Multi-Good Consumption Dynamics: Necessities and Frictions

Jim Been* and Eduard Suari-Andreu†

January 23, 2026

Abstract

We develop and estimate a multi-good extension of the permanent–transitory income model, allowing consumption categories to differ in preferences and adjustment frictions. The data reveal large cross-category heterogeneity in responses to permanent income shocks—variation obscured by standard one-good formulations. This heterogeneity implies substantial dispersion in welfare costs and shows that food-based proxies drastically understate exposure to long-run income risk. Our framework separates households' *desired* adjustment from their *feasible* adjustment across goods, providing a richer mapping from income risk to consumption behavior.

JEL Codes: E21, D12, D31, H31, J22.

Keywords: Consumption, Dynamics, Income Shocks, Necessities, Adjustment

*Department of Economics at Leiden University and Netspar. E-mail: j.been@law.leidenuniv.nl
†Department of Economics at Leiden University and Netspar. E-mail:e.suari.andreu@law.leidenuniv.nl

1 Introduction

Understanding how households respond to income risk is a central question in empirical economics. A large literature documents that income varies substantially across individuals and over the life cycle, and that these fluctuations shape consumption behavior, inequality, and welfare (e.g. Meghir and Pistaferri, 2004; Kaplan and Violante, 2010; Blundell et al., 2016; Guvenen et al., 2014, 2021; Heathcote et al., 2020). Yet we lack a systematic understanding of how income shocks propagate into the *composition* of household spending. Most empirical approaches treat consumption as a single aggregate good—often proxied by food expenditures—thereby imposing a strong and rarely justified assumption that the welfare consequences of income shocks can be summarized by a one-dimensional response.

This simplification is convenient, but it obscures economically meaningful heterogeneity. Classical work by Deaton (1992) and Ziliak (1998) emphasizes that aggregation across heterogeneous goods can distort inference about both preferences and the income process. Prominent applied studies—such as Stephens (2004), Aguiar and Hurst (2005), and Hendren (2017)—use food expenditures as a proxy for total consumption; however, food is only a small and unusually smooth component of household spending. As Chetty and Szeidl (2007) show, ignoring within-period substitution and category-specific elasticities can bias inferences about risk exposure, precautionary motives, and welfare. Likewise, Aguiar and Hurst (2013) document large differences across nondurable spending categories over the life cycle, suggesting that aggregation masks important behavioral margins.

This paper extends the classic permanent–transitory income framework of Blundell et al. (2008)—itself rooted in the canonical analyses of Hall (1978), Hall and Mishkin (1982), and Flavin (1981)—to a multi-good environment in which each consumption category has its own loading on permanent and transitory income shocks. Using high-quality longitudinal microdata on disaggregated household expenditures, we estimate category-specific exposures to permanent and transitory income shocks. This approach embeds disaggregated consumption choices directly into a structural income-dynamics model and provides a unified framework

for measuring consumption smoothing across all components of household budgets.

Our methodological contribution is twofold. First, we show that interpreting income risk through the lens of a single nondurable good is deeply misleading. Food expenditures—the standard proxy for total consumption in the literature—respond only weakly to permanent income shocks. Because permanent shocks drive the bulk of welfare costs (Blundell et al., 2016), relying on food dramatically understates households' exposure to long-run income risk. Our estimates imply that welfare costs would be biased downward by nearly an order of magnitude when using food as the consumption measure, consistent with the concerns raised by Ziliak (1998) and with sizable implications for applications such as the design of optimal unemployment insurance in Hendren (2017).

Second, the multi-good structure uncovers substantial heterogeneity in consumption responses that is invisible in one-good models. Housing and other contractual or necessity goods exhibit near-zero sensitivity to permanent shocks, while discretionary categories—leisure, personal care, transport, donations, and especially insurance—display large, precisely estimated loadings. We further show that each category reflects a distinct pair of structural forces: a *desire* to adjust (captured by the preferences for a good) and an *ability* to adjust (captured by estimated adjustment frictions). This decomposition reveals that insurance expenditures, for example, combine exceptionally high long-run responsiveness with low adjustment frictions and negative short-run responses, making them a central margin of household adjustment to income risk.

Taken together, our findings reshape how income risk should be measured and interpreted. By documenting that permanent and transitory income shocks propagate through household budgets in highly uneven ways, we demonstrate that heterogeneity across consumption categories is crucial for both positive and normative analysis. The multi-good BPP framework provides a tractable and empirically grounded tool for future work on risk, inequality, household dynamics, and welfare measurement, and it offers a clear path for improving the empirical inputs used in quantitative macroeconomic models.

2 Model

2.1 Income Dynamics

Our empirical framework builds on the standard permanent–transitory decomposition of income and the partial-insurance model developed by Hall (1978), Flavin (1981), and, in its canonical modern form, by Blundell et al. (2008) (BPP). We adopt a BPP-style income process: log household income is the sum of a permanent component that follows a random walk and a purely transitory component.[1]

Let y_{it} denote log household income for individual i in period t. We write:

$$y_{it} = p_{it} + \varepsilon_{it}, \tag{1}$$

where p_{it} is permanent income and ε_{it} is a transitory shock. The permanent component follows:

$$p_{it} = p_{i,t-1} + \eta_{it}, \tag{2}$$

where η_{it} represents a permanent innovation. We assume that

$$\eta_{it} \sim i.i.d.\ (0, \sigma_\eta^2), \quad \varepsilon_{it} \sim i.i.d.\ (0, \sigma_\varepsilon^2),$$

with η_{it} and ε_{it} mutually independent and uncorrelated across households and over time. This process has become standard in the income dynamics literature (e.g. Blundell et al., 2008; Meghir, 2011; Guvenen et al., 2014; Blundell et al., 2016) and is supported by both survey and administrative data for several countries.

[1] Our methodology is easily extended to richer income dynamics. Any parametric specification that implies a tractable covariance structure for income growth—such as models with persistent transitory components, time-varying variances, or heterogeneous shocks—can be incorporated by replacing the BPP moment conditions with their generalized counterparts. Only the mapping from observed covariances of income and consumption growth to the underlying shock processes is modified.

Taking first differences yields:

$$\Delta y_{it} = y_{it} - y_{i,t-1} = \eta_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}.$$

This implies the well-known moment conditions:

$$\mathrm{Var}(\Delta y_{it}) = \sigma_\eta^2 + 2\sigma_\varepsilon^2, \tag{3}$$

$$\mathrm{Cov}(\Delta y_{it}, \Delta y_{i,t+1}) = -\sigma_\varepsilon^2. \tag{4}$$

From these, we can identify the variances of permanent and transitory income shocks:

$$\sigma_\varepsilon^2 = -\mathrm{Cov}(\Delta y_{it}, \Delta y_{i,t+1}), \tag{5}$$

$$\sigma_\eta^2 = \mathrm{Var}(\Delta y_{it}) + 2\mathrm{Cov}(\Delta y_{it}, \Delta y_{i,t+1}). \tag{6}$$

2.2 Consumption Dynamics

The standard BPP framework considers a single (log) nondurable consumption measure c_{it} and specifies consumption growth as

$$\Delta c_{it} = \phi\, \eta_{it} + \psi\, \varepsilon_{it} + \xi_{it}, \tag{7}$$

where ϕ and ψ capture the loading of consumption on permanent and transitory income shocks, and ξ_{it} is an idiosyncratic (taste) shock orthogonal to income innovations.

Our contribution is to extend this one-good framework to a setting with multiple consumption categories that match the empirical expenditure breakdown in our data. Let c_{kit} denote real expenditure in category $k \in \{1, \ldots, K\}$ for household i at time t (e.g. food at home, rent, clothing). We define log-differences for each category

$$\Delta c_{kit} \equiv \log c_{kit} - \log c_{k,i,t-1},$$

and specify the following multi-good analogue of (7):

$$\Delta c_{kit} = \phi_k \, \eta_{it} + \psi_k \, \varepsilon_{it} + \xi_{kit}. \tag{8}$$

Here:

- ϕ_k is the *loading* of category k on permanent income shocks. A value $\phi_k = 0$ corresponds to full insurance of category k against permanent shocks, while $\phi_k = 1$ corresponds to a one-for-one response of category k to a permanent shock.
- ψ_k is the *loading* of category k on transitory income shocks. Under full short-run insurance, consumption should not respond to purely transitory shocks ($\psi_k = 0$); deviations of ψ_k from zero quantify the degree of incomplete smoothing of transitory fluctuations.
- ξ_{kit} is an idiosyncratic preference or taste shock for category k, which captures movements in spending that are unrelated to income shocks.

