

SUPPLEMENT TO “EARNINGS AND CONSUMPTION DYNAMICS:
A NONLINEAR PANEL DATA FRAMEWORK”
(*Econometrica*, Vol. 85, No. 3, May 2017, 693–734)

BY MANUEL ARELLANO, RICHARD BLUNDELL, AND STÉPHANE BONHOMME

S1. CONSUMPTION RESPONSES IN A TWO-PERIOD MODEL

CONSIDER A STANDARD TWO-PERIOD SETUP, with a single risk-free asset. Let A_t denote beginning-of-period- t assets, and assume that $A_3 = 0$. Agents have CRRA utility. The Euler equation (assuming $\beta(1+r) = 1$ for simplicity) is

$$C_1^{-\gamma} = \mathbb{E}_1[(1+r)A_2 + Y_2]^{-\gamma},$$

where γ denotes risk aversion and the expectation is conditional on period-1 information. Here we have used the budget constraint $A_3 = (1+r)A_2 + Y_2 - C_2 = 0$. Equivalently,

$$C_1^{-\gamma} = \mathbb{E}_1[(1+r)^2 A_1 + (1+r)Y_1 - (1+r)C_1 + Y_2]^{-\gamma}. \quad (\text{S1})$$

Let $X_1 = (1+r)A_1 + Y_1$ denote “cash on hand” (as in Deaton (1991)). Let also $Y_2 = \mathbb{E}_1(Y_2) + \sigma W$. We will expand the Euler equation as $\sigma \rightarrow 0$. We denote the certainty equivalent consumption level as

$$\bar{C}_1 = \frac{(1+r)X_1 + \mathbb{E}_1(Y_2)}{2+r}.$$

Expanding in orders of magnitude of σ , we have

$$C_1 \approx \bar{C}_1 + a\sigma + b\sigma^2 + c\sigma^3. \quad (\text{S2})$$

It is easy to see that $a = 0$, since $\mathbb{E}_1(W) = 0$. Hence,

$$C_1^{-\gamma} \approx \bar{C}_1^{-\gamma} \left(1 - \frac{\gamma}{\bar{C}_1} b\sigma^2 - \frac{\gamma}{\bar{C}_1} c\sigma^3 \right). \quad (\text{S3})$$

Moreover, by (S1) and (S2),

$$C_1^{-\gamma} \approx \mathbb{E}_1[(\bar{C}_1 - (1+r)b\sigma^2 - (1+r)c\sigma^3 + \sigma W)^{-\gamma}],$$

from which it follows that

$$\begin{aligned} C_1^{-\gamma} \approx \bar{C}_1^{-\gamma} \mathbb{E}_1 \left[1 + \frac{\gamma}{\bar{C}_1} (1+r)b\sigma^2 + \frac{\gamma}{\bar{C}_1} (1+r)c\sigma^3 - \frac{\gamma}{\bar{C}_1} \sigma W \right. \\ \left. + \frac{\gamma(\gamma+1)}{2} \left(\frac{1}{\bar{C}_1} \right)^2 \sigma^2 W^2 - \gamma(\gamma+1) \left(\frac{1}{\bar{C}_1} \right)^2 (1+r)b\sigma^3 W \right. \\ \left. - \frac{\gamma(\gamma+1)(\gamma+2)}{6} \left(\frac{1}{\bar{C}_1} \right)^3 \sigma^3 W^3 \right]. \quad (\text{S4}) \end{aligned}$$

Finally, equating the coefficients of σ^2 and σ^3 in (S3) and (S4), using that $\mathbb{E}_1(W) = 0$, and denoting as $R = (1+r)X_1 + \mathbb{E}_1(Y_2)$ the expected period-2 resources, we obtain

$$b = -\frac{\gamma+1}{2R}\mathbb{E}_1(W^2), \quad c = \frac{(2+r)(\gamma+1)(\gamma+2)}{6R^2}\mathbb{E}_1(W^3).$$

This yields the following expression for period-1 consumption:

$$\begin{aligned} C_1 \approx & \underbrace{\frac{(1+r)X_1 + \mathbb{E}_1(Y_2)}{2+r}}_{\text{certainty equivalent}} - \underbrace{\frac{\gamma+1}{2R}\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2)}_{\text{precautionary (variance)}} \\ & + \underbrace{\frac{(2+r)(\gamma+1)(\gamma+2)}{6R^2}\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3)}_{\text{precautionary (skewness)}}. \end{aligned} \tag{S5}$$

Note that $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2)$ is the conditional variance of Y_2 , and $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3)$ is its conditional third-order moment.

Example: A Simple Nonlinear Earnings Process. To illustrate the effect of earnings shocks on consumption in this model, we consider the following simple earnings process (in levels):

$$Y_2 = Y_2^D + \rho(Y_1^P, V_2)Y_1^P + V_2 + Y_2^T,$$

where Y_2^D is the deterministic component, $Y_2^P = \rho(Y_1^P, V_2)Y_1^P + V_2$ is the persistent component, and Y_2^T is the transitory component. We set $\rho(Y_1^P, V_2) = 1 - \delta$ if $(Y_1^P < -c, V_2 > b)$ or $(Y_1^P > c, V_2 < -b)$, and $\rho(Y_1^P, V_2) = 1$ otherwise. Moreover, $\Pr(V_2 > b) = \Pr(V_2 < -b) = \tau$, with $\tau < 1/2$, and we assume that V_2 and Y_2^T are symmetrically distributed with zero mean. This earnings process has the following properties:

- If $|Y_1^P| \leq c$, then the process coincides with the “canonical” earnings model (in levels). So $\mathbb{E}_1(Y_2) = Y_2^D + Y_1^P$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2) = \text{Var}(V_2) + \text{Var}(Y_2^T)$, and $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) = 0$.
- If $|Y_1^P| > c$, $\mathbb{E}_1(Y_2) = Y_2^D + (1 - \delta\tau)Y_1^P$ (“state-dependent persistence”).
- If $|Y_1^P| > c$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2) = \tau(1 - \tau)\delta^2(Y_1^P)^2 + 2\tau\delta\mathbb{E}(V_2|V_2 > b)|Y_1^P| + \text{Var}(V_2) + \text{Var}(Y_2^T)$ (“state-dependent risk”).
- Last, if $Y_1^P < -c$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) > 0$, and if $Y_1^P > c$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) < 0$ (“state-dependent skewness”). For example, if $Y_1^P < -c$, we have

$$\begin{aligned} \mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) = & -\tau(1 - 2\tau)\delta Y_1^P [(1 - \tau)\delta^2(Y_1^P)^2 - 3\delta Y_1^P \mathbb{E}(V_2|V_2 > b) \\ & + 3(\mathbb{E}(V_2^2|V_2 > b) - \mathbb{E}(V_2^2|V_2 \leq b))] > 0. \end{aligned}$$

Discussion. State-dependent persistence implies that low- and high-earnings households respond less to variations in Y_1^P than middle-earnings households. Low-earnings households save less than in the canonical linear model, while high-earnings households save more.

State-dependent risk implies that both low- and high-earnings households save more than in the canonical model because of higher variability of earnings. As shown by (S5), the effect is increasing in risk aversion and higher for low-assets households. Note that the effect is scaled by expected resources R . Compared to the canonical linear earnings

model, this effect will tend to increase savings for high-earnings households and decrease savings for low-earnings households.

Last, *state-dependent skewness* implies that, compared to the canonical model, high-earnings households save more, and low-earnings households save less.

Overall, the comparative statics for high-earnings households are unambiguous, while the combined effect for low-earnings households is ambiguous.¹

S2. SIMULATIONS OF A LIFE-CYCLE MODEL

Here we present an illustrative simulation to show some possible implications on consumption and assets of nonlinearity in income in the standard model outlined in Section 3. We start by comparing the linear canonical earnings model with a simple nonlinear earnings model which features the presence of “unusual” earnings shocks. We then report simulations based on the nonlinear earnings process we estimated on the PSID.

S2.1. Simulation Exercise

To simulate the model, we follow Kaplan and Violante (2010). Each household enters the labor market at age 25, works until 60, and dies with certainty at age 95. After retirement, families receive Social Security transfers Y_i^{ss} from the government, which are functions of the entire realizations of labor income. Income is assumed not to be subject to risk during retirement. Agents’ utility is CRRA with risk aversion $\gamma = 2$. The interest rate is $r = 3\%$ and the discount factor is $\beta = 1/(1+r) \approx .97$. We consider the following process for η_{it} :

$$\eta_{it} = \rho_t(\eta_{i,t-1}, v_{it})\eta_{i,t-1} + v_{it}, \quad (\text{S6})$$

and we compare two specifications. In the first specification, $\rho_t = 1$ (and v_{it} is normally distributed), which corresponds to the “canonical” earnings model used by Kaplan and Violante.² In the second specification, nonlinear persistence in income is approximated through a simple switching process:

$$\begin{aligned} \rho_t(\eta_{i,t-1}, v_{it}) = & 1 - \delta(\mathbf{1}\{\eta_{i,t-1} < -d_{t-1}\}\mathbf{1}\{v_{it} > b_t\} \\ & + \mathbf{1}\{\eta_{i,t-1} > d_{t-1}\}\mathbf{1}\{v_{it} < -b_t\}), \end{aligned} \quad (\text{S7})$$

where, at each age t , d_t is set so that $|\eta_{it}| > d_t$ with probability π , and b_t is set so that $|v_{it}| > b_t$ with probability π . In model (S6)–(S7), the persistence of the η process is equal to 1 unless an “unusual” positive shock v hits a low-income household or an “unusual” negative shock v hits a high-income household, leading persistence to drop to $1 - \delta = 0.8$. The latter happens with probability $\pi = 0.15$ in every period. Details on the simulation are provided below.

The simple parametric process (S7) is designed to roughly approximate the earnings process that we estimate on PSID data; see Section 6. It is worth noting that, since our flexible, quantile-based process is first-order Markov, it is easy to take the estimated η process to simulate, and possibly estimate, life-cycle models of consumption and saving such as the one we focus on here. In Section S2.3 of this Supplemental Material, we show

¹Note that here A_1 is taken as exogenous. In a complete model of the life-cycle, household assets will be different when facing a nonlinear or a linear (“canonical”) earnings process.

