Measuring Consumption Insurance*

Agustin Diaz† Justin Franco Lam‡ Sean McCrary§ Kharis Sokolov¶

April 20, 2024

Abstract

We explore the properties of the consumption pass-through estimator proposed by *Heathcote, Storesletten and Violante (2014)* (HSV) using simulated data from a standard life cycle incomplete-markets model with labor supply decision. We show that, without savings, the HSV wage-to-consumption pass-through estimator accurately captures consumption insurance via labor supply adjustment. On the other hand, when there is savings, HSV has an upward bias and that bias gets more severe when savings accumulate or when the persistence of the wage process deviates from the unit root assumption that underlies the method. Compared to the estimate of *Blundell, Pistaferri and Preston (2008)* (BPP), HSV performs better when there is little savings in the model, while BPP performs better when there is ample savings. The findings suggest that combining the two methods, such as using HSV for agents close to the borrowing constraint while applying BPP to agents away from the borrowing constraint, may yield more comprehensive consumption insurance estimates.

*We are grateful to our advisors José-Víctor Ríos-Rull, Iourii Manovskii, Dirk Krueger and Joachim Hubmer for their insightful advice and continuous support. We also thank Harold Cole, Guillermo Ordonez, Sara Casella, Jincheng Huang, Xincheng Qiu and seminar participants at Penn Macro Lunch for helpful comments. The views expressed herein are solely those of the author and do not necessarily reflect the views of the Central Bank of Chile. We are responsible for all errors.

†Central Bank of Chile Email: adiaz@bcentral.cl.

‡Department of Economics, University of Pennsylvania. Email: jflam@sas.upenn.edu.

§Department of Economics, Ohio State University. Email: mccrary.65@osu.edu.

¶Long-Range Planning Data Science Team, Microsoft Corporation. Email: ksokolov@microsoft.com.

1 Introduction

The amount of consumption insurance in an economy is one of the most important topic in modern macroeconomics since it determines agents' welfare in the presence of income risk and the role of government when the market can't provide insurance. The success of welfare programs and tax systems, for instance, depends on understanding the existing degree of risk-sharing among individuals in an economy. For this reason, in this paper, we evaluate and compare the reliability of two of the most used methods in the literature: Heathcote, Storesletten and Violante (2014) (HSV hereafter) and Blundell, Pistaferri and Preston (2008) (BPP hereafter). Through this exercise, we highlight the weaknesses and strengths of each estimator and propose remedies to obtain a more comprehensive measure.

The goal of both methods is to estimate the income shock pass-through coefficient to consumption, which measures how sensitive consumption is to income shocks. Less sensitive consumption implies that individuals have a higher capacity to protect themselves against adverse shocks and hence higher consumption insurance. Blundell, Pistaferri and Preston (2008) provided a non-structural method to estimate consumption pass-through, which depends on the orthogonality of consumption growth and a linear combination of future and past income shocks. As a result, this method requires that current consumption doesn't depend on future or previous income shocks. In the presence of borrowing constrained individuals, this assumption is violated, causing an upward bias to the estimator as shown in Kaplan and Violante (2010).

To overcome these difficulties, Heathcote, Storesletten and Violante (2014) provides a fully structural approach to back out the consumption pass-through coefficient in a model-consistent way. Their method is appealing as it allows direct identification of consumption insurance through the variance and covariance of consumption, wage, and working hours.

The importance and versatility[1] of HSV's framework motivate us to explore the reliability of the estimator in the spirit of Kaplan and Violante (2010) and Wu and Krueger (2021), where they explore the reliability of the Blundell, Pistaferri and Preston (2008) and Blundell, Pistaferri and Saporta-Eksten (2016) estimator respectively. The three questions we aim to answer are: (i) How accurate is the HSV method? (ii) How does the performance compare with BPP? (iii) Can we combine the methods to improve the estimator?

We proceed by using the standard life cycle incomplete markets model as a laboratory to test the estimators' performance. The presumption is that the model can capture key mechanisms of the real-world data. By simulating datasets from the model, we will know the true data-generating process, which will allow us to compare the true consumption pass-through coefficient

[1]For instance, Heathcote, Storesletten and Violante (2017), Boerma and Karabarbounis (2020) and Boerma and Karabarbounis (2021), who extended Heathcote, Storesletten and Violante (2014) to infer optimal tax progressivity, inequality with home production and to account for trends in leisure time allocation respectively.

with the HSV estimates.

Our main findings are that the HSV estimates accurately capture the pass-through coefficient when there is no savings motive or when agents are near the borrowing constraint, but it is upwardly biased when agents save. When agents accumulate savings as they age or when we depart from the unit root wage process assumption, the upward bias in HSV grows more severe. Combined with the fact that the BPP estimator is upwardly biased for borrowing-constrained individuals and unbiased for agents with sufficient savings, our findings suggest that applying HSV to agents close to the borrowing constraint while applying BPP to agents away from it may provide more accurate consumption insurance estimates.