We assume that ξ_{kit} is mean-zero, serially uncorrelated, and orthogonal to the income innovations for all i, k, t:

$$\mathbb{E}[\xi_{kit}] = 0, \tag{9}$$

$$\mathrm{Cov}(\xi_{kit}, \eta_{is}) = 0, \tag{10}$$

$$\mathrm{Cov}(\xi_{kit}, \varepsilon_{is}) = 0, \tag{11}$$

$$\mathrm{Cov}(\xi_{kit}, \xi_{k'is'}) = 0 \quad \text{for } (k,t) \neq (k',s'). \tag{12}$$

We denote the variance of the taste shock for category k by

$$Var(\xi_{kit}) = \sigma^2_{\xi,k}. \tag{13}$$

Given the income process outlined in Section 2.1 and the consumption equation (8), the key

second moments linking income and category-level consumption growth are:

$$\begin{aligned}\mathrm{Cov}(\Delta c_{kit}, \Delta y_{it}) &= \mathrm{Cov}\big(\phi_k \eta_{it} + \psi_k \varepsilon_{it} + \xi_{kit}, \eta_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}\big) \\ &= \phi_k \sigma_\eta^2 + \psi_k \sigma_\varepsilon^2, \end{aligned} \tag{14}$$

$$\begin{aligned}\mathrm{Cov}(\Delta c_{kit}, \Delta y_{i,t+1}) &= \mathrm{Cov}\big(\phi_k \eta_{it} + \psi_k \varepsilon_{it} + \xi_{kit}, \eta_{i,t+1} + \varepsilon_{i,t+1} - \varepsilon_{it}\big) \\ &= -\psi_k \sigma_\varepsilon^2, \end{aligned} \tag{15}$$

$$\begin{aligned}\mathrm{Var}(\Delta c_{kit}) &= \mathrm{Var}\big(\phi_k \eta_{it} + \psi_k \varepsilon_{it} + \xi_{kit}\big) \\ &= \phi_k^2 \sigma_\eta^2 + \psi_k^2 \sigma_\varepsilon^2 + \sigma_{\xi,k}^2. \end{aligned} \tag{16}$$

Equations (14)–(15) show that the permanent-loading ϕ_k and the transitory-loading ψ_k are identified from the covariances of category-level consumption growth with current and future income growth, once the income shock variances $(\sigma_\eta^2, \sigma_\varepsilon^2)$ have been identified from (5)–(6). Specifically,

$$\psi_k = -\frac{\mathrm{Cov}(\Delta c_{kit}, \Delta y_{i,t+1})}{\sigma_\varepsilon^2}, \tag{17}$$

$$\phi_k = \frac{\mathrm{Cov}(\Delta c_{kit}, \Delta y_{it}) + \mathrm{Cov}(\Delta c_{kit}, \Delta y_{i,t+1})}{\sigma_\eta^2}. \tag{18}$$

Given (ϕ_k, ψ_k), the variance of taste shocks is then identified from (16) as

$$\sigma_{\xi,k}^2 = \mathrm{Var}(\Delta c_{kit}) - \phi_k^2 \sigma_\eta^2 - \psi_k^2 \sigma_\varepsilon^2. \tag{19}$$

Equations (17)–(19) provide a direct multi-good extension of the BPP moment inversion. In our empirical implementation, we construct these moments from first differences of residual log income and residual log consumption for each category,[2] and solve the system implied by (5), (6), (17), (18), and (19) jointly for $(\sigma_\eta^2, \sigma_\varepsilon^2, \phi_k, \psi_k, \sigma_{\xi,k}^2)$.

[2]Here, we take into account a *predictable* part of the income process defined by age categories, gender, educational level, household size, number of children, and period-specific effects. For details of those estimates, we refer to Tables 6 and 7 in the Appendix.

2.3 Assumptions and Relation to the Literature

Our framework generalizes the one-good permanent–transitory model of Blundell et al. (2008) to a multi-good setting while remaining conceptually close to their identification strategy. The key assumptions can be summarized as follows:

1. **Permanent–transitory income structure.** Household income follows a permanent–transitory decomposition with a random-walk permanent component and an i.i.d. transitory component, as in Blundell et al. (2008) and related work (e.g. Blundell et al., 2016; Guvenen et al., 2014). This structure delivers the simple moment conditions in (5)–(6), which we exploit to identify $(\sigma_\eta^2, \sigma_\varepsilon^2)$.

2. **Orthogonality and serial independence of shocks.** Permanent shocks η_{it}, transitory shocks ε_{it}, and taste shocks ξ_{kit} are mutually orthogonal across time and households. This assumption underlies the mapping between observed covariances of income and consumption growth and the structural parameters, and follows the standard partial-insurance literature.

3. **Consumption responds to both permanent and transitory shocks.** In contrast to a pure "permanent-shocks-only" specification, we follow Blundell et al. (2008) in allowing consumption to respond to both permanent and transitory income innovations through (ϕ_k, ψ_k). Under full smoothing of transitory fluctuations, we would have $\psi_k = 0$ for all k. Deviations from zero therefore quantify incomplete insurance against transitory shocks at the category level.

4. **Idiosyncratic preference shocks are orthogonal to income.** Taste shocks ξ_{kit} capture fluctuations in category-level consumption that are unrelated to income. We assume they are orthogonal to η_{it} and ε_{it} and uncorrelated across time and categories. This ensures that $\sigma_{\xi,k}^2$ can be recovered as a residual variance as in (19).

5. **Residualization of predictable components.** The theoretical model abstracts

from predictable life-cycle patterns and observable heterogeneity. In the data, we first regress log income and log consumption in each category on age, calendar year, household composition and other controls. For details of those estimates, we refer to Tables 6 and 7 in the Appendix. All moments are computed using the resulting residuals. Under the assumption that controls span the predictable component of income and consumption, the residual-based moments are equal to the moments implied by (8).

Relative to the canonical one-good analyses of Blundell et al. (2008) and subsequent contributions, our multi-good framework allows the degree of partial insurance to vary across detailed expenditure categories. The parameters $(\phi_k, \psi_k, \sigma^2_{\xi,k})$ trace out how permanent and transitory income shocks are absorbed (or amplified) in the composition of the consumption basket, not only in its aggregate level. This richer structure is directly aligned with our empirical strategy: we implement the BPP-style moment-based identification separately for each consumption category using the same underlying income process, thereby delivering a coherent multi-good generalization of the original BPP model.

2.4 Estimation Procedure

This section describes the construction of the empirical moments, the mapping from moments to the structural parameters of the multi-good BPP model, and the computation of standard errors. The procedure extends the original single-good estimator of Blundell et al. (2008) to a setting with multiple consumption categories, permanent and transitory income shocks, and category-specific taste shocks.

Throughout, all variables are first residualized with respect to age, calendar year, household composition, and other demographic controls. The residualization step removes predictable components that are outside the BPP model and ensures that first differences identify innovations rather than systematic life-cycle movements.

Step 1: Constructing First Differences

For each household i observed in consecutive years t and $t+1$, we compute first differences of residual log income and residual log consumption in each category k:

$$\Delta y_{it} = y_{it} - y_{i,t-1}, \qquad \Delta c_{kit} = c_{kit} - c_{k,i,t-1}.$$

The theoretical counterparts of these objects are

$$\Delta y_{it} = \eta_{it} + \varepsilon_{it} - \varepsilon_{i,t-1},$$

and

$$\Delta c_{kit} = \phi_k \eta_{it} + \psi_k \varepsilon_{it} + \xi_{kit}.$$

These moment conditions form the basis for identification of both the income process $(\sigma_\eta^2, \sigma_\varepsilon^2)$ and the consumption parameters $(\phi_k, \psi_k, \sigma_{\xi,k}^2)$.

Step 2: Estimating Income-Shock Variances

We construct the variance and first-order autocovariance of Δy_{it}:

$$\widehat{\mathrm{Var}}(\Delta y) = \frac{1}{N} \sum_{i,t} (\Delta y_{it} - \overline{\Delta y})^2,$$

$$\widehat{\mathrm{Cov}}(\Delta y_{it}, \Delta y_{i,t+1}) = \frac{1}{N} \sum_{i,t} (\Delta y_{it} - \overline{\Delta y})(\Delta y_{i,t+1} - \overline{\Delta y}).$$

Under the income process in Section 2.1, the population moments satisfy

$$\mathrm{Var}(\Delta y_{it}) = \sigma_\eta^2 + 2\sigma_\varepsilon^2, \qquad \mathrm{Cov}(\Delta y_{it}, \Delta y_{i,t+1}) = -\sigma_\varepsilon^2.$$

Thus the income-shock variances are identified via the BPP inversion:

$$\widehat{\sigma}_{\varepsilon}^2 = -\,\widehat{\mathrm{Cov}}(\Delta y_{it}, \Delta y_{i,t+1}), \qquad \widehat{\sigma}_{\eta}^2 = \widehat{\mathrm{Var}}(\Delta y_{it}) + 2\,\widehat{\mathrm{Cov}}(\Delta y_{it}, \Delta y_{i,t+1}).$$

Step 3: Estimating Consumption Loadings on Permanent and Transitory Shocks

Given $(\widehat{\sigma}_{\eta}^2, \widehat{\sigma}_{\varepsilon}^2)$, we use covariances between category-level consumption growth and income growth to identify each category's exposure to permanent and transitory income shocks.