²Kaplan and Violante (2010) also considered a more general AR(1) log-earnings process.

the results of using the nonlinear dynamic quantile earnings model we estimated on PSID data as an input in this simple life-cycle simulation model.

The simulation results are presented in Figure S35. In the simulation, we use a natural borrowing limit.³ Graphs (a) and (b) show that a qualitative implication of the nonlinear earnings process is to reduce consumption among those on higher incomes. A negative shock for those on higher incomes reduces the persistence of the past and consequently is more damaging in terms of expected future incomes. This induces higher saving and lower consumption at younger ages. Conversely, we see that consumption is (slightly) higher for the nonlinear process for those on lower income. Graphs (c) and (d) show that the nonlinear model also results in a higher consumption variance among older households, and steeper accumulation and subsequent decumulation of assets over the life-cycle. In addition, in graph (e), we report estimates of the average derivative of the conditional mean of log-consumption with respect to log-earnings, holding assets and age fixed at percentiles indicated on the two horizontal axes. We see that in the simulated economy, consumption responses to changes in earnings tend to decrease with age and, to a lesser extent, the presence of assets. These simulation results provide further motivation for the use of a nonlinear earnings model to study consumption dynamics.

S2.2. Details on the Simulations

Agents live for T periods, and work until age T_{ret} , where both T and T_{ret} are exogenous and fixed. ξ_t is the unconditional probability of surviving to age t , where $\xi_t = 1$ before retirement, and $\xi_t < 1$ after retirement. Households have expected life-time utility

$$\mathbb{E}_0 \sum_{t=1}^T \beta^{t-1} \xi_t u(C_{it}).$$

During working years $1 \leq t < T_{\text{ret}}$, agents receive after-tax labor income Y_{it} , which is decomposed into a deterministic experience profile κ_t , a permanent component η_{it} , and a transitory component ε_{it} :

$$\begin{aligned} \ln Y_{it} &= \kappa_t + y_{it}, \\ y_{it} &= \eta_{it} + \varepsilon_{it}. \end{aligned}$$

We consider the following process for η_{it} :

$$\eta_{it} = \rho_t(\eta_{i,t-1}, v_{it})\eta_{i,t-1} + v_{it},$$

where η_{i0} is drawn from an initial normal distribution with mean zero and variance $\sigma_{\eta_0}^2$. The shocks ε_{it} and v_{it} have mean zero, are normally distributed with variances σ_{ε}^2 and $\sigma_{v_t}^2$. The persistence of the η process is approximated as (S7) where, at each age t , d_t is set so that $|\eta_{it}| > d_t$ with probability π , and b_t is set so that $|v_{it}| > b_t$ with probability π . We set $\pi = 0.15$ and $1 - \delta = 0.8$.

Define gross labor income as \tilde{Y}_{it} , with $\tilde{Y}_{it} = G(Y_{it})$, where G function is the inverse of a tax function estimated by Gouveia and Strauss (1994).⁴ After retirement, agents receive

³Results for a zero borrowing limit are given in Figure S36.

⁴The tax function is

$$\tau(\tilde{Y}_{it}) = 0.258 \times [\tilde{Y}_{it} - (\tilde{Y}_{it}^{-0.768} + \tau^s)^{-\frac{1}{0.768}}],$$

after-tax Social Security transfers Y^{ss} , which are a function of lifetime average individual gross earnings $P(\frac{1}{T_{ret}-1} \sum_{t=1}^{T_{ret}-1} \tilde{Y}_{it})$.

Last, throughout their lifetime, households have access to a single risk-free, one-period bond whose constant return is $1 + r$, and face a period-to-period budget constraint

$$A_{i,t+1} = (1 + r)A_{it} + Y_{it} - C_{it}, \quad \text{if } t < T_{ret},$$

$$\left(\frac{\xi_t}{\xi_{t+1}}\right)A_{i,t+1} = (1 + r)A_{it} + Y^{ss} - C_{it}, \quad \text{if } t \geq T_{ret}.$$

For the calibration, we use Kaplan and Violante's (2010, KV hereafter) preferred parameters. Specifically we use the following:

Demographics. The model period is one year. Agents enter the labor market at age 25, retire at age 60, and die with certainty at age 95. So we set $T_{ret} = 35$ and $T = 70$. The survival rates ξ_t are obtained from the National Center for Health Statistics (1992).

Preferences. We assume the utility function is of the CRRA form $u(C) = C^{1-\gamma}/(1-\gamma)$, where the risk aversion parameter is set to $\gamma = 2$.

Discount Factor and Interest Rate. The interest rate r is assumed to equal 3%. We set the discount factor β to match an aggregate wealth-income ratio of 2.5, which is the average wealth to average income ratio computed from the 1989 and 1992 Survey of Consumer Finances.

Income Process. We use KV's choice of deterministic age profile κ_t , which is estimated from the PSID. The estimated profile peaks after 21 years of labor market experience at roughly twice the initial value, and then it slowly declines to about 80% of the peak value. For the stochastic components of the income process, we set the initial variance of the permanent shocks $\sigma_{\eta_0}^2 = 0.15$ to match the dispersion of household earnings at age 25. We set the variance of transitory shocks $\sigma_{\epsilon}^2 = 0.05$, which is the value estimated by Blundell, Pistaferri, and Preston (2008). The permanent component follows the nonlinear switching process (S7). In order to make its variance comparable with the random-walk process in KV, we calibrate $\sigma_{v_t}^2$ to match the dispersion of age-specific variance of η_{it} in KV.

Initial Wealth and Borrowing Limit. Households' initial assets are set to 0. We impose two alternative borrowing limits: either a natural borrowing limit, in which the agent cannot die with debt, or a zero borrowing limit, in which the agent's net worth cannot fall below zero.

Social Security Benefits. This setup follows KV exactly. Social Security benefits are a function of lifetime average individual gross earnings $Y^{ss} = P(\frac{1}{T_{ret}-1} \sum_{t=1}^{T_{ret}-1} \tilde{Y}_{it})$. The P function is designed to mimic the actual U.S. system. This is achieved by specifying that benefits are equal to 90% of average past earnings up to a given bend point, 32% from this first bend point to a second bend point, and 15% beyond that. The two bend points are set at, respectively, 0.18 and 1.10 times cross-sectional average gross earnings, based on the U.S. legislation and individual earnings data for 1990. Benefits are then scaled proportionately so that a worker earning average labor income each year is entitled to a replacement rate of 45%.

Discretization of the Earnings Process. In order to use the switching process (S7) in the life-cycle model, we compute age-specific Markov transition probabilities using a simulation approach, as follows:

where τ^s is chosen so that the ratio of total personal current tax receipts on labor income (not including Social Security contributions) to total labor income is the same as for the U.S. economy in 1990, that is, roughly 25%.

1. Based on the process (S7), we simulate one million households throughout their working years.
2. At each age, we rank households by their earnings and put them into 40 different bins. We then count the transitions between any two bins from two neighboring ages, and estimate transition probabilities.

S2.3. Simulations Based on the Estimated Nonlinear Earnings Model

In this subsection, we report the results of a simulation exercise closely related to the one shown in the first part of this section, except that it is based on the nonlinear quantile-based earnings process that we estimated on the PSID. Calibration and simulation details are similar to the ones above, with a few differences. In order to ensure comparability with the data, we let each household enter the model at age 25, work until 61, and die at age 93. As the PSID is biennial, we set each period in the model to be equal to two years. There is one risk-free asset with a constant interest rate $r = 1.03^2 - 1 \approx 6\%$, and the discount factor is set to $\beta = (1 + r)^{-1}$. We impose a natural borrowing limit in the simulation.

We simulate life-cycle profiles for one million households. When discretizing persistent and transitory earnings components, we use 100 bins for the former and 80 bins for the latter. We checked that the discretized process fits the nonlinear persistence in Figure 1 well. We remove age-specific means in the persistent component (which may be different from zero for the particular cohort for which we are performing the simulation). In order to ensure comparability with a canonical linear earnings process, in the nonlinear process we compute age-specific variances of transitory shocks and target the age profile of the variance of the persistent component, and we set the parameters of the linear process (i.e., a random walk plus an independent shock) using the resulting values.

The results of the simulation are presented in Figure S37. As in Figure S35, we see that households on higher incomes tend to consume less when exposed to the nonlinear process than to the linear one. Conversely, consumption is slightly higher for the nonlinear process for those on lower income. Consumption and asset variance are lower for the nonlinear process, similarly as in Figure S35, although the differences between the two processes are stronger in Figure S37, which is based on the process estimated on PSID.⁵ Last, the marginal propensity to consume out of earnings in panel (e) tends to decrease with age and the presence of assets, similarly to what we found in the PSID.

S3. EXTENSIONS

Here we consider four extensions of the model: to allow for unobserved heterogeneity in earnings, dependent ε , advance earnings information, and consumption habits, respectively. The estimation strategy can be modified to handle each of these extensions.

S3.1. Unobserved Heterogeneity in Earnings

It is possible to allow for unobserved heterogeneity in earnings, in addition to heterogeneity in the initial condition η_{it} . Specifically, let η_{it} be a first-order Markov process

⁵Differences in scale with Figure S35 are due to the different period considered in this paper relative to Kaplan and Violante (2010) and different parameter choices. Also, note that consumption data in the PSID waves which we use only capture a share of consumption expenditures, and similarly for assets. This can explain the differences between the levels in Figure S37 and the descriptive statistics reported in Table C.I.

conditional on another latent component ζ_i :

$$\eta_{it} = Q_t(\eta_{i,t-1}, \zeta_i, u_{it}), \quad (\text{S8})$$

where u_{it} is i.i.d. standard uniform, independent of η_i^{t-1} and ζ_i . ε_{it} is independent over time, independent of η_{is} for all s , and independent of ζ_i .

With a vector-valued ζ_i , (S8) would nest linear earnings models with slope heterogeneity as in Guvenen (2007) and Guvenen and Smith (2014), for example. A simpler case is our baseline model (1)–(2) augmented with a household-specific fixed-effect, that is, equation (20) in the paper.