Road Map. The rest of the paper is organized as follows. Section 2 reviews the consumption pass-through coefficient and the methodology of HSV. Section 3 describes the standard life cycle incomplete markets model from which we simulate our dataset. Section 4 presents the performance of HSV in the simulated dataset. Section 5 concludes.

2 Consumption pass-through and HSV Review

In this section, we review the consumption pass-through coefficient and the methodology proposed by Heathcote, Storesletten and Violante (2014) to recover it from the data. First, we defined the consumption pass-through. Second, in order to understand the method assumptions, we review the HSV environment and market structure. Finally, we review the HSV estimator and discuss its identification.

2.1 Consumption pass-through

The consumption pass-through coefficient summarizes how much consumption changes with respect to an income shock. The less sensitive consumption is, the more insurance it implies. Following the literature, the consumption pass-through coefficient for shock x_{it} at age t is defined as:

$$\phi_t = \frac{cov(\Delta \log c_{it}, x_{it})}{var(x_{it})} \tag{1}$$

where c_{it} is the consumption for household i at age t, and the variance and covariance are taken cross-sectionally over households at age t.

An appealing characteristic of this metric is its intuitive interpretation as the univariate OLS regression coefficient of change in log consumption on the (unobservable) income shock. The micro-foundation of this metric is presented in BPP, where the authors showed that a linearization of the Euler equation in a version of the Permanent Income Hypothesis with CRRA preferences approximates the following equation:

$$\Delta \log c_{it} = \lambda_{it}\varepsilon_{it} + \psi_{it}\zeta_{it} + \xi_{it}$$

where ε_{it}, ζ_{it}, and ξ_{it} refer to the permanent income shock, transitory income shock, and measurement error, respectively, while λ_{it} and ψ_{it} denote the loading factors of the shocks. Hence, the pass-through coefficient can be rationalized as an estimate for the theoretical loading factor of income shocks.

Unfortunately, income shocks are unobservable to econometricians, so such OLS regression cannot be estimated. To overcome this difficulty, we need the methods studied in this paper to recover the pass-through coefficient. We will now turn to explain the methodology of HSV.

2.2 HSV Environment

HSV uses a structural equilibrium model with a particular asset structure to derive closed-form solutions to the equilibrium allocations. This allows them to use cross-sectional variance and covariance of the joint distribution of wages, worked hours, and consumption to identify the extent of risk sharing in the economy. In this subsection, we will focus on the main components of their model structure.

Demographics. The authors adopt a Yaari, perpetual youth model where agents are born at age zero and survive from age a to $a+1$ with probability $\delta < 1$. A new generation of mass $(1-\delta)$ enters the economy at each date t.

Wage process and Information. In order to address Deaton (1997)'s critique that structural-based estimates will be sensitive to the details of the risk-sharing mechanisms considered in the model, HSV models the insurability of income shock flexibly and focuses on letting the data identify the insurability nature of the shock rather than specifying the mechanism of how it is achieved. The key innovation of HSV's framework is to model the shocks to hourly wages to be of two types: one that is completely insurable, as in incomplete market models, and one that is uninsurable due to the lack of state-contingent claims, as in the incomplete-markets model. In particular, the wage process is specified as follows:

$$\log\ w_t = \alpha_t + \varepsilon_t$$
$$\alpha_t = \alpha_{t-1} + \omega_t \quad \text{with } \omega_t \sim F_{\omega,t}$$
$$\varepsilon_t = \kappa_t + \theta_t \quad \text{with } \theta_t \sim F_{\theta,t}$$
$$\kappa_t = \kappa_{t-1} + \eta_t \quad \text{with } \eta_t \sim F_{\eta,t}$$

where ε_t denotes the fully insurable shock since HSV specifies the market structure such that agents can and are willing to trade (η_t, θ_t) contingent assets. On the other hand, α_t is the "uninsurable" shock as the specified environment makes it such that trades in ω_t contingent assets cannot happen. The quote-and-quote for uninsurable is because α_t shocks are still allowed to be smoothed via labor supply adjustments, via borrowing and lending in the risk-free bond, or via government redistribution, achieving partial insurance. Agents take as given the sequence distribution of $\{F_{\alpha^0,t}, F_{\kappa^0,t}, F_{\omega,t}, F_{\theta,t}, F_{\eta,t}\}$ having perfect foresight over future wage distribution.

Market Structure. All assets in the economy are in zero net supply, and asset markets are competitive. Agents are born with zero financial wealth. Individuals born in year b draw values for α^0 and θ^0 before any markets open. HSV economy is composed of an island structure and individuals are allocated to an island, which is defined by a sequence of ω_t common to all island members. Within each island, individuals are allowed to trade a complete set of insurance contracts indexed to $s_{t+1} = \{\omega_{t+1}, \eta_{t+1}, \theta_{t+1}\}$. However, between islands, individuals can only trade insurance contracts indexed $z_{t+1} = (\eta_{t+1}, \theta_{t+1})$, ruling out inter-island contracts contingent on island-level shocks. In this structure, agents can trade risk-free bonds by buying one contract that pays one unit of consumption for every pair $(\eta_{t+1}, \theta_{t+1})$.

