RESERVE BANK
OF NEW ZEALAND

DP2011/05

Time-Varying Returns, Intertemporal Substitution and Cyclical Variation in Consumption

Emmanuel De Veirman and Ashley Dunstan

August 2011

JEL classification: C22, C32, E21

www.rbnz.govt.nz/research/discusspapers/

Discussion Paper Series

ISSN 1177-7567

DP2011/05

Time-Varying Returns, Intertemporal Substitution and Cyclical Variation in Consumption*

Emmanuel De Veirman and Ashley Dunstan†

Abstract

This paper studies the importance of intertemporal substitution in consumption for the cyclical co-movement of consumption, net worth and income in New Zealand. We can largely explain the empirical hump-shaped consumption response to a transitory wealth increase by allowing for time-varying returns in an otherwise standard Permanent Income Hypothesis (PIH) model. At the net worth peak, households bring consumption forward in anticipation of low returns on saving. The PIH model fully explains the empirical response when households initially expect the net worth shock to be permanent, but gradually learn that it is in fact transitory.

* The Reserve Bank of New Zealand's discussion paper series is externally refereed. The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Reserve Bank of New Zealand. We thank Craig Burnside, Phil Briggs, Satyajit Chatterjee, Andrew Coleman, Felix Delbrueck, Karen Dynan, Oleg Korenok, Andreas Lehnert, Thomas Lubik, Pierre-Daniel Sarte, Jirka Slacalek, as well as participants at the 2011 Econometric Society Australasian Meetings, the 2011 annual meeting of the New Zealand Economic Association, and seminar participants at the Federal Reserve Bank of Richmond, the Banque de France, the College of William and Mary, Virginia Commonwealth University, the University of Alabama, the University of Amsterdam, the University of Leuven, and the Reserve Bank of New Zealand for comments and discussion.

† Corresponding author. *email address:* emmanuel.deveirman@rbnz.govt.nz, Reserve Bank of New Zealand, 2 The Terrace, Wellington 6011, New Zealand. De Veirman is a research associate of the Centre for Applied Macroeconomic Analysis (CAMA), Canberra, Australia.

ISSN 1177-7567

1 Introduction

This paper studies the role of intertemporal re-allocation in consumption for cyclical fluctuations in consumption.

Much of the literature and policy analysis on the relationship between household net worth and consumption focuses on the wealth effect of net worth on consumer spending.[1] In traditional Permanent Income Hypothesis (PIH) theory, the wealth effect implies that in response to a surprise increase in net worth, households choose to consume more at all times by a value that equals the annuity value of the wealth increase.[2]

Our paper's starting point is the intuition that exogenous net worth changes also affect households' decisions about the timing of consumption. The reason is that empirical changes in net worth often reflect wealth returns. When wealth returns are higher than usual, households have an incentive to hold more wealth by postponing consumption.

In this paper, we allow for time-varying returns in an otherwise standard PIH model. In this model, the wealth effect co-exists with the intertemporal substitution effect, where the latter is the incentive for households to re-allocate consumption across time in response to changes in returns. We define returns as all changes in net worth that are exogenous to the household's consumption decisions and do not reflect labor income. The aggregate return on net worth includes capital gains and losses on assets, rental income on real estate, dividend earnings, as well as interest earnings and payments.

Our theoretical contribution is to derive an explicit consumption function in log deviations from steady-state from the Euler equation and budget

[1] For some recent macroeconomic evidence on wealth (and collateral) effects, see Lettau and Ludvigson (2004), Ludwig and Slok (2004), Case, Quigley, and Shiller (2005), Carroll, Otsuka and Slacalek (2006), Rudd and Whelan (2006), Chen (2006), Fernandez-Corugedo, Price and Blake (2007), Slacalek (2009), and Fisher, Otto and Voss (2009). For recent microeconomic evidence, see Lehnert (2004), Campbell and Cocco (2007), Attanasio, Blow, Hamilton and Leicester (2009), and Bostic, Gabriel and Painter (2009).

[2] The permanent income hypothesis originated with Friedman (1957), and the lifecycle model with Modigliani and Brumberg (1954). For overviews of PIH/lifecycle theory, see Deaton (1992), Muellbauer (1994), Attanasio (1999) and Davis and Palumbo (2001).

constraint of a PIH model with time-varying returns. This consumption function makes explicit the intuition that unlike labor income or wealth windfalls, deviations in returns from steady-state affect consumption through two channels: a wealth effect and an intertemporal substitution effect.

By means of that consumption function, we investigate the PIH model's ability to explain the cyclical co-movement in consumption, net worth and income that we observe in New Zealand data. We compare the theoretical and empirical consumption paths implied by an empirical transitory net worth and income shock. Our paper relates to the literature that estimates the intertemporal elasticity of substitution from the Euler equation.[3] However, our study is different because we do not estimate the Euler equation itself as a relation between one period's return and the next period's consumption growth, but we explicitly track the dynamic response of consumption to a persistent shock to returns and income. By means of these impulse-responses to transitory shocks, we gain the following insights about the role of intertemporal substitution in consumption cycles.

The PIH model with intertemporal substitution is consistent with the data in that it implies a hump-shaped consumption response to a transitory wealth increase. The intuition is that when net worth reaches the peak of its cycle, households that anticipate the subsequent wealth decline bring consumption forward because they anticipate low returns on saving. This result suggests that intertemporal substitution increases the positive response of consumption to a wealth increase that households expect to be transitory.[4]

While the pure PIH model with time-varying returns mirrors the data in its hump-shaped consumption response, it differs from the data in that it implies a higher level of consumption in the first few years after a transitory wealth increase. However, the PIH model replicates the empirical consumption response when households mispredict future wealth returns,

[3] See Hall (1988), Mankiw (1981, 1985), Campbell and Mankiw (1989), Ogaki and Reinhart (1998), Basu and Kimball (2002), Vissing-Jorgensen (2002) and Guvenen (2006).

[4] This finding relates to papers such as Hall (1979), Barro (1981) and King (1989) that investigate whether temporary government purchases imply a large consumption response because of a high intertemporal elasticity of substitution.

in the sense that they initially expect the net worth increase to be permanent, but gradually learn that it is in fact transitory. We obtain this result by solving for successive paths of expected returns that rationalize the empirical consumption response within the PIH model.

In the traditional PIH model with constant returns, there is no transitory variation in consumption because consumption adjusts instantly to all changes in expectations. In this paper, we find that accounting for intertemporal substitution goes a long way in explaining empirical cyclical co-movement in consumption, net worth and income. Therefore, we can explain consumption cycles without introducing departures from the traditional PIH model that are common in the contemporary literature, such as allowing for credit market imperfections and incorporating a precautionary saving motive.[5] This does not constitute direct evidence against the importance of credit constraints or precautionary saving for explaining the response of consumption to transitory net worth shocks. Investigating whether credit market imperfections are consistent with the empirical response of consumption to a transitory shock, and if so, quantifying the relative role of intertemporal substitution and credit constraints is a topic for future research.

Finally, our paper relates to studies in financial economics on portfolio allocation and consumption with time-varying returns, and on the pricing of risk contained in time-varying asset returns.[6] Our paper is different in that we do not study the choice among different asset classes, but focus on the trade-off between consuming more today and accumulating more overall net worth.

[5] For lifecycle models with collateral constraints, see Japelli and Pagano (1994), Kiyotaki, Michaelides and Nikolov (2006), Ortalo-Magne and Rady (2006) and Coleman (2007). For general equilibrium models with household collateral constraints, see Aoki, Proudman and Vlieghe (2004), Iacoviello (2005), Davis and Heathcote (2005), Campbell and Hercowitz (2005), Calza, Monacelli and Stracca (2007), Monacelli (2009), and Iacoviello and Neri (2010). On the relation between collateral and risk sharing, see Lustig and Van Nieuwerburg (2010). For buffer stock saving models, see Zeldes (1989), Kimball (1990), Deaton (1991), Carroll (1997), Attanasio, Banks, Meghir and Weber (1999) and Gourinchas and Parker (2002). In a general equilibrium setting, see Krusell and Smith (1999). Aiyagari (1994) and Carroll and Kimball (2001) account for both liquidity constraints and precautionary saving.

[6] For instance, see Hansen and Singleton (1983), Attanasio and Weber (1989), Epstein and Zin (1991), Campbell (1996), and Campbell and Viceira (1999).

Section 2 derives the log-linearized consumption function from a Permanent Income Hypothesis model with time-varying returns. Section 3 simulates the effect of persistent income and return shocks in the PIH model. Section 4 decomposes empirical variation into a permanent and a transitory component, and produces empirical impulse-responses to a transitory shock. Section 5 compares empirical and theoretical responses to a transitory shock. Section 6 solves for a succession of return expectations that equates the theoretical and empirical consumption paths. Section 7 concludes.

2 Theory

In this section, we present the theoretical implications of a Permanent Income Hypothesis (PIH) model with time-varying returns. We refer to Appendix I for the derivations.

The representative consumer maximizes the expected utility of future consumption C_{t+i}, discounted by factor β:

$$\max \ E_t \sum_{i=0}^{\infty} \beta^i u(C_{t+i}) \tag{1}$$

subject to the flow budget constraint, according to which the household has net worth A_{t-1} at the end of period $t-1$, earns labor income Y_t and consumes C_t in period t, and transfers its net worth to the end of period t at the rate of return r_t applicable in period t:

$$A_t = (1 + r_t)\ (A_{t-1} + Y_t - C_t) \tag{2}$$

Equations (1) and (2) are standard except for the fact that the return on net worth r_t is time-varying. Assuming constant relative risk aversion (CRRA) utility such that $u(C_t) = \left[C_t^{1-\rho}/(1-\rho)\right]$, and applying the approximations $\log(1+r_t) \approx r_t$ and $\log(1+r^*) \approx r^*$ for r_t and r^* near zero, we derive the log-linearized Euler equation:

$$E_t\ \widehat{c}_{t+1} \approx \widehat{c}_t + \frac{1}{\rho}\ \overline{r}_t \tag{3}$$

where $\widehat{c}_t \equiv \log\ C_t - \ \log\ C^*$, $\overline{r}_t \equiv r_t - r^*$, and C^* and r^* are steady-state consumption and steady-state returns, respectively. This equation implies that when returns in t are above steady-state such that $\overline{r}_t > 0$, the consumer anticipates positive consumption growth from period t to $t+1$. The reason is that high returns to saving have an intertemporal substitution effect on consumption, which tends to lead the household to postpone consumption. As equation (3) reveals, the magnitude of this effect is determined by the elasticity of intertemporal substitution $1/\rho$.

We also generalize the Euler equation for expected future consumption growth:

$$E_t\ \widehat{c}_{t+i} \approx \widehat{c}_t + \frac{1}{\rho}\sum_{j=1}^{i} E_t\ \overline{r}_{t+j-1} \quad \forall i \geq 1 \tag{4}$$

which implies that expected consumption growth between period t and any future period $t+i$ depends on the unweighted sum of returns from period t to $t+i-1$, with the extent of re-allocation of consumption across periods determined by the intertemporal elasticity of substitution.

We log-linearize the flow budget constraint, equation (2), around the steady-state and obtain:

$$\widehat{a}_t \approx \overline{r}_t + (1+r^*)\ (\widehat{a}_{t-1} + \frac{Y^*}{A^*}\widehat{y}_t - \frac{C^*}{A^*}\widehat{c}_t) \tag{5}$$

where $\widehat{a}_t \equiv \log\ A_t - \log\ A^*$, $\widehat{y}_t \equiv \log\ Y_t - \log\ Y^*$, and A^* and Y^* are steady-state net worth and income, respectively. The term $\overline{r}_t$ in equation (5) captures the income effect of returns. By virtue of the income effect, a time-t increase in returns above steady-state such that $\overline{r}_t > 0$ relaxes the budget constraint, which tends to allow for more consumption in every period from t onwards.

