Precautionary Saving for Consecutive Life-Cycle Risks*

Richard Blundell†
Ben Etheridge‡
Tom Stoker§

November 4, 2014

Abstract

We examine whether households combine (or 'complement') precautionary saving for near-term risks with saving for long-term risks. In a realistic life-cycle model, we find this complementarity effect accounts for 8-16% of precautionary savings. Almost all this effect is driven by permanent shocks. We then obtain analytical results from a 3-period model. We find permanent shocks induce complementarity for a general class of preferences, including those with constant relative risk aversion (CRRA). However, for most preferences in this class, the interaction of transitory shocks amplifies the precautionary motive. We interpret these results in terms of the structure of risks and the pattern of prudence over the wealth spectrum.

JEL Classification: D91, E21, D81.

Keywords: Precautionary saving, risk, life-cycle models.

*We thank Hamish Low and Louis Eeckhoudt for comments on an earlier draft. We also thank seminar audiences at Essex and at the 2014 ESEM meetings in Toulouse. We are grateful to the financial support of the ESRC Centre for the Micro-Economic Analysis of Fiscal Policy at IFS. All errors are our own.

†UCL and Institute for Fiscal Studies
‡University of Essex
§MIT Sloan School of Management

1 Introduction

Economists have long recognized the variety of motives for saving.[1] More recently, researchers have debated the relative importance of the precautionary motive for emergencies versus the life-cycle motive.[2] [3] Clearly, households' life-cycle motive has, by definition, a long-term focus. In contrast, it is tempting to consider precautionary saving as driven by short-term possible needs. However, households face risks of different types and magnitudes over the whole life-cycle. Within the variety of motives for general saving, therefore, households have competing motives for precautionary saving, for near-term as well as for far-off emergencies.

Assessing the importance of each saving motive prompts a further question. Can savings be decomposed cleanly according to motive or do the motives overlap? Liquid savings, for example, can presumably meet many motives simultaneously: a dollar saved in the bank can be spent in retirement, can be used in a near-term emergency, or can be used eventually in a far-off emergency. And presumably households need save less when they can overlap motives in this way. Understanding these overlaps is therefore important for understanding household, and hence aggregate, wealth accumulation.

In this paper we examine these overlaps in saving. Using more formal language, we consider these overlaps as 'complementary' saving. In particular, we focus on precautionary saving for consecutive, distinct future income risks. On one intuition, such precautionary saving seems particularly likely to be complementary because of its contingent nature: precautionary savings are only rarely needed because emergencies only rarely occur. Therefore the accumulated savings for the first risk can presumably be put to other ends, and specifically as rolled-over precautionary savings for subsequent income risk. On the other hand, the theoretical literature on saving under background

[1]see Browning and Lusardi (1996) for an exhaustive list.

[2]See Kimball (1990) and Caballero (1990) for early theoretical formulations of the precautionary motive in terms of prudence. Deaton (1991) and Carroll (1997) provide classic analyses focusing on explaining patterns of life-cycle consumption and wealth formation.

[3]Carroll and Samwick (1997) and Gruber and Yelowitz (1999), for example, emphasize the importance of precautionary saving in total wealth, while Dynan (1993), for example, finds precautionary saving to be less important.

risk (discussed, for example, in Gollier, 2004) emphasizes that the presence of multiple risks amplifies risk aversion. On this intuition, precautionary savings for consecutive risks may not be complementary but may in fact be amplified.

We formalize the intuitive notion of complementarity in the following way: first we quantify benchmark precautionary saving when the household faces risk both in mid-life and then late-life consecutively. We then quantify precautionary saving when the household faces either risk in isolation. We label the sum of precautionary saving for these two isolated periods of risk the 'total' precautionary effect of the risks. If the total precautionary effect is greater than the initial precautionary saving benchmark, we say that saving is complementary, or exhibits complementarity.

Our results are both quantitative and analytical. On the quantitative side, we simulate a standard life-cycle consumption and saving model with both permanent and transitory shocks. In a range of realistic parametrizations, the total precautionary effect of the risks is 8-16% higher than the precautionary saving benchmark. We find that this complementarity effect is driven entirely by the permanent shocks. In contrast, as these shocks get less persistent eventually complementarity is overturned. We find that purely transitory risks amplify precautionary saving, although the effect is small. We conclude that the complementarity effect quantitatively dominates.

To understand how these results arise we then simplify the model and examine saving behaviour analytically. We use a 3-period framework and examine various special cases of the risk environment. First, and because it is quantitatively the most important case, we admit small permanent risks alone. We find that complementarity depends on the shape of relative prudence over the wealth spectrum.[4] Utility functions with constant relative prudence (those in the CRRA class) always permit complementary saving. More generally, so do all utility functions with harmonic absolute risk aversion (HARA).[5] The HARA class includes not only CRRA but also constant absolute risk

[4] Relative prudence is defined as $-\frac{xu'''(x)}{u''(x)}$ for utility function $u(x)$ and consumption level x. This can be interpreted as the strength of the desire to have higher wealth when facing a given proportional gamble (such as $\pm 10\%$). See Kimball (1990).

[5] HARA preferences have absolute risk aversion of the form $\frac{1}{ax+b}$, where a and b are constant, x is the consumption level and risk aversion is defined to be $-\frac{u''(x)}{u'(x)}$.

aversion (CARA) and quadratic functions.[6]

We build intuition for this behaviour by examining related environments. In our second special case, we examine saving when the agent faces consecutive small transitory risks alone. In line with the simulations, first-period saving is not complementary but is amplified for standard preferences, such as CRRA. We term this behaviour 'excessive' saving. Again we derive more general conditions for this behaviour in terms of the shape of prudence. In this second case the conditions are more interpretable: saving is excessive if absolute prudence is declining and convex.[7] The intuition is as follows: when the prudence function is convex then average prudence is greater the more independent risks are added, an application of Jensen's inequality. And because average prudence is greater, precautionary saving is larger. Convexity and a negative slope are attractive properties of the prudence function.[8]

In fact we gain the most intuition for these results by considering preferences of the CARA form and transitory risks, our third special case. This third case can be solved analytically and simply. Here saving is neutral; it is neither excessive nor complementary, but can be additively decomposed. We understand this result by considering the following informal argument: From the viewpoint of the middle period, the precautionary motive for the final-period risk does not depend on the outcome of the most recent shock. This is because innovations to wealth from transitory risks are independent across time and because precautionary saving with CARA preferences is independent of wealth. Returning to the first period, the household must save for both middle- and final-period risks. And because the middle-period outcome will not affect behaviour, the size of the middle risk does not affect the precautionary motive for the final risk. In short, the risks don't interact in the initial saving decision.[9]

[6]CRRA preferences can be represented by utility function $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$ for some parameter γ. These preferences have both constant relative risk aversion and constant relative prudence. CARA preferences can be represented by utility function $u(x) = \mathrm{Exp}(-\gamma x)$ for some parameter γ. These preferences have both constant *absolute* risk aversion and constant *absolute* prudence.

[7]Absolute prudence is defined as $-\frac{u'''(x)}{u''(x)}$. This differs from relative prudence by employing an *absolute* gamble in the definition, such as $\pm\$10$.

[8]All HARA preferences except CARA and quadratic preferences have a convex and declining absolute prudence.

[9]See Caballero (1990) for the full mathematical treatment.

We can understand the more complex environments by considering where this argument breaks down. For example, consider the first, main special case, with two consecutive permanent risks. In this case the variances of the innovations to life-time wealth are no longer independent. In fact the final-period variance is increasing quadratically in the middle-period outcome: in particular, the variance of subsequent wealth innovations is reduced following a bad income shock. This dependence arises even though the raw income shocks are independent. The risk process itself therefore provides a kind of insurance that limits the need for precautionary saving for standard preferences. Permanent risks therefore permit complementarity much more readily than do transitory risks. This is true not only for CARA preferences but for all general HARA preferences. It is worth emphasizing that this feature of permanent risks arises not because *mean* income is dependent over time. It arises because of the dependence over time of the *variance*.

We provide some further interesting results. For example, even with permanent shocks alone, saving might not be complementary whenever relative prudence is locally strongly declining. This occurs whenever households have minimum consumption needs, such as with Stone-Geary preferences, and a household is on the 'breadline'. With these preferences, households are so averse to consecutive bad shocks that the interaction of the risks amplifies precautionary saving. Such excessive saving is only local, however. When the household is sufficiently far from the breadline, precautionary saving is complementary again.

After the early works cited above, the theoretical and empirical literature on precautionary savings has developed steadily in, for example, Carroll and Kimball (2001) and Carroll (2004). These papers emphasize that when households discount the future strongly, display prudence and face income risk, they save to meet a target wealth holding, the buffer stock. This buffer stock intuitively conforms to the notion that a fixed level of wealth can meet all subsequent risks. The models in these papers, however, include constant income risk over the planning horizon and so do not explicitly explore the distinction between near-term and far-term risk.

Aside from the life-cycle literature on precautionary savings, this paper contributes to an established theoretical literature on the interaction of risks. Classic theoretical works by Kimball (1993), Pratt and Zeckhauser (1987), Gollier and Pratt (1996) characterize conditions on the utility function for the introduction of background risks to affect the desirability of risk bearing. More recently, Franke et al. (2011) examine multiplicative as well as additive risks, similarly to our focus on transitory versus permanent shocks. These papers focus on risk bearing and so characterize results in terms of the coefficient of risk aversion. Our work, on the other hand, relates to precautionary saving so characterizes the results in terms of the coefficient of prudence. In a similar vein, our paper relates to work on the time-diversification of risks using 3-period settings, such as in Blundell and Stoker (1999) and Eeckhoudt et al. (2005).

There is also a growing empirical literature on background risks, and particularly the effect of risk at the end of the life cycle on prior behaviour. For example, Goldman and Maestas (2013) look at the effect of medical expenditure risk on portfolio decisions earlier in retirement. De Nardi et al. (2010) look at the effect of medical expenditure and mortality risk on the precautionary saving motive earlier in retirement. Guiso et al. (2009) look at the effect of pension risks on portfolio allocation earlier in the life-cycle. All these papers find that the background risks have an important impact on household behaviour.