Preferences and Agent Problem. Lifetime utility for an agent born in cohort b is given by:

$$\sum_{t=b}^{\infty} (\beta\delta)^{t-b} \max_{c_t, h_t, B_t, B^*} \int_{t,b,\mathbb{S}} \left(\frac{c_t^{1-\sigma} - 1}{1-\sigma} - e^{\varphi} \frac{h_t^{1+\gamma}}{1+\gamma} \right) d_{t,b,\mathbb{S}}$$

, where c is consumption, h worked hours, σ the intertemporal elasticity of substitution for consumption, γ the convexity of the disutility of labor supply, and φ is the risk aversion to work. The agent's budget constraint at each period is given by:

$$w(s^t)h_t(s^t) + B_{t-1}(s_t; s^{t-1}) + B^*_{t-1}(z_t; s^{t-1}) = c_t(s^t) +$$
$$\int_{\mathbb{S}} Q_t(s_{t+1}; s^t) B_t(s_{t+1}; s^t) d_{s_{t+1}} + \int_{\mathbb{Z}} Q^*_t(z_{t+1}; s^t) B^*_t(z_{t+1}; s^t) d_{z_{t+1}}$$

, where $s^t = (s^{t-1}, s_t)$ denotes the individual history from birth to period t. $Q_t(s_{t+1}; s^t)$ denotes the price of insurance claims purchased at date t from within-island insurers by an agent with history s^t that delivers one unit of consumption at $t+1$ if and only if s_{t+1} is realized. $B_t(s_{t+1}; s^t)$ denotes the quantity of claims purchases that are paid in individual states s_{t+1}. Finally, Q^* and B^* denote the price and quantity of claims purchased from outside island insurer that pays at realizations z_{t+1}.

No-trade result. The key result in Heathcote, Storesletten and Violante (2014) is that in equilibrium, agents choose not to use the risk-free bond to smooth the uninsurable shock. This result extends the logic in Constantinides and Duffie (1996). Due to this no-trade result, in stark contrast to the standard life cycle incomplete-markets model, we can derive closed-form solutions to the allocation of consumption and hours worked in the economy:

$$\log\ c_t(\varphi, \alpha, \varepsilon) = -(1-\tau)\hat{\varphi} + (1-\tau)\left(\frac{1+\hat{\gamma}}{\hat{\gamma}+\sigma}\right)\alpha_t$$
$$\log\ h_t(\varphi, \alpha, \varepsilon) = -\hat{\varphi} + \frac{1-\sigma}{\hat{\gamma}+\sigma}\alpha_t + \frac{1}{\hat{\gamma}}\varepsilon_t$$

where τ denotes the progressivity of the tax system, σ denotes the risk aversion parameter of the CRRA preference, $\hat{\gamma}$ denotes a function of the elasticity of labor supply and $\hat{\varphi}$ denotes a function of the disutility of work. From the equations, we can see that consumption only reacts to changes in α_t but not ε_t, showcasing that one shock is completely insurable while the other is only partially insurable. These closed-form solutions enable the computation of variance and covariances of the joint equilibrium distribution of wages, worked hours, and consumption, which in turn, allows the authors to prove the identification of the model's parameters and hence back out the consumption pass-through coefficient.

2.3 HSV Consumption Pass-Trough Methodology and Identification.

The connection between the consumption pass-through coefficient and the model parameters is as follows:

$$\phi_t^{HSV} \equiv \frac{cov(\Delta \log c_t, \omega_t + \eta_t)}{var(\omega_t + \eta_t)} = (1-\tau) \cdot \frac{1+\hat{\gamma}}{\hat{\gamma}+\sigma} \cdot \frac{v_{\omega t}}{v_{\omega t} + v_{\eta t}} \tag{2}$$

where $v_{\omega t}$ and $v_{\eta t}$ are the variance of the respective shocks. This result shows that the HSV pass-through coefficient effectively measures the portion of the permanent shock variance that is uninsurable ($\frac{v_{\omega t}}{v_{\omega t}+v_{\eta t}}$). To identify these underlying parameters, HSV proceeds in the following sequential way[2].

[2]The exposition closely follows that in *Appendix C* of Heathcote, Storesletten and Violante (2014).