Solving this equation forward, and imposing the transversality condition

$\lim_{n\to\infty}\left[1/\left(1+r^*\right)\right]^{n+1}\widehat{a}_{t+n}=0$, we obtain:

$$\widehat{a}_{t-1}\approx\sum_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\left(\frac{C^*}{A^*}E_t\;\widehat{c}_{t+i}\;-\frac{Y^*}{A^*}E_t\;\widehat{y}_{t+i}-\frac{1}{1+r^*}\;E_t\;\overline{r}_{t+i}\right) \quad (6)$$

Substituting equation (4) into equation (6), re-arranging, and assuming that $r^*>0$ such that $\sum\limits_{i=0}^{\infty}\left[1/(1+r^*)\right]^i=\left[(1+r^*)/r^*\right]$, we derive the log-linearized consumption function:

$$\widehat{c}_t\approx\frac{r^*}{1+r^*}\left[\begin{array}{c}\frac{A^*}{C^*}\widehat{a}_{t-1}+\frac{Y^*}{C^*}\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i E_t\;\widehat{y}_{t+i}\\ +\left(\frac{A^*}{C^*}\frac{1}{1+r^*}-\frac{1}{\rho}\;\frac{1}{r^*}\right)\;\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i E_t\;\overline{r}_{t+i}\end{array}\right] \quad (7)$$

Because of the factor $r^*/(1+r^*)$, this equation reflects the traditional implication of PIH theory that consumers spend the annuity value of their lifetime resources. In terms of deviations from steady-state, consumption depends positively on end-of-last period net worth, as well as on human wealth, which is the expected present discounted value of current and future labor income. The net response of consumption to changes in the present discounted value of current and expected future returns depends on the parameter values. The income effect of returns, captured by the term $(A^*/C^*)[1/(1+r^*)]$, tends to imply that consumption depends positively on the present discounted value of returns dated t and later. The intertemporal substitution effect, captured by $-(1/\rho)\;(1/r^*)$, tends to imply a negative response of current consumption to contemporaneous and expected future returns.

3 Simulation

To gain intuition for the effects of changes in income and returns in the PIH model with time-varying returns, we use equations (5) and (7) to simulate impulse-responses to exogenous changes in income and returns.

Periods are quarters. We assume a time discount factor $\beta=0.99$. The steady-state Euler equation corresponding to equation (A2) in the appendix is $C^{*-\rho} = \beta(1 + r^*)\ C^{*-\rho}$, which implies that $\beta = 1/(1 + r^*)$. Correspondingly, we assume that in steady-state, net worth yields a quarterly real return close to 1 percent, i.e. $r^* = 0.01$, and an annual return close to 4 percent.[7] The return reflects changes in net worth that are exogenous to the household's optimization problem and are not related to labor income. Empirically, the return reflects interest earned and paid, as well as dividends, rental income, and capital gains and losses on assets. The theoretical return captures the aggregate return on net worth, and abstracts from differences in returns across asset classes.

We calibrate steady-state ratios by setting them equal to the ratios of sample averages. Therefore, we do not explicitly account for steady-state growth in the PIH model. We discuss data in the next section. In our sample, $A^*/C^* = 23.81$, implying that net worth is about twenty-four times a quarter's consumption, and $Y^*/C^* = 1.08$. Those values are consistent with $Y^*/A^* = 0.05$ and $C^*/A^* = 0.04$.

We use three values for the intertemporal elasticity of substitution $1/\rho$. In the baseline, we set the coefficient of relative risk aversion $\rho = 2$. In this case, $(A^*/C^*)[1/(1+r^*)] - (1/\rho)\ (1/r^*)$ is negative, implying that according to equation (7), current consumption reacts negatively to expected wealth returns. The reason is that the intertemporal substitution effect dominates the income effect of returns on current consumption. Secondly, we consider the case where substitution and income effects of returns on current consumption exactly offset each other such that $(A^*/C^*)[1/(1+r^*)] - (1/\rho)\ (1/r^*) = 0$. In this case, $\rho = 4.20$. Finally, we consider the case without any intertemporal re-allocation, $\rho = \infty$.

Both $\rho = 2$ and $\rho = 4.20$ are within the range of plausible values. $\rho = \infty$ is implausible, but by comparing its implications with the two other cases we can show the effects of intertemporal substitution.

Figure 1 shows the impulse-responses to a transitory but moderately persistent shock in labor income that occurs in time t. The economy is in steady-state before the shock. The shock implies that income is 5 percent above steady-state at time t. After that, the log deviation from steady-state decays gradually according to $E_t\ \widehat{y}_{t+i+1} = 0.5\ \widehat{y}_{t+i}$ for $i \geq 0$. This

[7] Since the quarterly return r^* and annual return r_a^* are related as $(1+r^*)^4 = 1+r_a^*$, our assumption that $r^* = 0.01$ implies that $r_a^* = 0.04$.

shock implies that $\sum_{i=0}^{\infty}[1/(1+r^*)]^i E_t\ \widehat{y}_{t+i} = 9.90$. That is, human wealth is 9.90 percent above steady-state on impact.[8] The PIH model implies that in response to the income surprise, the household increases consumption instantly by a small amount that leaves total expected lifetime resources intact, as would be the case if the household only consumed the interest revenue earned from depositing the surprise income. To keep total lifetime resources intact, the household saves a large enough fraction of the income increases (or pays off enough debt) so that the eventual dollar increase in asset net worth exactly equals the initial dollar increase in human wealth. Figure 1 reflects the common intuition that according to the Permanent Income Hypothesis, consumption only responds to transitory income surprises to the extent that the shocks affect permanent income. The impulse-response is the same for any value of ρ.

Figure 2 displays the impulse-responses to a surprise increase in returns of 5 percentage points above steady-state on impact, which subsequently decays back to steady-state according to $E_t\ \overline{r}_{t+i+1} = 0.5\ \overline{r}_{t+i}$. In the consumption and net worth panels, the solid blue line graphs the response when $\rho = 2$, while the dotted green line plots the response when $\rho = 4.20$ and the dashed red line when $\rho = \infty$. In each case, the period of above-equilibrium returns expands the present discounted value of lifetime resources, and therefore allows for a permanent increase in consumption by virtue of a wealth effect. However, any non-zero degree of intertemporal substitution implies that consumers also have a desire to reduce consumption as long as the returns to saving are above steady-state.

When we set $\rho = \infty$, consumption instantly jumps to its new value, in analogy with the response to the income shock.[9] On the other hand, when $\rho = 2$, consumption declines on impact because the intertemporal substitution effect on current consumption dominates the income effect. As time progresses, consumption increases. The reason is that the intertemporal

[8] The simulated income impulse implies that for any future date $t+i$, $E_t\ \widehat{y}_{t+i} = 0.5^i\ \widehat{y}_t$. Therefore, it implies that $\sum_{i=0}^{\infty}[1/(1+r^*)]^i E_t\ \widehat{y}_{t+i} = \sum_{i=0}^{\infty}[1/(1+r^*)]^i\ 0.5^i\ \widehat{y}_t = 1.98\ \widehat{y}_t$, which equals 9.90 since $\widehat{y}_t = 5$.

[9] The difference is that the consumption response to the graphed increase in returns is on the order of 20 times larger than the consumption response to the income shock in Figure 1. This is because $Y^*/A^* = 0.05$, implying that the dollar value of a 5 percent increase in income above steady-state is much smaller than the dollar value of a 5 percentage point increase in returns.

substitution effect gradually weakens as expected return deviations, and the present discounted value of return deviations $\sum_{i=0}^{\infty}[1/(1+r^*)]^i E_t\, \overline{r}_{t+i}$ in equation (7), decay to zero. Meanwhile, the saved fraction of realized returns accumulates in the net worth term $\widehat{a}_{t-1}$ in equation (7), and therefore continues to affect consumption through a wealth effect. With $\rho = 4.20$, consumption stays in steady-state on impact because in this case, the income effect and intertemporal substitution effect on current consumption exactly offset each other. Afterwards, consumption increases as the intertemporal substitution effect weakens over time.

4 Permanent-Transitory Decomposition

This section's aim is to decompose empirical variation in consumption, labor income, and net worth into a trend and cyclical component. We first discuss the data, before estimating empirical long-run and short-run relations. Finally, we use those estimates to disentangle permanent and transitory variation.

We use aggregate New Zealand data on real per-capita household consumption, household wealth, and after-tax labor income.[10] All data are for the period 1990Q1-2009Q1.

Household consumption C_t includes spending on durable goods, non-durable goods and services, but excludes housing consumption, which is the imputed gross rental value of owning a home. Our income measure Y_t is after-tax labor income, which differs from disposable income in that it excludes property income. This is consistent with the theoretical model, where property income enters the return r_t. Our measure of household net worth A_t combines the aggregate value of the housing stock with the total value of financial assets, and subtracts the value of debt outstanding. Net worth excludes durable goods. As in our theoretical model, A_t measures net worth at the end of period t. In Section 5, we will construct a measure for the return r_t implied by the data.[11]

[10] We deflate all nominal variables by the Consumer Price Index (CPI) with base year 2006.

[11] Our data definitions avoid double counting on various levels. First, durable goods

We refer to Appendix II for detailed information on data sources and construction.

Figure 3 plots our series for consumption, income and net worth in real per capita terms. The recession of the early 1990s reflects a disinflation associated with the first years of inflation targeting by the Reserve Bank of New Zealand. Household net worth increases dramatically from the early 2000s through 2007Q3, and declines after that. This reflects a housing boom and subsequent downturn. As De Veirman and Dunstan (2010) discuss, changes in the value of the housing stock account for most changes in New Zealand households' net worth. This is in line with the fact that on average over our sample period, housing accounts for 68% of total household assets, which is more than double the share in the United States (31% over our sample period).[12]

We now proceed to decompose empirical variation into changes that reflect trends and changes that reflect transitory deviations from trend. As Table 1 shows, Augmented Dickey-Fuller tests suggest that all log series are non-stationary. From Table 2, we conclude that Johansen L-max and Trace tests indicate the existence of a single cointegration relationship between consumption, income and net worth. The co-trending relationship is:

$$c_t - \kappa - \gamma_a \, a_t - \gamma_y \, y_t = \eta_t \tag{8}$$

where c_t, a_t and y_t are log real per-capita consumption, net worth and income, κ is a constant, and η_t is the cointegrating residual. As usual, we normalize the coefficient on log consumption to one. We estimate the cointegrating relation because it is a prerequisite for performing permanent-transitory decomposition. For this purpose, we do not need to interpret the coefficients γ_a and γ_y as elasticities that capture the long-run response of consumption to exogenous wealth and income changes. Instead, we interpret equation (8) as reflecting the empirical long-run correlation between the three variables. This statistical long-run relation likely reflects

enter consumption but do not enter net worth. Therefore, we treat the purchase of durable goods as consumption rather than as wealth accumulation. Second, owning a house enters net worth but not consumption. Third, property income does not enter our labor income measure, but enters the implicit return.