Consider the scalar- ζ_i case for concreteness, and take $T = 5$. In this model, (y_{i1}, y_{i2}) , y_{i3} , and (y_{i4}, y_{i5}) are conditionally independent given (η_{i3}, ζ_i) .⁶ By Hu and Schennach’s (2008) theorem, for bivariate latent (η_{i3}, ζ_i) , and under suitable injectivity conditions, the marginal distribution of ε_{i3} is thus identified given five periods of earnings data. As a result, the joint density of η ’s is identified by a similar argument as in Section 4. Identification of the densities of ζ_i and of η_{it} given $(\eta_{i,t-1}, \zeta_i)$ can then be shown along the lines of Hu and Shum (2012), under a suitable “scaling” condition. A scaling condition is implicit in equation (20), which is the model we implement.

S3.2. Dependence in the Transitory Earnings Component

In the baseline model, ε_{it} are independent over time. It is possible to allow for serial dependence while maintaining identification. To see this, consider the setup where ε_{it} is an m -dependent process with $m = 1$ (e.g., an $MA(1)$ process), and consider a panel with $T \geq 5$ periods. Then it is easy to see that y_{i1} , y_{i3} , and y_{i5} are conditionally independent given η_{i3} . As a result, identification arguments based on “Hidden Markov” structures (Hu and Schennach (2008), Wilhelm (2015)) can be applied.

S3.3. Advance Information

If households have advance information about future earnings shocks, the consumption rule (10) takes future earnings components as additional arguments; see Blundell, Pistaferri, and Preston (2008). For example, consider a model where households know the realization of the one-period-ahead persistent component, in which case

$$c_{it} = g_t(a_{it}, \eta_{it}, \eta_{i,t+1}, \varepsilon_{it}, \nu_{it}), \quad t = 1, \dots, T - 1. \quad (\text{S9})$$

Identification can be established using similar arguments as in the baseline model. To see this, consider first period’s consumption. We have

$$f(c_1|a_1, y) = \int \int f(c_1|a_1, \eta_1, \eta_2, y_1) f(\eta_1, \eta_2|a_1, y) d\eta_1 d\eta_2.$$

It can be shown that $f(\eta_1, \eta_2|a_1, y)$ is identified under completeness in (y_{i2}, \dots, y_{iT}) of the distribution of $(\eta_{i1}, \eta_{i2}|y_i)$, using the earnings process and first period’s assets.

⁶Indeed,

$$\begin{aligned} f(y_1, y_2, y_3, y_4, y_5|\eta_3, \zeta) &= f(y_1, y_2|\eta_3, \zeta) f(y_3|\eta_3, \zeta, y_1, y_2) f(y_4, y_5|\eta_3, \zeta, y_3, y_2, y_1) \\ &= f(y_1, y_2|\eta_3, \zeta) f(y_3|\eta_3, \zeta) f(y_4, y_5|\eta_3, \zeta). \end{aligned}$$

If the distribution of $(\eta_{i1}, \eta_{i2}|a_{i1}, y_i)$ is complete in (y_{i2}, \dots, y_{iT}) , it thus follows that $f(c_1|a_1, \eta_1, \eta_2, y_1)$ is identified. In this case, we need at least two “excluded instruments” in y_i for (η_{i1}, η_{i2}) . The other steps in the identification arguments of Section 4 can be similarly adapted.

Last, similar arguments can be used to show identification in models where households have advance information about future transitory shocks $\varepsilon_{i,t+s}$, as well as in models where the consumption rule depends on lags of η 's or ε 's, for example in models where η_{it} follows a higher-order Markov process.

S3.4. Consumption Habits

In the presence of habits, the consumption rule takes the form

$$c_{it} = g_t(c_{i,t-1}, a_{it}, \eta_{it}, \varepsilon_{it}, v_{it}), \quad t = 2, \dots, T. \quad (\text{S10})$$

Identification can be shown under similar conditions as in Section 4. For example, in the second period, equation (18) becomes

$$f(c_2|c_1, a_2, a_1, y) = \int f(c_2|c_1, a_2, \eta_2, y_2) f(\eta_2|c_1, a_2, a_1, y) d\eta_2.$$

Provided the distribution of $(\eta_{i2}|c_{i1}, a_{i2}, a_{i1}, y_i)$ is identified, and complete in $(a_{i1}, y_{i1}, y_{i3}, \dots, y_{iT})$, it thus follows that the density $f(c_2|c_1, a_2, \eta_2, y_2)$ is identified. Intuitively, in the presence of habits, the first lag of consumption cannot be used as an “excluded instrument” since it affects second period’s consumption directly.

S4. SUMMARY OF THE ARGUMENT IN WILHELM (2015)

We consider model (1)–(2) with $T = 3$. We omit i subscripts for conciseness. Let $L^2(f)$ denote the set of squared-integrable functions with respect to a weight function f . We define $\mathcal{L}_{y_2|y_1}$ as the linear operator such that $\mathcal{L}_{y_2|y_1} h(a) = \mathbb{E}[h(y_2)|y_1 = a] \in L^2(f_{y_1})$ for every function $h \in L^2(f_{y_2})$. Similarly, let $\mathcal{L}_{\eta_2|y_1}$ be such that $\mathcal{L}_{\eta_2|y_1} h(a) = \mathbb{E}[h(\eta_2)|y_1 = a] \in L^2(f_{y_1})$ for every function $h \in L^2(f_{\eta_2})$. We denote as $\mathcal{R}(\mathcal{L}_{y_2|y_1})$ the range of $\mathcal{L}_{y_2|y_1}$, that is,

$$\mathcal{R}(\mathcal{L}_{y_2|y_1}) = \{k \in L^2(f_{y_1}), \text{ s.t. } k = \mathcal{L}_{y_2|y_1} h \text{ for some } h \in L^2(f_{y_2})\}.$$

We assume the following, in addition to the Markovian and independence assumptions on η 's and ε 's.

ASSUMPTION S1:

- (i) $\mathcal{L}_{y_2|y_1}$ and $\mathcal{L}_{\eta_2|y_1}$ are injective.
- (ii) There exists a function $h \in L^2(f_{y_3})$ such that

$$\mathbb{E}[h(y_3)|y_1 = \cdot] \in \mathcal{R}(\mathcal{L}_{y_2|y_1}), \quad \text{and} \quad (\text{S11})$$

$$\mathbb{E}[y_2 h(y_3)|y_1 = \cdot] \in \mathcal{R}(\mathcal{L}_{y_2|y_1}). \quad (\text{S12})$$

Thus, there exist s_1 and s_2 in $L^2(f_{y_2})$ such that

$$\mathbb{E}[h(y_3)|y_1 = \cdot] = \mathcal{L}_{y_2|y_1} s_1, \quad \text{and} \quad \mathbb{E}[y_2 h(y_3)|y_1 = \cdot] = \mathcal{L}_{y_2|y_1} s_2.$$

(iii) Let $\tilde{s}_1(y) = ys_1(y)$. The Fourier transforms $\mathcal{F}(s_1)$, $\mathcal{F}(\tilde{s}_1)$, and $\mathcal{F}(s_2)$ (where $\mathcal{F}(h)(u) = \int h(x)e^{iux} dx$) are ordinary functions. Moreover, $\mathcal{F}(s_1)(u) \neq 0$ for all $u \in \mathbb{R}$.

Part (i) is an injectivity/completeness condition. Part (ii) is not standard. It is related to the existence problem in nonparametric instrumental variables. Horowitz (2012) proposed a test for (S11) in the case where $\mathcal{L}_{y_2|y_1}$ is a compact operator. Part (iii) is a high-level assumption. See Wilhelm (2015) for more primitive conditions.

By Assumption S1(ii) we have, almost surely in y_1 ,

$$\begin{aligned}\mathbb{E}[h(y_3)|y_1] &= \mathbb{E}[s_1(y_2)|y_1], \\ \mathbb{E}[y_2h(y_3)|y_1] &= \mathbb{E}[s_2(y_2)|y_1].\end{aligned}$$

Moreover, s_1 and s_2 are the unique solutions to these equations by Assumption S1(i).

Hence, given the model's assumptions,

$$\mathbb{E}[\mathbb{E}(h(y_3)|\eta_2)|y_1] = \mathbb{E}[\mathbb{E}(s_1(y_2)|\eta_2)|y_1] \quad \text{a.s.}$$

It thus follows from the injectivity of $\mathcal{L}_{\eta_2|y_1}$ in Assumption S1(i) that, almost surely in η_2 ,

$$\mathbb{E}[h(y_3)|\eta_2] = \mathbb{E}[s_1(y_2)|\eta_2]. \quad (\text{S13})$$

Likewise, $\mathbb{E}[y_2h(y_3)|\eta_2] = \mathbb{E}[s_2(y_2)|\eta_2]$. Hence

$$\eta_2 \mathbb{E}[h(y_3)|\eta_2] = \mathbb{E}[s_2(y_2)|\eta_2] \quad \text{a.s.} \quad (\text{S14})$$

Combining (S13) and (S14), we obtain

$$\eta_2 \mathbb{E}[s_1(y_2)|\eta_2] = \mathbb{E}[s_2(y_2)|\eta_2] \quad \text{a.s.}$$

That is, almost surely in η_2 ,

$$\eta_2 \int s_1(y)f_{\varepsilon_2}(y - \eta_2) dy = \int s_2(y)f_{\varepsilon_2}(y - \eta_2) dy. \quad (\text{S15})$$

The functional equation (S15) depends on s_1 and s_2 , which are both uniquely determined given the data generating process, and on the unknown f_{ε_2} . By Assumption S1(iii), we can take Fourier transforms and obtain

$$i\mathcal{F}(s_1)(u) \frac{d\psi_{\varepsilon_2}(-u)}{du} + \mathcal{F}(\tilde{s}_1)(u)\psi_{\varepsilon_2}(-u) = \mathcal{F}(s_2)(u)\psi_{\varepsilon_2}(-u), \quad (\text{S16})$$

where $\psi_{\varepsilon_2}(u) = \mathcal{F}(f_{\varepsilon_2})(u)$ is the characteristic function of ε_2 .

Noting that $\psi_{\varepsilon_2}(0) = 1$, (S16) can be solved in closed form for $\psi_{\varepsilon_2}(\cdot)$, because $\mathcal{F}(s_1)(u) \neq 0$ for all u by Assumption S1(iii). This shows that the characteristic function of ε_2 , and hence its distribution function, are identified.