Getting $v_{\omega t}$ is the most important step of the identification proof and it relies on the within-cohort changes in moments, the subscript a denotes the age of that cohort and Δ denotes the first difference with its lag variable within that cohort:

$$\frac{\Delta cov_t^a(\log w, \log c)^2}{\Delta var_t^a(\log c)} = v_{\omega t} \tag{3}$$

After obtaining $v_{\omega t}$, $(v_{\eta t} + \Delta v_{\theta t})$ is identified by:

$$\Delta var_t^a(\log w) = v_{\omega t} + (v_{\eta t} + \Delta v_{\theta t}) \tag{4}$$

Given $v_{\omega t}$ and $(v_{\eta t} + \Delta v_{\theta t})$, the tax-modified Frisch elasticity $\hat{\gamma}$ is identified by:

$$\Delta var_t^a(\log h) = \left(\frac{\Delta cov_t^a(\log h, \log c)}{\Delta cov_t^a(\log w, \log c)}\right)^2 v_{\omega t} + \frac{1}{\hat{\gamma}^2}(v_{\eta t} + \Delta v_{\theta t}) \tag{5}$$

Given $\hat{\gamma}$, parameter σ is identified by:

$$\frac{\Delta cov_t^a(\log h, \log c)}{\Delta cov_t^a(\log w, logc)} = \frac{1-\sigma}{\hat{\gamma}+\sigma} \tag{6}$$

Finally, we can back out $v_{\theta t}$ by differencing between the dispersion in growth rates and the growth rate of within-cohort dispersion:

$$\begin{aligned} cov_t^a(\Delta \log w, \Delta \log h) + var_t^a(\Delta \log h) - \Delta cov_t^a(\log w, \log h) \\ -\Delta var_t^a(\log h) = \frac{2(1+\hat{\gamma})}{\hat{\gamma}^2} v_{\theta, t-1} \end{aligned} \tag{7}$$

With $(v_{\eta t} + \Delta v_{\theta t})$ identified from (4) and $v_{\theta t}$ identified from (7). Finally, $v_{\eta t}$ will be identified as the residual. As a result, we get all the parameters to compute the consumption pass-through coefficient.

In our exercises below, we would make use of equations (3) to (7) to back out the consumption pass-through coefficient implied by the HSV method and compare it to the true coefficient.

3 Standard Life Cycle Incomplete Markets Model

This section describes the standard life cycle incomplete markets model which we use to simulate the data and evaluate the accuracy of the HSV estimate. We consider a life cycle partial equilibrium model of the style of Deaton (1991). The key differences between this environment and the HSV environment are: i) the market structure, individuals are only allowed to trade risk-free bonds, and ii) the economy allows a positive supply of risk-free assets.

3.1 Model Environment

The model environment is cast in discrete time. Agents live T periods, dying with certainty at the end of the last period, and discount the future at rate β.

Technology. There is no aggregate uncertainty, and the economy is populated with ex-ante identical households. Workers face idiosyncratic shocks to their wages and make consumption-savings decisions as well as labor supply choices, taking interest rates as given. For assets, there is only a risk-free bond, and households can self-insure by saving at an exogenous risk-free rate of r. In our baseline specification, we consider the zero borrowing constraint.

Preferences. Households maximize the discounted sum of time-separable expected utility by choosing the amount of consumption and saving as well as the number of worked hours each period, subject to the budget constraint. The instantaneous utility function is as follows:

$$u(c_t) - v(h_t)$$

where $u(c_t)$ refers to the utility of consumption and $v(h_t)$ refers to the disutility of hours worked.

Wage Process.

Following the literature, our baseline specification considers a permanent-transitory structure for the wage process:

$$\log w_t = \alpha_t + \zeta_t, \quad \zeta_t \sim \mathcal{N}(0, \sigma_\zeta) \tag{8}$$

$$\alpha_t = \rho\alpha_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon) \tag{9}$$

, where $\rho \leq 1$.

Household Problem. The environment specified above will give rise to a recursive formulation of the life-cycle consumption-saving problem. In the last period, agents know that they will die with certainty in the next period, and they don't derive utility from dying with savings. As a result, at age T, their value function becomes:

$$
\begin{aligned}
V(a,w,T) &= \max_{c,a',h} u(c) - v(h) \\
\text{s.t.} \qquad & c + a' = wh + (1+r)a \\
\log w_T &= \alpha_T + \zeta_T, \quad \zeta_T \sim \mathcal{N}(0,\sigma_\zeta) \\
\alpha_T &= \rho\alpha_{T-1} + \varepsilon_T, \quad \varepsilon_T \sim \mathcal{N}(0,\sigma_\varepsilon) \\
& c \geq 0, h \geq 0, a^{'} \geq 0
\end{aligned}
$$

where c, a, h, T denotes consumption, assets, hours worked, and the final age, respectively. It is trivial to see that agents consume all their income wh and their savings $(1+r)a$. For the rest of the periods, agents' value function is as follows :

$$
\begin{aligned}
V(a,w,t) &= \max_{c,a',h} u(c) - v(h) + \beta E[V(a',w',t+1)|w] \\
\text{s.t.} \qquad & c + a' = wh + (1+r)a \\
\log w_t &= \alpha_t + \zeta_t, \quad \zeta_t \sim \mathcal{N}(0,\sigma_\zeta) \\
\alpha_t &= \rho\alpha_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,\sigma_\varepsilon) \\
& c \geq 0, h \geq 0, a^{'} \geq 0
\end{aligned}
$$

The problems faced by households differ from the last period as they enjoy utility from consuming in period $t+1$, creating incentives to save.