[12] US data are from Haver. For comparability, we express both the New Zealand and US ratios as a share of household assets excluding durable goods. Both figures are averages for 1990Q1-2009Q1.

many other factors, including the effects of productivity growth on consumption, income and wealth.[13]

We also estimate the following Vector Error Correction Model (VECM) for stationary log-differenced data:

$$\Delta X_t = A_0 + A_1 \, \Delta X_{t-1} + \alpha\gamma' X_{t-1} + \varepsilon_t \tag{9}$$

where X_t is a three-by-one vector stacking c_t, y_t and a_t. A_0 is a vector of constants, and A_1 is a matrix of coefficients on lagged growth rates.[14] The vector γ' contains the estimated cointegrating coefficients such that equation (8) implies $\gamma' X_{t-1} = \eta_{t-1}$. The vector α captures every variable's response to the cointegration residual. When cointegration exists, η_{t-1} is stationary. Stationarity in η_{t-1} can only be achieved if at least one of the variables error-corrects, i.e. it tends to undo deviations from the long-run relationship as captured by η_{t-1}.

We find that $\gamma_a = 0.12$ and $\gamma_y = 0.86$, and both are statistically significant at the 1 percent level.[15] Table 3 reveals that net worth and income error-correct at the 5% significance level or better. For instance, when consumption is above its long-run relation with income and net worth such that η_{t-1} is positive, this tends to anticipate an increase in income and net worth that restores the long-run relationship. Consumption error-corrects marginally, at the 10% significance level.

We now perform permanent-transitory decomposition, using the technique developed by Gonzalo and Ng (2001) and applied by Lettau and Ludvigson (2004). Gonzalo and Ng (2001) show that, by multiplying the vector of reduced-form VECM residuals ε_t by a matrix that depends on the estimated short-run adjustment parameters α as well as on the estimated

[13] Lantz and Sarte (2001) show that in general equilibrium, technological improvements imply an increase in consumption as well as an increase in the market value of the firm. This matters in our setting because the latter implies an increase in households' stock market wealth.

[14] The Bayesian Information Criterion selects a lag length of order two in a VAR specified in levels of X_t, suggesting that a lag length of order one is appropriate for the VECM (9). (The Akaike Information Criterion suggests lag length three for the VAR in levels, and therefore two for the VECM.)

[15] The t-statistic for γ_a is 3.17, while the t-statistic for γ_y is 9.64. These t-statistics are adjusted for serial correlation using Hayashi (2000, p. 654 ff.). We estimate equation (8) using Dynamic Ordinary Least Squares (DOLS) with three leads and lags of the differenced log variables.

cointegrating vector γ, one obtains a set of permanent and transitory shocks u_t. Explicitly writing ε_t and u_t in terms of their entries, this means in our case:

$$\begin{bmatrix} u_{P1,t} \\ u_{P2,t} \\ u_{T,t} \end{bmatrix} = \begin{bmatrix} \alpha_{\perp}^{'} \\ \gamma \end{bmatrix} \begin{bmatrix} \varepsilon_{c,t} \\ \varepsilon_{a,t} \\ \varepsilon_{y,t} \end{bmatrix} \tag{10}$$

where $\varepsilon_{c,t}$, $\varepsilon_{a,t}$ and $\varepsilon_{y,t}$ are the reduced-form residuals corresponding to the VECM equation for consumption, net worth and income, respectively. The matrix $\alpha_{\perp}$ is orthogonal to α such that $\alpha_{\perp}^{'}\alpha = 0$. In our application, permanent-transitory decomposition yields two permanent shocks $u_{P1,t}$ and $u_{P2,t}$ and one transitory shock $u_{T,t}$.

The permanent shocks encompass any shocks that induce non-stationary variation in the variables, such as productivity shocks or permanent changes in consumer preferences. In a system with three non-stationary variables and one common trend, there are two separate permanent shocks. The single transitory shock reflects mean-reverting fluctuations around the single long-run relationship.

The shocks in u_t are in principle correlated. To obtain uncorrelated shocks, we orthogonalize the permanent and transitory shocks by applying a Cholesky decomposition to the vector u_t. We order the transitory shock last. Thus, we assume that transitory shocks do not affect permanent shocks contemporaneously.

Table 4 decomposes variation in consumption, net worth and income growth into variation caused by transitory and permanent shocks. The upper part of the table shows the fraction of the forecast error variance explained by transitory shocks, while the lower part displays the share of permanent variation. Over virtually all forecast horizons, about one quarter of variation in consumption, wealth and income is stationary, while the non-stationary component accounts for the remaining three quarters. For our purposes, it is important to note that although most variation in consumption, wealth and income is permanent, there is a substantial cyclical component to each of the variables.

We now document empirical co-movement in consumption, income and wealth conditional on a transitory shock. In Figure 4, the bold black lines in the panels ‘consumption’, ‘income’, and ‘net worth’ plot the impulse-responses to a one-standard deviation orthogonalized transitory shock implied by equations (8) through (10).

Income is 5.81 percent higher on impact than it would have been without the shock. Net worth is 7 percent higher on impact, and reaches 11.96 percent above steady-state two quarters after the shock. Consumption declines by 2.51 percent on impact, but then starts rising, reaching 2.33 percent above steady-state three quarters after the shock. Since the empirical shock is transitory by construction, it implies that eventually, consumption, income and wealth all return to their steady-state growth path. These impulse-responses form the basis for our analysis in the next section.

5 Consumption Response to Empirical Transitory Shocks

We now use Section 2's Permanent Income Hypothesis (PIH) model with time-varying returns to analyze the theoretical consumption response to empirical transitory shocks. We compare the theoretical and empirical consumption responses before analyzing the theoretical response in more detail.

5.1 Theoretical vs. Empirical Responses

By means of equations (5) and (7), we analyze the theoretical consumption response to empirical transitory shocks. As in Section 3, we calibrate the steady-state value for r^* to a value close to 1 percent[16] and set steady-state ratios involving C^*, A^* and Y^* equal to their sample averages. We consider the same three cases for the coefficient of relative risk aversion, i.e. $\rho = 2$, 4.20, and ∞.

We consider a transitory shock that occurs at time t, and assume that all variables are in steady-state before the shock. In this section, we assume that households form their income and return expectations as of time t based on the empirical impulse-responses to a transitory shock.

[16] $\beta = 0.99$ implies $r^* = 0.0101$, or 1.01 percent.

The path for households' expectations of income deviations $E_t\ \widehat{y}_{t+i}$, for $i \geq 0$, directly corresponds to the empirical impulse-response for income as plotted in the upper right panel of Figure 4. The lower right panel graphs the corresponding human wealth deviation $\sum_{i=0}^{\infty}[1/(1+r^*)]^i E_t\ \widehat{y}_{t+i}$. Human wealth increases 27.95% on impact, but gradually declines over time as an ever smaller part of income surprises lies in the future.

Implicitly, changes in net worth reflect the aggregate return r_t to the extent that these changes are not due to changes in income and consumption. To compute a path for return deviations $E_t\ \overline{r}_{t+i}$ implied by the empirical impulse-responses, we re-arrange the flow budget constraint (5), write it explicitly in terms of times $t+i$ where i indicates the number of quarters since the shock, and take expectations as of time t:

$$E_t\ \overline{r}_{t+i}\ \approx E_t\ \widehat{a}_{t+i} - (1+r^*)\ (E_t\ \widehat{a}_{t+i-1} + \frac{Y^*}{A^*} E_t\ \widehat{y}_{t+i} - \frac{C^*}{A^*} E_t\ \widehat{c}_{t+i}) \quad (11)$$

For this purpose, we fill the right-hand side of equation (11) with the empirical impulse-response paths for $E_t\ \widehat{c}_{t+i}$, $E_t\ \widehat{a}_{t+i}$ and $E_t\ \widehat{y}_{t+i}$. The upper middle panel of Figure 4 graphs the implied empirical return deviations.

The return is 6.62 percentage points above equilibrium on impact and stays above equilibrium in $t+1$. The return is a mere 0.02 percent below steady-state in $t+2$. After that, return deviations turn substantially negative, with returns reaching their trough in $t+4$ at 2.27 percentage points below steady-state.[17] The return profile largely reflects the empirical net worth path. High returns in t and $t+1$ correspond to sharp net worth increases in those quarters. The negative return deviations after $t+2$ correspond to net worth declining from that point on. This indicates that a transitory net worth increase largely reflects a period of above-equilibrium returns followed by a period of low returns.

The lower middle panel of Figure 4 graphs the present discounted value of expected return deviations $\sum_{i=0}^{\infty}[1/(1+r^*)]^i E_t\ \overline{r}_{t+i}$. The present discounted

[17] Since $\overline{r}_t \equiv r_t - r^*$, return deviations are in percentage point differences from steady-state. Since we assume a steady-state quarterly real return of 1.01 percent, the trough for the real return is -1.25 percent.

value is near-zero (albeit slightly negative) on impact, reflecting the fact that the period of positive returns roughly offsets the period with negative returns. After that, return deviations turn markedly negative in present discounted value terms, with a trough of -12.23% in $t+3$. The present discounted value of expected return deviations is negative after the shock because the contribution of the positive returns episode reduces over the course of the response horizon, and quickly goes to zero.

Given the empirical expected income and return deviations, we compute the paths for theoretical intended consumption deviations $E_t\ \widehat{c}_{t+i}$ and implied net worth deviations $E_t\ \widehat{a}_{t+i}$. We do so by iterating on equations (7) and (5), rolling these equations 1 period forward at every step, and taking expectations as of time t.

The two left panels of Figure 4 graph the theoretical consumption and net worth paths, along with the empirical consumption and net worth impulse-responses that we discussed in the previous section. In the remainder of the paper, bold black lines indicate empirical impulse-responses while colored lines indicate implications from PIH theory.

The solid blue line graphs the consumption response for the case where $\rho = 2$. Consumption increases marginally on impact, but increases substantially in the next few quarters, attaining 6.06 percent above steady-state in $t+2$. Consumption flattens out in $t+3$, and then starts declining. Consumption declines to a level slightly below the original steady-state. As will become clear later in this section, this is because households have brought consumption forward to the immediate aftermath of the shock and therefore have fewer resources available later on. In this respect, also note that theoretical net worth converges to a level 1.10 percent below its pre-shock value.

The overall finding at this stage is that for a plausible degree of intertemporal re-allocation, the PIH model implies a hump-shaped consumption response, consistent with the data. However, the theoretical response is about three times larger at its peak. Also, at this stage the theory does not explain the empirical drop in consumption on impact.

The green dotted line in the consumption panel of Figure 4 reveals that, when households are less prone to intertemporal re-distribution as captured by $\rho = 4.20$, consumption increases less markedly after the shock.

The consumption and net worth responses are both closer to the data than when $\rho = 2$. The response continues to be hump-shaped, but consumption reaches its peak at 2.94 percent in $t+2$. Consumption now converges to a value near steady-state. With a smaller temporary consumption increase in the aftermath of the shock than when $\rho = 2$, households accumulate somewhat more wealth, which means that they have more resources left for consuming in the long run.

The red dashed line in the same panel reveals that, if we shut down intertemporal re-allocation altogether by setting $\rho = \infty$, consumption increases slightly on impact, to 0.10 percent above steady-state, and subsequently stays constant at that new level. When $\rho = \infty$, return shocks only affect consumption through a wealth effect. By virtue of the wealth effect in the PIH model, households adjust consumption by an amount equal to the annuity value of the wealth change. Since Figure 4 considers a transitory shock, the annuity value is small, implying that consumption only adjusts by a small value.

Comparing the consumption responses across values for ρ, we find that the elasticity of intertemporal substitution matters greatly for whether the PIH model is able to account for the empirical co-movement in consumption, wealth and income induced by a transitory shock. The more prone households are to re-distribute consumption across time, the larger the temporary increase in consumption after the shock.