Define the empirical moments

$$\widehat{m}_{k,0} = \widehat{\mathrm{Cov}}(\Delta c_{kit}, \Delta y_{it}), \qquad \widehat{m}_{k,1} = \widehat{\mathrm{Cov}}(\Delta c_{kit}, \Delta y_{i,t+1}).$$

The theoretical moments from equations (14)–(15) imply

$$\widehat{m}_{k,1} = -\psi_k\,\widehat{\sigma}_{\varepsilon}^2, \qquad \widehat{m}_{k,0} = \phi_k \widehat{\sigma}_{\eta}^2 + \psi_k \widehat{\sigma}_{\varepsilon}^2.$$

Solving yields the BPP-style identification:

$$\widehat{\psi}_k = -\,\frac{\widehat{m}_{k,1}}{\widehat{\sigma}_{\varepsilon}^2}, \tag{20}$$

$$\widehat{\phi}_k = \frac{\widehat{m}_{k,0} + \widehat{m}_{k,1}}{\widehat{\sigma}_{\eta}^2}. \tag{21}$$

This provides a pair of structural loadings for every consumption category k.

Step 4: Recovering Taste-Shock Variances

The remaining structural parameter for each category k is the variance of taste shocks $\sigma_{\xi,k}^2$. From the variance decomposition,

$$\mathrm{Var}(\Delta c_{kit}) = \phi_k^2 \sigma_{\eta}^2 + \psi_k^2 \sigma_{\varepsilon}^2 + \sigma_{\xi,k}^2,$$

we obtain the estimator

$$\widehat{\sigma}^2_{\xi,k} = \widehat{\mathrm{Var}}(\Delta c_{kit}) - \widehat{\phi}^2_k \widehat{\sigma}^2_\eta - \widehat{\psi}^2_k \widehat{\sigma}^2_\varepsilon.$$

This step is implemented by computing the sample variance of Δc_{kit} for each k and subtracting the fitted contributions of the income shocks.

Step 5: Computing Standard Errors via Block Bootstrap

We compute standard errors using a nonparametric block bootstrap at the household level. Each bootstrap replication resamples entire household histories to preserve the temporal dependence structure of $(\Delta y_{it}, \Delta c_{kit})$. For each bootstrap sample, we:

1. reconstruct the empirical moments in Steps 2–4,
2. recompute $\widehat{\sigma}^2_\eta$ and $\widehat{\sigma}^2_\varepsilon$,
3. obtain new estimates of $(\widehat{\phi}_k, \widehat{\psi}_k, \widehat{\sigma}^2_{\xi,k})$.

The reported standard errors are the empirical standard deviations over 499 bootstrap replications.[3] This procedure is robust to serial correlation, heteroskedasticity, unbalanced panel structure, and cross-category dependence within the household.

3 Institutional Context

A central feature of the Dutch economic environment is the extensive system of public insurance that shapes both the distribution of income risk and its transmission into household consumption. This stands in sharp contrast to the institutional structure of the United

[3] We use 499 bootstrap replications to obtain stable inference while keeping computation tractable. Simulation studies show that 400–600 replications deliver reliable confidence intervals for estimators of comparable complexity (e.g. Davison and Hinkley, 1997; Efron and Tibshirani, 1994). Increasing replications further yields negligible improvements relative to the computational cost in our high-dimensional multi-equation estimation.

States—which has been studied most frequently in the context of consumption dynamics—where government-provided insurance is substantially more limited. These differences are essential for interpreting the consumption responses estimated in our study.

Recent evidence by De Nardi et al. (2021) documents earnings dynamics in both the Netherlands and the United States. They show that the propagation of primary income shocks into *household disposable income*—which is closely related to consumption—is fundamentally different across the two countries. In the Netherlands, government transfers, taxes, and income-dependent benefits substantially reduce the variance, negative skewness, and kurtosis of income changes. Disposable income is therefore much smoother than primary income. In the United States, by contrast, the public safety net plays a much smaller stabilizing role, leaving households more exposed to persistent and tail income risk.

Several institutional features contribute to the high degree of income smoothing in the Netherlands. Unemployment insurance replaces roughly 70–75% of prior earnings for up to two years,[4], two years of sickness benefits (with a 100% replacement rate in many cases) followed by disability insurance (replacement rate of 70-75%), means-tested social assistance ensures a minimum income floor (about 1,400 euros p/m), and universal health insurance eliminates out-of-pocket medical expenditure risk. Together, these programs substantially compress the distribution of disposable income shocks, particularly those in the left tail.

The Dutch system of government-provided insurance thus significantly limits the extent to which earnings risk translates into consumption volatility. These institutional aspects help explain the core empirical patterns we document in this paper.

[4]Been et al. (2024) report that job loss induces a significant but relatively minor drop of 5% in total consumption suggesting that Dutch households are largely insulated from unemployment risk in the short-run.

4 Data

4.1 LISS Panel

This study uses longitudinal microdata from the Longitudinal Internet Studies for the Social Sciences (LISS) panel, a probability-based sample of Dutch households administered by CentERdata. The LISS panel is uniquely well suited for studying the structural relationship between income risk and household consumption because it contains a rich set of annual expenditure categories reported at the household level, alongside detailed demographic and socioeconomic variables. In contrast to widely used US datasets such as the PSID or CEX, the LISS panel includes a comprehensive breakdown of household spending categories, enabling us to estimate category-specific responses to permanent income shocks.

Our analysis combines self-reported consumption data with administrative income variables, including total net household income. To ensure comparability over time, all nominal spending and income variables are deflated using the consumer price index (CPI) with base year 2015 (CPI = 100). We construct real income as

$$\text{income}_{it}^{\text{real}} = \frac{\text{income}_{it}^{nominal}}{\text{CPI}_{it}},$$

and analogously deflate each consumption category.

We focus on a set of consumption categories that span both nondurables (e.g. food, clothing), quasi-durables (e.g. transport expenditures), housing-related commitments (e.g. mortgage payments, rent, utilities), and discretionary categories (e.g. holidays, leisure, donations). This rich decomposition allows us to estimate the consumption response to income shocks across multiple margins of household adjustment, rather than relying on a single aggregate measure. We restrict our sample to households in which the household head is aged 25-60, as is common in studies of income dynamics.

4.2 Descriptive Statistics

Table 1 presents descriptive statistics and mean budget shares for the main consumption categories used in the analysis. All variables are reported in real 2015 euros. The table reveals substantial heterogeneity across both levels and composition of household expenditures. On average, households allocate the largest shares of their budgets to necessities such as food at home, mortgage or rent payments, utilities, and insurance, consistent with the importance of housing and essential services in household budgets. Transport, clothing, and leisure also represent meaningful expenditure categories but with greater variation across households. In contrast, categories such as donations, daycare, schooling, and personal care exhibit relatively low average spending and small budget shares, reflecting demographic differences in household composition as well as institutional features—such as subsidized childcare and regulated medical costs—that reduce out-of-pocket expenses. The wide dispersion in budget shares across categories highlights the importance of analyzing consumption responses at a disaggregated level rather than focusing solely on aggregate spending.

5 Results

This section presents the empirical estimates of the structural income and consumption model using Dutch longitudinal data from the LISS panel. We first discuss the estimated variances of permanent and transitory income shocks, before turning to the heterogeneous responses of consumption categories to permanent income shocks. Finally, we calculate the welfare costs from income risk.

5.1 Income Shock Estimates

We begin by discussing the estimated variances of permanent and transitory income shocks, $(\sigma_\eta^2, \sigma_\varepsilon^2)$, obtained from the covariance structure of residual income growth as described in Section 2. The estimates are reported in Table 2. Both variances are precisely estimated

Table 1: Descriptive Statistics and Mean Budget Shares

Category	Obs	Mean	Std. Dev.	Min	Max	Mean Budget Share
Total Consumption	30,979	1,995.84	1,130.38	0.93	8,335.04	—
Mortgage	28,805	352.50	438.02	0	2,046.32	0.1695
Rent	29,214	169.14	265.10	0	1,050.00	0.1133
Utilities	28,462	204.94	102.55	0	653.47	0.1281
Transport	29,143	124.22	117.79	0	700.00	0.0678
Insurance	28,193	235.12	151.75	0	905.88	0.1389
Debt Payments	29,624	23.59	76.42	0	552.85	0.0118
Cleaning	29,032	46.13	66.09	0	452.94	0.0244
Food at Home	29,455	325.27	200.20	0	1,116.17	0.1971
Daycare	29,528	19.31	98.63	0	943.33	0.0077
Holiday	22,664	126.55	192.88	0	1,500.00	0.0460
Food Out	22,636	53.26	66.79	0	436.73	0.0227
Clothing	22,617	85.96	90.75	0	624.74	0.0393
Personal Care	9,873	46.41	44.39	0	312.37	0.0094
Medical Care	30,211	45.26	79.21	0	650.00	0.0302
Schooling	22,390	8.27	30.92	0	260.31	0.0032
Leisure	13,139	62.02	70.94	0	520.62	0.0160
Donations	22,596	49.29	62.55	0	525.82	0.0258
Net income	30,979	2,759.13	4,525.29	0	315,965.30	—

Notes: All amounts are price-deflated using CPI (2015 = 100) and expressed in euros per month at the household level.

and statistically significant at the one-percent level.

Interpretation of the coefficients. In the permanent–transitory framework, σ_η^2 measures the variance of innovations to the *expected lifetime path of income*. A permanent shock η_{it} represents a revision to the household's permanent income and therefore affects the entire future income profile. Consequently, σ_η^2 captures the magnitude of long-run uncertainty faced by households. Our estimated permanent income variance of 0.032 implies that a 1-standard-deviation permanent shock raises (or lowers) household income by approximately 20 percent,[5] shifting the entire future income path.