This paper proceeds as follows. In section 2 we lay out the basic three-period model through which all the results can be understood. We also discuss our definition of complementarity. In section 3 we present the first, quantitative results. Here we use a more realistic life-cycle model which includes both transitory and permanent income risk and a longer-term planning horizon. In section 4 we present the analytic results. First we discuss savings behaviour for small permanent risks alone and standard classes of preferences. We then build intuition by considering transitory shocks alone. Meanwhile we extend intuition by constructing exotic preferences which induce different saving behaviour. Section 5 concludes and discusses other features which may affect complementarity of saving: liquidity constraints and holdings of illiquid assets such as housing and pensions.

2 The Model and Definition of Complementarity

In this section we lay out a standard consumption and saving problem in a multi-period framework. We focus on the three-period environment used for the more theoretical analysis in section 4. Section 3, in contrast, presents quantitative results for a longer planning period, but all the analysis is intelligible in terms of the shorter model.

2.1 The Budget Constraint and Income Process in the Basic 3-Period Model

Households must choose consumption in three periods (indexed by $t = 0, 1, 2$) subject to the following budget constraint:

$$c_0 + \frac{c_1}{R} + \frac{c_2}{R^2} = y_0 + \frac{y_1}{R} + \frac{y_2}{R^2} \tag{1}$$

where c_t is consumption at time t, y_t is income at time t. Households save in a risk-free bond with interest rate R.

Income follows a standard stochastic multiplicative permanent-transitory process:

$$\begin{aligned} y_t^P &= y_{t-1}^P G_t \psi_t \\ y_t &= y_t^P \xi_t \\ \mathbb{E}(\xi_t) = \mathbb{E}(\psi_t) = 1 \quad &, \quad \mathrm{Var}(\xi_t) = \sigma_{\xi_t}^2, \mathrm{Var}(\psi_t) = \sigma_{\psi_t}^2 \quad t = 1, 2 \end{aligned}$$

where G_t is deterministic growth, y_t^P is latent permanent income, ψ_t represents the permanent shock to income, and ξ_t a transitory shock to income.[10] These shocks are uncorrelated with each other and uncorrelated with other shocks across time. This process nests the standard lognormal process, in which case, for example, $\ln \psi_t \sim \mathcal{N}\left(-\frac{\sigma^2}{2}, \sigma^2\right)$, where $\sigma^2 = \ln\left(\sigma_{\psi_t}^2 + 1\right)$.

[10] In this 3-period model all shocks have a large impact on life-time wealth, so transitory shocks have greater impact on consumption than in a one-year-per-period model. The quantitative importance of transitory risks therefore cannot be determined easily from the 3-period model.

We have presented the income process in terms of raw income shocks. However, the intuition behind results depends on characterizing the variance of income levels and the risks to life-time wealth. With multiplicative shocks the relationship between these variables is a little involved. It is therefore worth spelling out these relationships in detail.

We first show how the variances of the income shocks affect the variances of income levels. Defining $\bar{y}_t \equiv \mathbb{E}_0\left(y_t\right) = y_0 \prod_{j=1}^{t} G_j$:

$$
\begin{aligned}
y_1 &= \bar{y}_1 \psi_1 \xi_1 \\
y_2 &= \bar{y}_2 \psi_1 \psi_2 \xi_2
\end{aligned}
$$

To simplify the exposition, we focus on the cases first with only transitory shocks, then only permanent shocks. The general environment is built up from these special cases. We define $\mathrm{Var}_t(y_2)$ to be the variance of period-2 income conditional on the period-t information set ($t \in \{0, 1\}$). First, when there are only transitory shocks:

$$
\begin{aligned}
\mathrm{Var}(y_1) &= && \left(\bar{y}_1\right)^2 \sigma^2_{\xi_1} \\
\mathrm{Var}_0(y_2) &= \mathrm{Var}_1(y_2) = && \left(\bar{y}_2\right)^2 \sigma^2_{\xi_2}
\end{aligned}
$$

in which case the variance of period-2 income is independent of the period-1 innovations, but depends on the size of expected income growth. When there are only permanent shocks:

$$
\begin{aligned}
\mathrm{Var}(y_1) &= \left(\bar{y}_1\right)^2 \sigma^2_{\psi_1} \\
\mathrm{Var}_1(y_2) &= \left(\bar{y}_2 \psi_1\right)^2 \sigma^2_{\psi_2} \\
\mathrm{Var}_0(y_2) &= \left(\bar{y}_2\right)^2 \left(\sigma^2_{\psi_1} + \left(1 + \sigma^2_{\psi_1}\right) \sigma^2_{\psi_2}\right)
\end{aligned}
$$

where the last line can be derived using the formula for the variance of products (as in Goodman (1960)).[11] In this case, the variance of period-2 incomes depends on the the period-1 environment.

[11] For uncorrelated random variables $\mathrm{Var}\left(xy\right) = \mathrm{Var}\left(x\right) \mathrm{E}\left(y\right)^2 + \mathrm{Var}\left(y\right) \mathrm{E}\left(x\right)^2 + \mathrm{Var}\left(x\right) \mathrm{Var}\left(y\right)$.

We now consider how the variances of income shocks affect the variances of innovations to life-time wealth. Here we focus on the most interesting case of pure permanent shocks. The 2nd-period innovation to life-time wealth is:

$$\zeta_2^* \equiv \frac{y_2 - \mathbb{E}_1 y_2}{R^2} = \frac{\bar{y}_2 \psi_1 (\psi_2 - 1)}{R^2}$$

and therefore:

$$\begin{aligned}\mathrm{Var}_1(\zeta_2^*) &= \frac{(\bar{y}_2\psi_1)^2}{R^4}\sigma_{\psi_2}^2 &(2)\\ \mathrm{Var}_0(\zeta_2^*) &= \frac{\bar{y}_2^2}{R^4}\left(1+\sigma_{\psi_1}^2\right)\sigma_{\psi_2}^2 &(3)\end{aligned}$$

Equation 2 shows that, from the viewpoint of period 1, the variance of the period-2 wealth innovations also depends on the realization of the period-1 income shock. It is worth emphasizing that this connection of innovations is an important feature of multiplicative risk. In consequence, equation 3 shows that, from the viewpoint of period 0, the variance of these period-2 wealth innovations depends on the variance of period-1 income risk. We want the ex-ante (i.e. seen from period-0) variance of period-2 innovations not to depend on the period-1 variance. Why? Because our analysis depends on varying period-1 and period-2 risks in isolation. We therefore generally work with a scaled and recentred income shock $\tilde{\psi}_2 = \frac{\psi_2 + \sqrt{1+\sigma_{\psi_1}^2} - 1}{\sqrt{1+\sigma_{\psi_1}^2}}$. For this recentred shock, $\mathbb{E}_0\left(\tilde{\psi}_2\right) = 1$ and $\mathrm{Var}\left(\tilde{\psi}_2\right) = \frac{\sigma_{\psi_2}^2}{1+\sigma_{\psi_1}^2}$. If $y_2 = \bar{y}_2\psi_1\tilde{\psi}_2$, and $\tilde{\zeta}_2^*$ is the wealth innovation for the scaled shock then:

$$\begin{aligned}\mathrm{Var}_1\left(\tilde{\zeta}_2^*\right) &= \frac{(\bar{y}_2\psi_1)^2}{R^4\left(1+\sigma_{\psi_1}^2\right)}\sigma_{\psi_2}^2 &(4)\\ \mathrm{Var}_0\left(\tilde{\zeta}_2^*\right) &= \frac{\bar{y}_2^2}{R^4}\sigma_{\psi_2}^2 &(5)\end{aligned}$$

Equations 2 and 4 show that the connection between period-1 income shocks and period-2 wealth innovations is the same for both $\tilde{\psi}_2$ and ψ_2. But, importantly, equation

5 shows that the ex-ante variance of scaled wealth innovations no longer depends on the size of period-1 income risks. In fact, using $\tilde{\psi}_2$ instead of ψ_2 does not affect the results substantially, but does indeed make the analysis more interpretable. See appendix A for an extensive discussion.

2.2 The Consumption Problem

The agent faces the following value function problem:

$$V_0(W_0) = \max_{\{c_t(W_t):t=0,1,2\}} u(c_0) + \mathbb{E}_0\left(\beta u(c_1) + \beta^2 u(c_2)\right) \tag{6}$$

subject to the budget constraint given in equation 1 and the process for wealth innovations. Here $u(x)$ is the per-period felicity function, $\mathbb{E}_t$ is the expectations operator at time t and β is the rate of time preference, assumed to be common across both future periods. $V_0(W_0)$ is the value of the programme to the household at time 0 at wealth level W_0. c_0 is period-0 saving. We focus on how initial-period saving responds to changes in future risk. Therefore we could rewrite programme 6 as:

$$V_0(W_0, \Gamma) = \max_{\{s_0(W_0,\Gamma)\}} u(W_0 - s_0) + \mathbb{E}_0\left(V_1\left(s_0, \Gamma\right)\right) \tag{7}$$

where $V_1()$ represents the value function at period 1, Γ represents the parameters of the problem, including, for example, time preference, interest rates and the distributions of future income risk, and s_0 represents period-0 saving.

Interest features solely on the effect of changes in income risk. It is well known that assigning a one-dimensional measure of riskiness is tricky. The standard approach is to use the notion of mean-preserving spreads, which provides a partial ordering of income distributions in terms of second-order stochastic dominance. However, in general, higher order features of the distribution will also affect saving.[12] In this paper we will refer to the size of income risk purely in terms of the variance or standard deviation.

[12] See Eeckhoudt and Schlesinger (2008) for an analysis in a 2-period environment.

This is justified by two main considerations. First, in the quantitative analysis in section 3 we perform numerical simulations using a log-normal distribution. Log-normal shocks (with unit mean) are characterized completely by their 2nd moment. Second, in the analytical analysis in section 4, we study behaviour following the introduction of small, mean zero risks. For small risks, again only the 2nd moment is relevant: the effect of higher-order moments vanishes.