5.2 Dissecting the Theoretical Response

So far in this section, we have discussed the combined effect of return and income surprises on consumption and wealth accumulation. To understand why these theoretical responses occur, we now analyze the response to three distinct components of the shocks: the transitory income shock; the episode with positive return deviations; and the episode with negative return deviations.

Figure 5 shows the theoretical impulse-response to the empirical transitory income shock, holding returns constant at steady-state. In this case, the consumption response does not depend on the value for ρ. In other words, income affects consumption purely through a wealth effect. The result is reminiscent of the simulated response to a transitory income shock in Section 3. Consumption instantly increases slightly to 0.30 percent above steady-state, and remains constant thereafter. Households consume their permanent income, as they would if they were to consume only the interest revenue on the surprise income. As income and human wealth decline back to steady-state, the household keeps life-time resources intact by accumulating net worth. Net worth converges to a value 1.27 percent above steady-state.

As we are about to see, the overall pattern of the theoretical consumption and net worth response is dominated by return surprises. To better understand the effect of return surprises, we now break Figure 4's transitory net worth increase down into a permanent net worth increase (Figure 6) and a subsequent permanent net worth decrease (Figure 7).

Figure 6 plots the consumption response to the two quarters, t and $t+1$, where empirical returns are above steady-state. Counterfactually, we hold subsequent returns at steady-state and abstract from the income shock. This return path implies a permanent increase in net worth.

When $\rho = \infty$, households instantly adjust consumption to 2.60 percent above steady-state and hold consumption unchanged from that point on. Since households do not intertemporally substitute consumption at all in this case, this consumption response captures the pure wealth effect of the positive return deviations on consumption. Since the wealth increase is permanent in this scenario, the wealth effect implies a relatively strong consumption response.

When $\rho = 4.20$, households also wish to save more in order to benefit from the higher returns on saving in t and $t+1$, and therefore consume less in t and $t+1$ than the pure wealth effect would imply. When $\rho = 2$, the intertemporal substitution effect more than offsets the wealth effect on current consumption, implying that consumption falls to 2.86 percent below steady state on impact. Consumption is moderately positive at $t+1$, and from $t+2$ stays constant at 2.68 percent above steady-state.

By comparing impulse-responses across different values for ρ, we find that households that are more prone to re-allocate consumption across periods tend to consume less during a permanent net worth increase. This documents that the intertemporal substitution effect tends to work in the opposite direction of the wealth effect in the case of a permanent wealth shock.

Figure 7 plots the response to empirical negative return deviations only. In this scenario, we fix returns at steady-state in t and $t+1$, but account for the empirical returns below steady-state from $t+2$ onwards. When we shut off intertemporal substitution, the expectation of these negative return deviations implies that consumption drops to 2.80 percent below steady-state at time t due to a wealth effect.

When we allow for intertemporal substitution, this tends to increase consumption in the aftermath of the shock, reflecting households' incentive to consume more today because of expected low future returns on saving. With $\rho = 4.20$, households hold consumption at steady-state until $t+2$, and very close to steady-state at $t+3$, as substitution and wealth effects offset each other. With $\rho = 2$, consumption is 3.08 percent above steady-state in period t through $t+2$, and very close to that level in $t+3$. In this case, intertemporal substitution more than offsets the wealth effect in the aftermath of the shock. As time progresses, consumption does decline, because the intertemporal substitution effect weakens as the proportion of wealth returns that is still in the future declines. Consumers who respond to the intertemporal incentive to dissave ($\rho = 2$, 4.20) accumulate less net worth, implying that they consume less from about three years after the shock than hypothetical consumers who never shift consumption across periods.

Adding up the impulse-responses in Figures 5 through 7 yields the theoretical responses in Figure 4. This yields the following insights. Focusing on the case where $\rho = 2$ in Figure 4, we distinguish three phases. On impact, consumption increases only slightly, as the wealth and substitution effects from episodes with positive and negative return deviations roughly cancel each other out. In a second phase, $t+1$ and $t+2$, consumption increases since the incentive to postpone consumption associated with the episode of returns above steady-state weakens and vanishes. At $t+2$, the incentive to bring consumption forward associated with the episode of low expected returns is still at full force. In a third phase, consumption decreases because the incentive to consume more today weakens as an ever larger proportion of negative return deviations materializes.

In this setting, transitory fluctuations in consumption are dominated by changes in the strength of the intertemporal substitution effect. On the other hand, the strength of the wealth effect stays constant throughout the impulse-response horizon. Therefore, the pure wealth effect cannot explain any transitory variation in consumption. With $\rho = \infty$, consumption adjusts instantly, such that all theoretical changes in consumption are permanent.

Finally, note that the overall wealth effect from a transitory shock is small since the positive wealth effect on consumption from the period with returns above steady-state roughly offsets the negative wealth effect on consumption from the period with returns below steady-state. Therefore, the positive response of consumption to a transitory wealth increase is almost uniquely due to intertemporal substitution in consumption.

6 Allowing for Mispredictions

The previous section assumed that PIH households set their expectations about future income and returns equal to the expectations implied by the empirical model of Section 4. The current section drops that assumption, and allows for the possibility that households mispredict returns. In particular, we solve for return expectations that imply that the PIH consumption response to a transitory shock is equal to the empirical response.

Consider the scenario of a time t transitory shock without any further surprises, such that the empirical impulse-responses turn out to be true. The intuition of the present section's exercise is that households can make inaccurate predictions about future returns, but learn history as it materializes. Households still use the empirical model to forecast income.

We first solve for a path of expected returns that exactly replicates the empirical impulse-response at time t. In particular, we set $\widehat{c}_t$, $\widehat{a}_t$ and $\overline{r}_t$ equal to their values for the empirical responses, and solve for a path of return expectations $E_t\ \overline{r}_{t+i}$ that warrants the empirical value for $\widehat{c}_t$ in the PIH model. Given these return expectations and given income expectations, we then compute the households' intentions as of time t about how much to spend in any future period $t+i$, i.e. $E_t\ \widehat{c}_{t+i}$ for $i \geq 0$, and associated wealth accumulation $E_t\ \widehat{a}_{t+i}$.

Conditional on the households' solution in t, we next investigate which return expectations as of time $t+1$ rationalize the empirical impulse-responses up to $t+1$. We solve for a path $E_{t+1}\ \overline{r}_{t+i+1}$ for return expectations as of time $t+1$ and the corresponding time $t+1$-intended consumption and savings profile $E_{t+1}\ \widehat{c}_{t+i+1}$ and $E_{t+1}\ \widehat{a}_{t+i+1}$.

Repeating this procedure, we eventually replicate the entire empirical consumption and net worth profile. This exercise produces successive expectations profiles $E_{t+j}\ \overline{r}_{t+i+j}$, $E_{t+j}\ \widehat{c}_{t+i+j}$ and $E_{t+j}\ \widehat{a}_{t+i+j}$, with i and $j \geq 0$, that at every step rationalize the empirical consumption and net worth outcomes up to $t+j$.

Before reporting quantitative results, we discuss the qualitative implications of the PIH model. In this section, we consider the case $\rho = 2$. The upper left panel of Figure 4 shows that empirical consumption drops on impact, while the PIH model calls for a slight increase in consumption. The subsequent hump-shaped response in consumption is substantially smaller in the data than in the PIH model. Therefore, theoretical consumption is 'too high' relative to its empirical counterpart when PIH households form return expectations in line with empirical model forecasts. When $\rho = 2$, the intertemporal substitution effect dominates, such that future expected returns have a negative effect on current consumption. Therefore, we need an increase in future expected returns above empirical expectations to bring down theoretical consumption so that it equals its empirical counterpart.

We now quantify the required increase in future expected returns. We solve for a path of expected returns that exactly produces the present discounted value of return deviations $\sum_{i=0}^{\infty} \left[1/\left(1+r^{*}\right)\right]^{i} E_t\ \overline{r}_{t+i}$ necessary to explain current empirical consumption $\widehat{c}_t$ according to the consumption function (7), given empirical net worth $\widehat{a}_{t-1}$ and given the empirical impulse-response for income $E_t\ \widehat{y}_{t+i}$.[18] Figure 8 plots the result. The bold black line graphs the empirical response while the blue line is the theoretical response. The PIH model can explain the drop in consumption on impact to 2.51 percent below steady-state if households expect positive

[18] An infinite number of paths yields the required present discounted value of returns. We solve for a particular path by assuming that expected returns are at steady-state after $t+25$, and by favoring a path where earlier expected returns that differ from the empirical return take on a common value.

return deviations up to $t+1$ equal to empirical return deviations, but from $t+2$ on expect returns near steady-state. This suggests that households expect that the return surprise implies a permanent increase in net worth. At time t, households intend to increase consumption after the shock, to 3.03 percent above steady-state in $t+2$. The peak increase in consumption is less strong than with empirical expectations and it is not hump-shaped. The reason is that there are no substantial negative expected return deviations that would cause households to re-allocate consumption from the future to the immediate aftermath of the shock.

Given that outcome for time t, we now ask how PIH households would have to update their expectation profiles over time in order to optimally decide to carry out the empirical consumption path. Figure 9 shows the results. The solid green line shows the results with expectations as of $t+1$. We find that a similar profile for expected returns can explain the empirical time-t drop as well as the empirical time-$t+1$ increase in consumption of 0.73 percent above steady-state.

The solid red line reveals that given empirical returns up to $t+2$, expectations as of $t+2$ that the return is 0.18 percent above steady-state from $t+3$ through $t+25$ are sufficient to explain empirical consumption of 2.13 percent above steady-state in $t+2$. This expectation involves net worth increasing to 16.65 percent above steady-state in the long run, which somewhat exceeds the 12.63 percent above steady-state expected as of $t+1$.

Dashed lines show expectation profiles as of $t+3$ through $t+5$, and dotted lines are for expectations as of $t+6$, $t+8$, and $t+12$. In the successive consumption paths, the short-lived decreases in intended consumption reflect intertemporal substitution.

We find that we can exactly match the empirical consumption and net worth profiles by assuming that households recursively update expectations about future returns and implied future values for intended consumption and net worth. On impact and one quarter after the shock, households mistakenly believe that the wealth increase will be permanent. When the wealth increase lasts for two quarters, households shift to thinking that wealth may increase somewhat more in the future. As wealth decreases, households gradually learn that net worth will return to steady-state.

7 Conclusion

We document that in New Zealand data, a substantial transitory increase in consumption tends to happen at almost the same time as a transitory wealth increase. The wealth effect cannot explain this fact. We document that in a PIH model, the wealth effect implies a very small and permanent consumption response to a transitory wealth shock. Consumption does not respond much because a transitory wealth increase only has a small effect on lifetime resources and permanent income.

However, we find that intertemporal substitution in consumption implies a positive hump-shaped response of consumption to a transitory wealth shock, which is much closer to the empirical consumption response. Within a PIH model with time-varying returns, we define a transitory wealth shock as an episode of returns above steady-state followed by a period with low returns. Consumption reaches its peak a few quarters after the transitory wealth shock, as households expect the initial wealth increase to be followed by low returns on saving and therefore choose to bring consumption forward.

Within the PIH model with time-varying returns, we interpret a permanent wealth increase as a single episode with returns above steady-state. We document that the wealth effect implies a substantial consumption increase in response to a permanent wealth increase. In this case, intertemporal substitution works in the opposite direction as the wealth effect. Households that are more prone to intertemporal substitution tend to consume less during a permanent wealth increase.

Overall, our paper finds that intertemporal substitution in response to changes in returns is an important determinant of cyclical fluctuations in consumption. Therefore, it may be instructive for policymakers to consider intertemporal substitution in addition to wealth effects in order to assess the likely effect of any present or future net worth change on consumption.