[5] A permanent income variance of $\sigma_\eta^2 = 0.032$ means that the innovation η_{it} follows a normal distribution $\eta_{it} \sim N(0, 0.032)$. The corresponding standard deviation is $\sigma_\eta = \sqrt{0.032} \approx 0.179$, so a one–standard-deviation permanent shock raises (or lowers) log income by about 0.18, which is equivalent to a change in income of approximately $\exp(0.18) - 1 \approx 20\%$.

By contrast, σ_ε^2 measures the variance of *purely transitory fluctuations* in earnings that do not alter expectations of future income. These shocks represent short-lived deviations such as bonuses, temporary job-loss spells, or short-run hours adjustments. Because they are mean-reverting by construction, they do not enter p_{it} and have only temporary effects on resources. Our estimated transitory income variance of 0.013 implies that a 1-standard-deviation transitory shock raises (or lowers) household income by approximately 12 percent.[6] Unlike permanent shocks, these fluctuations do not alter expected future income and therefore should be largely smoothed under full insurance.

The absolute magnitude of these variances therefore carries clear economic meaning: larger values of σ_η^2 correspond to greater long-run income risk, while larger values of σ_ε^2 reflect more volatile short-run income dynamics.

Comparison with U.S. evidence. Our estimates indicate that the *permanent income variance* is sizeable yet modest when compared to U.S. evidence. In particular, the Dutch estimate of σ_η^2 is substantially smaller than the benchmark estimates reported by Blundell et al. (2008), who find permanent-variance estimates in the range of 0.05–0.10 for U.S. households using the PSID.[7] In our setting, the permanent variance is closer to the lower end of the international range. This suggests that households in the Netherlands face comparatively muted long-run income risk, consistent with institutional features of the Netherlands.

The *transitory income variance*, on the other hand, is also precisely estimated but remains somewhat smaller than U.S. counterparts. Blundell et al. (2008) report transitory-variance estimates between 0.03 and 0.05, depending on specification. Our transitory variance estimates fall noticeably below this range, again consistent with institutional features of the Netherlands.

[6] A transitory income variance of σ_ε^2 implies that the innovation $\varepsilon_{it} \sim N(0, \sigma_\varepsilon^2)$ affects income only in the current period and partially in the adjacent periods through the differenced term $\Delta y_{it} = \eta_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}$. If, for example, $\sigma_\varepsilon^2 = 0.013$, then the standard deviation of the transitory shock is $\sqrt{0.013} \approx 0.114$, implying a temporary change in log income of about 0.11, or roughly $\exp(0.11) - 1 \approx 12\%$.

[7] See also Blundell et al. (2016) and Guvenen et al. (2014) for similar U.S. estimates and cross-cohort comparisons.

Taken together, the Dutch income process displays:

1. *Lower permanent income risk* than in the United States, implying that long-horizon income uncertainty is comparatively limited.

2. *Somewhat lower transitory income volatility*, suggesting that short-run earnings fluctuations are also smoother.

This pattern is consistent with earlier cross-country studies showing that continental European economies tend to exhibit lower income volatility than the United States.[8]

Implications. The smaller values of both σ_η^2 and σ_ε^2 have two important implications for the interpretation of our consumption results in Section 5.2. First, they imply that the absolute contribution of income risk to consumption variation is smaller in the Netherlands than in the U.S. context traditionally studied in the BPP literature. Second, because permanent shocks are less volatile, the relative heterogeneity in category-specific loadings ϕ_k becomes particularly informative about the mechanisms through which households smooth income risk across different components of their consumption basket.

In summary, our results confirm the presence of both long-run and short-run income risk, but at substantially reduced magnitudes relative to the United States. This difference is economically meaningful and aligns with institutional and labor-market features distinguishing the Dutch economy from the U.S. benchmark of Blundell et al. (2008).

5.2 Consumption Responses to Permanent Income Shocks

Table 3 reports the estimated loading of each consumption category on permanent income shocks (ϕ_k), transitory income shocks (ψ_k), and the variance of category-specific taste shocks ($\sigma^2_{\xi,k}$). Three patterns stand out: (i) substantial heterogeneity across categories in the degree

[8] See, for example, Venn (2011) (OECD), Pistolesi (2014) (France), Arellano et al. (2017) (Denmark), Ramos (2003); Blundell and Etheridge (2010); Cappellari and Jenkins (2014) (UK), and De Nardi et al. (2021) (Netherlands).

Table 2: Estimated Income Shock Parameters

	Estimate	Standard Error
Permanent income variance σ^2_η	0.0322***	0.0020
Transitory income variance σ^2_ε	0.0128***	0.0008

Notes: Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels.

of pass-through to permanent income shocks, (ii) limited and largely insignificant responses to transitory shocks, and (iii) pronounced dispersion in the idiosyncratic (taste) volatility across consumption categories. Each of these patterns provides insight into the mechanisms through which households smooth income risk and adjust the composition of their spending.

Responses to permanent income shocks. Permanent income shocks shift the expected lifetime budget constraint. In theory, forward-looking households should adjust expenditures in categories whose marginal utility responds strongly to long-term changes in resources. Our estimates reveal substantial heterogeneity in ϕ_k, ranging from near-zero for housing categories (mortgage, rent, utilities) to sizable elasticities for non-essential or discretionary categories such as leisure ($\phi_k = 0.480$), insurance ($\phi_k = 0.432$), transport ($\phi_k = 0.318$), and personal care ($\phi_k = 0.311$).

This pattern is economically intuitive. Housing expenditures in the Netherlands are largely predetermined in the short run due to long-term mortgage contracts, rental agreements, and regulatory constraints; such institutional rigidities curb the ability of households to adjust these components in response to permanent income revisions. By contrast, categories with a higher discretionary margin—recreation, holidays, out-of-home food consumption, and personal services—provide direct scope for adjustment and therefore exhibit strong pass-through. These responses are consistent with models of non-homothetic preferences, where luxury or quality-intensive goods react proportionally more to long-run changes in lifetime resources.

Responses to transitory income shocks. Transitory shocks should, under full insurance, have little effect on consumption. We indeed find that most ψ_k estimates are small and statistically insignificant. The exception—payments on insurances—likely reflect that temporary income declines increase short-run uncertainty thereby raising the perceived value of protection and reinforcing precautionary motives. Still, the broad pattern aligns with the permanent-income hypothesis: category-level consumption is well insulated from short-run income fluctuations.

Taste-shock heterogeneity. The estimated taste-shifter variances $\sigma^2_{\xi,k}$ exhibit large dispersion across categories, with particularly high values for medical spending, schooling, holidays, and leisure. These categories represent domains in which household-specific events (health shocks, educational needs, accumulated work stress) generate substantial idiosyncratic volatility that is orthogonal to income shocks. Conversely, highly routine and predictable categories such as mortgage payments and utilities display relatively low taste-shock variances.

Why heterogeneity matters. The heterogeneity we uncover is not merely descriptive but has important implications for economic behavior and welfare. First, variation in ϕ_k directly informs which components of the consumption bundle adjust to long-run income risk and therefore carry the welfare consequences of imperfect insurance. Categories with high permanent-loadings act as the primary "adjustment margin" in smoothing income shocks, while rigid categories (e.g. housing) amplify the sensitivity of other categories.

Second, heterogeneity in adjustment elasticities is central to understanding the effective curvature of the household utility function over different goods. As shown by Chetty and Szeidl (2007), differences in category-level elasticities encode information about the structure of household preferences and the degree of within-period substitutability across consumption goods. Our findings imply that households smooth income shocks not through uniform proportional adjustments but by reallocating spending toward and away from specific categories,

consistent with multi-good models of partial insurance.

Moreover, our results clearly indicate that food spending may not be representative for an approximation of total consumption as the loading of food spending (both at home and away) on permanent income shocks are much lower than the loading of total income on permanent income shocks: $\phi_{food} << \phi_{total}$. This implies that studies that use food spending as a proxy for total spending are likely to severely underestimate the consequences of income shocks and overestimate the degree of (partial) insurance from income shocks (Hendren, 2017).

Finally, heterogeneity in $\sigma^2_{\xi,k}$ matters because idiosyncratic volatility is a major driver of consumption inequality across households. Categories with high taste-shock variance—medical care, schooling, leisure—contribute disproportionately to cross-sectional dispersion even in the absence of income risk. This implies that policy interventions targeting these categories (e.g. education subsidies, health insurance design, or childcare support) may materially affect consumption inequality and risk exposure.

5.3 Welfare Costs of Income Risk: A Numerical Example

To quantify the welfare cost of income risk, we follow the standard BPP approach and consider a CRRA utility function (where the CRRA is measured by γ)[9] over a composite consumption index C_t,

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}, \qquad \gamma = 2.$$

Assume that (log) total consumption is approximately normally distributed,

$$\log C_t = \mu_C + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_C^2),$$

[9] $\gamma = 2$ is widely used in quantitative work (e.g. Blundell and Stoker, 1999; Carroll and Samwick, 1997).