S5. ADDITIONAL FIGURES

(a) PSID

(b) Norwegian administrative data





FIGURE S1.—Nonlinear earnings persistence in log-wages (PSID, males) and individual income (Norwegian population register data). *Note:* Estimates of the average derivative of the conditional quantile function of y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$, evaluated at percentile τ_{shock} and at a value of $y_{i,t-1}$ that corresponds to the τ_{init} percentile of the distribution of $y_{i,t-1}$. Age 25–60. Left: PSID, male log hourly wage residuals, 1999–2009; right: results from the Norwegian population register data, individual log-earnings residuals, years 2005–2006, see Appendix C; are provided as part of the project on ‘Labour Income Dynamics and the Insurance from Taxes, Transfers and the Family’.

(a) Three sub-intervals

(b) Five sub-intervals





FIGURE S2.—Nonlinear persistence, piecewise-linear specification quantile regression. *Note:* See the notes to Figure 2. Piecewise-linear quantile regression with three and five equally-spaced sub-intervals.


(a) Persistent component η_{it} , nonlinear model

FIGURE S3.—Nonlinear persistence of η_{it} , averaged over ages. *Note:* Estimates of the average derivative of the conditional quantile function of η_{it} on $\eta_{i,t-1}$ with respect to $\eta_{i,t-1}$, based on estimates from the nonlinear earnings model, averaged over all ages of household heads.



(a) Log-earnings residuals y_{it} (b) Persistent component η_{it} 

FIGURE S4.—Conditional dispersion of log-earnings residuals (fit) and η component. *Note:* Conditional dispersion $\sigma(y, \tau) = Q(\tau|y_{i,t-1} = y) - Q(1 - \tau|y_{i,t-1} = y)$ and $\sigma(\eta, \tau)$, for $\tau = 11/12$. Log-earnings residuals (left) and η component (right). On the left graph, dark is data and light is model fit. The x -axis shows the conditioning variable, the y -axis shows the corresponding value of the conditional dispersion measure.


(a) Log-earnings residuals y_{it} 

FIGURE S5.—Conditional skewness of log-earnings residuals, fit. *Note:* Conditional skewness $sk(y, \tau)$, see equation (6), for $\tau = 11/12$. Log-earnings residuals. The x -axis shows the conditioning variable, the y -axis shows the corresponding value of the conditional skewness measure. Dark is PSID data, light is nonlinear model.


(a) Skewness
Nonlinear model

FIGURE S6.—Skewness of log-earnings residuals growth at various horizons, fit. *Note:* Robust skewness estimate of $y_{it} - y_{i,t-s}$, at various horizons s (reported on the x -axis), computed according to the formula $sk(\tau, \alpha) = \frac{Q(\tau) + Q(1-\tau) - 2Q(1/2)}{Q(\tau) - Q(1-\tau)}$, where Q denote unconditional quantiles, and $\tau = 11/12$. Dark is PSID data, light is nonlinear model.


(a) Kurtosis
Nonlinear model

FIGURE S7.—Kurtosis of log-earnings residuals growth at various horizons, fit. *Note:* Robust kurtosis estimate of $y_{it} - y_{i,t-s}$, at various horizons s (reported on the x -axis), computed according to the formula $kur(\tau, \alpha) = \frac{Q(1-\alpha) - Q(\alpha)}{Q(\tau) - Q(1-\tau)}$, where Q denote unconditional quantiles, $\tau = 10/12$, and $\alpha = 1/12$. Dark is PSIS data, light is nonlinear model, and the horizontal line denotes the values of $kur(\tau, \alpha)$ for a normal distribution.

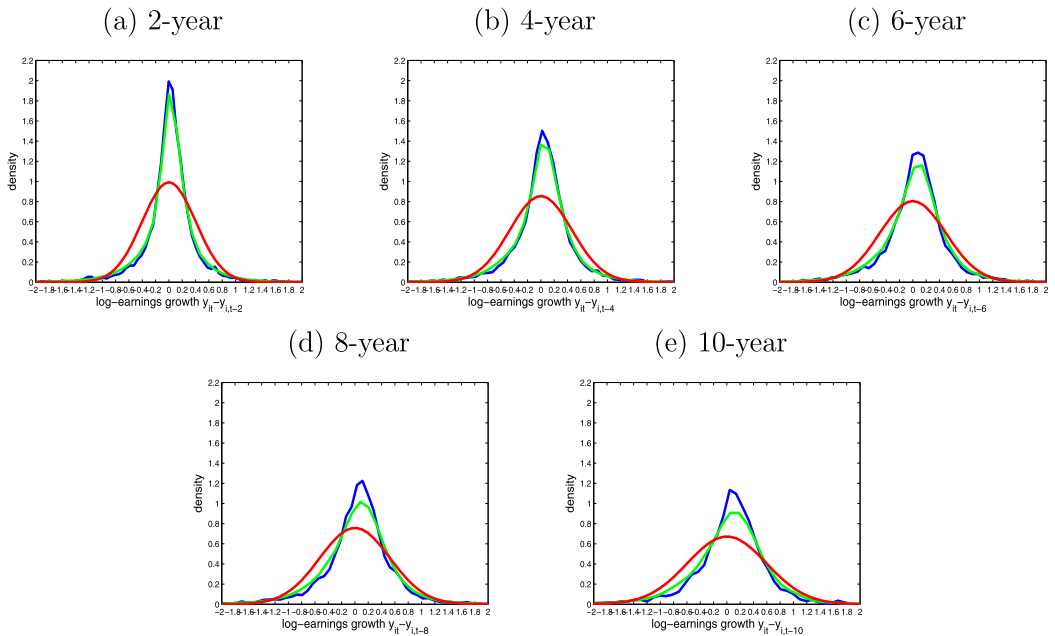




FIGURE S8.—Densities of log-earnings growth at various horizons. *Note:* Estimated density of $y_{it} - y_{i,t-s}$, at various horizons s . Dark is PSIS data, light is nonlinear model. Added to each graph is the Gaussian density with zero mean and the same variance as in the data.

(a) PSID data

(b) Nonlinear model

Parametric bootstrap



Nonparametric bootstrap





FIGURE S9.—Nonlinear persistence in earnings, 95% pointwise confidence bands. *Note:* See notes to Figure 2. Pointwise 95% confidence bands. 500 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

(a) PSID data

(b) Nonlinear model

Nonparametric bootstrap





FIGURE S10.—Nonlinear persistence in earnings, 95% uniform confidence bands. *Note:* See notes to Figure 2. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 500 replications.

(a) Parametric bootstrap

(b) Nonparametric bootstrap





FIGURE S11.—Nonlinear persistence in η_{it} , 95% pointwise confidence bands. *Note:* See notes to Figure 2. Pointwise 95% confidence bands. 500 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

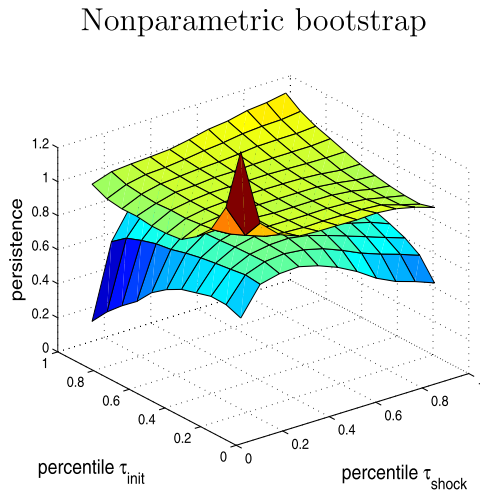




FIGURE S12.—Nonlinear persistence in η_{it} , 95% uniform confidence bands. *Note:* See notes to Figure 2. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 500 replications.

(a) Log-earnings residuals y_{it} (b) Persistent component η_{it}

Parametric bootstrap



Nonparametric bootstrap





FIGURE S13.—Conditional skewness of log-earnings residuals and η component, 95% pointwise confidence bands. *Note:* See notes to Figure 4. Pointwise 95% confidence bands. 500 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

(a) Log-earnings residuals y_{it} (b) Persistent component η_{it}

Nonparametric bootstrap




FIGURE S14.—Conditional skewness of log-earnings residuals and η component, 95% uniform confidence bands. *Note:* See notes to Figure 2. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 500 replications.

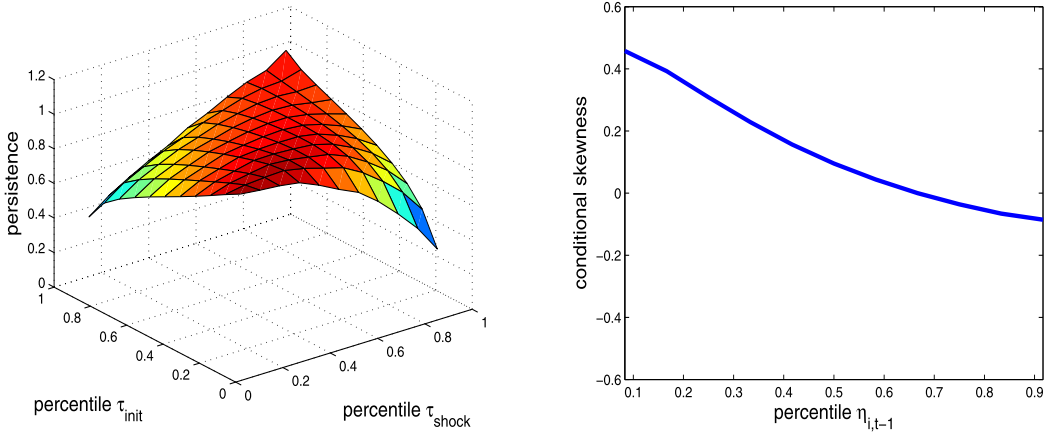
(a) Nonlinear persistence of η_{it} (b) Conditional skewness of η_{it} 

FIGURE S15.—Household heterogeneity in earnings. *Note:* (a) Estimates of the average derivative of the conditional quantile function of η_{it} on $\eta_{i,t-1}$ with respect to $\eta_{i,t-1}$, based on estimates from the nonlinear earnings model with an additive household-specific effect. (b) Conditional skewness $sk(\eta, \tau)$, see equation (6), for $\tau = 11/12$, based on the same model.