Figures and Tables

Figure 1: Response to Simulated Transitory Income Shock

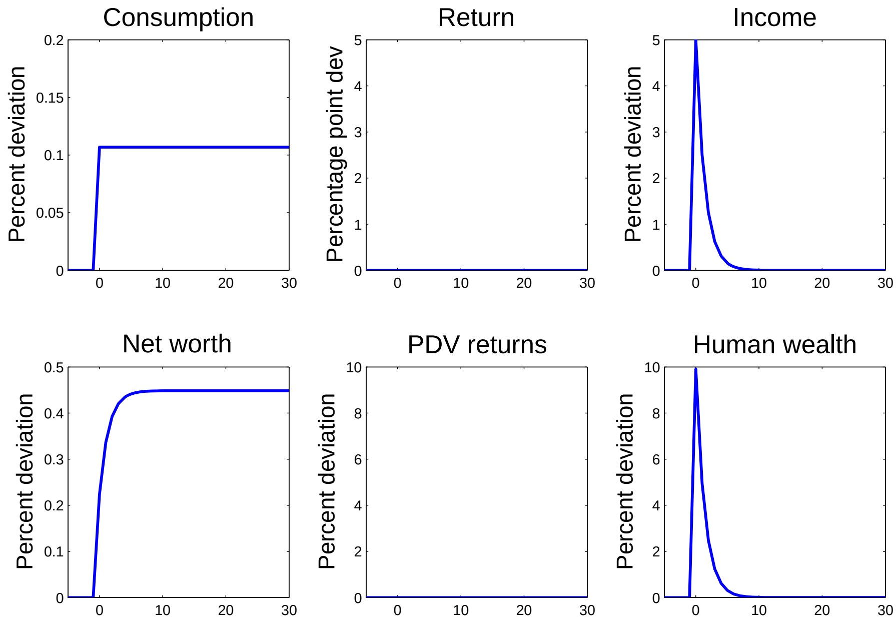

Note: this figure plots impulse-responses to an income shock implied by the Permanent Income Hypothesis (PIH) model of Section 2. Before the time-t shock, the economy is in steady-state. We simulate a persistent income shock such that on impact, the percentage deviation of income from steady-state $\widehat{y}_t = 5$ and afterwards, income deviations decay according to $E_t\ \widehat{y}_{t+i+1} = 0.5\ \widehat{y}_{t+i}$ for $i \geq 0$. In response, consumption instantly increases by a small amount that leaves total expected lifetime resources intact. All variables are graphed in percent deviations from steady-state, except the return which is in percentage point deviations from steady-state. The horizontal axis plots quarters since the shock.

Figure 2: Response to Simulated Return Shock

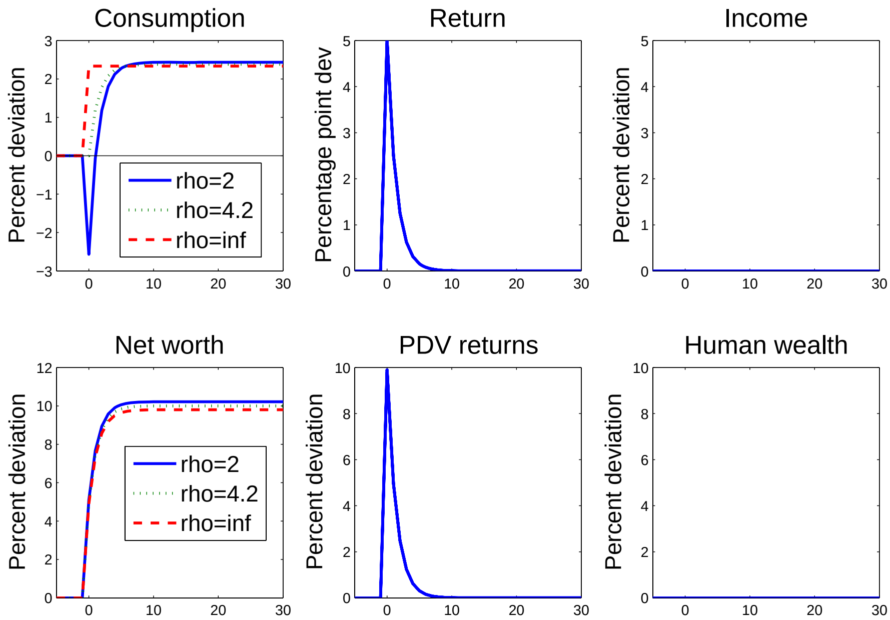

Note: This figure plots impulse-responses to a return shock implied by Section 2's PIH model. Before the time-t shock, the economy is in steady-state. We simulate a persistent return shock such that on impact, the percentage point deviation of returns from steady-state $\overline{r}_t = 5$ and afterwards, return deviations decay according to $E_t\ \overline{r}_{t+i+1} = 0.5\ \overline{r}_{t+i}$ for $i \geq 0$. The solid blue line graphs the consumption response when the coefficient of relative risk aversion $\rho = 2$, the dotted green line plots the response when $\rho = 4.20$ and the dashed red line when $\rho = \infty$. With $\rho = \infty$, there is no intertemporal substitution, and consumption increases instantly according to the wealth effect. With $\rho = 2$ and $\rho = 4.20$, consumers also desire to postpone consumption as they expect a period of above-equilibrium returns to saving. All variables are graphed in percent deviations from steady-state, except the return which is in percentage point deviations from steady-state. The horizontal axis plots quarters since the shock.

Figure 3: Data: Real Per Capita Consumption, Income and Net Worth

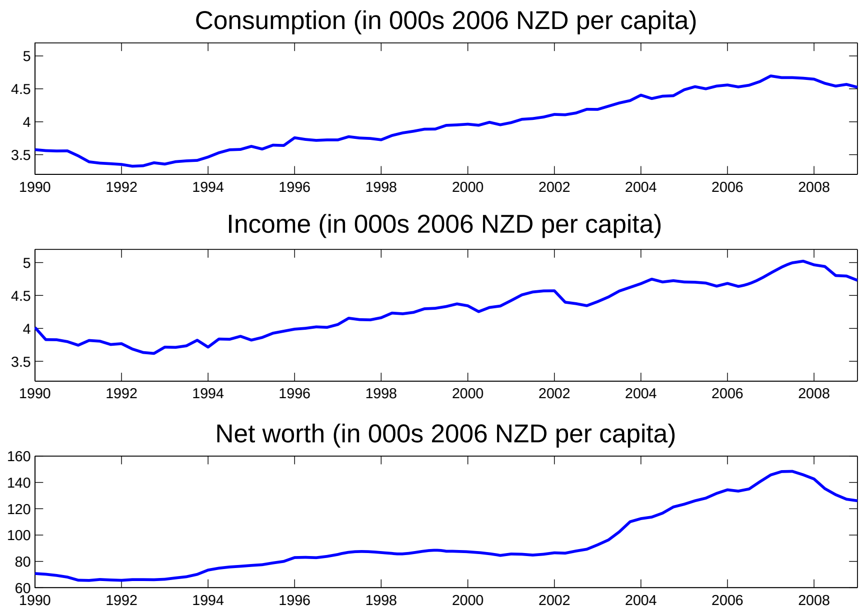

Note: This figure plots New Zealand consumption, income, and net worth, for 1990Q1-2009Q1. All data are in per-capita terms, and deflated by the Consumer Price Index with base year 2006. Consumption is quarterly household consumption of durable goods, non-durable goods and services, but excludes housing consumption. Income is quarterly after-tax labor income. We construct the end-of-quarter stock of net worth as housing plus financial assets minus household debt, excluding durable goods.

Figure 4: Empirical and Theoretical Responses to a Transitory Shock

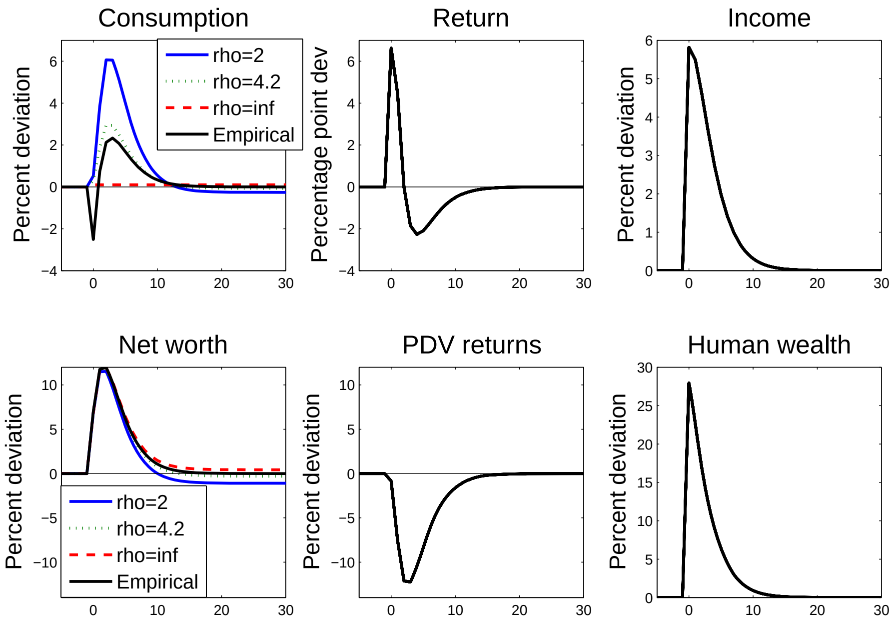

Note: This figure compares the empirical consumption response to a transitory shock implied by Section 4's VECM with the theoretical response implied by Section 2's PIH model. All variables are graphed in percent deviations from steady-state, except the return which is in percentage point deviations from steady-state. The economy is in steady-state before the shock. Bold black lines in the consumption, net worth and income panels indicate empirical impulse-responses. Returns follow from the empirical impulse-responses through equation (11). Conditional on a transitory shock that implies increasing net worth and income, actual consumption declines on impact, but then increases, displaying a hump-shaped response. The red dashed lines document that the PIH model cannot replicate the hump-shaped response without intertemporal substitution, i.e. when $\rho = \infty$. The PIH model does generate transitory consumption variation when $\rho = 2$ and $\rho = 4.20$. However, the PIH model fails to account for the empirical drop in consumption on impact, and tends to call for a more pronounced consumption response in the first few years after the shock.

Figure 5: Reconstituting Theoretical Response: Transitory Income Shock

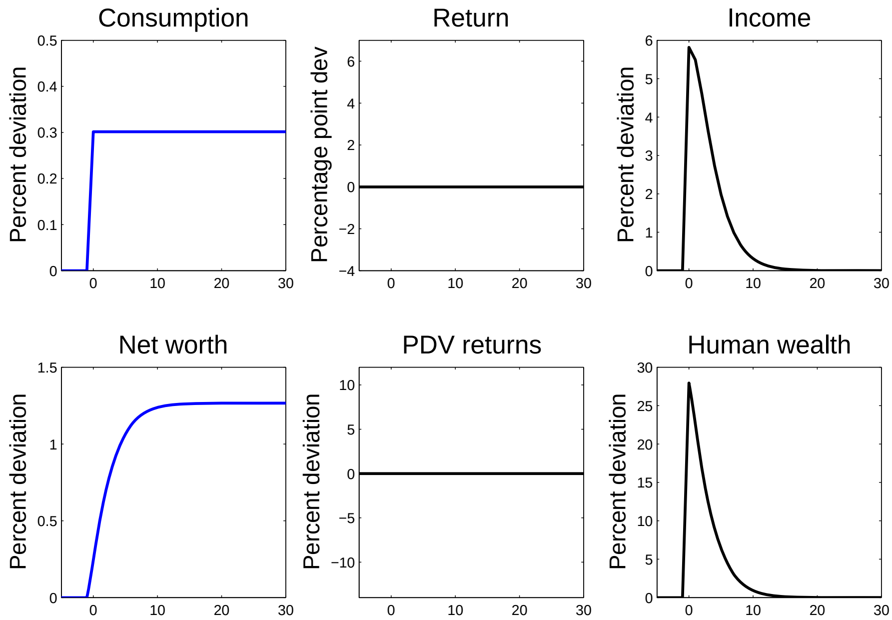

Note: This figure separately considers the empirical transitory income shock. It differs from Figure 4 in that it keeps returns constant at steady-state. Black lines are the empirical shocks. The blue lines graph the theoretical consumption response and its implications for net worth accumulation. Consumption increases on impact, but then stays constant at its new level. The consumption response is fairly small. The response is the same for any value of the coefficient of relative risk aversion ρ.