Table 3: Category-Level Responses to Permanent and Transitory Income Shocks

Category	ϕ_k	SE(ϕ_k)	ψ_k	SE(ψ_k)	$\sigma^2_{\xi,k}$	SE($\sigma^2_{\xi,k}$)
Total consumption	0.251***	0.055	0.098	0.239	0.382***	0.014
Mortgage	-0.002	0.033	0.237	0.168	0.138***	0.008
Rent	0.072	0.068	-0.117	0.195	0.151***	0.019
Utilities	0.013	0.048	0.045	0.150	0.228***	0.006
Transport	0.318***	0.071	-0.096	0.301	0.676***	0.013
Insurance	0.432***	0.063	-0.660**	0.274	0.631***	0.012
Debt payments	-0.184**	0.088	0.387	0.304	0.440***	0.020
Cleaning	0.224**	0.091	0.503	0.415	0.999***	0.016
Food at home	0.008	0.055	0.214	0.308	0.431***	0.010
Food away from home	0.141*	0.074	0.199	0.338	0.864***	0.016
Clothing	0.022	0.062	0.254	0.220	0.696***	0.014
Personal care	0.311***	0.103	0.083	0.429	0.738***	0.018
Medical care	0.197*	0.118	-0.334	0.494	1.576***	0.026
Daycare	-0.256*	0.140	0.493	0.709	0.647***	0.042
Holiday	0.263**	0.103	0.004	0.412	0.900***	0.021
Schooling	0.185	0.118	0.153	0.506	1.418***	0.041
Leisure	0.480***	0.114	-0.057	0.535	0.889***	0.021
Donations	0.268***	0.079	-0.024	0.333	0.841***	0.015

Notes: Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. Permanent shock loading = ϕ_k, transitory shock loading = ψ_k, taste-shifter variance = $\sigma^2_{\xi,k}$.

where σ^2_C is the variance of log consumption induced by permanent income shocks.[10] Under lognormality, expected utility is

$$\mathbb{E}[u(C_t)] = \frac{1}{1-\gamma} \exp\big((1-\gamma)\mu_C + \tfrac{1}{2}(1-\gamma)^2\sigma^2_C\big)\,.$$

The consumption-equivalent variation (CEV) λ is defined as the proportional reduction in deterministic consumption that leaves the household indifferent between a risky environment and a risk-free one:

$$\mathbb{E}[u(C_t)] = u\left((1-\lambda)\,\bar{C}\right),$$

[10] Consistent with our estimates, we abstract from transitory shocks.

where $\bar{C} = \exp(\mu_C)$ is the mean consumption level. Under CRRA and lognormality, this yields the closed-form CEV

$$1 - \lambda = \exp\left(\tfrac{1}{2}(1-\gamma)\sigma_C^2\right), \qquad \lambda = 1 - \exp\left(\tfrac{1}{2}(1-\gamma)\sigma_C^2\right).$$

In our setting, we can relate σ_C^2 to the underlying income risk using the estimated loading of total consumption on permanent income shocks, ϕ_{total}, and the permanent income variance, σ_η^2. Ignoring any additional sources of consumption risk, we approximate

$$\sigma_{C,\text{total}}^2 \approx \phi_{\text{total}}^2 \sigma_\eta^2.$$

With our final estimates, $\phi_{\text{total}} = 0.251$ and $\sigma_\eta^2 = 0.0322$, this implies

$$\sigma_{C,\text{total}}^2 \approx 0.251^2 \times 0.0322 \approx 0.0020.$$

Plugging this into the CEV formula with $\gamma = 2$,

$$\lambda_{\text{total}} = 1 - \exp\left(-\tfrac{1}{2}\sigma_{C,\text{total}}^2\right) \approx 1 - \exp(-0.0010) \approx 0.0010,$$

so the welfare cost of permanent income risk in the one-good benchmark is about 0.10% of consumption. Put differently, a representative household would be willing to give up roughly one-tenth of one percent of its deterministic consumption each year to eliminate the permanent income risk that we estimate for the Netherlands.

A Multi-Good Welfare Framework

The one-good calculation abstracts from the rich heterogeneity we have documented across consumption categories. Households in our data choose a vector of goods $(c_{1t}, \ldots, c_{Kt})$, and the pass-through of permanent income shocks differs sharply by category, as captured by the ϕ_k parameters in Table 3. To highlight the welfare implications of this heterogeneity, we

consider a simple multi-good extension in which the welfare cost of income risk is computed separately for each category.

For category k, we approximate the variance of log consumption due to permanent income risk by

$$\sigma^2_{C,k} \approx \phi^2_k \sigma^2_\eta,$$

where σ^2_η is the common permanent income variance estimated in Section 5.1. Under CRRA utility with $\gamma = 2$ and lognormality,[11] the category-specific CEV is then

$$\lambda_k = 1 - \exp\left(-\tfrac{1}{2}\phi^2_k \sigma^2_\eta\right), \qquad \gamma = 2. \tag{22}$$

This formula mirrors the one-good case, but now applies to the volatility of each individual consumption category induced by permanent income shocks.

Empirical CEVs by Category

Table 4 reports the empirical CEVs by category, computed from the final estimates of ϕ_k in Table 3, the permanent variance $\sigma^2_\eta = 0.0322$, and equation (22). The table also reports the implied variance $\sigma^2_{C,k} = \phi^2_k \sigma^2_\eta$ of log consumption in each category due to permanent income risk.

Several patterns emerge. First, the welfare cost of permanent income risk for *total* consumption is modest at about 0.10%, only slightly larger than the illustrative one-good calculation above. Second, the cross-category dispersion in welfare costs is substantial. Leisure (0.37%), insurance (0.30%), transport (0.16%), personal care (0.16%), donations (0.12%), holidays (0.11%), and daycare (0.11%) exhibit welfare losses that are roughly two to three times larger than the one-good benchmark. By contrast, essential and rigid categories such as mortgage payments, rent, utilities, and food at home display negligible CEVs: their effective exposure to permanent income risk is extremely low.

[11] $\gamma = 2$ is widely used in quantitative work (e.g. Blundell and Stoker, 1999; Carroll and Samwick, 1997).

Table 4: Consumption-Equivalent Variation (CEV λ_k) by Consumption Category

Category	ϕ_k	$\sigma^2_{C,k} = \phi_k^2 \sigma_\eta^2$	λ_k (%)
Leisure	0.480	0.00741	0.370
Insurance	0.432	0.00599	0.299
Transport	0.318	0.00325	0.162
Personal care	0.311	0.00311	0.156
Donations	0.268	0.00231	0.115
Holiday	0.263	0.00223	0.112
Daycare	-0.256	0.00211	0.106
Total consumption	0.251	0.00203	0.101
Cleaning	0.224	0.00161	0.080
Medical care	0.197	0.00125	0.062
Schooling	0.185	0.00110	0.055
Debt payments	-0.184	0.00108	0.054
Food away from home	0.141	0.00064	0.032
Rent	0.072	0.00017	0.008
Clothing	0.022	0.00002	0.001
Utilities	0.013	0.00001	0.000
Food at home	0.008	0.00000	0.000
Mortgage	-0.002	0.00000	0.000

Notes: CEVs are computed under CRRA utility with $\gamma = 2$, using $\sigma_\eta^2 = 0.0322$ and the category-specific loadings ϕ_k from Table 3. λ_k is expressed as a percentage of deterministic consumption in category k.

Comparing One-Good and Multi-Good Welfare Costs

Two findings stand out when comparing the one-good and multi-good perspectives:

1. The one-good CEV of roughly 0.10% *masks substantial heterogeneity* in how income risk affects different components of the consumption bundle. Several categories—especially

leisure and insurance—carry welfare costs that are multiples of the aggregate estimate.

2. The distribution of welfare losses across categories is highly uneven. Rigid, necessity-like categories (housing, utilities, basic food) contribute little to welfare costs from income risk, whereas discretionary and quality-intensive categories (leisure, holidays, personal services) bear a disproportionate share of the adjustment. These categories effectively act as "shock absorbers" in the household budget.

This heterogeneity is important for welfare analysis. It implies that policies aimed at stabilizing disposable income or insuring specific expenditure risks (e.g. health, childcare, unemployment) may have effects that are highly non-uniform across the consumption basket.

Size of Welfare Costs

Our baseline CEV for permanent income risk—about 0.10 percent of consumption for the one-good case, and at most 0.37 percent across categories in the multi-good model—is small relative to the welfare costs typically found in calibrated heterogeneous-agent models. For example, Storesletten et al. (2004) estimate welfare losses from incomplete markets between 1 and 3 percent of lifetime consumption, while Krebs (2007) reports welfare costs of roughly 1 percent for plausible calibrations of permanent productivity shocks.

Our relatively modest numbers are consistent with the lower income risk documented for the Netherlands in Section 5.1. At the same time, the multi-good perspective shows that even small aggregate welfare costs can conceal sizeable losses in specific categories that play a central role in household well-being (such as leisure and insurance), underscoring the importance of modeling the full consumption basket rather than a single nondurable aggregate.

6 A Multi-Good Environment with Good-Specific Preferences and Frictions

To give a more structural meaning to our empirical findings, we embed the estimated category-level responses to permanent income shocks into a parsimonious multi-good environment. We allow this environment to have good-specific preferences–in particular the necessity of the good–and frictions–in terms of adjustment costs. In this way, we take a more formal approach in identifying goods that households *want* to adjust versus goods that households *can* adjust.