Computation and Simulation. We solve the household problem via backward induction. Then, the artificial data set is simulated using the policy function obtained in the previous stage. We assume that all agents are born with zero assets at age 0. The model period is set to be a year and final age $T = 70$. We simulate the dataset with 150,000 agents.

3.2 Parametrization

To facilitate comparison, we parametrize the preference function in the same way as HSV: $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$ and $v(h) = e^{\varphi}\frac{h^{1+\gamma}}{1+\gamma}$. We chose the parameter values by following Wu and Krueger (2021). Table 1 shows the parameter values in our baseline case.

Table 1: Baseline parameters for standard life cycle incomplete markets model

	Description	Value
σ	Risk Aversion	1.73
γ	Elasticity of labor supply	1.89
φ	Disutility of labor	0.931
r	Interest rate	0.02
ρ	Persistence of permanent component in wage process	1
σ_ϵ	Standard deviation of permanent wage innovation	0.174
σ_ζ	Standard deviation of transitory wage innovation	0.166
β	Discount rate	0.995

Notes: This table reports the baseline parameters for the standard life cycle incomplete markets model.

4 HSV in a Standard Life Cycle Incomplete Market Model

In this section, we present the results. First, we show how the HSV estimator performs at different levels of savings. Second, we exhibit how the HSV estimator performs when we deviate from the permanent shock unit root assumption in Heathcote, Storesletten and Violante (2014). Finally, we compare the HSV with the BPP estimator, finding that HSV outperforms BPP if agents have little savings.

4.1 HSV's Performance with and without Savings

We first show the performance of the HSV methodology under the baseline model with precautionary savings and contrast it with the case where the precautionary savings motive is absent. We calculate the HSV coefficient by first identifying the HSV model parameters using equations (3) to (7) for each age and then plug them into equation (2). The true pass-through coefficient is obtained using (1) since we know the data-generating process and observe the income shocks.

The left panel in figure 1 plots the HSV estimates and the true coefficients for each age[3] under the baseline model. As can be seen, the HSV pass-through coefficient has an upward bias for all ages compared to the true coefficients, and the bias increases as individuals age in the model. At age 2, the bias is 19.8%, increasing to 86.1% at age 70. The reason behind the decrease in the true consumption pass-through over age is because of the accumulation of assets. In the life cycle incomplete markets model, as time passes by, agents accumulate wealth as a

[3] Note that the reported coefficients start from age 2 because the HSV method uses first difference of variables.

Figure 1: HSV pass-through coefficient with and without precautionary savings

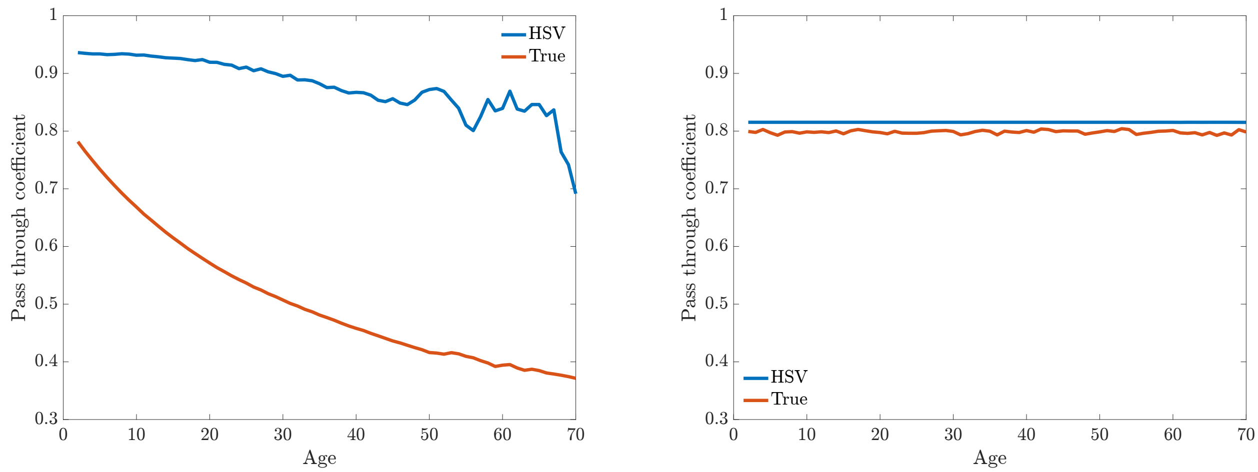

(a) Baseline

(b) Without precautionary savings

Notes: The figure plots the HSV wage-to-consumption pass-through coefficient with our baseline parameters (left) and when we set $r = -1$ to switch off precautionary savings motive while keeping everything else the same (right).

buffer so that they can draw down their savings to cope with negative income shocks. The fact that the downward trend of the HSV estimates in age is not as strong as the true coefficients implies that it misses the consumption insurance channel via precautionary savings.