With this in mind we rewrite programme 7 further as:

$$V_0(W_0, \sigma_1, \sigma_2|\Gamma) = \max_{\{s_0(W_0,\sigma_1,\sigma_2|\Gamma)\}} u(W_0 - s_0) + \mathbb{E}_0\left(V_1\left(s_0, \sigma_1, \sigma_2|\Gamma\right)\right)$$

for some parameters σ_1 and σ_2 governing the 2nd moment of risks in periods 1 and 2. These could govern permanent risk, transitory risk or some combination of the two. A solution to this programme is given by the function $s_0\left(W_0, \sigma_1, \sigma_2|\Gamma\right)$, or more simply $s_0\left(\sigma_1, \sigma_2\right)$.

2.3 Formalization of Complementarity

Let $s_0(\sigma_1, \sigma_2)$ be period-0 saving as a function of future risks, with standard deviations given by σ_1 and σ_2. To repeat, we will assume that, given other features of the environment (such as the shape of the distribution of shocks), we can classify and order risks purely in terms of the standard deviations of income shocks. We further assume that initial-period saving is twice differentiable in these standard deviations. Given a counterfactual risk profile $\langle\sigma_1^*, \sigma_2^*\rangle$, we consider changes in saving from the base case:

$$\Delta s_0(\sigma_1^*, \sigma_2^*) \equiv s_0(\sigma_1^*, \sigma_2^*) - s_0(\sigma_1, \sigma_2)$$

Similarly we denote changes in saving from the base case for each risk in isolation by $\Delta s_0(\sigma_1^*, \sigma_2)$, and $\Delta s_0(\sigma_1, \sigma_2^*)$, (so, for example, $\Delta s_0(\sigma_1^*, \sigma_2) = s_0(\sigma_1^*, \sigma_2) - s_0(\sigma_1, \sigma_2)$). We say that savings exhibit complementarity for this environment if, for $\sigma_1^* > \sigma_1$ and $\sigma_2^* > \sigma_2$:

$$\Delta s_0(\sigma_1^*, \sigma_2^*) < \left(\Delta s_0(\sigma_1^*, \sigma_2) + \Delta s_0(\sigma_1, \sigma_2^*)\right) \tag{8}$$

In terms of the earlier discussion, the right hand side of 8 gives the total precautionary effect of the risks. Taking a Taylor-series expansion of these terms gives:

$$\Delta s_0(\sigma_1^*, \sigma_2^*) - (\Delta s_0(\sigma_1^*, \sigma_2) + \Delta s_0(\sigma_1, \sigma_2^*)) = \frac{1}{2}\Delta\sigma_1\Delta\sigma_2\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2} + \mathcal{O}\left((\Delta\sigma_1)^3, (\Delta\sigma_2)^3\right) \tag{9}$$

where $\Delta\sigma_j \equiv \left(\sigma_j^* - \sigma_j\right)$ Therefore, for small changes in risk, this complementarity depends on the sign of the cross partial derivative $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$. This gives the following definition:

Definition. Savings (s_0) display *complementarity* for utility function $u\,()$, for future risks parametrized by $\langle\sigma_1, \sigma_2\rangle$ and for environment Γ, if $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2} < 0$. If $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2} > 0$ then savings display *excessiveness.*

During the rest of the paper we sometimes refer to $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$ as the complementarity measure.

Further discussion of the definition is merited. First, the definition is in terms of the standard deviation of future risks rather than the variance. Of course $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2} = 4\sigma_1\sigma_2\frac{\partial^2 s_0}{\partial\sigma_1^2\partial\sigma_2^2}$, and so $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$ and $\frac{\partial^2 s_0}{\partial\sigma_1^2\partial\sigma_2^2}$ always have the same sign when σ_1 and σ_2 are positive. Therefore the distinction is inconsequential.

Finally, it is helpful to compare the definition to other possible formalizations of the intuitive notion of complementarity. An obvious alternative is to wonder whether saving is higher or lower if a given amount of future risk is 'bunched' in one particular period or spread out more evenly. Using the notation above, and allowing ϵ to denote a small deviation in a risk parameter we might try to study

$$s_0\left(\sigma + \epsilon, \sigma - \epsilon\right) - s_0\left(\sigma, \sigma\right) \tag{10}$$

In this case, we are bunching risk in the middle period. Corresponding to our intuitive notion of complementarity, we might say that saving is complementary if agents save less when risk is spread out, ie expression 10 is negative.[13]

[13] We must take care in this definition that total ex-ante wealth is held constant, and specifically that risk sequence $\langle\sigma + \epsilon, \sigma - \epsilon\rangle$ yields the same risk to lifetime wealth as does $\langle\sigma, \sigma\rangle$.

Similarly as before, taking a Taylor-series expansion of expression 10 gives

$$\epsilon\left(\frac{\partial s_0}{\partial \sigma_1}(\sigma,\sigma)-\frac{\partial s_0}{\partial \sigma_2}(\sigma,\sigma)\right)+\frac{1}{2}\epsilon^2\left(\frac{\partial^2 s_0}{(\partial \sigma_1)^2}(\sigma,\sigma)-2\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}(\sigma,\sigma)+\frac{\partial^2 s_0}{(\partial \sigma_2)^2}(\sigma,\sigma)\right)+O\left(\epsilon^3\right)$$

We immediately see that this definition is, in fact, less interpretable. It combines first-order and second-order derivatives of the savings function: it is not clear whether or not $\left(\frac{\partial s_0}{\partial \sigma_1}(\sigma,\sigma)-\frac{\partial s_0}{\partial \sigma_2}(\sigma,\sigma)\right)$ is zero and, if not, what sign it takes. Its sign might depend on other features of the environment. In contrast, we will see that $\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2}$ can be characterized more cleanly. This latter definition is more useful for related but distinct questions, such as what is the arrangement of risk such that precautionary saving is minimized? This may be an interesting question if the amount of precautionary saving provides a good approximation to the welfare cost of risk.

3 Quantifying Complementarity of Precautionary Saving

We begin the analysis by numerically simulating a realistic life-cycle consumption and savings model. In the following section (section 4) we return to the three period model to examine the conditions for complementarity in more theoretical detail.

We choose the following basic parametrization: the household is 'born' at age 20, and receives labour income for 45 years (from 20 to 65). The household then retires for another 15 years, in which time it finances consumption out of savings. We set $R = 1.02$ and $\beta R = 1$. Initial income is 1 unit. Income grows at 1% per year over working age. The household has CRRA preferences with coefficient of relative risk aversion of 2. This is close to the standard coefficient estimated using microdata such as in Attanasio and Weber (1995) and is lower than estimates used in the macro literature (see Barro (2006)).

In keeping with the 3-period framework, we divide the planning horizon into three tranches of fifteen years, corresponding to ages 21-35, 36-50 and 51-65.[14] We assume no income risk in the first tranche and then a constant yearly variance of permanent and transitory shocks in the final two tranches: in our base ('high') scenario, the yearly variance of permanent shocks is 0.02 (standard deviation of 14%); the yearly variance of transitory shocks is 0.04 (standard deviation of 20%). Income fluctuations are lognormally distributed. We exclude risk in the first 15 periods because we are interested in precautionary saving for risk in the medium and long term. By excluding risk in the first tranche we are able to identify pure life-cycle saving when we switch off risk in these latter periods. Allowing risk in the first tranche does not affect results but obscures their interpretation somewhat.

We then run the following experiments. First we switch off all risk over all periods. Second we keep risk (both permanent and transitory shocks) in the middle tranche alone. Third we allow for risk in the final tranche alone. Finally we allow for risk in both the middle and late tranches as standard. As discussed in section 2, when running these experiments we are careful to account for the change in the variance of life-time wealth precisely. In short, we reweight the shocks when the household faces both tranches of risk. Using unadjusted shocks, the variance of life-time wealth would be higher when facing both tranches of risk than the total variance from each risk in isolation because the variances of a multiplicative process do not sum precisely. Further details are given in appendix A.

Table 1 shows the results from these simulations. It shows saving at the end of the 15th period, just before income risk kicks in. The first column shows results for the baseline ('high') environment. The first row shows savings when there is no income risk. This therefore represents pure life-cycle saving caused by the pattern of life-cycle income and the retirement period. The second row shows accumulated saving in the standard model with risk in both middle and final tranches. This accounts for both life-cycle and precautionary saving. The third row subtracts life-cycle saving (row 1) to leave precautionary saving alone. The fourth row shows saving when there is just risk in the middle tranche. The fifth row shows saving when there is risk in the final

[14] The household lives for another 15 years after retiring but faces no risk, so faces a trivial planning problem.

tranche alone. Note that precautionary saving is much larger in the 4th row than the 5th, because permanent shocks in mid life persist until retirement and so have a greater effect on life-time wealth. The sixth row shows the sum of the precautionary saving in these two scenarios (row 4 + row 5 - 2×row 1). This can be thought of as the total precautionary effect of the two tranches of risk. The seventh row shows the difference between the total precautionary effect and benchmark precautionary saving (row 6 - row 3). We interpret this as the complementarity effect of precautionary saving for this environment. The household can save the equivalent of 40% of its initial yearly income less because of the sequencing of risks. The final row represents this complementarity effect as a percentage of precautionary saving (in row 3): complementary saving is around 16% of total precautionary saving, a considerable sum.

The remaining columns of table 1 show the equivalent results when we vary the parameters. The second column (the 'low' scenario) shows saving and complementarity when we reduce the variance of permanent shocks to 0.01 and the variance of transitory shocks to 0.02. The complementarity effect is reduced to around 15% of initial income and a little over 9% of precautionary wealth. The lower variance of permanent shocks itself reduces precautionary saving by around 40%. The third column shows saving and complementarity when we remove the permanent shocks but keep a variance of transitory shocks of 0.04 per year. Precautionary saving is now excessive, it is greater than the total precautionary effect. The size of precautionary savings is tiny however because these transitory shocks pose little risk to life-time wealth.