We assume the following multi-good environment. Households derive period utility from a CES aggregate of category-level consumptions $\{c_{kt}\}_{k=1}^{K}$,

$$U_t = \frac{C_t^{1-\gamma}}{1-\gamma}, \qquad C_t = \left(\sum_{k=1}^{K} \theta_k^{\frac{1}{\sigma}} c_{kt}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{23}$$

where γ is the coefficient of relative risk aversion, σ is the intratemporal elasticity of substitution across categories, and $\theta_k > 0$ is a category-specific "luxury" parameter. A higher θ_k raises the marginal utility weight of good k and, holding budget shares fixed, makes that category more income-elastic in the sense of Chetty and Szeidl (2007).[12]

On the dynamic margin, we allow for convex adjustment costs in the growth rate of each category,

$$\mathcal{A}_t = \sum_{k=1}^{K} \frac{\kappa_k}{2} \left(\Delta \log c_{kt} \right)^2, \tag{24}$$

where $\kappa_k \geq 0$ indexes the tightness of contracts, habits, or other frictions that make category k difficult to adjust in response to shocks. In the absence of adjustment costs ($\kappa_k = 0$), first-order conditions for intertemporal optimization yield category-level responses to a permanent income shock that are governed by preferences alone; with $\kappa_k > 0$, the optimal response of

[12] Equation (23) is a Stone–Geary-type CES aggregator once we reinterpret θ_k as capturing both taste heterogeneity and necessity-versus-luxury differences across goods.

$\Delta \log c_{kt}$ is attenuated. In a log-linear approximation around the steady state, the loading of category k on a permanent income shock can be written schematically as

$$\phi_k \approx \frac{\tilde{\phi}_k(\theta_k, \sigma)}{1 + \gamma\sigma^2\kappa_k}, \tag{25}$$

where $\tilde{\phi}_k(\theta_k, \sigma)$ is the frictionless (static) response implied by the CES aggregator and $\gamma\sigma^2\kappa_k$ summarizes the dynamic friction. This expression clarifies the economic roles of our two key parameters: θ_k tells us *where households want to adjust* in response to permanent income shocks (the frictionless margin), while κ_k tells us *where they are able to adjust* given contractual and other adjustment frictions.

Calibration of θ_k and κ_k. We calibrate the luxury parameters θ_k and adjustment-cost parameters κ_k to match (i) observed budget shares s_k from our descriptive statistics, and (ii) the estimated pass-through coefficients ϕ_k from Table 3. For $\sigma = 0.5$ and $\gamma = 2$,[13] the frictionless income elasticity of category k with respect to a permanent income shock is

$$E_k^{\text{fricless}}(\theta_k, \sigma) \equiv \frac{\partial \log c_{kt}}{\partial \log Y_t} = \frac{\theta_k^{1/\sigma} c_{kt}^{(\sigma-1)/\sigma}}{\sum_j \theta_j^{1/\sigma} c_{jt}^{(\sigma-1)/\sigma}},$$

which, evaluated at the steady-state budget shares $s_k = p_k c_{kt}/Y_t$, implies a one-to-one mapping between (θ_k, σ) and $(E_k^{\text{fricless}}, s_k)$. Using our estimate for total consumption, $\phi_{\text{total}} = 0.251$, we interpret

$$E_k \equiv \frac{\phi_k}{\phi_{\text{total}}}$$

as a category-specific income elasticity of permanent shocks. Given (E_k, s_k) and $\sigma = 0.5$, we invert the CES demand system to obtain θ_k for each category. Table 5 reports the resulting calibration.

[13] We follow much of the empirical consumption literature in setting $\gamma = 2$, a value widely considered in quantitative work (e.g. Blundell and Stoker, 1999; Carroll and Samwick, 1997). Our benchmark calibration $\sigma = 0.5$ lies well within the range of empirical estimates for nondurable consumption (e.g. Aguiar and Hurst, 2013)., which typically fall between 0 and 1 (e.g. Selvanathan et al., 2025).

Dynamic adjustment costs are then backed out from the observed attenuation of ϕ_k relative to its frictionless benchmark. Under the log-linear approximation in (25), and taking $\tilde{\phi}_k \approx E_k \phi_{\text{total}}$, we calibrate

$$\kappa_k = \frac{1}{\gamma \sigma^2} \left(\frac{\tilde{\phi}_k}{\phi_k} - 1 \right), \tag{26}$$

for all categories with $\phi_k > 0$.[14] Intuitively, κ_k is decreasing in ϕ_k: categories that move strongly with permanent income shocks must have low adjustment frictions, while categories that barely move require large κ_k to reconcile their muted response with a non-trivial frictionless elasticity.

Table 5 summarizes the calibrated (θ_k, κ_k) for $\sigma = 0.5$ and $\gamma = 2$.

Table 5: Calibration of Luxury and Adjustment Parameters by Category

Category	Budget share s_k	ϕ_k	Income elasticity E_k	Luxury θ_k	Adj. cost κ_k
Mortgage	0.170	-0.002	-0.008	0.023	∞
Rent	0.113	0.072	0.286	0.027	25.822
Utilities	0.128	0.013	0.050	0.019	156.558
Transport	0.068	0.318	1.266	0.115	4.293
Insurance	0.139	0.432	1.718	0.585	2.635
Debt payments	0.012	-0.184	-0.731	0.000	∞
Cleaning	0.024	0.224	0.890	0.020	6.945
Food at home	0.197	0.008	0.030	0.028	261.390
Daycare	0.008	-0.256	-1.021	0.000	∞
Holiday	0.046	0.263	1.049	0.051	5.591
Food away from home	0.023	0.141	0.561	0.009	12.185
Clothing	0.039	0.022	0.088	0.006	88.824
Personal care	0.009	0.311	1.239	0.015	4.426
Medical care	0.030	0.197	0.785	0.020	8.147
Schooling	0.003	0.185	0.737	0.002	8.811
Leisure	0.016	0.480	1.912	0.099	2.165
Donations	0.026	0.268	1.067	0.029	5.461

Notes: Budget shares s_k are from the descriptive statistics in Table 1. The pass-through coefficients ϕ_k are from Table 3. The income elasticity $E_k = \phi_k / \phi_{\text{total}}$ uses $\phi_{\text{total}} = 0.251$. Luxury parameters θ_k are inverted from the CES demand system with $\sigma = 0.5$. Adjustment-cost parameters κ_k are calibrated from equation (26) with $\gamma = 2$ and $\sigma = 0.5$.

[14] For categories with $\phi_k \leq 0$ the linear approximation in (25) is less informative; we treat these as having effectively infinite adjustment costs in the calibration.

Where households want and are able to adjust. Figure 1 plots the calibrated (θ_k, κ_k) pairs for all categories with $\phi_k > 0$. The horizontal axis measures the luxury parameter θ_k, while the vertical axis (on a log scale) measures the adjustment cost κ_k.