To further illustrate the effect of savings on the HSV estimates, on the right panel of figure 1, we report the HSV estimates and the true coefficients when the precautionary savings motive is switched off, by setting $r = -1$ while keeping all other parameters at the baseline value. In this environment without savings, the HSV estimates closely align with the true coefficients, having an average upward bias of 2.1% only, accurately capturing the consumption insurance channel via labor supply adjustment.

4.2 HSV's Performance with Deviation from Unit Root

The second result shows the performance of the HSV estimates when we deviate from the unit root assumption of the permanent component in the wage process. Note that due to mean reversion, when the persistence of the permanent component in wage is less than 1, the distribution of wage and hence of consumption and asset in our model start to display little change after age 40. Since the HSV methodology makes use of the change in variance or covariance of consumption and wage, this will lead to numerical instability of HSV, hence we only report the coefficient from age 2 to age 40. The results are shown in figure 2. In the unit root case, the average upward bias between age 2 and age 40 is 56.2%, and as persistence is reduced to 0.99 and 0.98, the average upward bias increases to 62.1% and 70%, respectively.

This reflects that the estimator is sensitive to the underlying persistence assumption of the wage process.

Figure 2: HSV pass-through coefficient with deviation from unit root assumption

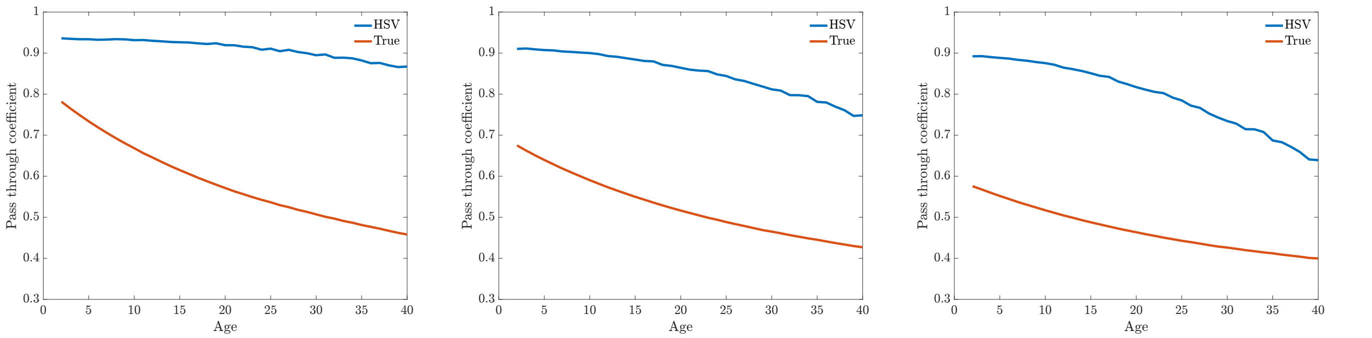

(a) $\rho = 1$ (b) $\rho = 0.99$ (c) $\rho = 0.98$

Notes: The figure plots the HSV, BPP and the true wage-to-consumption pass-through coefficients with our baseline parameters (left), $\rho = 0.99$ (middle) and $\rho = 0.98$ (right).

4.3 Comparison between HSV and BPP

As mentioned in the introduction, BPP is a widely used alternative method that offers a less structural approach to measure consumption insurance. In this section, we compare the two methodologies and highlight the differences to provide a benchmark for the HSV's performance.

BPP Refresher. The BPP estimator is essentially an instrumental variable regression for the pass-through coefficient. BPP uses the variable $\Delta \log w_{i,t-1} + \Delta \log w_{i,t} + \Delta \log w_{i,t+1}$ as an instrument for the unobservable permanent income shock, and the pass-through coefficient is defined as:

$$\phi_t^{BPP} = \frac{cov(\Delta \log c_{i,t}, \Delta \log w_{i,t-1} + \Delta \log w_{i,t} + \Delta \log w_{i,t+1})}{cov(\Delta \log w_{i,t}, \Delta \log w_{i,t-1} + \Delta \log w_{i,t} + \Delta \log w_{i,t+1})} \tag{10}$$

BPP assumed that the unexplained (from observable characteristics) log income process can be decomposed into a permanent component P_{it} and a transitory component v_{it}.

$$\log w_{it} = P_{i,t} + v_{i,t}$$

The permanent component follows a unit root process with shock ε_{it} being serially uncorrelated; on the other hand, the transitory component follows an MA(q) process with $\theta_0 \equiv 0$:

$$P_{it} = P_{i,t-1} + \varepsilon_{it}$$
$$v_{it} = \sum_{j=0}^{q} \theta_j \zeta_{i,t-j}$$