FIGURE S16.—Conditional skewness, Norwegian administrative data. *Note:* Conditional skewness of log-earnings measured as in (6) for $\tau = 1/10$. Age 25–60, years 2005–2006.

(a) Earnings, Norwegian data



(b) Earnings, nonlinear model



(c) Persistent component η_{it} , nonlinear model

FIGURE S17.—Nonlinear persistence, Norwegian data. *Note:* See the notes to Figure 2. Random subsample of 2873 households, from 2000 to 2005 Norwegian administrative data, non-immigrant residents, age 30 to 60.

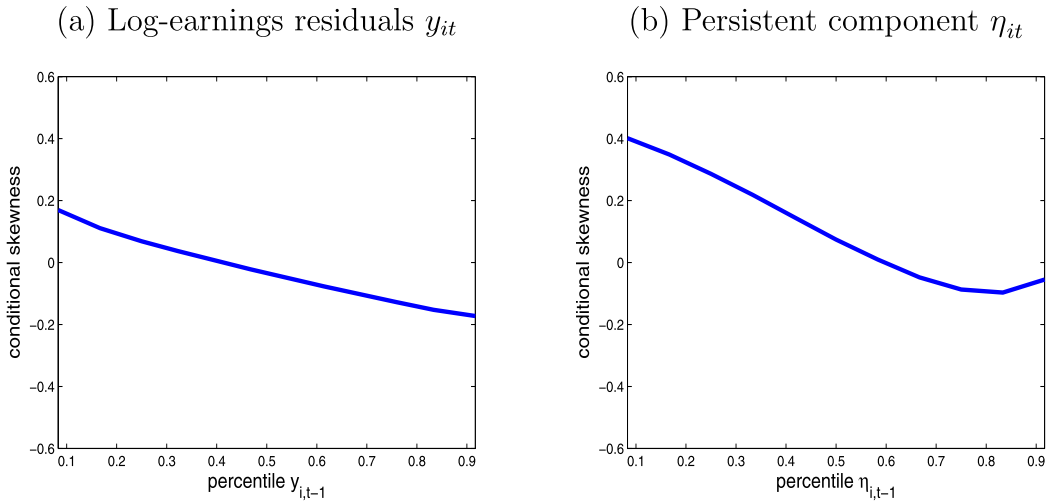


FIGURE S18.—Conditional skewness of log-earnings residuals and η component, Norwegian data. *Note:* See the notes to Figure 4. Random subsample of 2873 households, from 2000 to 2005 Norwegian administrative data, non-immigrant residents, age 30 to 60.

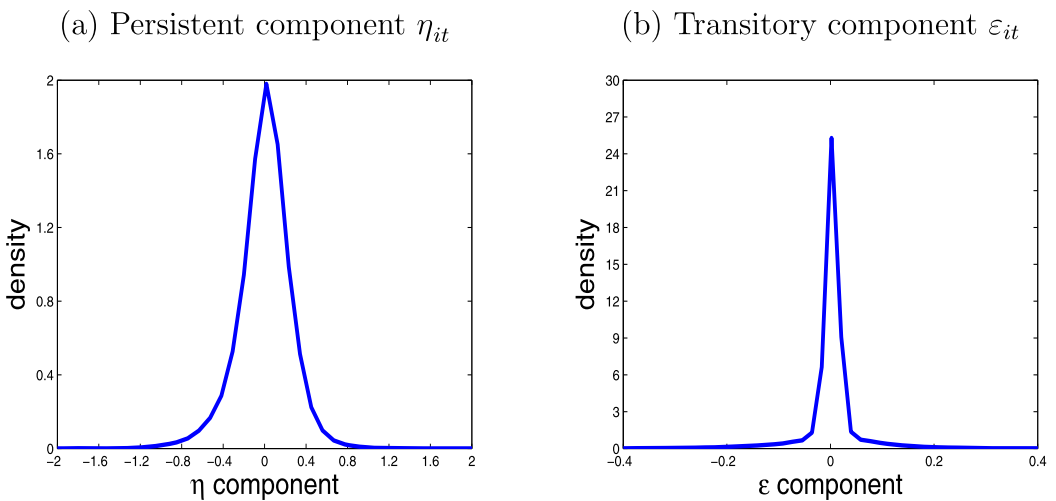


FIGURE S19.—Densities of persistent and transitory earnings components, Norwegian data. *Note:* See the notes to Figure 3. Random subsample of 2873 households, from 2000 to 2005 Norwegian administrative data, non-immigrant residents, age 30 to 60.

Density of log-consumption




FIGURE S20.—Density of log-consumption in the data and model, with and without unobserved heterogeneity. *Note:* Source PSID and nonlinear consumption model. Kernel density estimates. Dashed is Gaussian fit.

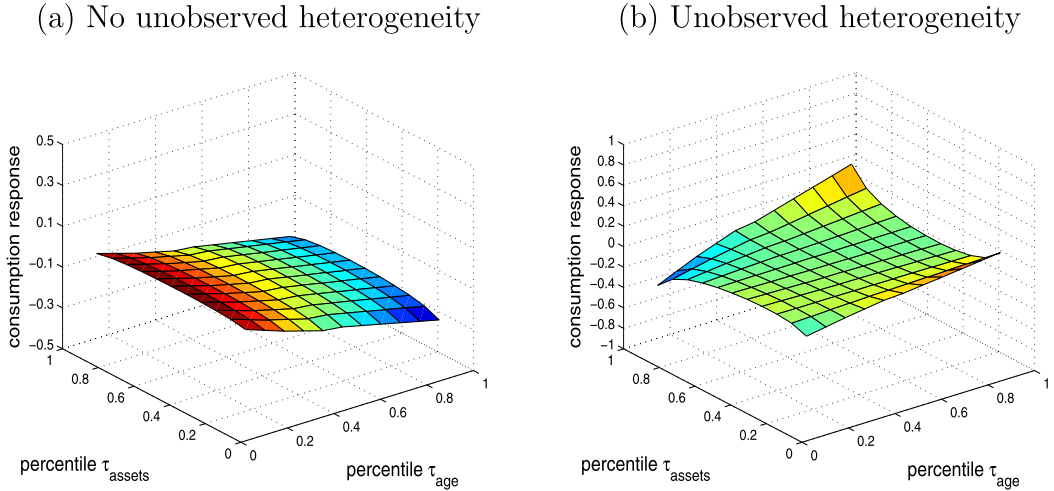




FIGURE S21.—Consumption responses to transitory earnings shocks, by assets and age. *Note:* Estimates of the average consumption responses to variations in ε_{it} , evaluated at τ_{assets} and τ_{age} .

(a) Consumption response to η_{it}
Nonlinear model

(b) Consumption response to ε_{it}
Nonlinear model

Parametric bootstrap



Nonparametric bootstrap





FIGURE S22.—Consumption responses to earnings shocks, by assets and age, 95% pointwise confidence bands. *Note:* See notes to Figures 5 and S21. Pointwise 95% confidence bands. 200 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

(a) Consumption response to η_{it}
Nonlinear model

(b) Consumption response to ε_{it}
Nonlinear model

Nonparametric bootstrap





FIGURE S23.—Consumption responses to earnings shocks, by assets and age, 95% uniform confidence bands. *Note:* See notes to Figures 5 and S21. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 200 replications.

(b) Response to earnings
Nonlinear model

(c) Response to η_{it}
Nonlinear model






FIGURE S24.—Consumption responses to earnings shocks, by assets and age, model with household-specific unobserved heterogeneity. *Note:* See the notes to Figure 5. Consumption and assets model with household-specific unobserved heterogeneity.

(b) Given y_{it}
Nonlinear model



(c) Given η_{it}
Nonlinear model





FIGURE S25.—Consumption responses to assets, model with unobserved heterogeneity. *Note:* See the notes to Figure 6. Model with unobserved heterogeneity.


Nonlinear model

$\tau_{init} = .1$

(a) $\tau_{shock} = .1$




(b) $\tau_{shock} = .9$




$\tau_{init} = .5$

(c) $\tau_{shock} = .1$




(d) $\tau_{shock} = .9$



$\tau_{init} = .9$

(e) $\tau_{shock} = .1$



(f) $\tau_{shock} = .9$





FIGURE S26.—Impulse responses, earnings, 95% pointwise confidence bands. *Note:* See notes to Figure 7. Point estimates and pointwise 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.


Nonlinear model

$$\tau_{init} = .1$$

(a) $\tau_{shock} = .1$




(b) $\tau_{shock} = .9$




$$\tau_{init} = .5$$

(c) $\tau_{shock} = .1$




(d) $\tau_{shock} = .9$



$$\tau_{init} = .9$$

(e) $\tau_{shock} = .1$



(f) $\tau_{shock} = .9$




FIGURE S27.—Impulse responses, earnings, 95% uniform confidence bands. *Note:* See notes to Figure 7. Point estimates and uniform 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.