Figure 6: Reconstituting Theoretical Response: Returns Above Steady-State

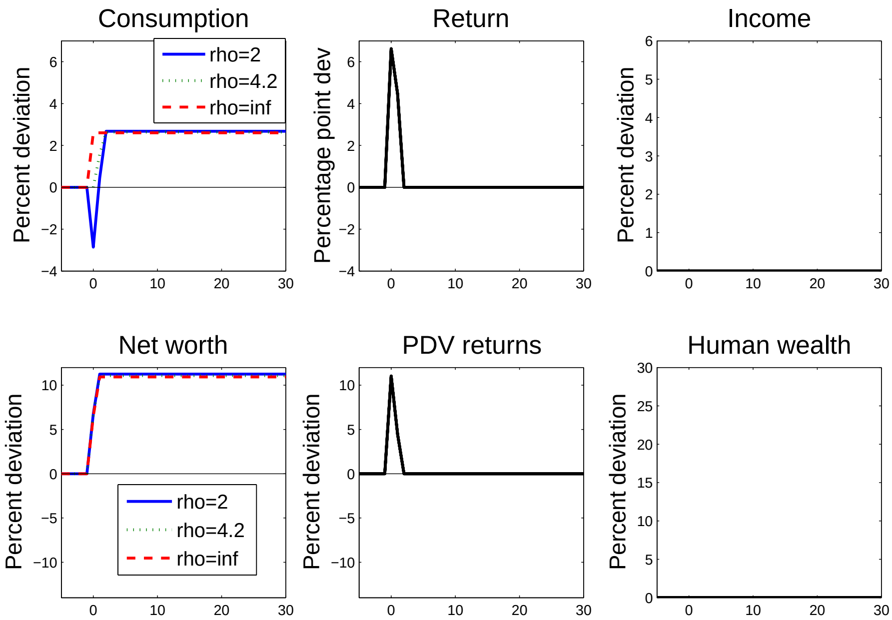

Note: This figure separately considers the part of the empirical returns impulse where returns are above steady-state. Returns are above steady-state on impact and one quarter after the shock. Unlike in Figure 4, this figure keeps subsequent returns at steady-state, and also keeps income at steady-state. Black lines are the empirical shocks. The solid blue line graphs the consumption response when the coefficient of relative risk aversion $\rho = 2$, the dotted green line plots the response when $\rho = 4.20$ and the dashed red line when $\rho = \infty$. In the latter case, consumption increases on impact and stays constant from that point on, reflecting the wealth effect. In the other two cases, intertemporal substitution implies that consumption is initially lower as consumers postpone consumption in order to collect high wealth returns.

Figure 7: Reconstituting Theoretical Response: Returns Below Steady-State

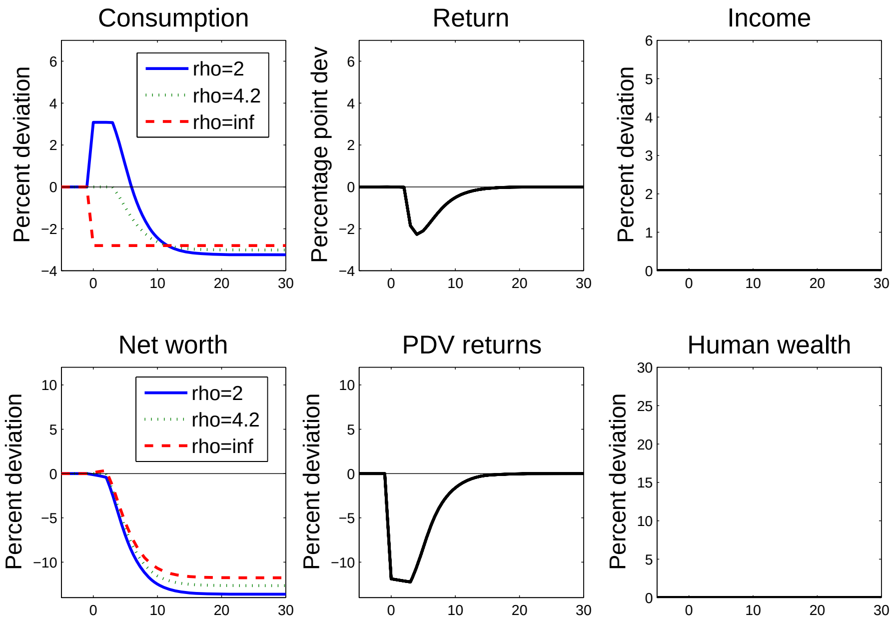

Note: This figure separately considers returns below steady-state. It differs from Figure 4 in that it keeps returns in steady-state at the time of the shock and one quarter after the shock, and keeps income at steady-state. Black lines are the empirical shocks. The solid blue line graphs the consumption response when the coefficient of relative risk aversion $\rho = 2$, the dotted green line plots the response when $\rho = 4.20$ and the dashed red line when $\rho = \infty$. In the latter case, consumption decreases on impact and stays constant from that point on, reflecting the wealth effect. In the other two cases, intertemporal substitution implies that consumption is initially higher as households bring consumption forward in anticipation of low returns on wealth.

Figure 8: Allowing for Mispredictions: Expected Paths as of Time of the Shock

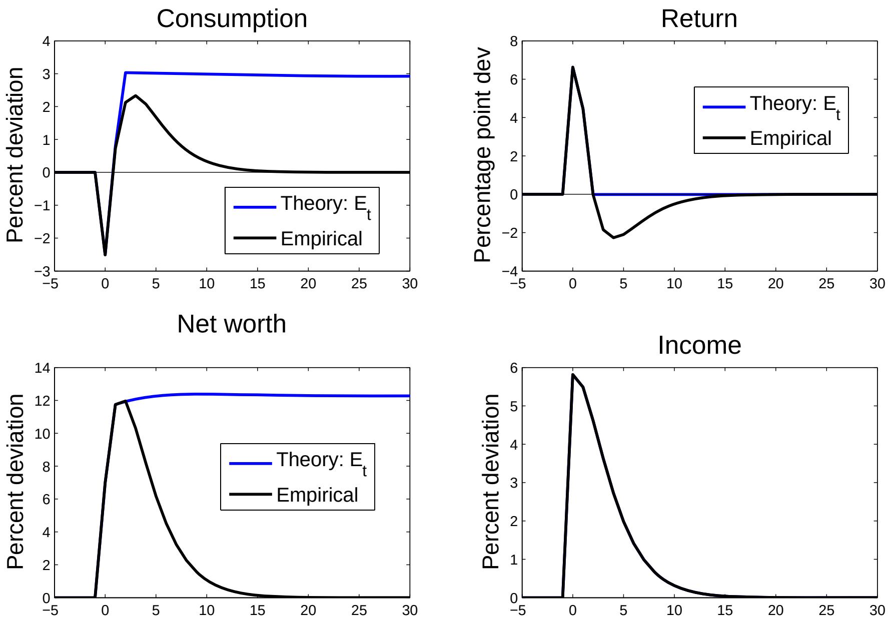

Note: The bold black lines show the empirical impulse-responses to a time-t transitory shock. The blue lines in the returns panel show a path of expectations for future returns that implies that households in the PIH model choose to consume the empirical level of consumption at the time of the shock. The blue lines in the consumption and net worth panels show the implied paths for intended consumption and expected net worth. Expectations are as of time t. The PIH model explains the consumption response on impact if at that time, households expect the net worth increase to be permanent.

Figure 9: Allowing for Mispredictions: Updating Expectations

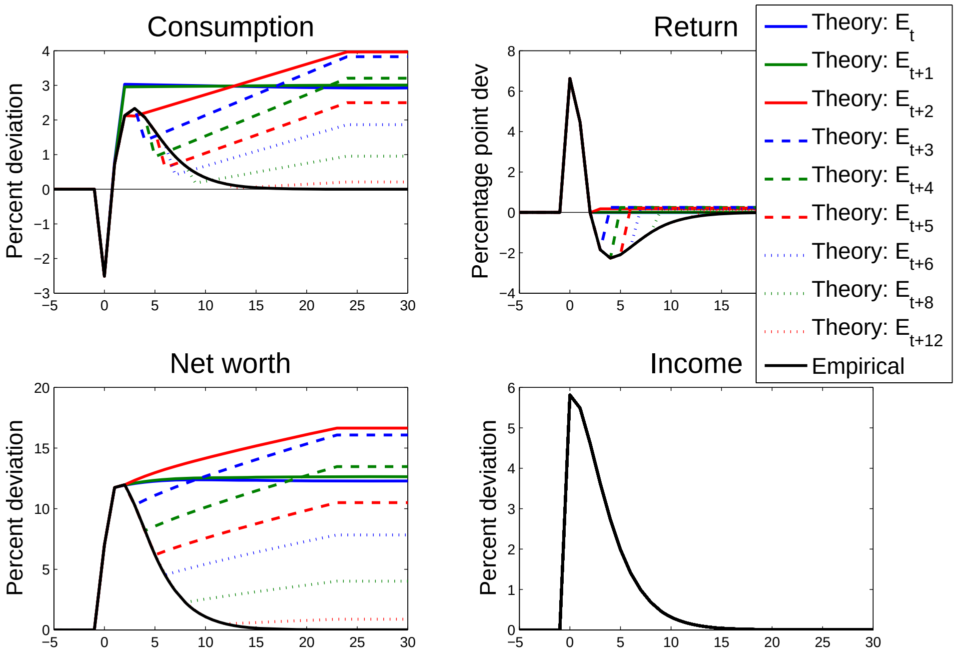

Note: The bold black lines show the empirical impulse-responses to a time-t transitory shock. The colored lines graph successive paths for expectations about future returns, and implied consumption and net worth paths. Solid colored lines indicate expectations as of times t through $t+2$. Dashed lines indicate expectations as of $t+3$ through $t+5$. Dotted lines pertain to expectations as of $t+6$, $t+8$ and $t+12$. The PIH model explains the empirical consumption and net worth path when households initially expect the net worth increase to be permanent, but gradually learn that it is in fact transitory.

Table 1: Augmented Dickey-Fuller Tests for Stationarity

	Lag length 0	1	2	3	4	5
Consumption	$0.17^{A,B}$	-0.009	-0.22	-0.31	-0.54	-0.81
Net worth	0.07	$-0.76^{A,B}$	-0.76	-1.01	-1.15	-1.07
Income	-0.30^{B}	-0.42	-0.71	-1.01	-0.56	-0.19^{A}
5% critical value	-2.86					

Note: Every variable is in log real per capita terms. For every variable, we report the ADF test statistic for the null hypothesis that the series has a unit root. 'Lag length' refers to the order of augmentation of the testing regression. Superscripts A and B refer to the lag length selected by the Akaike and Bayesian Information Criteria, respectively. In every test regression, we include a constant but no trend. We cannot reject the null hypothesis of a unit root for any variable at any of the reported orders of augmentation. *** indicates significance at 1% level, ** at 5% level, and * at 10% level.