The unexplained income growth can thus be written as:

$$\Delta \log w_{it} = \varepsilon_{i,t} + \Delta v_{i,t}$$

For the consumption process, BPP assumed that the unexplained change in log consumption can be written as:

$$\Delta \log c_{it} = \lambda \varepsilon_{i,t} + \psi \zeta_{i,t}$$

Finally, the following identifying assumptions are in place:

$$cov(\Delta \log c_{it}, \varepsilon_{i,t+1}) = cov(\Delta \log c_{it}, \zeta_{i,t+1}) = 0 \qquad \text{(No Foresight)}$$
$$cov(\Delta \log c_{it}, \varepsilon_{i,t-1}) = cov(\Delta \log c_{it}, \zeta_{i,t-2}) = 0 \qquad \text{(Short Memory)}$$

If these assumptions are satisfied, then the BPP estimator for permanent income shock ε_{it} (10) will be equivalent to the true consumption pass-through coefficient (1). The intuition behind this assumption is that current consumption is independent of previous and future unexpected income shocks.

4.3.1 HSV vs BPP: Savings Margin

Guided by our result above, we know that savings is an important margin for HSV's performance. Coincidentally, as documented by Kaplan and Violante (2010), this is also true for BPP. Since savings is a crucial margin for both estimators, a comparison of the two methods along the savings margin is presented in figure 3 to draw insights.

On the left panel of figure 3, we show the comparison between HSV and BPP under our baseline parameters. In this environment where agents have considerable savings, BPP outperforms HSV as it almost exactly identifies the true wage to consumption pass-through coefficients.

On the right panel of figure 3, we show the same comparison but in an environment with little savings. We achieve this by setting $r = -0.012$. Consistent with Kaplan and Violante (2010) and our results above, BPP has an upward bias due to the violation of the short memory

Figure 3: HSV vs BPP: baseline and little savings

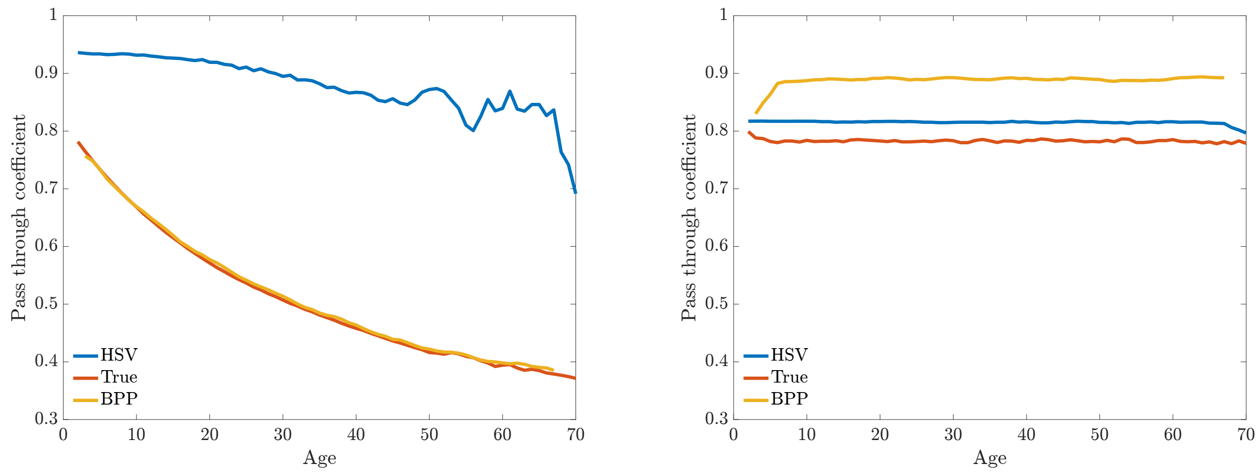

(a) Baseline (b) Little savings

Notes: The figure plots the HSV, BPP and the true wage-to-consumption pass-through coefficient with our baseline parameters (left) and when we set $r = -0.012$ while keeping all other parameters the same (right).

condition, whereas HSV closely aligns with the true coefficients. The average upward bias for BPP is 13.4%, while that for HSV is just 4.12%. As explained in Kaplan and Violante (2010), the intuition behind the bias in BPP is best illustrated when we consider a household close to the borrowing constraint and receiving a negative transitory shock at age $t-2$. Should they dissave and reduce their wealth to smooth consumption, they will have a large expected utility cost because of the chance of becoming constrained in the future. Therefore, households will choose to optimally decrease consumption rather than dissave. Importantly, due to the transitory nature of the shock, changes in consumption are expected to be positive in the next and next next period, as income gradually returns to its normal value through mean reversion. This causes $cov(\Delta \log c_{it}, \zeta_{i,t-1}) < 0$ and crucially the short memory condition to be violated ($cov(\Delta \log c_{it}, \zeta_{i,t-2}) < 0$).