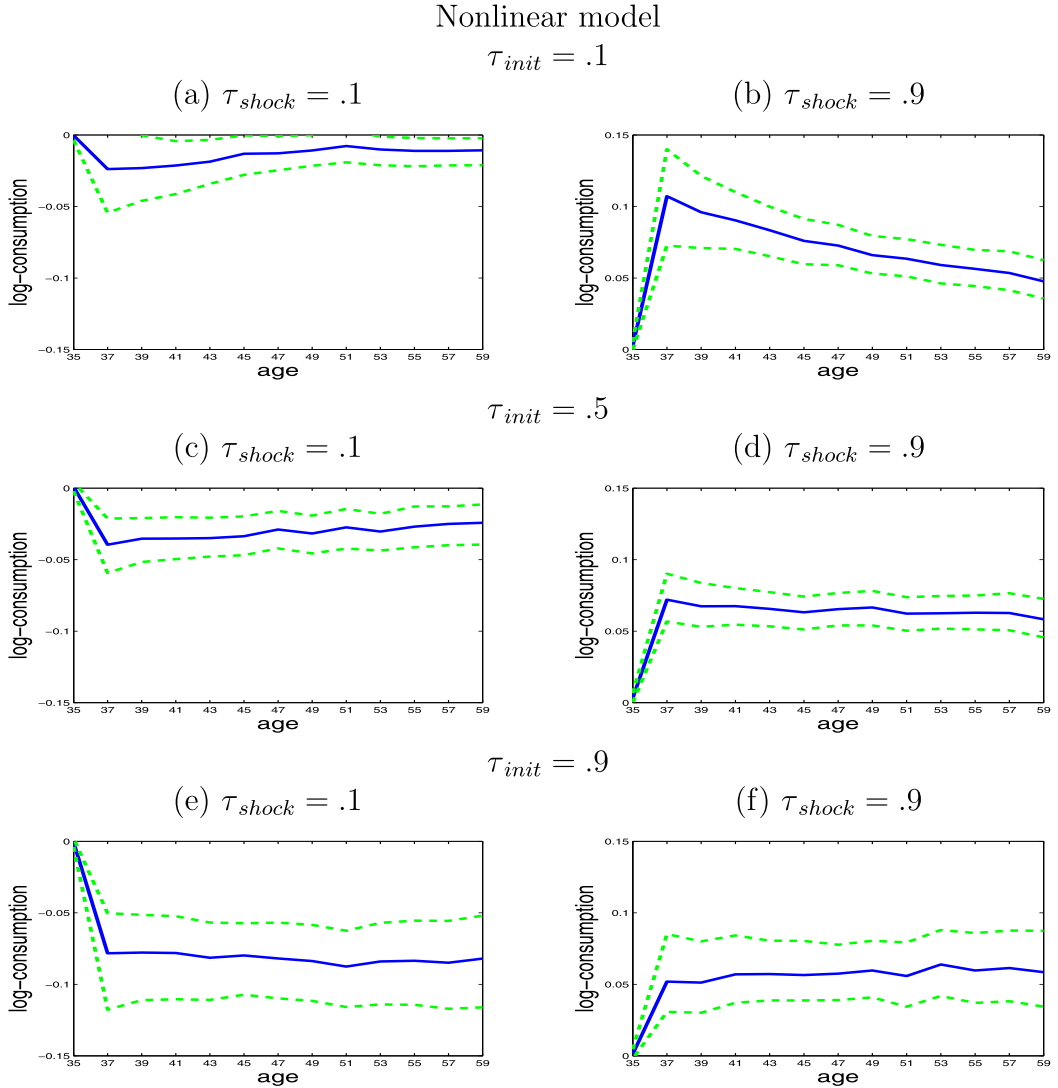


FIGURE S28.—Impulse responses, consumption, 95% pointwise confidence bands. *Note:* See notes to Figure 8. Linear assets accumulation rule (7), $r = 3\%$. $a_{it} \geq 0$. Point estimates and pointwise 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.

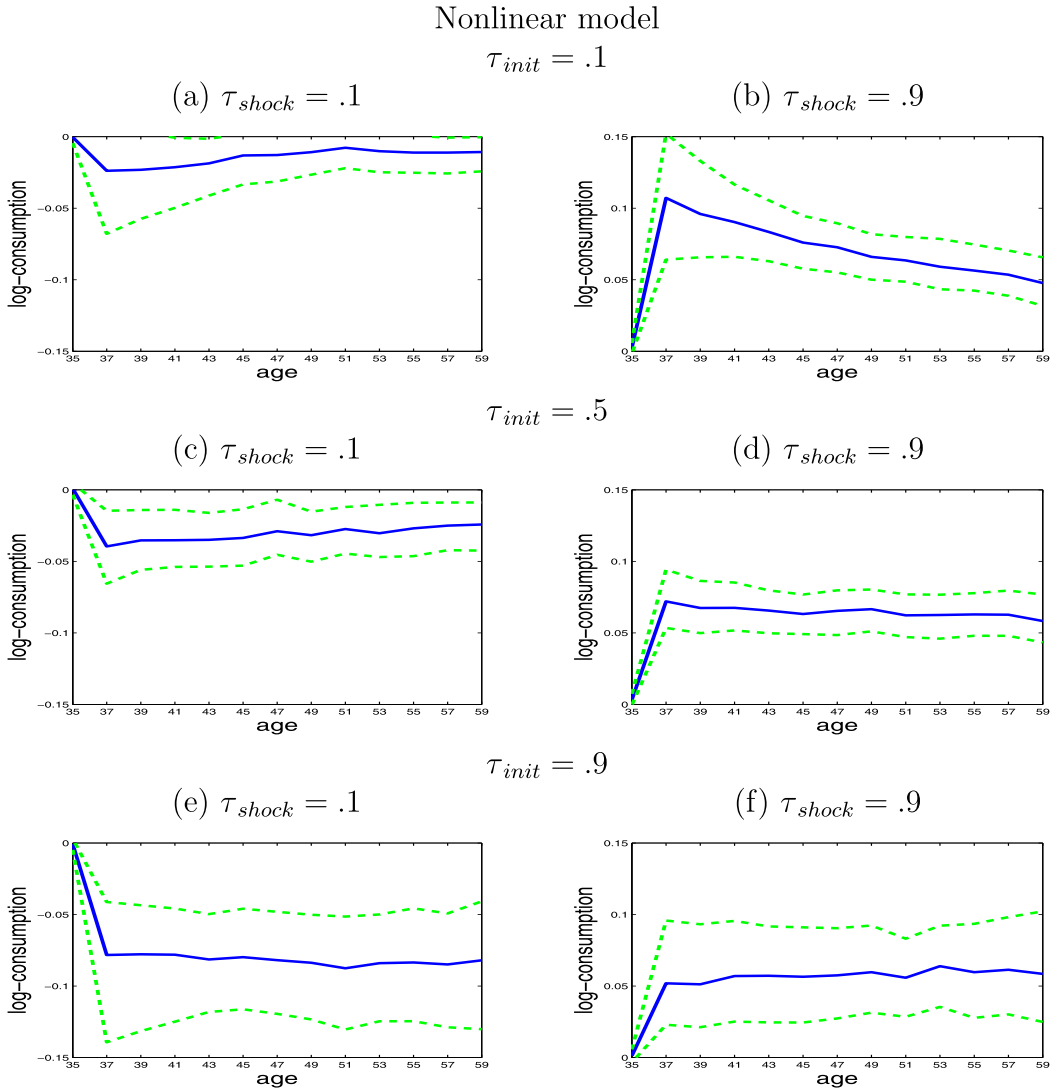


FIGURE S29.—Impulse responses, consumption, 95% uniform confidence bands. *Note:* See notes to Figure 8. Linear assets accumulation rule (7), $r = 3\%$. $a_{it} \geq 0$. Point estimates and uniform 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.

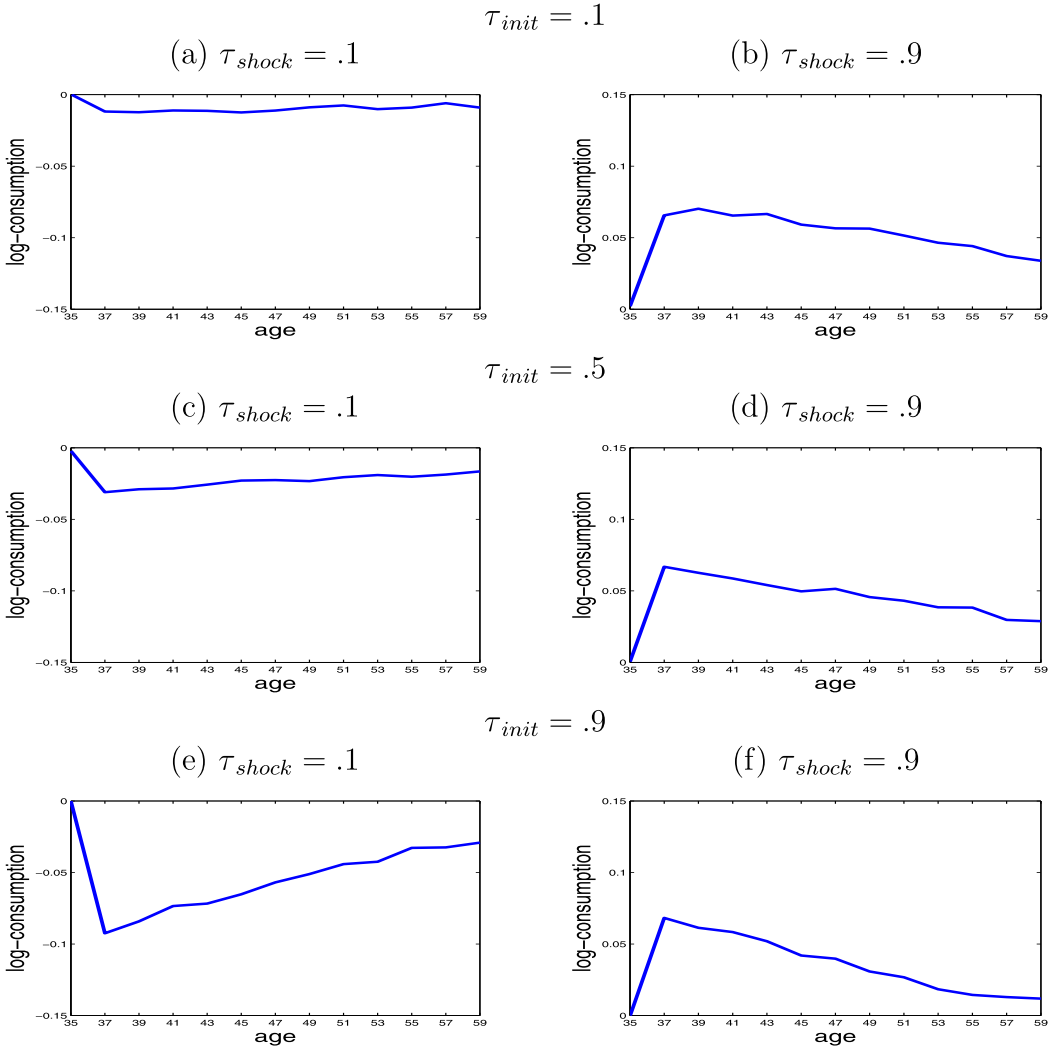






FIGURE S30.—Impulse responses, consumption, model with household unobserved heterogeneity in consumption. *Note:* See notes to Figure 8. Nonlinear model with household unobserved heterogeneity in consumption. Linear assets accumulation rule (7), $r = 3\%$. $a_{it} \geq 0$.

