5).

savings), whose current income impacts their consumption level one for one, rather than by that of unconstrained precautionary savers. Figure 2h displays the total contributions of constrained and unconstrained households, and shows that much of the consumption dynamics is driven by the behaviour of precautionary savers. In fact, the total consumption of constrained households rises mildly, rather than falls, after the shock. This occurs for two reasons. First, constrained households rise in number after the shock has taken place. Second, some of the constrained households actually consume more than they would have had the shock not occurred: those who liquidate their assets (i.e. *ue* households), whose precautionary wealth has increased since the beginning of the crisis following rising uncertainty about employment prospects. Figure 2i divides the total consumption of precautionary savers (i.e. *ue* and *ee* households) by the total number of such savers (i.e. $\omega_t^{ee} + \omega_t^{ue}$). This again establishes the leading role of the precautionary motive in determining the dynamics of aggregate consumption in the model.

Figure 3 carries out a sensitivity check with respect to the (exogenous) job-finding rate ρ, which is now set at 0.3 instead of 0.5. This focus on the

job-finding rate is motivated by the empirical observation that such rates vary substantially across OECD countries, while average job-loss rates do not (see Hobjin and Sahin 2009). One implication of this finding is that the higher unemployment rate experienced many European countries relative to the US is primarily due to fewer transitions out of the unemployment pool in Europe than in the USA; we thus wish to investigate what a smaller job-finding rate implies for the transmission of labour reallocation shocks. The figure show that shocks to the job-loss rate have stronger and more persistent effects on aggregate consumption when the job-finding rate is smaller. This is due to two competing effects. On the one hand, a smaller job-finding rate raises the share of constrained (i.e. unemployed) agents in the economy and lowers that of unconstrained (i.e. employed) precautionary savers; this goes towards lowering the reaction of aggregate consumption to the shock. On the other hand, for a given path of the job-loss rate, a smaller job-finding rate magnifies the employment response to the shock (as a higher proportion of employed agents turn out to lose their job), which reduces the aggregate endowment available for consumption and saving—see the $w_t N_t$ term in (15). This second effect turns out to be dominant (at least under our choice of calibration for the other parameters), leading to stronger and more persistent responses of most aggregates.

5 Concluding remarks

This article has built on a now old, but not mainstream, tradition in macroeconomic analysis that emphasizes the importance of uninsurable income risk and borrowing constraints in determining aggregate outcomes. Early work in this literature, including those of Bewley (1983), Scheinkman and Weiss (1986), Kehoe et al. (1991) and Imrohoroglu (1992), focused on the precautionary demand for money under incomplete income insurance. Aiyagari (1994) and Huggett (1997) extended early models by analysing the effects of the precautionary motive when assets are claims to the capital stock. The latest wave of articles including Heathcote (2005) Krusell and Smith (1997, 1998), and Storesletten et al. (2007) introduce aggregate shocks (together with idiosyncratic uncertainty) into this framework, thereby contributing to the extension of this approach to a variety of macroeconomic issues, including business cycles, fiscal policy, and asset pricing. Heathcote et al. (2009) provide a survey of these models, commonly referred to as 'heterogeneous agents' models.

The purpose of our current research, of which the model developed in Sections 3 and 4 is a representative example, is to offer simple class of models/equilibria that incorporate a variety of channels related to the

precautionary motive and the demand for 'aggregate liquidity'. Indeed, it seems to us that the scope of analysis of incomplete market models is often limited by the computationally difficult exercise of approximating the dynamics of the equilibrium wealth distribution. In contrast, our model provides an illustration of what can be achieved by confining oneself to the simplest deviation from the complete markets assumption; essentially, the outcome is a model that can accommodate uninsured individual risk, agents' heterogeneity, and the precautionary motive for saving, while at the same time taking the form of a standard dynamic stochastic general equilibrium model, for which a variety of workable solution techniques are available. Ultimately, which of a fully fledged computational model or a limited-heterogeneity approach should be used depends on the question under scrutiny. Clearly, our framework generates overly simplistic wealth distributions, and is thus poorly suited to the matching and characterisation of their empirical analogues, which is the focus of many articles in the incomplete-markets literature. However, our approach flexibly incorporates the precautionary motive for holding assets into dynamic equilibrium and can thus be extended in several directions while maintaining tractability, e.g. by considering endogenous interest rate adjustments (such as those which would occur in a closed or a large open economy), more realistic labour market frictions and adjustments, a larger number of liquidation periods for the unemployed, etc.