The results in figure 3 show that HSV's performance is the opposite of BPP's along the savings margin, and there may be room to combine the two methods to achieve a more comprehensive consumption insurance measure.

4.3.2 HSV vs BPP: Wage Persistence Margin

Guided again by our results above, we know that the HSV estimate is sensitive to the wage persistence assumption. This section thus compares the performance of HSV and BPP along the wage persistence margin. The result is presented in figure 4.

Figure 4: HSV vs BPP: wage persistence

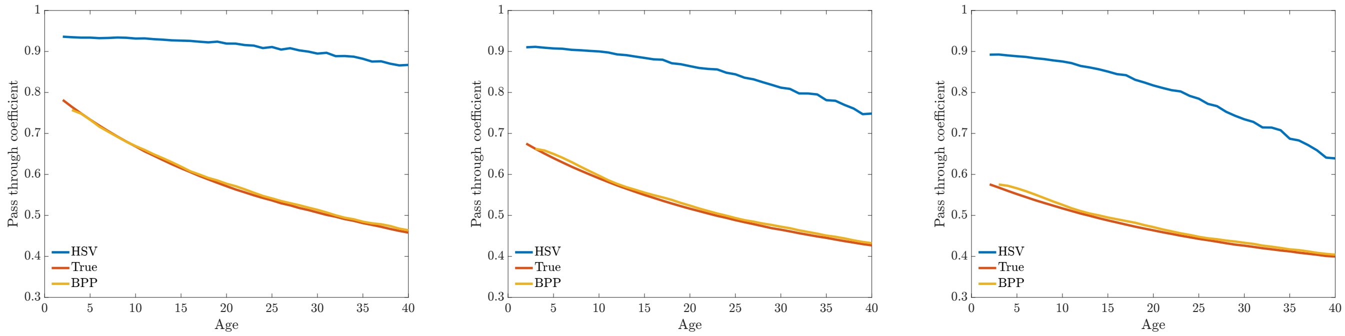

(a) Baseline, $\rho = 1$ (b) $\rho = 0.99$ (c) $\rho = 0.98$

Notes: The figure plots the HSV, BPP and the true wage-to-consumption pass-through coefficients with our baseline parameters (left), $\rho = 0.99$ (middle) and $\rho = 0.98$ (right).

Consistent with the findings of Kaplan and Violante (2010), the BPP estimate is insensitive to the persistence assumption of the permanent shock, outperforming the HSV along this margin.

5 Conclusion

In this paper, we explored the properties of the HSV consumption pass-through estimator using simulated data from a standard incomplete-markets model. There are four takeaways from this exercise. First, the HSV method accurately captures the consumption pass-through coefficient when agents do not have savings. Second, the HSV pass-through coefficient has an upward bias when precautionary savings are present. Third, the upward bias of HSV grows more severe as savings accumulate or when the wage process deviates further from the unit-root assumption. Fourth, the performance of HSV is the opposite of BPP along the savings margin, where BPP has an upward bias if agents are close to the borrowing constraints but capture the pass-through coefficient well if agents are away from the borrowing constraints. The results of this paper suggest that combining the two methods: HSV for agents close to the borrowing constraints and BPP for agents away from the borrowing constraints, can be a promising direction in providing a more comprehensive measure of consumption insurance for the economy as a whole.

References

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption inequality and partial insurance." *American Economic Review*, 98(5): 1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption inequality and family labor supply." *American Economic Review*, 106(2): 387–435.

Boerma, Job, and Loukas Karabarbounis. 2020. "Labor market trends and the changing value of time." *Journal of Economic Dynamics and Control*, 115: 103885.

Boerma, Job, and Loukas Karabarbounis. 2021. "Inferring inequality with home production." *Econometrica*, 89(5): 2517–2556.

Constantinides, George M, and Darrell Duffie. 1996. "Asset pricing with heterogeneous consumers." *Journal of Political Economy*, 104(2): 219–240.

Deaton, Angus. 1991. "Saving and liquidity constraints." *Econometrica*, 59(5): 1221–1248.

Deaton, Angus. 1997. *The analysis of household surveys: a microeconometric approach to development policy.* World Bank Publications.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2014. "Consumption and labor supply with partial insurance: An analytical framework." *American Economic Review*, 104(7): 2075–2126.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2017. "Optimal tax progressivity: An analytical framework." *The Quarterly Journal of Economics*, 132(4): 1693–1754.

Kaplan, Greg, and Giovanni L Violante. 2010. "How much consumption insurance beyond self-insurance?" *American Economic Journal: Macroeconomics*, 2(4): 53–87.

Wu, Chunzan, and Dirk Krueger. 2021. "Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation." *American Economic Journal: Macroeconomics*, 13(1): 79–113.