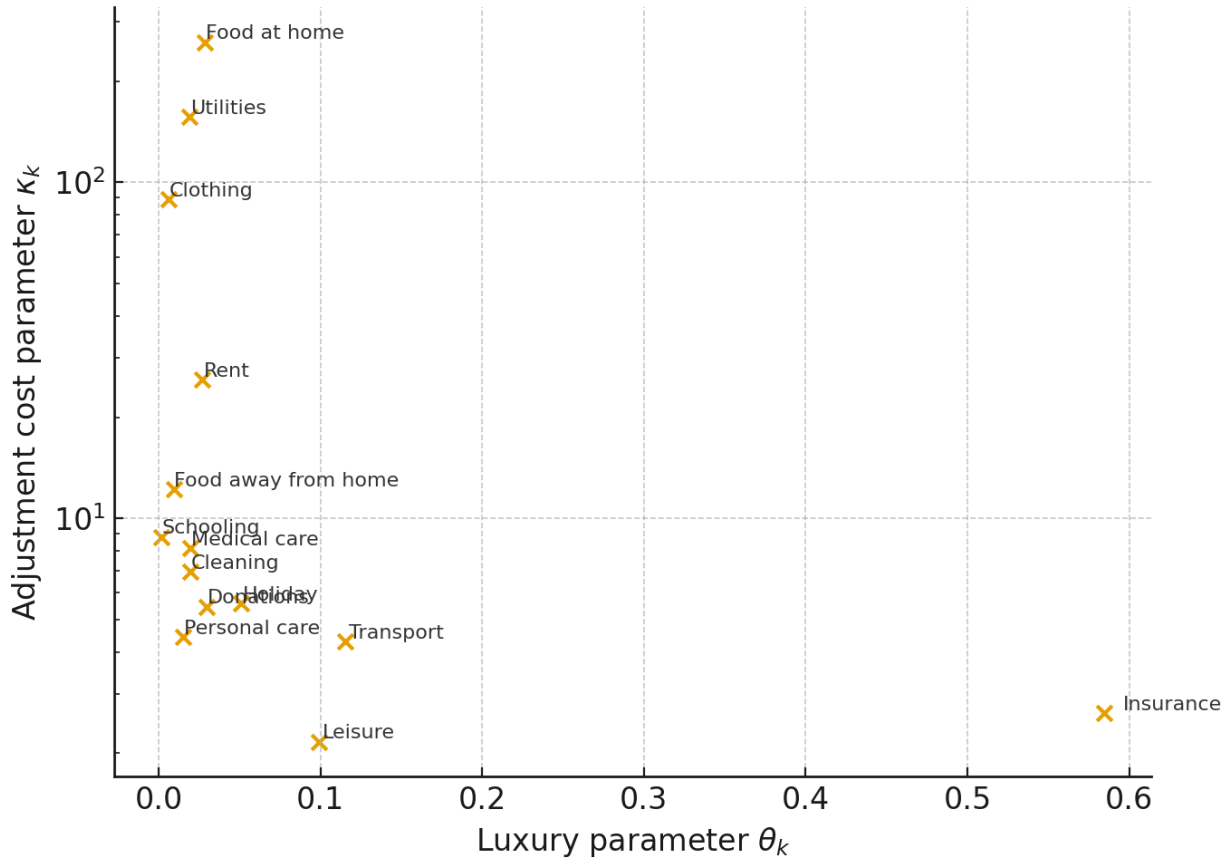


Figure 1: Luxury Parameters and Adjustment Costs by Category

Notes: Categories in the upper-left quadrant (high κ_k, low θ_k) are "hard-to-adjust necessities" (e.g. food at home, utilities), while those in the lower-right quadrant (high θ_k, low κ_k) are "easy-to-adjust luxuries" (e.g. leisure, insurance, transport).

Two patterns are striking. First, goods such as food at home, utilities, and clothing cluster in the region with very low θ_k and very high κ_k, indicating that these categories are both necessity-like in preferences and subject to substantial adjustment frictions, consistent with contractual arrangements and subsistence needs. Second, leisure, insurance, and transport lie in the opposite region, with high luxury parameters and low adjustment costs. These are precisely the categories in which we observe the largest ϕ_k in the data: households both

want to adjust these goods when permanent income changes and are *able* to do so.

The role of insurances. Our calibration places insurance expenditures in a distinctive position within household consumption dynamics. First, insurance displays the highest luxury parameter θ_k among all categories, implying that in the frictionless benchmark households *desire* to adjust insurance spending strongly in response to permanent income innovations. This is consistent with the idea that perceptions of background risk, precautionary motives, and financial slack co-move with long-run income. Second, insurance also exhibits one of the lowest adjustment-cost parameters κ_k, indicating that households can adjust insurance outlays more flexibly than contractual or necessity goods. The combination of a high desired sensitivity (θ_k) and low friction (κ_k) produces the large permanent-shock loading we estimate in the data. Third, the negative loading on transitory shocks suggests an intertemporal substitution mechanism: when short-run disposable resources tighten, households temporarily reduce insurance premia, even though long-run preferences favor maintaining coverage. Taken together, these patterns identify insurance as a central margin of adjustment. Because insurance products directly intermediate financial risk, heterogeneity in insurance responsiveness has first-order implications for welfare losses from income volatility and for the design of policies that subsidize, mandate, or otherwise influence insurance participation.

7 Conclusion

This paper extends the canonical permanent–transitory framework of Blundell et al. (2008) by embedding disaggregated consumption directly into a structural model of income and consumption dynamics. Using high-quality longitudinal expenditure data, we estimate category-specific loadings of consumption on permanent and transitory income shocks, together with the dispersion of preference shocks at the category level. This generalized multi-good approach unlocks a set of insights that cannot be obtained from the one-good representations that dominate the empirical literature (e.g. Stephens, 2004; Aguiar and Hurst, 2005; Ziliak,

1998; Hendren, 2017) and provides a sharper lens on household risk exposure and welfare.

Three substantive findings emerge. First, income risk is considerably smaller than that typically estimated for the United States. The variance of permanent shocks, $\sigma_\eta^2 = 0.032$, and the variance of transitory shocks are both close to half of U.S. benchmark estimates (e.g. Meghir and Pistaferri, 2004; Guvenen et al., 2014; Blundell et al., 2016). This pattern is consistent with institutional features such as the generosity and coverage of Dutch social insurance programs, and it complements recent administrative evidence documenting substantial cross-country heterogeneity in income dynamics (Guvenen et al., 2021).

Second, we uncover pronounced heterogeneity in the pass-through of permanent income shocks across consumption categories. Contractual and necessity goods—such as housing, utilities, and food at home—are effectively unresponsive to long-run income movements. In contrast, discretionary categories—especially leisure, insurance, personal care, donations, and transport—respond strongly to the same shocks. These results provide structural evidence for the type of within-period substitution emphasized by Chetty and Szeidl (2007) and highlight the limitations of studies that proxy total consumption with a single subcomponent, most notably food spending (Stephens, 2004; Aguiar and Hurst, 2005; Hendren, 2017). Our estimates show that food spending understates permanent-income sensitivity by an order of magnitude, implying that prior welfare analyses relying on food as a proxy for nondurable consumption likely suffer from substantial downward bias.

Third, translating category-level elasticities into welfare reveals similarly large dispersion. Although the aggregate consumption-equivalent cost of income risk is modest—about 0.10% of consumption under conventional preference parameters—several categories exhibit welfare losses two to three times larger. Leisure and insurance, for example, generate welfare costs of roughly 0.37% and 0.30%, respectively, while necessity goods contribute almost nothing. These findings show that welfare costs are concentrated precisely in the goods where both desired responsiveness and adjustment capacity are high—something a one-good model cannot reveal.

Taken together, our results demonstrate that moving beyond a single composite consumption measure is essential for accurately characterizing how households absorb income risk and how income fluctuations translate into welfare losses. By structurally integrating disaggregated expenditure into the permanent–transitory model, we provide a framework capable of quantifying both where households *want* to adjust consumption and where they are *able* to do so. This distinction matters for positive and normative questions alike, and it opens the door to a new class of models and empirical strategies in which income risk, consumption composition, and welfare are analyzed jointly rather than through a single aggregate index.

References

Aguiar, M. and Hurst, E. (2005). Consumption versus expenditure. *Journal of Political Economy*, 113(5):919–948.

Aguiar, M. and Hurst, E. (2013). Deconstructing life-cycle expenditure. *Journal of Political Economy*, 121(3):437–492.

Arellano, M., Blundell, R., and Bonhomme, S. (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica*, 85(3):693–734.

Been, J., Suari-Andreu, E., Knoef, M., and Alessie, R. J. M. (2024). Consumption and time use responses to unemployment: Implications for the lifecycle model. *Economica*, 91(361):1–32.

Blundell, R. and Etheridge, B. (2010). Consumption, income and earnings inequality in britain. *Review of Economic Dynamics*, 13(1):76–102.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption inequality and family labor supply. *American Economic Review*, 106(2):387–435.

Blundell, R. and Stoker, T. M. (1999). Consumption and the timing of income. *European Economic Review*, 43:475–507.

Cappellari, L. and Jenkins, S. P. (2014). Earnings and labour market volatility in britain, with a transatlantic comparison. *Labour Economics*, 30:201–211.

Carroll, C. D. and Samwick, A. A. (1997). The nature of precautionary wealth. *Journal of Monetary Economics*, 40(1):41–71.

Chetty, R. and Szeidl, A. (2007). Consumption commitments and risk preferences. *Quarterly Journal of Economics*, 122(2):831–877.

Davison, A. and Hinkley, D. (1997). *Bootstrap Methods and their Application*. Cambridge University Press.

De Nardi, M., Fella, G., Knoef, M., Paz-Pardo, G., and Van Ooijen, R. (2021). Family and government insurance: Wage, earnings, and income risks in the Netherlands and the U.S. *Journal of Public Economics*, 193:104323.

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

Efron, B. and Tibshirani, R. (1994). *An Introduction to the Bootstrap*. CRC Press.

Flavin, M. (1981). The adjustment of consumption to changing expectations about future income. *Journal of Political Economy*, 89(5):974–1009.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of US workers reveal about life-cycle earnings risk? *Econometrica*, 89(5):2303–2339.

Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6):971–987.

Hall, R. E. and Mishkin, F. S. (1982). The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica*, 50(2):461–481.

Heathcote, J., Storesletten, K., and Violante, G. L. (2020). Wealth inequality, earnings dynamics, and heterogeneity. *Annual Review of Economics*, 12:67–100.

Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review*, 107(7):1778–1823.

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

Krebs, T. (2007). Job displacement risk and the cost of business cycles. *American Economic Review*, 97(3):664–686.

Meghir, C. (2011). Earnings, consumption and life cycle choices. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics, Volume 4B*, pages 185–265. Elsevier, Amsterdam / Oxford.

Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32.

Pistolesi, N. (2014). Income and consumption risk: Evidence from France. *Annales d'Économie et de Statistique*, 113/114:347–377.

Ramos, X. (2003). The covariance structure of earnings in great britain, 1991–1999. *Economica*, 70(278):353–374.

Selvanathan, E. A., Clements, K. W., and Selvanathan, S. (2025). Estimating the elasticity of substitution: An index-number approach. *Journal of International Trade and Economic Development*, 34(6):1378–1393.

Stephens, M. (2004). Job loss expectations, realizations, and household consumption behavior. *Review of Economics and Statistics*, 86(1):253–269.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633.