The 4th and 5th columns show results with the same configurations of income risk but an enhanced life-cycle motive. In this scenario the household lives for 20 years after retirement (dies at 85). Moreover we set income to be flat in the final third of working life (income still grows at 1% pa until age 50). With the higher variance of income shocks the complementarity effect is slightly reduced to 14.2%, indicating that the higher life-cycle saving crowds out the complementarity effect from precautionary saving. Comparing the 5th to the 2nd columns we see a similar effect: complementarity is reduced from 9.3% to 8.4%.

The final two columns of table 1 show results when the baseline variance of permanent shocks is held at 0.02 and the variance of transitory shocks at 0.04, but some other feature of the environment is changed. In the 6th column we reduce the coefficient of relative risk aversion to 1. Total precautionary saving is reduced by 25% and the complementarity effect is reduced to 12.5% of precautionary saving. The final column shows results when we don't reweight permanent shocks. The results are very similar here to the base case, indicating that the results are not caused by how we account for the variance of life-time wealth.

Figure 1 links our results to the formal definition of complementarity. This figure shows numerical calculations of the complementarity measure $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$ at various levels of risk. The horizontal axes display the variances of permanent shocks. We tie the variances of transitory shocks to be double those of the permanent shocks. The vertical axis shows the amount of complementarity to small changes in risk around these levels. We term this 'local' complementarity. This figure is computed by solving the model numerically on a 30-by-30 grid of income risks. A negative amount here shows local complementarity. A positive amount would show local excessiveness. Here there is local complementarity at all levels of risk except at zero. Of course, the results in table 1 could be obtained by integrating this surface over larger changes in risk.

Table 1: Simulated Household Wealth at Age 35 as % of Initial Income

	Environment	Baseline			High life-cycle		Low risk aversion	No re-weighting
	Variance of perm. shocks	High	Low	Zero	High	Low	High	High
(1)	No risk	68	68	68	176	176	68	68
(2)	Risk throughout career	322	225	71	388	306	261	325
(3)	Precautionary wealth*	254	157	3.39	212	131	193	257
(4)	Risk in mid-career only	306	207	70	381	296	244	306
(5)	Risk in late-career only	124	100	70	212	197	109	124
(6)	Total prec. effect**	294	172	3.37	242	141	217	294
(7)	Complementarity effect***	41	15	-0.012	30	11	24	37
(8)	as % of prec. wealth****	16.0	9.3	-0.4	14.2	8.4	12.5	14.4

Notes: * = row(2)-row(1)
** = (4)-(1) + (5)-(1)
*** = (6)-(3)
**** = 100×(7)/(3)
- 'High' variance of permanent shocks is 0.02 per year. 'Low' is 0.01 per year. In 'zero' scenario we remove permanent shocks completely but keep transitory shocks with a variance of log of 0.04 per year.
- In 'baseline' scenario expected income grows at 1% pa; households are 'born' at age 20 and work for 45 years, then live for 15 years in retirement. In 'high life-cycle' scenario income growth is flat in the last third of working life and retirement lasts 20 years. In 'no re-weighting' scenario, the variances of permanent shocks are not reweighted.
See text for more details.

Figure 1: Precautionary Saving is Complementary for Overall Life-Cycle Risks

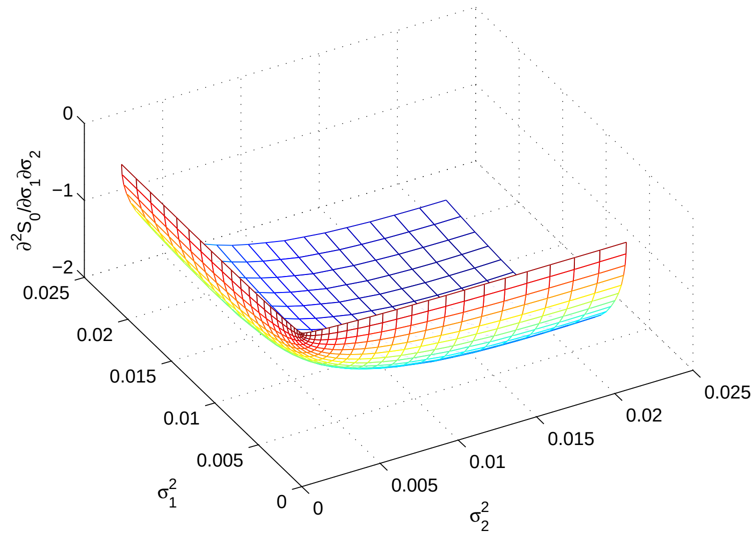

Notes: $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$ on the vertical axis is the complementarity measure. Horizontal axes show baseline variances of permanent shocks in mid-career (σ_1^2) and in late career (σ_2^2).

We set $\beta = 0.98$ and $\beta R = 1$. Initial income is 1 unit and income growth is 1% per year over the whole life time. Income fluctuations are lognormally distributed. The coefficient of relative risk aversion is 2. Households work for 45 years then live for 15 years in retirement. Initial-period saving is used. See text for more details.

4 Analytical Characterization of Complementarity

In this section we study the complementarity effect further by obtaining an analytical characterization of household saving in the three-period version of the model.

4.1 Complementarity for Permanent Shocks

We first consider the 3-period consumption and saving model given in section 2 when the agent faces only permanent risk. We also restrict attention to small risks in period 1 and 2 with standard deviations evaluated at $\sigma_1 = \sigma_2 = \sigma$, and for $\beta = R = 1$. When preferences are of the CRRA form we find that:

$$\begin{aligned}\frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2} &\approx -\mathcal{A}\sigma^2 c\left(p_r\left(c\right)+\frac{5}{6}p_r\left(c\right)^2\right) \\ &< 0\end{aligned} \tag{11}$$

where $\frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2}$ is our complementarity measure, $\mathcal{A}$ is some positive constant, σ^2 is the variance of the risks, c is first period consumption (which, for small risks, approximately equals flow income), and $p_r\left(c\right)$ is the coefficient of relative prudence, defined as $-\frac{cu'''(c)}{u''(c)}$.

We know that CRRA preferences must take the form $u\left(x\right)=\frac{x^{(1-\gamma)}}{1-\gamma}$, in which case the coefficient of relative prudence is $1+\gamma$. This result therefore indicates that CRRA preferences must permit complementarity ($\frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2}<0$) and that complementarity is increasing in the prudence parameter, alongside total precautionary saving. To link this result to the earlier simulations, figure 2 shows the complementarity measure for permanent risks and for CRRA preferences, obtained from simulations from the 3-period model. This figure extends the analytical result by showing that complementarity holds (i.e. the surface is negative) for small changes of risk, not just around the perfect certainty baseline, but also at higher levels of initial risk. The figure also shows that complementarity holds when the baseline variances of innovations (σ_1^2 and σ_2^2) are not equal.

We now explain how we arrive at condition 11, while also analysing saving for preferences other than CRRA. For general preferences, we obtain the following result, derived in appendix B:

$$\begin{aligned}\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2} &= \\ &-\sigma^2\frac{\mathcal{A}_1 c_0^3\left(-c_0u^{(3)}(c_0)^3-2c_0u^{(4)}(c_0)u^{(3)}(c_0)u''(c_0)+12u^{(4)}(c_0)u''(c_0)^2+3c_0u^{(5)}(c_0)u''(c_0)^2\right)}{u''(c_0)^3} \\ &+\mathcal{O}\left(\left(\sigma\right)^3\right)\end{aligned} \tag{12}$$

where $\mathcal{A}_1$ is another positive constant $u^{(n)}\left(\right)$ denotes the n^{th} derivative of $u\left(\right)$ and $c_0=W_0-s_0$ is period-0 consumption. Again σ_1,σ_2 are the standard deviations of wealth innovations in each subsequent period, and we look at $\sigma_1=\sigma_2=\sigma$. $\left(\sigma\right)^3$

Figure 2: Precautionary Saving is Complementary for Permanent Risks Alone

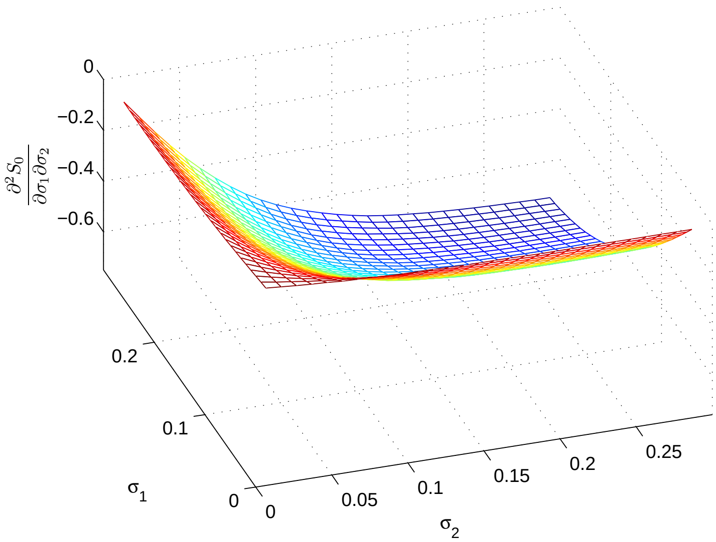

Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$: the complementarity measure. On the horizontal axes are standard deviations of wealth innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Preferences are CRRA with coefficient of relative risk aversion of 2.

denotes the cube of the standard deviation (and not some parameter of the third moment of the income distribution). To emphasize, we re-weight the risks as described in section 2 and appendix A.

It turns out that we can still reformulate the complementarity measure in equation 12 in terms of the coefficient of relative prudence, $p_r(c)$. Substituting terms we obtain:

$$\frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2} \approx \mathcal{A}_1\sigma^2 c\left(c^2 p_r''(c) - \frac{7}{3}cp_r'(c)\,p_r(c) + 2cp_r'(c) - 2p_r(c) - \frac{5}{3}p_r(c)^2\right) \tag{13}$$

where we have dropped the 0 subscripts on s and on c.