Because incomplete-markets economies typically generate aggregate inefficiencies, they leave ample room for welfare-improving policy interventions. For example, in the context of the model presented above, rising idiosyncratic uncertainty leads to a much larger fall in total private demand than is (first-best) efficient. This suggests that this model may provide a rationale for aggregate demand management of the 'Keynesian' kind, inasmuch as it may indirectly provide labour income insurance when direct insurance markets are missing.[8] Here again, it seems important to carry out theoretical and quantitative experiments using models in which the transmission channels of the different policy options are transparent.

More generally, it is our contention that generalizing the incorporation of uninsurable risk into the various strands of economic analysis may considerably improve our understanding of the macroeconomy, both along the business cycle and in time of major crisis such as that we are currently experiencing. To give substance to this view, let us think loud about some potential applications of our framework for a moment.

---

[8] For example, Challe and Ragot (2010) construct a model in which debt-financed expansionary fiscal policy may, by altering the stock of aggregate liquidity, boost aggregate consumption and welfare.

One typical example of behaviour akin to 'precautionary saving' or 'aggregate liquidity hoarding' is that adopted by financial intermediaries in the early stages of the crisis, during which they hoarded base money or liquid assets to avoid any form of credit risk.[9] This behaviour is clearly detrimental to the normal functioning of financial intermediation, and may even trigger the collapse of some financial markets (e.g. the interbank market). Take, as another example, the behaviour of firms facing uncertainty about idiosyncratic factors such as productivity, demand, etc. In this situation, high individual uncertainty leads to delayed investment in order to avoid commitment to future production.

To summarize, our answer to the question 'What's wrong with modern macroeconomics?' is that the latter has (thus far) considerably underestimated the extent and cyclicality of the individual uncertainty faced by economic agents. One implication is that macroeconomics has largely overlooked the role of precautionary savings in the business cycle and the way it affects aggregate demand. While the current article has focused on the effect of unemployment risk on aggregate consumption, the same argument applies, mutatis mutandis, to other economic agents or institutions such as firms and financial intermediaries.

Acknowledgements

References

Acharya, V.V. and O. Merrouche (2009), "Precautionary Hoarding of Liquidity and Inter-bank Markets: Evidence from the Sub-prime Crisis", Working Paper, July.

Aiyagari, S.R. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics* **109**, 659–684.

Algan, Y., E. Challe and X. Ragot (2009), "Incomplete Markets and the Output-inflation Tradeoff", *Economic Theory* 28 September 2009 [Epub ahead of print; doi:10.1007/s00199-009-0499-0].

---

[9] See, i.e. Acharya and Merrouche (2009) for direct evidence about banks' precautionary hoarding of liquidity during the money market freeze that started on 9 October, 2007.

Andolfatto, D. (1996), "Business Cycles and Labor Market Search", *American Economic Review* **86**, 112–132.

Bewley, T.F. (1983), "A Difficulty with the Optimum Quantity of Money", *Econometrica* **51**, 1485–1504.

Carroll, C.D. (1992), "The Buffer-stock Theory of Saving: Some Macroeconomic Evidence", *Brookings Papers on Economic Activity* **1992**, 61–156.

Cochrane, J.H. (1991), "A Simple Test of Consumption Insurance", *Journal of Political Economy* **99**, 957–976.

Challe, E. and X. Ragot (2010), "Fiscal Policy in a Tractable Liquidity-Constrained Economy", PSE Working Paper.

Challe, E., F. Le Grand and X. Ragot (2010), "Incomplete Markets, Liquidation Risk, and the Term Structure of Interest Rates", PSE Working Paper.

Den Haan, W.J., G. Ramey and J. Watson (2000), "Job Destruction and Propagation of Shocks", *American Economic Review* **90**, 482–498.

Gourinchas, P.-O. and J. Parker (2002), "Consumption Over the Life Cycle", *Econometrica* **70**, 47–89.