Venn, D. (2011). Earnings volatility and its consequences for households. Technical Report 125, OECD. DELSA/ELSA/WD/SEM(2011)10.

Ziliak, J. P. (1998). Does the choice of consumption measure matter? An application to the permanent-income hypothesis. *Journal of Monetary Economics*, 41(1):201–216.

Appendix A: Income regression

Table 6: Determinants of Log (Net) Household Income (p/m).

Variable	Coefficient	Std. Error	95% CI
Age group (ref: age group 2)			
Age 3	0.009	0.014	[-0.017; 0.036]
Age 4	0.015	0.013	[-0.012; 0.041]
Age 5	0.039**	0.013	[0.013; 0.064]
Age 6	0.061***	0.013	[0.036; 0.087]
Age 7	0.102***	0.013	[0.077; 0.127]
Age 8	0.103***	0.013	[0.079; 0.128]
Age 9	0.063***	0.019	[0.026; 0.100]
Age 10	-0.025	0.082	[-0.185; 0.136]
Age 11	0.282	0.240	[-0.187; 0.752]
Age 14	0.386*	0.208	[-0.021; 0.793]
Year (ref: 2009)			
2010	-0.043***	0.012	[-0.067; -0.019]
2012	-0.104***	0.013	[-0.128; -0.079]
2015	-0.187***	0.012	[-0.211; -0.163]
2017	-0.164***	0.013	[-0.189; -0.138]
2019	-0.179***	0.014	[-0.206; -0.152]
2020	-0.174***	0.013	[-0.201; -0.148]
2021a	-0.170***	0.013	[-0.196; -0.144]
2021b	-0.178***	0.013	[-0.203; -0.152]
Education (ref: primary/none)			
VMBO	0.017	0.017	[-0.017; 0.050]
HAVO/VWO	0.187***	0.019	[0.151; 0.223]
MBO	0.108***	0.016	[0.076; 0.140]
HBO	0.296***	0.016	[0.264; 0.328]
WO	0.415***	0.017	[0.381; 0.449]
Household characteristics			
Household size	0.494***	0.010	[0.474; 0.513]
Female (1=yes)	-0.040***	0.007	[-0.053; -0.028]
Partner status (ref: no partner)			
Partner, cohabiting	0.189***	0.038	[0.115; 0.263]
Partner, non-cohabiting	0.088**	0.038	[0.014; 0.163]
Single, no steady partner	0.041	0.037	[-0.033; 0.114]
Children			
Number of children	-0.449***	0.011	[-0.471; -0.428]
Constant			
Constant	2.199***	0.043	[2.114; 2.284]
Observations		17,209	
R-squared		0.443	
Root MSE		0.414	

Notes: Robust standard errors in parentheses.
* p<0.10, ** p<0.05, *** p<0.01.

Appendix B: Consumption regression

Table 7: Determinants of Log Total Consumption (p/m).

Variable	Coefficient	Std. Error	95% CI
Age group (ref: age group 2)			
Age 3	0.368***	0.022	[0.325; 0.412]
Age 4	0.400***	0.022	[0.357; 0.444]
Age 5	0.410***	0.022	[0.367; 0.453]
Age 6	0.373***	0.022	[0.331; 0.415]
Age 7	0.359***	0.021	[0.317; 0.400]
Age 8	0.337***	0.021	[0.296; 0.377]
Age 9	0.326***	0.031	[0.265; 0.388]
Age 10	0.256*	0.133	[-0.005; 0.517]
Age 11	0.485	0.381	[-0.262; 1.232]
Age 14	-1.246***	0.331	[-1.894; -0.598]
Year (ref: 2009)			
2010	-0.012	0.021	[-0.052; 0.028]
2012	-0.052**	0.021	[-0.093; -0.011]
2015	-0.127***	0.020	[-0.166; -0.087]
2017	-0.170***	0.021	[-0.211; -0.128]
2019	-0.156***	0.023	[-0.201; -0.111]
2020	-0.258***	0.022	[-0.302; -0.214]
2021a	-0.190***	0.022	[-0.233; -0.146]
2021b	-0.188***	0.022	[-0.232; -0.145]
Education (ref: primary/none)			
VMBO	0.190***	0.029	[0.134; 0.247]
HAVO/VWO	0.383***	0.031	[0.323; 0.444]
MBO	0.347***	0.027	[0.293; 0.401]
HBO	0.539***	0.027	[0.485; 0.592]
WO	0.630***	0.029	[0.573; 0.687]
Household members	-0.022	0.016	[-0.054; 0.010]
Female	-0.023**	0.011	[-0.044; -0.002]
Partner status (ref: no partner)			
Partner, cohabiting	1.923***	0.250	[1.434; 2.412]
Partner, not cohabiting	1.650***	0.250	[1.160; 2.140]
Single (no steady partner)	1.560***	0.250	[1.071; 2.050]
Children in household	0.032*	0.018	[-0.003; 0.068]
Constant	0.524**	0.253	[0.029; 1.020]
Observations		15,873	
R-squared		0.163	
Root MSE		0.659	

Notes: Robust standard errors in parentheses.
* p<0.10, ** p<0.05, *** p<0.01.

Appendix C: Robustness

Table A1: Estimated Income Shock Parameters (Gross Income)

	Estimate	Standard Error
Permanent income variance σ^2_η	0.0369***	0.0025
Transitory income variance σ^2_ε	0.0181***	0.0011

Notes: Estimates are based on gross household income. Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels.

Table A2: Category-Level Responses to Permanent and Transitory Income Shocks (Gross Income)

Category	ϕ_k	SE(ϕ_k)	ψ_k	SE(ψ_k)	$\sigma^2_{\xi,k}$	SE($\sigma^2_{\xi,k}$)
Total consumption	0.243***	0.052	0.089	0.177	0.382***	0.015
Mortgage	0.009	0.030	0.166	0.123	0.138***	0.008
Rent	0.099	0.065	-0.033	0.163	0.151***	0.020
Utilities	0.016	0.044	0.122	0.130	0.228***	0.006
Transport	0.298***	0.061	-0.129	0.215	0.676***	0.013
Insurance	0.337***	0.057	-0.373**	0.175	0.636***	0.012
Debt payments	-0.255***	0.082	0.249	0.227	0.439***	0.021
Cleaning	0.198**	0.078	0.243	0.303	1.000***	0.018
Food at home	-0.044	0.046	0.219	0.236	0.428***	0.011
Food away from home	0.064	0.060	0.272	0.251	0.863***	0.016
Clothing	0.065	0.058	0.125	0.170	0.697***	0.014
Personal care	0.365***	0.089	-0.069	0.319	0.733***	0.019
Medical care	0.270**	0.108	-0.120	0.376	1.576***	0.025
Daycare	-0.067	0.130	-0.209	0.602	0.651***	0.036
Holiday	0.366***	0.097	0.090	0.335	0.897***	0.021
Schooling	0.338***	0.100	-0.527	0.497	1.418***	0.039
Leisure	0.484***	0.103	0.363	0.339	0.889***	0.021
Donations	0.342***	0.075	0.213	0.304	0.840***	0.015

Notes: Estimates are based on gross household income. Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. Permanent shock loading = ϕ_k, transitory shock loading = ψ_k, taste-shifter variance = $\sigma^2_{\xi,k}$.

Table A3: Estimated Income Shock Parameters (Ages 25–85)

	Estimate	SE
Permanent income variance σ^2_η	0.0338***	0.0023
Transitory income variance σ^2_ε	0.0115***	0.0007

Notes: Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels.

Table A4: Category-Level Responses to Income Shocks (Ages 25–85)

Category	ϕ_k	SE(ϕ_k)	ψ_k	SE(ψ_k)	σ^2_ξ	SE
Total consumption	0.184***	0.034	-0.100	0.151	0.346***	0.009
Mortgage	0.035	0.030	-0.033	0.199	0.189***	0.008
Rent	0.015	0.031	-0.028	0.097	0.095***	0.011
Utilities	0.026	0.026	0.035	0.114	0.224***	0.004
Transport	0.262***	0.047	0.151	0.234	0.718***	0.009
Insurance	0.226***	0.041	-0.434***	0.164	0.656***	0.009
Debt payments	-0.215***	0.074	0.090	0.206	0.426***	0.017
Cleaning	0.160***	0.059	0.106	0.260	1.000***	0.018
Food at home	-0.048	0.041	0.145	0.216	0.431***	0.009
Food away from home	0.049	0.053	0.245	0.235	0.862***	0.014
Clothing	0.024	0.047	0.091	0.139	0.693***	0.012
Personal care	0.273***	0.070	-0.127	0.295	0.732***	0.016
Medical care	0.212***	0.082	0.080	0.339	1.572***	0.020
Daycare	-0.193*	0.112	0.054	0.526	0.669***	0.034
Holiday	0.298***	0.082	0.020	0.287	0.896***	0.018
Schooling	0.239***	0.087	-0.289	0.433	1.415***	0.034
Leisure	0.436***	0.095	0.224	0.309	0.891***	0.018
Donations	0.297***	0.063	0.127	0.273	0.838***	0.013

Notes: Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. Permanent shock loading = ϕ_k, transitory shock loading = ψ_k, taste–shifter variance = $\sigma^2_{\xi,k}$.