When the coefficient of relative prudence is constant, as in CRRA preferences, then approximation 13 reduces to expression 11. But approximation 13 yields some further, more general, results. First, notice that complementarity is linear in the variance of risk, σ^2, in this general setting. Second, condition 13 indicates that the complementarity measure is negative as long as the relative prudence function is near flat. In fact, it is easy to show that the general HARA class of preferences, of which CRRA is an example, induces complementarity. To see this this, note that HARA functions have coefficient of relative prudence $p_r(x) = \frac{x}{ax+b}$ for $b \geq 0$ and $0 \leq a < 1$. Therefore the complementarity measure in approximation 13, with HARA preferences, reduces to:

$$\begin{aligned} \frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2} &\approx -\frac{c_0^3(12(a+1)b + a(6a+5)c_0)}{3(ac_0+b)^3} \\ &< 0 \end{aligned}$$

Intuitively, HARA preferences have relative prudence which is weakly increasing and weakly concave, properties which fit many of the components of condition 13. As a final point, note that, for HARA preferences, complementarity is stronger the smaller are a and b.

4.1.1 Preferences for Which Complementarity Breaks Down

It is worth considering preferences which don't yield complementarity. Condition 13 implies that this might happen when prudence is (locally) very convex (i.e. $p_r''(c)$ is +ve). A good example of such behaviour comes from Stone-Geary preferences of the form $u(c) = \ln(c - \underline{c})$, where $\underline{c}$ is a minimum consumption need. For these preferences the relative prudence term is $p_r(c) = \frac{2c}{c-\underline{C}}$ and the complementarity measure given by equation 13 simplifies to $-\frac{\mathcal{A}c^3(4c-9\underline{C})}{(c-\underline{C})^3}$.[15] This term is positive for any $\underline{c} < c < \frac{9}{4}\underline{c}$ and hence savings are in fact excessive in this range. By way of explanation, note that, for these preferences, relative prudence rises to infinity as c approaches $\underline{c}$ from above. Clearly, households are particularly prudent when they are near the consumption floor. It also seems they are especially averse to a series of consecutive negative shocks, and so the consecutive risks amplify precautionary behaviour.

[15] Notice that the prudence term, $p_r(c) = \frac{2c}{c-\underline{C}}$ is similar to that for the HARA class, except that the additive term in the denominator is negative.

4.2 Excessiveness for Small Transitory Risks

We now turn to the characterization of period-0 saving when income shocks in periods 1 and 2 are transitory. For small risks with standard deviations $\sigma_1 = \sigma_2 = \sigma$, and for $\beta = R = 1$ we obtain the following result, also derived in appendix B:

$$\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2} = -\frac{\mathcal{B}\sigma^2\left(u^{(3)}(c_0)^3 + 2u^{(4)}(c_0)u^{(3)}(c)u''(c_0) - 3u^{(5)}(c_0)u''(c_0)^2\right)}{u''(c_0)^3} + \mathcal{O}\left((\sigma)^3\right) \quad (14)$$

where $\mathcal{B}$ is another positive constant and other terms are as before.

We can most conveniently reformulate equation 14 in terms of absolute prudence concepts. Letting $p_a(c)$ denote the coefficient of absolute prudence, given by $-u'''(c)/u''(c)$, and dropping the 0 subscript, and substituting terms, then we obtain:

$$\frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2} \approx \mathcal{B}_1\sigma^2\left(p''_a(c) - \frac{7}{3}p'_a(c).p_a(c)\right) \quad (15)$$

for another positive constant B_1.

For CRRA preferences the coefficient of absolute prudence is $\frac{1+\gamma}{c}$. For these preferences $p''_a(c) = \frac{2(1+\gamma)}{c^3} > 0$ and $p'_a(c) = -\frac{(1+\gamma)}{c^2} < 0$. Therefore, $\frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2} > 0$ and savings exhibit not complementarity in this case but excessiveness. Again to link to the earlier simulations, figure 3 shows $\frac{\partial^2 s}{\partial\sigma_1\partial\sigma_2}$ for transitory risks and for the CRRA utility function, obtained from simulating the 3-period model. The figure shows that saving is excessive at all levels of baseline risk and even when the risks are not identical. It further shows that the absolute height of the function is comparable to the absolute magnitude of the complementary measure for permanent shocks given in figure 2. However, transitory risks in the 3-period model have a far greater effect on life-time wealth than they do in a model with a longer time horizon. As discussed in section 2 we emphasize that the effect of transitory shocks are actually quantitatively small for realistic parametrizations.

By way of interpretation, approximation 15 shows that saving is excessive if prudence is declining and convex, because then $p'_a(c) < 0$ and $p''_a(c) > 0$. An intuitive reason for this sufficient condition is as follows: the wealth innovations caused by transitory

Figure 3: Precautionary Saving is Excessive for Transitory Risks Alone

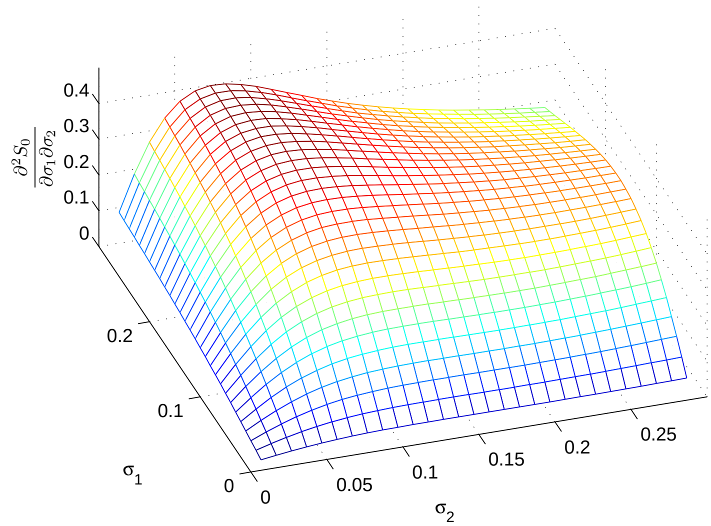

Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$: the complementarity measure. On the horizontal axes are standard deviations of wealth innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Preferences are CRRA with coefficient of relative risk aversion of 2.

income shocks are independent. If prudence is declining and convex, average prudence is greater the more independent risks are added (a consequence of Jensen's inequality). Therefore the interaction of risks amplifies the precautionary saving motive. It seems intuitively plausible that prudence be declining and convex because this implies that wealthier households have less need to avoid a gamble (of constant variance) but that the rate of decline of prudence is diminishing with wealth. Indeed this condition on prudence is satisfied by all preferences in the HARA class for which risk tolerance is not constant, i.e. notably excluding CARA and quadratic preferences.[16]

[16]The result for HARA preferences is easily shown by differentiating the coefficient of risk tolerance $-\frac{u'(c)}{u''(c)}$ with respect to wealth. Re-arranging we find that $p(c) = k.r(c)$ for some constant k and where $r\,(c)$ is the coefficient of absolute risk aversion $= -\frac{u''(c)}{u'(c)}$. Note that $r(c)$ is decreasing and convex for HARA functions (except those with constant risk tolerance), therefore so must be $p(c)$.

4.3 Further Intuition

The obvious question is why CRRA preferences induce excessive saving for transitory shocks but complementary saving for permanent shocks. Paradoxically we gain the best intuition for CRRA preferences by considering saving under CARA preferences. For CARA preferences, the first and second derivatives of absolute prudence are zero. Therefore, expression 15 implies that, for transitory risks, saving is neutral: it is neither excessive nor complementary. It is easily checked that saving is also neutral for larger risks (see Caballero (1990)). Figure 4 shows this result graphically. This result is obvious when we remember two facts. First, transitory risks are independent across time. Second, the wealth level is irrelevant under CARA preferences when saving for a given risk. From the viewpoint of the middle period therefore, the precautionary motive for the final-period risk does not depend on the outcome of the most recent shock. Returning to the initial period, the household must save for both middle- and final-period risks. However, because the middle-period outcome will not change behaviour, the size of the middle risk does not affect the precautionary motive for the final risk. Symmetrically, the size of the final risk does not affect the precautionary motive for the middle risk. In short, the risks don't interact in the initial saving decision.

In contrast, and just like CRRA preferences, saving is complementary for permanent shocks. Figure 5 shows this result. To see why, notice that the argument above for why CARA preferences induce neutral savings no longer holds. In this case the variance of final-period shocks now depends on the middle-period outcome: the risks are no longer independent. In fact the variance of middle-period shocks now decreases, the lower is the outcome in the first period. A prudential saver is most concerned about the lowest possible outcomes and places most weight on these outcomes for decision making. Therefore, precautionary savings to cover both risks combined need not to be as high as the sum of savings to match each risk on its own. In short, it seems the risk pattern itself provides a kind of insurance.

Figure 4: CARA Preferences: Saving is Neutral for Transitory Risks

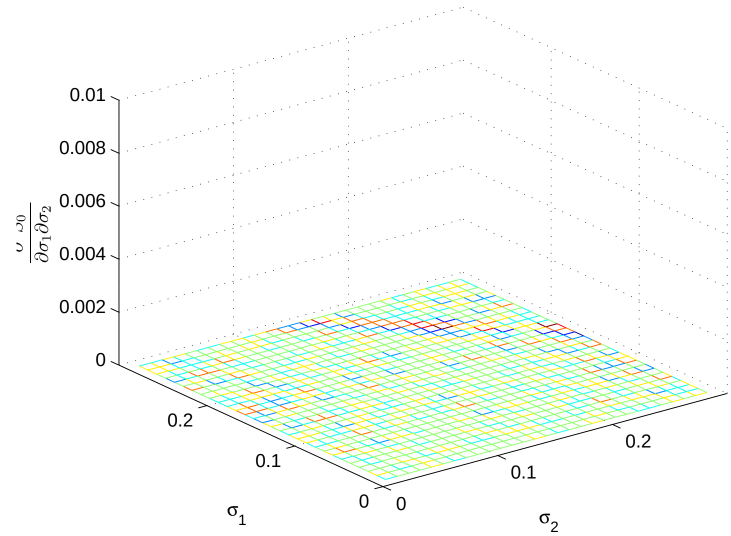

Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$: the complementarity measure. On the horizontal axes are standard deviations of wealth innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Coefficient of absolute risk aversion is 3.