Nonlinear model

$$\tau_{init} = .1$$

(a) $\tau_{shock} = .1$ (b) $\tau_{shock} = .9$ 

$$\tau_{init} = .5$$

(c) $\tau_{shock} = .1$ (d) $\tau_{shock} = .9$ 

$$\tau_{init} = .9$$



(e) $\tau_{shock} = .1$ (f) $\tau_{shock} = .9$ 

FIGURE S31.—Impulse responses, consumption, model without household unobserved heterogeneity in consumption, nonlinear assets rule. *Note:* See the notes to Figure 8. Estimated nonlinear assets rule.

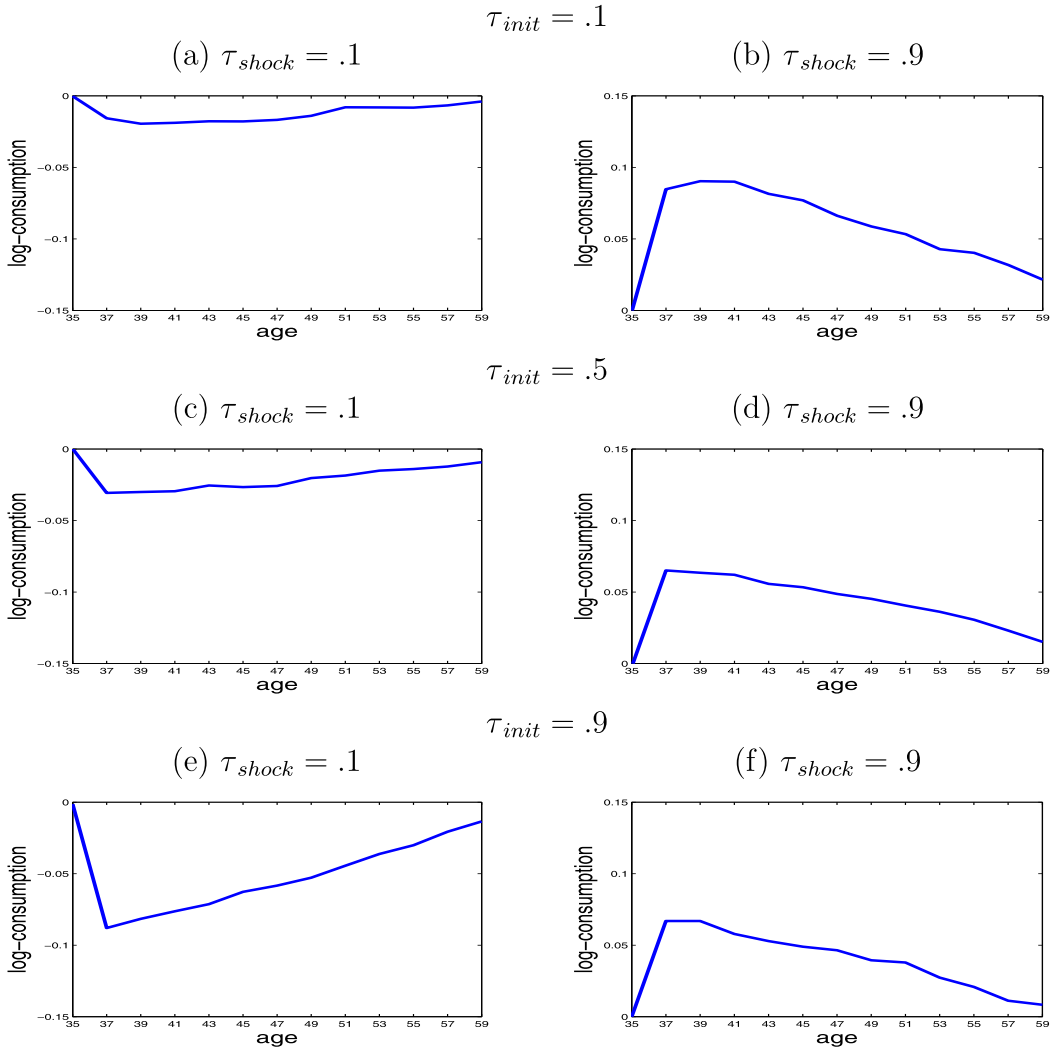


FIGURE S32.—Impulse responses, consumption, model with household unobserved heterogeneity in consumption, nonlinear assets rule. *Note:* See notes to Figure S31. Nonlinear model with household unobserved heterogeneity in consumption. Estimated nonlinear assets rule.

Consumption, model with unobserved heterogeneity





FIGURE S33.—Impulse responses by age and initial assets. *Note:* See notes to Figures 8 and S30. Initial assets at age 35 (for “young” households) or 51 (for “old” households) are at percentile 0.10 (dashed curves) and 0.90 (solid curves). Linear assets accumulation rule (7), $r = 3\%$. $a_{it} \geq 0$. In the simulation of the model with unobserved heterogeneity, ξ_i is set to zero.

Consumption, model without unobserved heterogeneity



Consumption, model with unobserved heterogeneity




FIGURE S34.—Impulse responses by age and initial assets, nonlinear assets rule. *Note:* See notes to Figure S31. Initial assets at age 35 (for “young” households) or 51 (for “old” households) are at percentile 0.10 (dashed curves) and 0.90 (solid curves). Estimated nonlinear assets rule. In the simulation of the model with unobserved heterogeneity, ξ_i is set to zero.

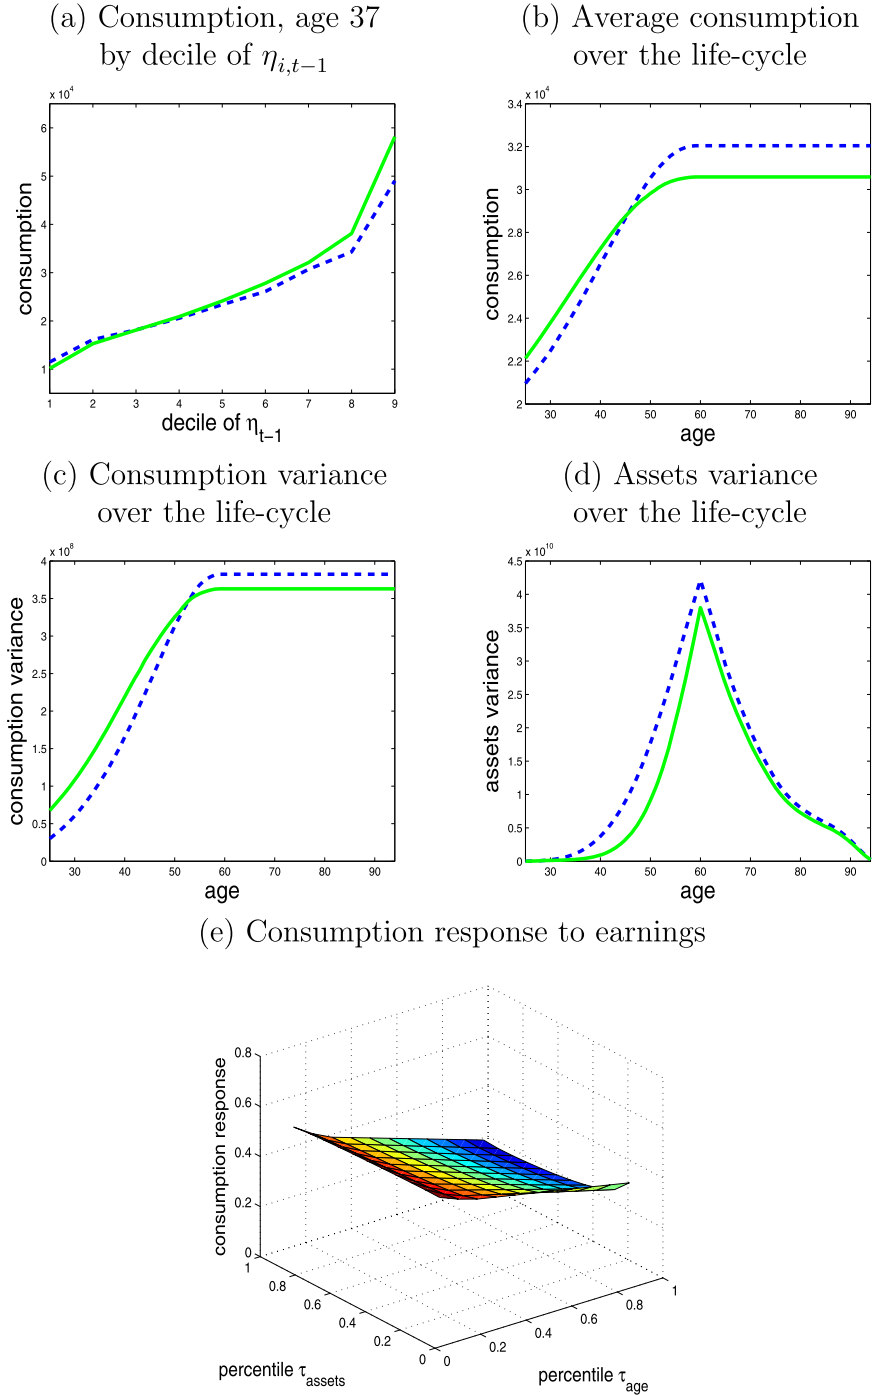


FIGURE S35.—Simulation exercise. *Notes:* In the top four panels, dashed is based on the nonlinear earnings process (S6)–(S7), solid is based on the canonical earnings process (3). Panel (e): estimate of the average derivative of the conditional mean of log-consumption with respect to log-earnings, given earnings, assets, and age, evaluated at values of assets and age that correspond to their τ_{assets} and τ_{age} percentiles, and averaged over the earnings values. Quantile regression on polynomials, see Section 6 for a description.