Table 2: Johansen L-Max and Trace Tests for Cointegration

				Lag length			
Test	H_0	H_A	5% crit. val.	1	2	3	4
L-max	r=0	r=1	21.13	24.38**	36.40***	37.52***	29.91***
	r=1	r=2	14.26	4.24	4.50	3.44	5.57
Trace							
	r=0	r$\geq$ 1	29.80	28.66*	40.99***	41.15***	35.71***
	r=1	r$\geq$ 2	15.49	4.28	4.59	3.64	5.80

Note: This table presents Johansen L-max and Trace test statistics for the number of cointegrating vectors in a three-variable system containing consumption, net worth and income. 'Lag length' refers to the number of lagged differences included in the VECM for the purpose of conducting the L-max and Trace tests. We include a constant, but no trend in either the VECM or the long-run relationship. For most lag specifications, we can reject the null hypothesis of no cointegration at the 5 percent level. We cannot reject the null hypothesis of one cointegrating vector. *** indicates significance at 1% level, ** at 5% level, and * at 10% level.

Table 3: Estimated Vector Error Correction Model

	Δc_t	Δa_t	Δy_t
constant	-0.06*	0.18***	0.15**
	(-1.68)	(3.13)	(2.61)
η_{t-1}	-0.11*	0.31***	0.26**
	(-1.73)	(3.09)	(2.60)
Δc_{t-1}	-0.17	-0.39**	-0.04
	(-1.47)	(-2.13)	(-0.22)
Δa_{t-1}	0.28***	0.69***	0.15
	(4.74)	(7.35)	(1.59)
Δy_{t-1}	-0.02	0.27**	0.12
	(-0.25)	(2.08)	(0.96)
R^2	0.27	0.58	0.17
$\overline{R}^2$	0.23	0.55	0.12

Note: This table reports estimates for the VECM in equation (9). t-statistics are in parentheses. η_{t-1} is the error-correction term. The results suggest that income and net worth error-correct at the 5 percent level or better, while consumption error-corrects at the 10 percent level. *** indicates significance at 1% level, ** at 5% level, and * at 10% level.

Table 4: Permanent-Transitory Decomposition

	Consumption	Net Worth	Income
Transitory			
1	24.56	37.34	24.32
4	25.60	27.96	23.49
8	25.45	27.76	23.64
20	25.47	27.82	23.66
∞	25.47	27.82	23.66
Permanent			
1	75.44	62.66	75.68
4	74.40	72.04	76.51
8	74.55	72.24	76.36
20	74.53	72.18	76.34
∞	74.53	72.18	76.34

Note: For consumption, net worth and income, respectively, this table provides the fraction of the forecast error variance, in percent, explained by transitory and permanent shocks. For consumption, the variance pertains to the forecast error $\Delta c_{t+h-1} - E_{t-1}\Delta c_{t+h-1}$ associated with consumption growth $h-1$ quarters after a shock in quarter t. Analogous definitions apply for net worth and income. We compute transitory and permanent shocks by Gonzalo-Ng (2001) decomposition as in equation (10). Next, we orthogonalize the permanent and transitory shocks by applying a Cholesky decomposition to the vector of transitory and permanent shocks, assuming that transitory shocks do not contemporaneously affect permanent shocks. At most forecast horizons, we find that transitory shocks account for about one quarter of the variation in consumption, net worth and income.

Appendix I

This appendix derives the equations for Section 2.

1 Deriving log-linearized Euler equation

This subsection derives the log-linearized Euler equation from the representative consumer's optimization problem with objective function (1) and flow budget constraint (2). From the first order conditions for consumption in two consecutive periods, we obtain:

$$u'(C_t) = \beta(1 + r_t)\ E_t\ u'(C_{t+1}) \tag{A1}$$

This is the standard Euler equation except for the fact that returns are time-varying. The equation indicates that given our specification of the optimization problem, the return in period t influences expected consumption growth from period t to the next period.

Assuming Constant Relative Risk Aversion (CRRA) utility, i.e. $u(C_t) = \left[C_t^{1-\rho}/(1-\rho)\right]$, the Euler equation becomes:

$$C_t^{-\rho} = \beta(1 + r_t)\ E_t\ C_{t+1}^{-\rho} \tag{A2}$$

Taking log deviations from steady-state, we obtain:

$$E_t\ \widehat{c}_{t+1} \approx \widehat{c}_t + \frac{1}{\rho}\ \overline{r}_t \tag{A3}$$

where $\widehat{c}_t \equiv \log\ C_t - \log\ C^*$, $\overline{r}_t \equiv r_t - r^*$, C^* is steady-state consumption and r^* is the steady-state return. To derive (A3), we used the approximation $\log(1 + r_t) \approx r_t$ which follows from a first-order Taylor expansion around $r_t = 0$, and the corresponding approximation for the steady-state return $\log(1+r^*) \approx r^*$. Equation (A3) is the log-linearized Euler equation (3) in the text.

Repeatedly applying the law of iterated expectations and repeatedly rolling (A3) one period forward, we obtain an expression for future expected log deviations from steady-state $E_t\ \widehat{c}_{t+i}$:

$$E_t\ \widehat{c}_{t+i} \approx \widehat{c}_t + \frac{1}{\rho}\sum_{j=1}^{i} E_t\ \overline{r}_{t+j-1} \quad \forall i \geq 1 \tag{A4}$$

This is equation (4) in the text. Note that for $i = 0$, $E_t\ \widehat{c}_{t+i} = \widehat{c}_t$ such that for $i = 0$ one replaces the term $\sum_{j=1}^{i} E_t\ \overline{r}_{t+j-1}$ by 0.

2 Deriving log-linearized intertemporal budget constraint

This subsection log-linearizes the flow budget constraint (2) and solves it forward to obtain the log-linearized intertemporal budget constraint. The flow budget constraint reads $A_t = (1 + r_t)\ (A_{t-1} + Y_t - C_t)$. Taking log deviations from steady-state, and defining $Z_{t-1} \equiv A_{t-1} + Y_t - C_t$, we obtain:

$$\widehat{a}_t \approx \overline{r}_t + \widehat{z}_{t-1} \tag{A5}$$

where $\widehat{a}_t \equiv \log\ A_t - \log\ A^*$ and $\widehat{z}_t \equiv \log\ Z_t - \log\ Z^*$. A^* is steady-state net worth, $Z^* = A^* + Y^* - C^*$, and Y^* is steady-state income. We apply the same approximations for returns as under equation (A3).

Log-linearizing the equation $Z_{t-1} = A_{t-1} + Y_t - C_t$ around steady-state, and substituting the resulting expression for $\widehat{z}_{t-1}$ into (A5), we obtain:

$$\widehat{a}_t \approx \overline{r}_t + (1 + r^*)\left(\widehat{a}_{t-1} + \frac{Y^*}{A^*}\widehat{y}_t - \frac{C^*}{A^*}\widehat{c}_t\right) \tag{A6}$$

Which is the log-linearized flow budget constraint, equation (5) in the text. Re-arranging that equation, iterating forward, and imposing the transversality condition $\lim_{n\to\infty}\left[1/(1+r^*)\right]^{n+1}\widehat{a}_{t+n}=0$, we obtain:

$$\widehat{a}_{t-1}\approx\sum_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\left(\frac{C^*}{A^*}E_t\ \widehat{c}_{t+i}\ -\frac{Y^*}{A^*}E_t\ \widehat{y}_{t+i}-\frac{1}{1+r^*}\ E_t\ \overline{r}_{t+i}\right)\quad \text{(A7)}$$

Which is equation (6) in the text.

3 Deriving log-linearized consumption function

Subtituting equation (A4) into equation (A7), re-arranging, and assuming that $r^*>0$ such that we can apply the formula for the infinite sum $\sum_{i=0}^{\infty}[1/(1+r^*)]^i=(1+r^*)/r^*$, we obtain:

$$\widehat{c}_t=\frac{r^*}{1+r^*}\left[\begin{array}{c}\frac{A^*}{C^*}\widehat{a}_{t-1}+\frac{A^*}{C^*}\frac{1}{1+r^*}\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\ E_t\ \overline{r}_{t+i}\\ +\frac{Y^*}{C^*}\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i E_t\ \widehat{y}_{t+i}-\frac{1}{\rho}\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\left(\sum_{j=1}^{i}E_t\ \overline{r}_{t+j-1}\right)\end{array}\right]\quad \text{(A8)}$$

Two terms in equation (A8) depend on returns, one of which involves a double summation. To simplify the expression, we do further work on the term $-\frac{1}{\rho}\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\left(\sum_{j=1}^{i}E_t\ \overline{r}_{t+j-1}\right)$. Recalling from the discussion under (A4) that for $i=0$ one replaces $\sum_{j=1}^{i}E_t\ \overline{r}_{t+j-1}$ by 0, we can express this term explicitly as:

$$-\frac{1}{\rho}\left[\begin{array}{c}0+\left(\frac{1}{1+r^*}\right)\overline{r}_t+\left(\frac{1}{1+r^*}\right)^2\left(\overline{r}_t+E_t\ \overline{r}_{t+1}\right)\\ +\left(\frac{1}{1+r^*}\right)^3\left(\overline{r}_t+E_t\ \overline{r}_{t+1}+E_t\ \overline{r}_{t+2}\right)+...\end{array}\right]\quad \text{(A9)}$$

Regrouping terms, this is equivalent to:

$$-\frac{1}{\rho}\left[\begin{array}{c}\left(\frac{1}{1+r^*}\right)\left(\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\right)\overline{r}_t+\left(\frac{1}{1+r^*}\right)^2\left(\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\right)E_t\ \overline{r}_{t+1}\\ +\left(\frac{1}{1+r^*}\right)^3\left(\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\right)E_t\ \overline{r}_{t+2}+...\end{array}\right]\quad \text{(A10)}$$

Still assuming that $r^* > 0$, such that $\sum\limits_{i=0}^{\infty} [1/(1+r^*)]^i = (1+r^*)/r^*$, we obtain:

$$-\frac{1}{\rho}\frac{1}{r^*}\left[\overline{r}_t + \left(\frac{1}{1+r^*}\right) E_t\ \overline{r}_{t+1} + \left(\frac{1}{1+r^*}\right)^2 E_t\ \overline{r}_{t+2} + ...\right] \tag{A11}$$

which is:

$$-\frac{1}{\rho}\frac{1}{r^*}\sum_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\ E_t\ \overline{r}_{t+i} \tag{A12}$$

Now replacing the term $-\frac{1}{\rho}\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i\left(\sum_{j=1}^{i} E_t\ \overline{r}_{t+j-1}\right)$ in (A8) by this new equivalent expression, we obtain:

$$\widehat{c}_t \approx \frac{r^*}{1+r^*}\left[\begin{array}{c} \frac{A^*}{C^*}\widehat{a}_{t-1} + \left(\frac{A^*}{C^*}\frac{1}{1+r^*} - \frac{1}{\rho}\ \frac{1}{r^*}\right)\ \sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i E_t\ \overline{r}_{t+i} \\ +\frac{Y^*}{C^*}\sum\limits_{i=0}^{\infty}\left(\frac{1}{1+r^*}\right)^i E_t\ \widehat{y}_{t+i} \end{array}\right] \tag{A13}$$

Which is equation (7) in the text.