4.4 Classifying Utility Functions

We can classify utility functions according to whether they induce complementary or excessive saving more generally. Comparing equations 14 and 12, the complementary function for permanent shocks is lower than for transitory shocks by a term in $u^{(4)}(c)u''(c)^2$. We therefore classify utility functions as follows:

Proposition. *As long as $u^{(4)}(c)$ is negative then if a utility function induces excessive saving for permanent shocks then it also induces excessive saving for transitory shocks.*

Saving is excessive for Stone-Geary preferences and permanent risks when consumption is near the 'breadline'. Therefore saving is excessive for these preferences and transitory shocks also. On the other hand, HARA preferences lie in the middle ground: they induce complementarity for permanent shocks but excessiveness for transitory shocks. To complete the classification it is instructive to construct a utility function that is complementary for both permanent and transitory shocks. According to the discussion

Figure 5: ...While Saving is Complementary for Permanent Risks

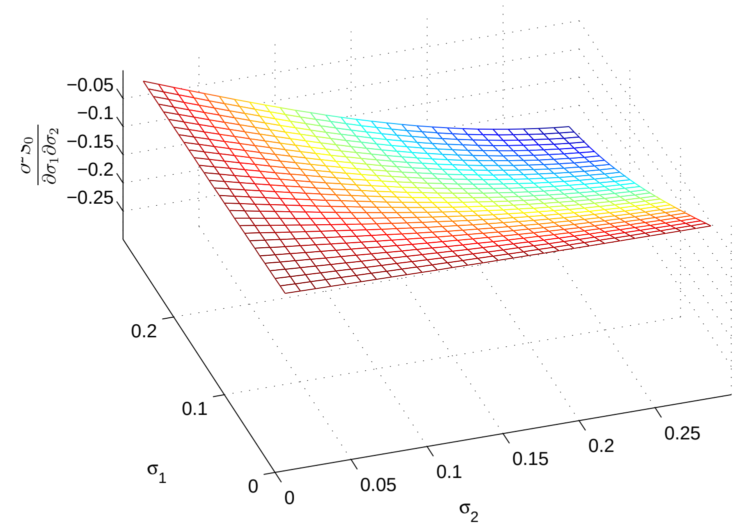

Notes: Figure shows results from the 3-period model. The vertical axis shows $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$: the complementarity measure. On the horizontal axes are standard deviations of wealth innovations in the middle and final periods. These innovations are binary and symmetric. Initial wealth is 1 unit. Coefficient of absolute risk aversion is 3.

in section 4.2, a utility function induces complementarity with transitory shocks if it has increasing absolute prudence. An example is the following:

$$u(c) = -e^{-\frac{c+1^2}{2}} - \sqrt{\frac{\pi}{2}}(c+1) \times \mathrm{erf}\left(\frac{c+1}{\sqrt{2}}\right) + 2(c+1)$$

where $\mathrm{erf}(x)$ represents the error function, the anti-derivative of $\frac{2}{\sqrt{\pi}}e^{-x^2}$. This utility function has positive third derivative and negative fourth derivative and coefficient of absolute prudence of c. Such a utility function has many undesirable features.[17] Nevertheless, a household with these preferences has complementary saving even for transitory shocks because it is little affected by a negative shock in the middle period. In fact, the household has less desire to save following a bad shock, because absolute

[17]For example, it does not obey Inada conditions because $u(0) = 2$. Moreover at high levels of wealth it displays high prudence but low risk aversion. The high prudence arises because, even though households do not lose much expected utility from risk, their elasticity of substitution is so high that they equally lose little from allocating consumption to the future. The latter affect dominates so they precautionary save and increasingly so at higher levels of wealth.

prudence is lower with lower wealth. Returning to the initial period, the household therefore need not save much more for consecutive risks than for just a single risk. It is happier to save little and to have more equal consumption across time in expectation.

4.5 Concluding Remarks

To our knowledge, no such similar conditions have been derived before. However the results for transitory shocks relate to the literature on multiple risk bearing in static settings in papers by Pratt and Zeckhauser (1987), Kimball (1993) and Gollier and Pratt (1996), summarized neatly in the last paper. These papers elucidate the related and intuitively attractive notions of standard risk aversion (Kimball), proper risk aversion (Pratt and Zeckhouser) and risk vulnerability (Gollier and Pratt). All are formalizations of the idea that background risks should make agents more averse to new risks; for example an agent should be more averse to investing in equities if exposed to high labour-market risk.[18] While this literature concerns risk aversion and portfolio choice, it seems intuitive that such effects carry over to precautionary saving when risks are independent (i.e. transitory). The results for permanent shocks of course relate to non-independent risks. Nevertheless it is striking that for these two standard specifications of risk the results should be so contrasting.

5 Conclusion

In this paper we define and examine the concept of complementarity in precautionary saving. Intuitively, precautionary saving should be complementary because of its contingent nature. Emergencies occur only rarely and precautionary savings are only rarely needed. Therefore the accumulated stock can presumably be put to other ends, specifically as rolled-over precautionary savings against subsequent income risk.

[18]These restrictions on utility functions can be summed up neatly in notation related to that above. If $r(c) \equiv -u''(c)/u'(c)$ is the coefficient of absolute risk aversion at consumption/wealth level c, then a necessary and sufficient condition for standard risk aversion is that both $r(c)$ and $p(c)$ be decreasing.

On the quantitative side, we simulate the standard life-cycle consumption and saving model with both permanent and transitory income risks. In a range of realistic parametrizations, the total precautionary effect of the risks is 8-16% higher than precautionary saving for the consecutive risks. This complementarity effect is driven entirely by the permanent shocks. In contrast, as these shocks get less persistent eventually complementarity is overturned and the risks amplify precautionary saving, although the effect is small. We conclude that the complementarity effect quantitatively dominates.

We then study the complementarity effect in more detail by analyzing a stripped-down version of the model with only 3 periods. We show that permanent shocks admit complementary savings for a general class of preferences, notably including the standard CRRA form. However, we find two instances which depart from our basic intuition. First, consecutive transitory shocks amplify precautionary savings for standard preferences. Second, even consecutive permanent risks can amplify precautionary savings when households have minimum consumption needs. The effect from transitory shocks is small, however, for empirically-plausible risk sizes. Moreover, the 'breadline' effect is only local: saving is complementary for wealthier households. We conclude that complementary saving is the norm.

The present study could be extended in several ways. First, it would be interesting to see how liquidity constraints affect the results. In general, constraints exacerbate the precautionary motive (Carroll and Kimball (2001)). However, constraints will not bind in expectation. The intuition above therefore carries through: households facing a constraint should not need to save much more for consecutive risks than for just a single risk.

A related question is how the distribution of shocks affects the results. As a specific example, what if households can receive zero (or very small) income in any period, for example due to unemployment? The analysis presented in section 4 cannot answer this question because it applies only to small, local risks. This question could, however, be studied quantitatively.

Finally, this paper suggests a much broader research agenda. Recalling the introductory comments we may ask the following questions: how do other savings motivations complement each other? How does the presence of different savings technologies affect this complementarity? What are the implications for household saving if pension and housing wealth become more liquid, perhaps because of financial innovation?

References

Attanasio, O. and G. Weber (1995): "Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey," *Journal of Political Economy*, 103, 1121–1157.

Barro, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, 823–866.

Blundell, R. and T. M. Stoker (1999): "Consumption and the Timing of Income Risk," *European Economic Review*, 43, 475–507.

Browning, M. and A. Lusardi (1996): "Household saving: Micro theories and micro facts," *Journal of Economic Literature*, 34, 1797–1855.

Caballero, R. (1990): "Consumption puzzles and precautionary savings," *Journal of monetary economics*, 25, 113–136.

Carroll, C. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis*," *Quarterly Journal of Economics*, 112, 1–56.

——— (2004): "Theoretical foundation of buffer stock saving," *NBER Working Paper*.

Carroll, C. and M. Kimball (2001): "Liquidity constraints and precautionary saving," *NBER Working Paper*, wp8496.

Carroll, C. and A. Samwick (1997): "The nature of precautionary wealth," *Journal of monetary Economics*, 40, 41–71.

De Nardi, M., E. French, and J. Jones (2010): "Why do the elderly save? the role of medical expenses," *Journal of Political Economy*, 118, 39–75.

Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 1221–1248.

Dynan, K. (1993): "How prudent are consumers?" *Journal of Political Economy*, 101, 1104–1113.

Eeckhoudt, L., C. Gollier, and N. Treich (2005): "Optimal consumption and the timing of the resolution of uncertainty," *European Economic Review*, 49, 761–773.

Eeckhoudt, L. and H. Schlesinger (2008): "Changes in Risk and the Demand for Saving," *Journal of Monetary Economics*, 55, 1329–1336.

Franke, G., H. Schlesinger, and R. Stapleton (2011): "Risk taking with additive and multiplicative background risks," *Journal of Economic Theory*, 146, 1547–1568.

Goldman, D. and N. Maestas (2013): "Medical expenditure risk and household portfolio choice," *Journal of Applied Econometrics*, 28, 527–550.

Gollier, C. (2004): *The economics of risk and time*, The MIT Press.

Gollier, C. and J. W. Pratt (1996): "Risk vulnerability and the tempering effect of background risk," *Econometrica*, 1109–1123.

Goodman, L. (1960): "On the exact variance of products," *Journal of the American Statistical Association*, 55, 708–713.

Gruber, J. and A. Yelowitz (1999): "Public health insurance and private savings," *Journal of Political Economy*, 107, 1249–1274.

Guiso, L., T. Jappelli, and M. Padula (2009): "Pension Risk, Retirement Saving and Insurance," *EUI Working Paper*, EUI ECO; 2.

Kimball, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica*, 53–73.

——— (1993): "Standard risk aversion," *Econometrica*, 589–611.

Pratt, J. W. and R. J. Zeckhauser (1987): "Proper risk aversion," *Econometrica*, 143–154.

A Reweighting The Permanent Shock Process

This appendix gives more detail on the reweighting of permanent income shocks discussed in section 3.