Guiso, L., T. Jappelli and D. Terlizzese (1996), "Income Risk, Borrowing Constraints, and Portfolio choice", *American Economic Review* **86**, 158–172.

Hall, R.E. (1978), "Stochastic Implications of the Life Cycle – Permanent Income Hypothesis", *Journal of Political Economy* **86**, 971–987.

Heathcote, J. (2005), "Fiscal Policy with Heterogeneous Agents and Incomplete Markets", *Review of Economic Studies* **72**, 161–188.

Heathcote, J., K. Storesletten and G. Violante (2009), "Quantitative Macroeconomics with Heterogeneous Households", *Annual Review of Economics* **1**, 319–354.

Heer, B. (2001), "Wealth Distribution and Optimal Inheritance Taxation in Life-cycle Economies", *Scandinavian Journal of Economics* **103**, 445–65.

Hobijn, B. and A. Sahin (2009), "Job-finding and Separation Rates in the OECD", *Economics Letters* **104**, 107–111.

Huggett, M. (1993), "The Risk-free Rate in Heterogeneous-Agent, Incomplete-insurance Economies ", *Journal of Economic Dynamics and Control* 17, 953–969.

Huggett, M. (1996), "Wealth Distribution in Life-cycle Economies", *Journal of Monetary Economics* **38**, 469–494.

Huggett, M. (1997), "The One-sector Growth Model with Idiosyncratic Shocks: Steady States and Dynamics", *Journal of Monetary Economics* **39**, 385–403.

Imrohoroglu, A. (1992), "The Welfare Cost of Inflation Under Imperfect Insurance", *Journal of Economic Dynamics and Control* **16**, 79–91.

Jappelli, T. (1990), "Who is Credit Constrained in the U.S. Economy?" *Quarterly Journal of Economics* **105**, 219–234.

Kehoe, T.J., D.K. Levine and M. Woodford (1992), "The Optimum Quantity of Money Revisited", in P. Dasgupta, D. Gale, O. Hart and E. Maskin, eds. *Economic Analysis of Markets and Games*, MIT Press, Cambridge, MA.

Krebs, T. (2007), "Job Displacement Risk and the Cost of Business Cycles", *American Economic Review* **97**, 664–686.

Krusell, P. and A.A. Smith (1997), "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns", *Macroeconomic Dynamics* **1**, 387–422.

Krusell, P. and A.A. Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy* **106**, 867–896.

Merz, M. (1995), "Search in the Labor Market and the Real Business Cycle", *Journal of Monetary Economics* **36**, 269–300.

Scheinkman, J.A. and L. Weiss (1986), "Borrowing Constraints and Aggregate Economic Activity", *Econometrica* **54**, 23–45.

Shimer, R. (2005), "The Cyclical Behavior of Equilibrium Unemployment and Vacancies", *American Economic Review* **95**, 25–49.

Walsh, C.E. (2005), "Labor Market Search, Sticky Prices, and Interest Rate Policies", *Review of Economic Dynamics* **8**, 829–849.

Woodford, M. (1990), "Public Debt as Private Liquidity", *American Economic Review* **80**, 382–388.

Appendix A

A.1 Existence conditions

This appendix works out the sufficient conditions for the equilibrium analysed in the body of the paper to exist. We must check that our equilibrium conditions are satisfied in the steady state, as characterized by the deterministic counterparts of (4–7), (9), and (10–13). There are two sets of conditions for the equilibrium to exist. First, the ranking of consumption levels must be consistent with the conjectured ranking under which the

equilibrium was derived, i.e. $c^{uu} < c^{eu} < c^* < c^{ue} < c^{ee}$. From (4) and (5), it is necessarily the case that $c^{ue} < c^{ee}$. Using (9), the condition that $c^{ue} = w - a > c^*$ gives

$$w > c^* \left(1 + \frac{\beta(1-\alpha)}{1 - \beta R\alpha}\right). \tag{A1}$$

From (6) and (9), we have that $c^{eu} < c^*$ if and only if $R\beta < 1$, which holds by assumption (3). Finally, from (6) and (7) the condition that $c^{uu} = \lambda < c^{eu} = aR$ gives