Appendix II

This appendix discusses data sources and construction for consumption, net worth and income.

Our series for household **consumption** is from the quarterly national accounts published by Statistics New Zealand (SNZ). Household consumption excludes the consumption of private non-profit organizations. We include consumption on durable and non-durable goods and services, but exclude housing consumption, which is the imputed gross rental value of owning a home.

Household **net worth** is the value of the housing stock plus financial assets minus debt.

We construct a quarterly series for housing wealth from SNZ data on the total number of dwellings and the Quotable Value New Zealand (QVNZ) house price index. As an input into its house price index, QVNZ obtains capital values from local authorities which conduct periodic revaluations for the purpose of levying rates.

Our quarterly series on financial assets reflects Reserve Bank of New Zealand estimates from 1995 onwards. The Reserve Bank does not compute a quarterly series for earlier quarters. Holdings of assets other than equity evolve gradually over time, such that we construct a pre-1995 quarterly measure for each of these assets by interpolating its respective annual series with a cubic spline. To capture higher-frequency variation in stock prices, we construct a quarterly measure of direct equity holdings before 1995 by ensuring that the quarterly growth rate of the interpolated series matches the growth rate of a weighted average of the New Zealand and Australian capital price indices. At any quarter, we set the weight on the New Zealand capital price index equal to the proportion of direct equity that is domestic, while the weight for the Australian index equals the proportion of direct equity that is held abroad.

Our series for household debt reflects Reserve Bank of New Zealand estimates.

As a caveat, note that in New Zealand, measured household net worth does not include assets held in farms and other unincorporated businesses, nor assets in privately held corporations. We did not impute wealth held in unincorporated businesses and privately held firms.

We compute after-tax labor **income** as follows: pre-tax labor income minus tax payments plus transfer income.

SNZ produces an annual series for labor income. However, there is no directly available measure for quarterly labor income in New Zealand. We construct quarterly labor income by multiplying average hourly earnings (including overtime payments) from the Quarterly Employment Survey (QES) by hours worked from the Household Labour Force Survey (HLFS). The latter survey includes hours worked by agricultural workers and workers that are otherwise self-employed, but the former study excludes earnings of workers in those sectors. Assuming that a typical farmer or entrepreneur earns the same hourly labor income as the average hourly income in other sectors, our measure of labor income captures the compensation of employees as well as entrepreneurial income.

We construct a series of quarterly tax payments using our measure of labor income as well as data on tax rates by income bracket. For the latter, we interpolated an annual series on implied effective tax rates from SNZ's annual national accounts.

As the final input for our labor income measure, we compute a quarterly measure of transfer income. For unemployment and pension benefits, the Ministry of Social Development provides benefit rates, from which we estimated the number of beneficiaries. For other transfer receipts, there are no available estimates on benefit rates. We account for these other payments by interpolating the corresponding annual data from Work and Income New Zealand. Unlike unemployment benefits, these other payments tend to vary gradually over time, such that it is unlikely that the interpolated series omit a substantial degree of quarterly fluctuation.

References

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

Aoki, K, J Proudman and G Vlieghe (2004), "House Prices, Consumption, and Monetary Policy: a Financial Accelerator Approach", *Journal of Financial Intermediation*, 13, 414-435.

Attanasio, O P (1999), "Consumption", in *Handbook of Macroeconomics*, Volume 1B, eds. John B. Taylor and Michael Woodford, Elsevier.

Attanasio, O P, J Banks, C Meghir and G Weber (1999), "Humps and Bumps in Lifetime Consumption", *Journal of Business and Economic Statistics*, 17, 22-35.

Attanasio, O P, L Blow, R Hamilton and A Leicester (2009), "Booms and Busts: Consumption, House Prices and Expectations", *Economica*, 76, 20-50.

Attanasio, O P and G Weber (1989), "Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption", *Economic Journal*, 99, 59-73.

Barro, R J (1981), "Output Effects of Government Purchases", *Journal of Political Economy*, 89, 1086-1121.

Basu, S and M S Kimball (2002), "Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption", mimeo, University of Michigan.

Bostic, R, S Gabriel and G Painter (2009), "Housing Wealth, Financial Wealth, and Consumption: New Evidence from Micro Data", *Regional Science and Urban Economics*, 39, 79-89.

Calza, A, T Monacelli and L Stracca (2007), "Mortgage Markets, Collateral Constraints, and Monetary Policy: Do Institutional Factors Matter?", CEPR Discussion Paper 6231.

Campbell, J R and Z Hercowitz (2005), "The Role of Collateralized Household Debt in Macroeconomic Stabilization", NBER Working Paper W11330.

Campbell, J Y (1996), "Understanding Risk and Return", *Journal of Political Economy*, 104, 298-345.

Campbell, J Y and J Cocco (2007), "How do House Prices Affect Consumption? Evidence from Micro Data", *Journal of Monetary Economics*, 54, 591-621.

Campbell, J Y and N G Mankiw (1989), "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence", in *NBER Macroeconomics Annual* 1989, University of Chicago Press.

Campbell, J Y and L M Viceira (1999), "Consumption and Portfolio Decisions when Expected Returns are Time Varying", *Quarterly Journal of Economics*, 114, 433-495.

Carroll, C D (1997), "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis", *Quarterly Journal of Economics*, 112, 1-55.

Carroll, C D and M S Kimball (2001), "Liquidity Constraints and Precautionary Saving", NBER Working Paper 8496.

Carroll, C D, M Otsuka and J Slacalek (2006), "How Large is the Housing Wealth Effect? A New Approach", NBER Working Paper 12746.

Case, K E, J M Quigley and R J Shiller (2005), "Comparing Wealth Effects: the Stock Market versus the Housing Market", *B.E. Journal of Macroeconomics (Advances)*, 5, article 1.

Chen, J (2006), "Re-evaluating the Association between Housing Wealth and Aggregate Consumption: New Evidence from Sweden", *Journal of Housing Economics*, 15, 321-348.

Coleman, A (2007), "Credit Constraints and Housing Markets in New Zealand", Reserve Bank of New Zealand Discussion Paper 2007/11.

Davis, M A and J Heathcote (2005), "Housing and the Business Cycle", *International Economic Review*, 46, 751-784.

Davis, M A and M G Palumbo (2001), "A Primer on the Economics and Time Series Econometrics of Wealth Effects", Finance and Economics Discussion Series 2001-09, Board of Governors of the Federal Reserve System.

Deaton, A (1991), "Saving and Liquidity Constraints", *Econometrica*, 59, 1221-1248.

Deaton, A (1992), "*Understanding Consumption*", Oxford University Press.

De Veirman, E and A Dunstan (2010), "Debt Dynamics and Excess Sensitivity of Consumption to Transitory Wealth Changes", Reserve Bank of New Zealand Discussion Paper 2010/09.

Epstein, L G and S E Zin (1991), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: an Empirical Analysis", *Journal of Political Economy,* 99, 263-286.

Fernandez-Corugedo, E, S Price and A P Blake (2007), "The Dynamics of Aggregate UK Consumers' Non-Durable Expenditure", *Economic Modelling*, 24, 453-469.

Fisher, L A, G Otto and G M Voss (2010), "The Response of Australian Consumption to Housing Wealth", *Journal of Macroeconomics*, 32, 284-299.

Friedman, M (1957), *"A Theory of the Consumption Function"*, Princeton University Press.

Gonzalo, J and S Ng (2001), "A Systematic Framework for Analyzing the Dynamic Effects of Permanent and Transitory Shocks", *Journal of Economic Dynamics and Control,* 25, 1527–46.

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

Guvenen, F (2006), "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: a Macroeconomic Perspective", *Journal of Monetary Economics*, 53, 1451-1472.

Hall, R E (1979), "Labor Supply and Aggregate Fluctuations", NBER Working Paper 385.

Hall, R E (1988), "Intertemporal Substitution in Consumption", *Journal of Political Economy*, 96, 339-357.

Hansen, L P and K J Singleton (1983), "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns", *Journal of Political Economy*, 91, 249-265.

Hayashi, F (2000), *"Econometrics"*, Princeton University Press.

Iacoviello, M (2005), "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle", *American Economic Review*, 95, 739-764.

Iacoviello, M and S Neri (2010), "Housing Market Spillovers: Evidence from an Estimated DSGE Model", *American Economic Journal: Macroeconomics,* 2, 125-164.

Japelli, T and M Pagano (1994), "Saving, Growth, and Liquidity Constraints", *Quarterly Journal of Economics,* 109, 83-109.

Kimball, M S (1990), "Precautionary Saving in the Small and in the Large", *Econometrica,* 58, 53-73.

King, R G (1989), "Value and Capital in the Equilibrium Business Cycle Program", Rochester Center for Economic Research Working Paper 207.

Kiyotaki, N, A Michaelides and K Nikolov (2006), "From Shirtsleeves to Shirtsleeves in a Long Lifetime", mimeo, Princeton University.

Krusell, P and A A Smith, Jr (1999), "On the Welfare Effects of Eliminating Business Cycles", *Review of Economic Dynamics*, 2, 245-272.

Lantz, C D and P G Sarte (2001), "Consumption, Savings, and the Meaning of the Wealth Effect in General Equilibrium", Economic Quarterly 87, 53-71, Federal Reserve Bank of Richmond.

Lehnert, A (2004), "Housing, Consumption, and Credit Constraints", Finance and Economics Discussion Series 2004-63, Federal Reserve Board.

Lettau, M and S Ludvigson (2001), "Consumption, Aggregate Wealth, and Expected Stock Returns", *Journal of Finance*, 56, 815–849.

Lettau, M and S Ludvigson (2004), "Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption", *American Economic Review*, 94, 276-299.

Ludwig, A and T Slok (2004), "The Relationship between Stock Prices, House Prices and Consumption in OECD Countries", *B.E. Journal of Macroeconomics (Topics)*, 4, Article 4.

Lustig, H and S Van Nieuwerburgh (2010), "How Much Does Household Collateral Constrain Regional Risk Sharing?", *Review of Economic Dynamics*, 13, 265-294.

Mankiw, N G (1981), "The Permanent Income Hypothesis and the Real Interest Rate", *Economics Letters*, 7, 307-311.

Mankiw, N G (1985), "Consumer Durables and the Real Interest Rate", *Review of Economics and Statistics*, 67, 353-362.

Modigliani, F and R Brumberg (1954), "Utility Analysis and the Consumption Function: an Interpretation of Cross-Section Data", in *Post Keynesian Economics*, ed. Kenneth K. Kurihara, Rutgers University Press.

Monacelli, T (2009), "New Keynesian Models, Durable Goods, and Collateral Constraints", *Journal of Monetary Economics*, 56, 242-254.

Muellbauer, J (1994), "The Assessment: Consumer Expenditure", *Oxford Review of Economic Policy,* 10, 1-41.

Ogaki, M and C M Reinhart (1998), "Measuring Intertemporal Substitution: the Role of Durable Goods", *Journal of Political Economy*, 106, 1078-1098.

Ortalo-Magne, F and S Rady (2006), "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints", *Review of Economic Studies*, 73, 459-485.

Rudd, J and K Whelan (2006), "Empirical Proxies for the Consumption-Wealth Ratio", *Review of Economic Dynamics*, 9, 34–51.

Slacalek, J (2009), "What Drives Personal Consumption? The Role of Housing and Financial Wealth", *B.E. Journal of Macroeconomics (Topics)*, 9, Article 37.

Vissing-Jorgensen, A (2002), "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution", NBER Working Paper 8896.

Zeldes, S P (1989), "Consumption and Liquidity Constraints: an Empirical Investigation", *Journal of Policitical Economy*, 97, 305-346.