We consider permanent shocks ψ_i for $i = 1...\mathrm{T}$ with mean 1 and variance $\sigma^2_{\psi_i}$. By Goodman's rule[19] it is easy to see that $\mathrm{Var}\left(y_t\right) = \mathrm{Var}\left(\Pi_i^t \psi_i\right) = \Pi_{i=1}^t \left(\sigma^2_{\psi_i} + 1\right) - 1$.[20] In this discussion we ignore the first, risk-free tranche of the life-cycle and consider two tranches: $1..\mathrm{T}_0$ and $\mathrm{T}_0..\mathrm{T}$, both containing permanent risk.

The problem is the following. When the household faces risk in the first tranche of working life alone then at time $\mathrm{T}_0 + 1$, we have $\sigma^2_{\psi_{T_0+1}} = 0$ and so:

$$\mathrm{Var}\left(y_{T_0+1}\right) = \mathrm{Var}\left(\Pi_{i=1}^{T_0} \psi_i\right)$$

Similarly when the household faces risk in the second tranche alone then

$$\mathrm{Var}\left(y_{T_0+1}\right) = \sigma^2_{\psi_{T_0+1}}$$

When the household faces both risks combined then

$$\begin{aligned}
\mathrm{Var}\left(y_{T_0+1}\right) &= \mathrm{Var}\left(\Pi_{i=1}^{T_0+1} \psi_i\right) \\
&= \Pi_{i=1}^{T_0+1} \left(\sigma^2_{\psi_i} + 1\right) - 1 \\
&= \Pi_{i=1}^{T_0} \left(\sigma^2_{\psi_i} + 1\right) \left(\sigma^2_{\psi_{T_0+1}} + 1\right) - 1 \\
&= \Pi_{i=1}^{T_0} \left(\sigma^2_{\psi_i} + 1\right) - 1 + \Pi_{i=1}^{T_0} \left(\sigma^2_{\psi_i} + 1\right) \sigma^2_{\psi_{T_0+1}} \\
&> \Pi_{i=1}^{T_0} \left(\sigma^2_{\psi_i} + 1\right) - 1 + \sigma^2_{\psi_{T_0+1}} \\
&= \mathrm{Var}\left(\Pi_{i=1}^{T_0} \psi_i\right) + \sigma^2_{\psi_{T_0+1}}
\end{aligned}$$

Therefore income risk in period $\mathrm{T}_0 + 1$ is greater when the household faces both risks

[19] For uncorrelated random variables $\mathrm{Var}\left(xy\right) = \mathrm{Var}\left(x\right) \mathrm{E}\left(y\right)^2 + \mathrm{Var}\left(y\right) \mathrm{E}\left(x\right)^2 + \mathrm{Var}\left(x\right) \mathrm{Var}\left(y\right)$.

[20] We ignore mean income growth here without loss of generality.

combined. By a simple induction argument we can show that the same holds for risks in time t for any t>T_0.

To solve this problem we reweight the risks to make the variance of lifetime wealth the same in all scenarios. To do this we create an alternative sequence of income shocks $\tilde{\psi}_i$ with variances $\sigma^2_{\tilde{\psi}_i}$ for $i > T_0$. We do this iteratively from period T_0 onwards as follows: when facing risk in the first tranche only, the variance of income $t > T_0$ is given as before by $\text{Var}\left(\Pi_{i=1}^{T_0}\psi_i\right)$. When facing risk only in the second tranche then income risk at time t is given again as before by $\text{Var}\left(\Pi_{i=T_0+1}^{t}\psi_i\right)$. When facing both risks combined and with the alternative income process the variance of income is

$$\begin{aligned}\text{Var}\left(\tilde{y}_t\right) &= \text{Var}\left(\Pi_{i=1}^{t}\tilde{\psi}_i\right) \\ &= \text{Var}\left(\Pi_{i=1}^{t-1}\tilde{\psi}_i\right)\sigma^2_{\tilde{\psi}_t} + \sigma^2_{\tilde{\psi}_t} + \text{Var}\left(\Pi_{i=1}^{t-1}\tilde{\psi}_i\right) \qquad (16)\end{aligned}$$

Setting this equal to the sum of the two isolated risks added together gives:

$$\text{Var}\left(\Pi_{i=1}^{t-1}\tilde{\psi}_i\right)\sigma^2_{\tilde{\psi}_t} + \sigma^2_{\tilde{\psi}_t} + \text{Var}\left(\Pi_{i=1}^{t-1}\tilde{\psi}_i\right) = \text{Var}\left(\Pi_{i=1}^{T_0}\psi_i\right) + \text{Var}\left(\Pi_{i=T_0+1}^{t}\psi_i\right)$$

solving for $\sigma^2_{\tilde{\psi}_t}$:

$$\sigma^2_{\tilde{\psi}_t} = \frac{\text{Var}\left(\Pi_{i=1}^{T_0}\psi_i\right) + \text{Var}\left(\Pi_{i=T_0+1}^{t}\psi_i\right) - \text{Var}\left(\Pi_{i=1}^{t-1}\tilde{\psi}_i\right)}{1 + \text{Var}\left(\Pi_{i=1}^{t-1}\tilde{\psi}_i\right)}$$

B Derivation of the Approximations for $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$

This appendix gives more details of the derivation of the formulae in section 4

B.1 Translating the Budget Constraint

For the 3 period model with $\beta = R$ and with constant expected income we have household problem

$$V_0(W_0) = \max_{\{c_t(W_t):t=0,1,2\}} u(c_0) + \mathbb{E}_0 \left(u(c_1) + u(c_2)\right)$$

subject to

$$c_0 + c_1 + c_2 = y_0 + y_0\psi_1\xi_1 + y_0\psi_1\psi_2\xi_2$$

It is convenient to rewrite the budget constraint in terms of relative wealth innovations.

In the case of transitory risk the budget constraint is:

$$c_0 + c_1 + c_2 = y_0 + y_0\xi_1 + y_0\xi_2$$

We make the transformation that $a_0 = 3y_0$, $\epsilon_1 = \frac{1}{3}(\xi_1 - 1)$, and $\epsilon_2 = \frac{1}{3}(\xi_2 - 1)$.[21] We can then write the budget constraint as

$$c_0 + c_1 + c_2 = a_0 + a_0\epsilon_1 + a_0\epsilon_2$$

such that ϵ_1 and ϵ_2 have mean zero and $\text{Var}(\epsilon_1) = \frac{1}{9}\sigma^2_{\xi_1}$ and $\text{Var}(\epsilon_2) = \frac{1}{9}\sigma^2_{\xi_2}$.

In the case of permanent shocks alone we can write the constraint

$$c_0 + c_1 + c_2 = y_0 + y_0\psi_1 + y_0\psi_1\tilde{\psi}_2$$

We make the transformations $a_0 = 3y_0$, $\epsilon_1 = \frac{2}{3}(\psi_1 - 1)$, and $\epsilon_2 = \frac{1}{3}\frac{(\psi_2-1)}{\sqrt{1+\sigma^2_{\psi_1}}}$.[22] We can then write the budget constraint as:

$$c_0 + c_1 + c_2 = a_0 + a_0\epsilon_1 + a_0\frac{(1 + \frac{3}{2}\epsilon_1)}{\sqrt{1 + \frac{9}{4}\text{Var}(\epsilon_1)}}\epsilon_2$$

[21] More generally, we can allow for non-constant expected income over the life cycle by allowing for $\epsilon_1 = \frac{y_1}{\sum y_i}(\xi_1 - 1)$, $\epsilon_2 = \frac{y_2}{\sum y_i}(\xi_2 - 1)$.

[22] Again more generally, we can allow for non-constant expected income over the life cycle by allowing for $\epsilon_1 = \frac{y_1+y_2}{\sum y_i}(\psi_1 - 1)$, $\epsilon_2 = \frac{y_2}{\sum y_i}\frac{(\psi_2-1)}{\sqrt{1+\sigma^2_{\psi_1}}}$.

such that ϵ_1 and ϵ_2 have zero mean and $\mathrm{Var}\left(\epsilon_1\right) = \frac{4}{9}\sigma^2_{\psi_1}$ and $\mathrm{Var}\left(\epsilon_2\right) = \frac{\sigma^2_{\psi_2}}{9\left(1+\sigma^2_{\psi_1}\right)}$. We have weighted ϵ_1 and ϵ_2 appropriately so that the variance of life-time wealth is held constant when we sum risks.

And we nest the permanent and transitory cases by writing the budget constraint:

$$c_0 + c_1 + c_2 = a_0 + a_0\epsilon_1 + a_0\frac{(1+\frac{3}{2}\theta\epsilon_1)}{\sqrt{1+\frac{9}{4}\mathrm{Var}\left(\theta\epsilon_1\right)}}\epsilon_2$$

When $\theta = 0$ this collapses into the standard transitory process. When $\theta = 1$ it collapses to the case for permanent shocks. where e_1 and e_2 are random variables with mean zero and variances σ_1^2 and σ_2^2. For the remainder of the proof we consider variation in σ_1^2 and σ_2^2 rather than $\sigma^2_{\xi_1}$ and $\sigma^2_{\xi_2}$, and $\sigma^2_{\psi_1}$ and $\sigma^2_{\psi_2}$. The approximation can be adjusted for the prior variances by scaling up by a constant.

For the rest of this derivation I assume that the random variables e_1 and e_2 are symmetric and binary, with outcomes $\pm\sigma_1$ and $\pm\sigma_2$ each with probability $\frac{1}{2}$. This restriction is without loss of generality: for small mean-zero risks, only the second moment matters, other aspects of the distribution of risk are irrelevant.