$$\frac{c^*}{\lambda} > \frac{1 - \beta R\alpha}{\beta R(1-\alpha)}. \tag{A2}$$

Second, the borrowing constraint must be binding for eu and uu households. Since $c^{uu} < c^{eu}$ under condition (A2), while the probability of leaving unemployment is the same for the two types, we only need to check that it is binding for eu households, which is the case if and only if

$$u'(c^{eu}) > \beta R(\rho u'(c^{uu}) + (1 - \rho)u'(c^{ue})).$$

With c^{eu} ($= aR$) and c^{uu} ($= \lambda$) to the left of c^* and c^{ue} to the right of c^*, this inequality becomes

$$a\beta R^2 \left(\frac{\rho}{\lambda} + \frac{1-\rho}{c^*}\right) < 1 \tag{A3}$$

The stochastic equilibrium constructed in the article exists whenever conditions (A1–A3) are jointly satisfied and fluctuations around the steady state are small. They are satisfied for any plausible parameter configuration provided that γ and c^* are chosen appropriately.

Appendix B

B.1 National income accounting

Clearing of the capital markets operates through international capital movements. End-of-period assets, A_t, serve to form the capital stock at the beginning of the next period, K_{t+1}, as well as net foreign assets at the beginning of the next period, F_{t+1}, so that $A_t = K_{t+1} + F_{t+1}$. Firms maximization under constant returns to scale implies that market output, Y_t, is

$$Y_t = wN_t + (R - 1 + \delta)K_t,$$

while total domestic output, which incorporates home production, is $\tilde{Y}_t = Y_t + \lambda\omega_t^{uu}$. Solving the market output equation for wN_t, substituting

it into (15) and using the asset market clearing condition, we get the accounting identity

$$C_t + (K_{t+1} - (1 - \delta)K_t) + (F_{t+1} - RF_t) = \tilde{Y}_t,$$

where the expressions inside the first and second pairs of brackets are savings invested domestically and savings invested abroad, respectively. Since the latter are also equal to net exports, changes in net foreign assets are net exports plus capital income from abroad, i.e.

$$\Delta F_{t+1} = (X_t - M_t) + (R - 1)F_t.$$

Appendix C

C.1 The representative agent economy

To construct the equivalent open-economy, representative agent economy, we must first set this agent's subjective discount factor, β^{ra}, to $1/R$ (otherwise foreign borrowing or lending would be unlimited). We then endow this agent with the aggregate labour income $wN_t + \lambda\omega_t^{uu}$, and compute the implied optimal consumption and saving plans given the forcing sequence $\{\alpha_t\}_{t=0}^{\infty}$. We assume that $\alpha_t = \alpha$ until some date $T-1$. At $t = T$, a unanticipated shock occurs that sets α_t in motion; once the shock has occurred, the path for α_t is perfectly anticipated. With $\beta^{ra} = 1/R$, the Keynes–Ramsey rule implies that $C_{t+1} = C_t$ for all $t \geq T$. Iterating the period budget constraint (15) forward, yields the following intertemporal budget constraint (IBC)

$$\sum_{i=0}^{\infty}\left(\frac{C_{t+i}}{R^i}\right) = \sum_{i=0}^{\infty}\left(\frac{wN_t + \lambda\omega_t^{uu}}{R^i}\right) + RA_{t-1}.$$

Using the IBC and the optimality condition $C_{t+1} = C_t \ \forall t \geq T$, we find that

$$C_t = \Lambda(H_t + RA_{t-1}),$$

where $H_t = \sum_{i=0}^{\infty}(wN_t + \lambda\omega_t^{uu})/R^i$ is the representative agent's human wealth, RA_{t-1} his (begining-of-period) financial wealth and $\Lambda = (R - 1)/R$ is the propensity to consume out of wealth. With the economy being in the steady state before $t = T$, we have $A_{t-1} = A_{t-2}$, so that the change in aggregate consumption triggered by the shock is simply the shock to human wealth, i.e. $\Delta C_t = \Lambda\Delta H_t$. Then, using unindexed variables to denote steady-state values, we have

$$\Delta C_t = \Lambda(H_t - H) = \Lambda\sum_{i=0}^{\infty}\left(w\left(\frac{N_{t+i} - N}{R^i}\right) + \lambda\left(\frac{\omega_t^{uu} - \omega^{uu}}{R^i}\right)\right).$$