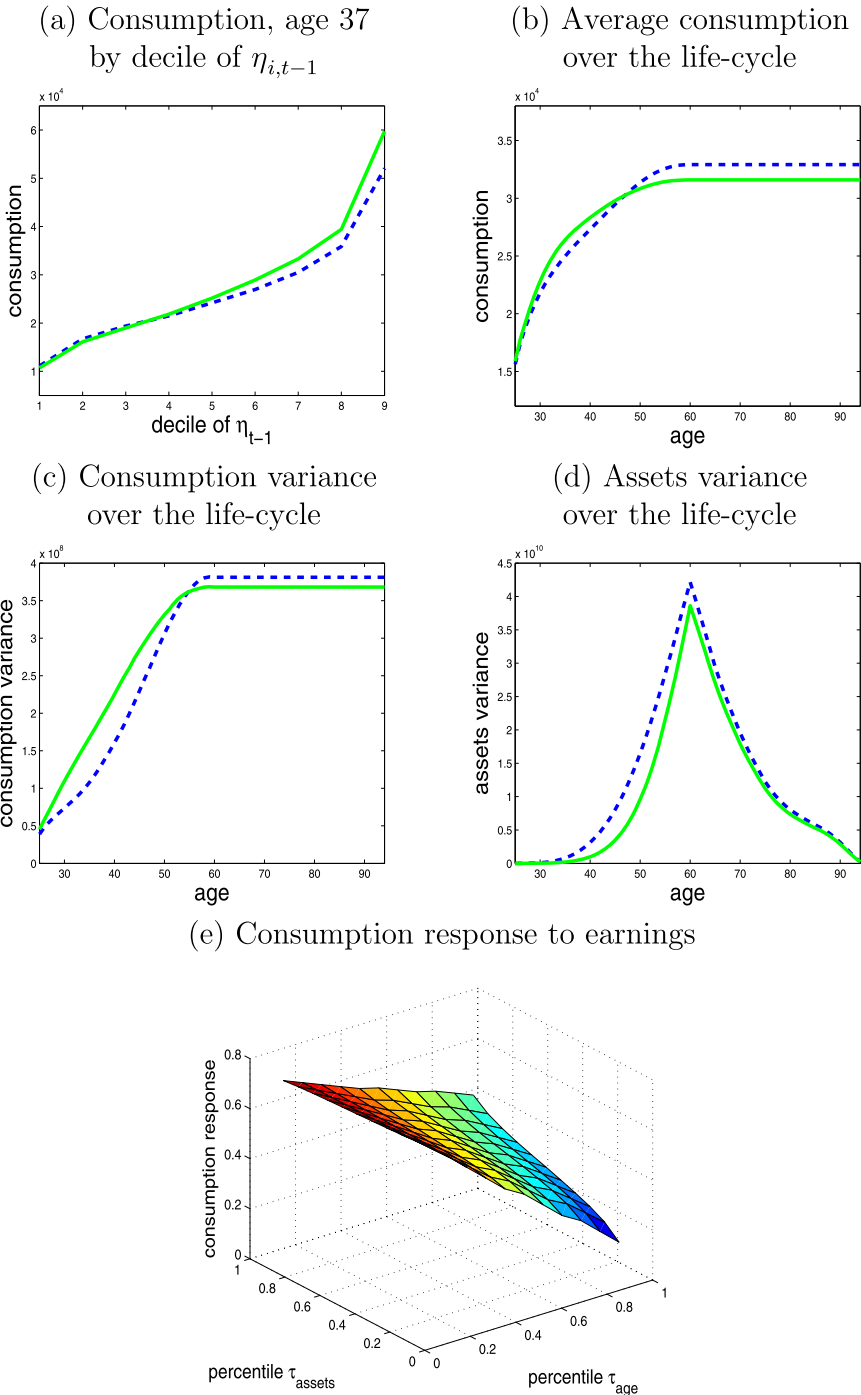


FIGURE S36.—Simulation exercise, no borrowing. *Notes:* In the top four panels, dashed is based on the non-linear earnings process (S6)–(S7), solid is based on the canonical earnings process (3). Panel (e): estimate of the average derivative of the conditional mean of log-consumption with respect to log-earnings, given earnings, assets, and age, evaluated at values of assets and age that correspond to their τ_{assets} and τ_{age} percentiles, and averaged over the earnings values. Assets are constrained to be nonnegative.

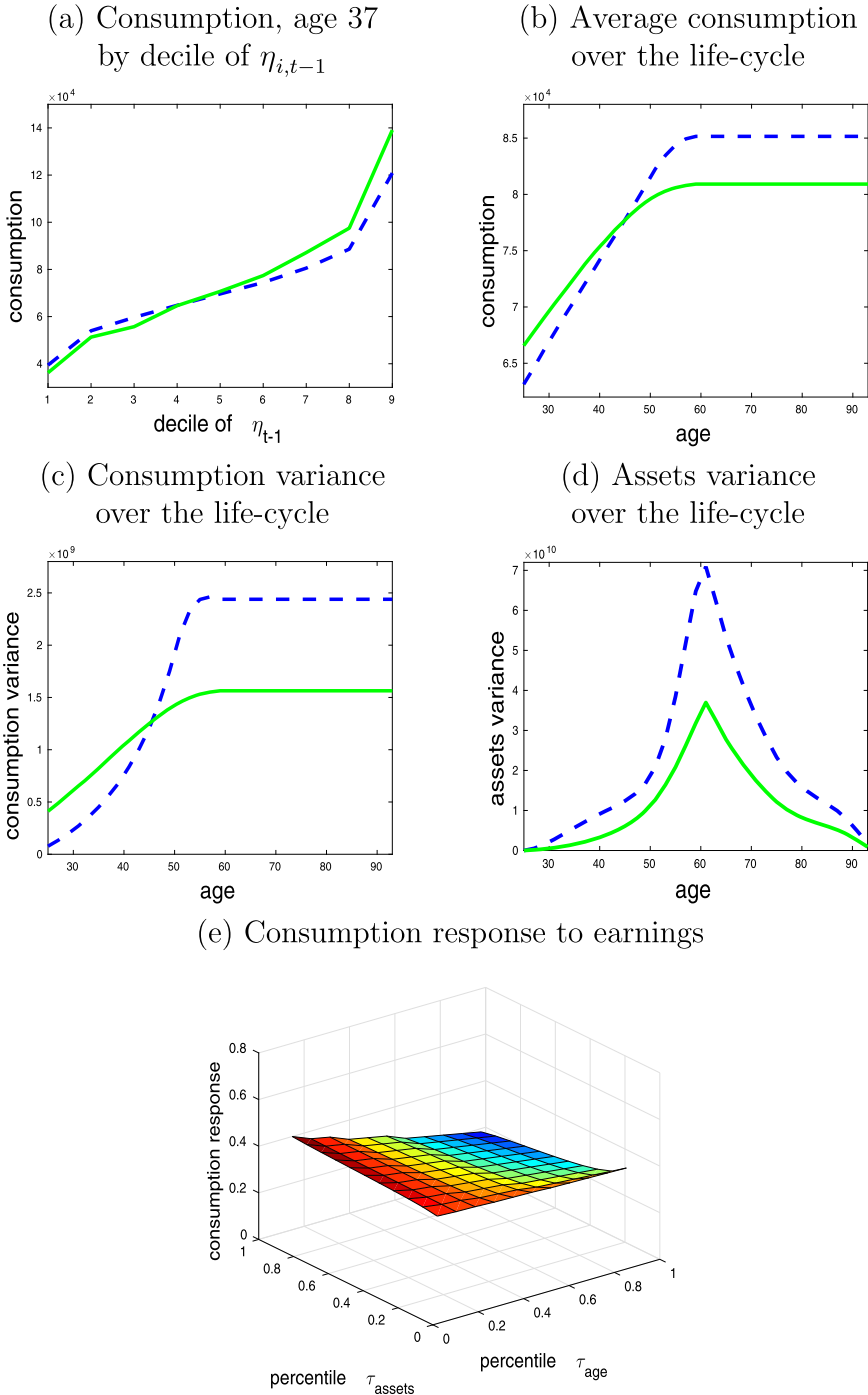


FIGURE S37.—Simulations based on the estimated nonlinear earnings model. *Notes:* In the top four panels, dashed is based on the nonlinear quantile-based earnings process estimated on the PSID, solid is based on a comparable canonical earnings process. Panel (e): estimate of the average derivative of the conditional mean of log-consumption with respect to log-earnings, given earnings, assets, and age, evaluated at values of assets and age that correspond to their τ_{assets} and τ_{age} percentiles, and averaged over the earnings values.

REFERENCES

- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98 (5), 1887–1921. [5,7]
- DEATON, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59 (5), 1221–1248. [1]
- GOUVEIA, M., AND R. P. STRAUSS (1994): "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis," *National Tax Journal*, 47 (2), 317–339. [4]
- GUVENEN, F. (2007): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review*, 97 (3), 687–712. [7]
- GUVENEN, F., AND A. SMITH (2014): "Inferring Labor Income Risk From Economic Choices: An Indirect Inference Approach," *Econometrica*, 82 (6), 2085–2129. [7]
- HOROWITZ, J. (2012): "Specification Testing in Nonparametric Instrumental Variables Estimation," *Journal of Econometrics*, 167, 383–396. [9]
- HU, Y., AND S. M. SCHENNACH (2008): "Instrumental Variable Treatment of Nonclassical Measurement Error Models," *Econometrica*, 76, 195–216. [7]
- HU, Y., AND M. SHUM (2012): "Nonparametric Identification of Dynamic Models With Unobserved State Variables," *Journal of Econometrics*, 171, 32–44. [7]
- KAPLAN, G., AND G. VIOLANTE (2010): "How Much Consumption Insurance Beyond Self-Insurance," *American Economic Journal*, 2 (4), 53–87. [3,5,6]
- WILHELM, D. (2015): "Identification and Estimation of Nonparametric Panel Data Regressions With Measurement Error," Cemmap Working Paper CWP34/15. [7-9]

*CEMFI, Casado del Alisal, 5, Madrid 28014, Spain; arellano@cemfi.es,
University College London, Gower Street, London WC1E 6BT, U.K. and Institute for Fiscal
Studies, 7 Ridgmount Street, London WC1E 7AE, U.K.; r.blundell@ucl.ac.uk,
and*

*Dept. of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637,
U.S.A.; sbonhomme@uchicago.edu.*

Co-editor Giovanni L. Violante handled this manuscript, which is based on a Presidential Address.

*Manuscript received 8 September, 2015; final version accepted 30 January, 2017; available
online 2 February, 2017.*