B.2 Obtaining the Approximation

In period 0, the agent therefore faces an optimization condition (Euler equation) of the following form:

$$\begin{aligned} u'\left(W_0 - s_0\right) - \mathbb{E}_0\left(u'\left(s_0 + e_1 - s_1\left(s_0 + e_1, W_0\frac{(1+\frac{3}{2}\theta\epsilon_1)}{\sqrt{1+\frac{9}{4}\mathrm{Var}\left(\theta\epsilon_1\right)}}\sigma_2\right)\right)\right) &= f\left(s_0\left(\sigma_1,\sigma_2\right),\sigma_1,\sigma_2\right) \\ &= 0 \end{aligned}$$

where $s_0\left(\sigma_1,\sigma_2\right)$ represents optimal savings as a function of income risk standard deviations, and the savings problem at $t = 1$, $s_1\left(s_0 + e_1, \sigma_2\right)$, is solved. Using the implicit function theorem twice, we can derive $\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2}$:

$$\frac{\partial^2 s_0}{\partial\sigma_1\partial\sigma_2} = \frac{-\frac{\partial^2 f}{\partial\sigma_1\partial\sigma_2} - \frac{\partial s_0}{\partial\sigma_2}\frac{\partial^2 f}{\partial s_0\partial\sigma_1} - \frac{\partial s_0}{\partial\sigma_1}(\frac{\partial^2 f}{\partial s_1\partial\sigma_2} + \frac{\partial s_0}{\partial\sigma_2}\frac{\partial^2 f}{\partial s_0^2})}{\frac{\partial f}{\partial s_0}} \tag{17}$$

We want to derive a Taylor-series approximation of this expression for $\sigma_1 = \sigma_2 = \sigma$ around 0. To break this expression up we should bear in mind the following algebra for such expansions:

If $\mathrm{Tayl}\left(f\left(x\right), n, x_0\right)$ represents the Taylor expansion of $f(x)$ to order n at x around x_0, i.e. $\mathrm{Tayl}\left(f\left(x\right), n, x_0\right) = f(x_0) + \sum_{k=1}^{n}\frac{(x-x_0)^k}{k!}f^{(k)}(x_0) + \mathcal{O}((x-x_0)^{n+1})$, then:

$$\mathrm{Tayl}\left(f(x)+g(x), n, x_0\right) = \mathrm{Tayl}\left(f(x), n, x_0\right) + \mathrm{Tayl}\left(g(x), n, x_0\right)$$

$$\mathrm{Tayl}\left(f(x)g(x), n, x_0\right) = \mathrm{Tayl}\left(f(x), n, x_0\right)\mathrm{Tayl}\left(g(x), n, x_0\right)$$

and if $f(x_0) \neq 0$, then:

$$\mathrm{Tayl}\left(\frac{1}{f\left(x\right)}, n, x_0\right) = \mathrm{Tayl}\left(\frac{1}{\mathrm{Tayl}\left(f(x), n, x_0\right)}, n, x_0\right)$$

Therefore, we can break the expression down and perform successive approximations on constituent parts to gain an accurate approximation to the whole.

Without displaying all calculations, we give an illustration of the approximations made, picking the first term in the numerator in equation 17. We now let a_{1h}, a_{1l}, s_{1h}, s_{1l}, c_{1h}, c_{1l} represent assets at the start of period 1, saving and consumption after resolution of high/low period-1 shocks.

We can expand each term in 17 in terms of savings functions and utility. For example:

$$\frac{\partial^2 f}{\partial s_0 \partial \sigma_2} = \frac{\left(1+\frac{3\mathrm{e}1\theta}{2}\right)}{2\sqrt{\frac{9\mathrm{e}1^2\theta^2}{4}+1}} W_0 \left(u^{(3)}(c_{1h})\frac{\partial s_{1h}}{\partial \sigma_2}\left(1-\frac{\partial s_{1h}}{\partial a_{1h}}\right)+u''(c_{1h})\frac{\partial^2 s_{1h}}{\partial a_{1h}\partial \sigma_2}\right)$$
$$+\frac{\left(1-\frac{3\mathrm{e}1\theta}{2}\right)}{2\sqrt{\frac{9\mathrm{e}1^2\theta^2}{4}+1}} W_0 \left(u^{(3)}(c_{1l})\frac{\partial s_{1l}}{\partial \sigma_2}\left(1-\frac{\partial s_{1l}}{\partial a_{1l}}\right)+u''(c_{1l})\frac{\partial^2 s_{1l}}{\partial a_{1l}\partial \sigma_2}\right) \quad (18)$$

We can further derive expressions for each component part in terms of the utility function:

$$\frac{\partial s_1}{\partial \sigma_2} = \frac{\frac{1}{2}(u''(c_{2h})-u''(c_{2l}))}{-u''(c_1)-E[u''(c_2)]}$$

$$\frac{\partial s_1}{\partial a_1} = \frac{u''(c_1)}{-u''(c_1)-E[u''(c_2)]}$$

$$\frac{\partial^2 s_1}{\partial a_1 \partial \sigma_2} = \frac{u^{(3)}(c_1)(1-\frac{\partial s_1}{\partial \sigma_2})\frac{\partial s_1}{\partial a_1}+u^{(3)}(c_{2h})(1-\frac{\partial s_1}{\partial \sigma_2})\frac{\partial s_1}{\partial a_1}-u^{(3)}(c_{2l})(1-\frac{\partial s_1}{\partial \sigma_2})\frac{\partial s_1}{\partial a_1}}{-u''(c_1)-E[u''(c_2)]}$$

We now begin the approximations. These are made in several stages. First, note that $c_{2h} = s_1+\sigma$ and $c_{2l} = s_1-\sigma$, so, to second order: $f(c_{2i}) \approx f(s_1) \pm \sigma.f'(s_1)+\frac{\sigma^2}{2}.f''(s_1)$. Using these approximations we get:

$$\frac{\partial s_1}{\partial \sigma_2} = -\sigma\frac{u'''(s_1)}{u''(c_1)+u''(s_1)} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial s_1}{\partial a_1} = \frac{u''(c_1)}{u''(c_1)+u''(s_1)} - \sigma^2\frac{u''(c_1)u''''(s_1)}{2(u''(c_1)+u''(s_1))^2} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial^2 s_1}{\partial a_1\partial \sigma_2} = \sigma\frac{(u''(s_1)u'''(c_1)u'''(s_1)+u''(c_1)u'''(s_1)^2-u''(c_1)^2u''''(s_1)-u''(c_1)u''(s_1)u''''(s_1))}{(u''(c_1)+u''(s_1))^3} + \mathcal{O}(\sigma^3)$$

Second, we now relate c_1 to s_1 for small σ. As shown by Kimball (1990), the growth in consumption depends on the coefficient of absolute prudence, $s_1 \approx c_1 + \frac{\sigma^2}{2}\frac{u'''(c_1)}{u''(c_1)}$.

Inserting this expression into our three formulae gives:[23]

$$\frac{\partial s_1}{\partial \sigma_2} = -\sigma \frac{u'''(c_1)}{2u''(c_1)} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial s_1}{\partial a_1} = \frac{1}{2} - \sigma^2 \frac{u''''(c_1)}{8u''(c_1)} + \mathcal{O}(\sigma^3)$$

$$\frac{\partial^2 s_1}{\partial a_1 \partial \sigma_2} = \sigma \frac{(u'''(c_1)^2 - u''(c_1)u''''(c_1))}{4u''\left(c_1\right)^2} + \mathcal{O}(\sigma^3)$$

Finally, we relate c_{1h} to c_{1l}, and apply to equation 18. Considering first-period savings/consumption as a function of assets:

$$s_{1i} = s_1(s_0 \pm \sigma) \approx s_1(s_0) \pm \sigma.s_1'(s_0) + \frac{\sigma^2}{2}.s_1''(s_0)$$

and similarly:

$$\begin{aligned} c_{1i} = c_1(s_0 \pm \sigma) &\approx c_1(s_0) \pm \sigma.c_1'(s_0) + \frac{\sigma^2}{2}.c_1''(s_0) \\ &= c_1(s_0) \pm \sigma.(1 - s_1'(s_0)) - \frac{\sigma^2}{2}.s_1''(s_0) \end{aligned}$$

Inserting these approximations into equation 18 and simplifying leads to the final expression. To illustrate these calculations we show the result for the last expression in equation 18:

$$\begin{aligned} u''(c_{1l})\frac{\partial^2 s_{1l}}{\partial a_{1l}\partial \sigma_2} &= \\ & -\frac{W_0\sigma\left(\sigma u^{(3)}(c_1)^3 + u''(c_1)^2\left(\sigma u^{(5)}(c_1) + u^{(4)}(c_1)(3\theta\sigma+2)\right)\right)}{16u''(c_1)^2} \\ & \quad + \frac{W_0\sigma\left(u^{(3)}(c_1)u''(c_1)\left(2\sigma u^{(4)}(c_1) + u^{(3)}(c_1)(3\theta\sigma+2)\right)\right)}{16u''(c_1)^2} + \mathcal{O}(\sigma^3) \end{aligned}$$

[23] These derivatives can all be expressed in terms of risk prudence and tolerance concepts. $\frac{\partial s_1}{\partial e_2} = \frac{e}{2}.p(c) + O(e^3)$, $\frac{\partial s_1}{\partial a_1} = \frac{1}{2} - \frac{e^2}{8}p(c)t(c) + O(e^3)$, and $\frac{\partial^2 s_1}{\partial a_1 \partial e_2} = \frac{e}{4}p'(c) + O(e^3)$, where $p(c)$ is the coefficient of absolute prudence, and $t(c)$ is the coefficient of absolute tolerance, defined to be $-\frac{u''''(c)}{u'''(c)}$.

Simplifying equation 18 we get:

$$\frac{\partial^2 f}{\partial s_0 \partial \sigma_2} = -\frac{1}{4} W_0 \sigma u^{(4)}(c_1) + \mathcal{O}(\sigma^3)$$

And finally:

$$\begin{aligned}\frac{\partial^2 s_0}{\partial \sigma_1 \partial \sigma_2} &= -\sigma^2 \frac{W_0^3 \left(-W_0 u^{(3)}(c)^3 + 36\theta u^{(4)}(c) u''(c)^2 - 2W_0 u^{(4)}(c) u^{(3)}(c) u''(c) + 3W_0 u^{(5)}(c) u''(c)^2\right)}{36 u''(c)^3} \\ &\quad + \mathcal{O}(\sigma^3)\end{aligned}$$

Letting $\theta = 1$ and noting that $c_0 = \frac{1}{3} W_0 + \mathcal{O}(\sigma^2)$ gives us the result for permanent shocks given in equation 13. Letting $\theta = 0$ gives us the result for transitory shocks given in equation 14.