Computing Price-Theoretic Decompositions in Heterogeneous-Agent Models*

Christian Bustamante
Bank of Canada

Heejeong Kim
Concordia University

September 2026

Abstract

We propose a general numerical method to compute price-theoretic decompositions of interest rate effects in heterogeneous-agent models with borrowing constraints and multiple assets. Using discrete-state transition operators under risk-adjusted probability measures, we implement backward recursions that isolate substitution, wealth, and precautionary channels in environments with endogenous policy rules and occasionally binding constraints. The algorithm accommodates stopping times induced by liquidity constraints and applies to multi-asset economies with illiquid assets and adjustment costs. In a benchmark Huggett model, the decomposition accurately replicates finite-difference responses to interest rate shocks. In a two-asset economy, liquidity constraints generate endogenous stopping behavior and amplify substitution effects, while precautionary channels remain quantitatively limited in deterministic interest rate environments. Our method provides a tractable tool for analyzing the transmission of interest rate changes in incomplete markets and for quantifying the roles of substitution, wealth, and precautionary channels in heterogeneous-agent models.

Keywords:

*The views expressed in this paper are our own and do not necessarily reflect the official views of the Bank of Canada. E-mails: mail@cbustamante.co (Bustamante), heejeong.kim.macro@gmail.com (Kim).

1 Introduction

Quantitative models of heterogeneous agents are now prevalent in the literature. They allow us to capture certain aspects of the data that the representative-agent framework abstracts from, and produce quantitatively rich predictions. However, as these models grow more complex, it becomes harder to disentangle the operating mechanisms and interpret them. For example, changes in the interest rate combine wealth and substitution effects along with portfolio revaluation and precautionary effects. And in richer environments, such as models with discrete choices, the number of these underlying effects increases. In this context, theory-driven decompositions are a useful way to measure and interpret these mechanisms following any changes in aggregate conditions.

Farhi, Olivi and Werning (2022) generalize results from the classical demand theory framework to settings with incomplete markets. Starting from the Slutsky equation, they provide analytical expressions to decompose consumption responses into wealth, substitution, and precautionary effects. Despite the simplicity of their resulting decomposition, it relies on taking expectations over endogenous households' histories using preference-adjusted probability measures, while accounting for binding borrowing constraints and sequences of marginal propensities to save. Even for simpler models, accounting for all these potential household histories over long periods becomes computationally costly. This cost increases exponentially in the horizon as well as model complexity and dimension.

In this paper, we develop a backward recursion method that addresses these computational difficulties and efficiently computes Farhi et al. (2022)'s price-theoretic decomposition in heterogeneous-agent models. Our main result is that the path-dependent objects used in the decomposition can be represented as backward recursions defined over the household state space. In a recursive setting, we can treat binding constraints as boundary conditions, adjusted probabilities as transition matrices or operators, while dealing with compounding products over histories as continuation values. In this way, our computational implementation does not require simulating and tracking households' histories over time.

The backward representation delivers substantial computational gains. Tracking path-dependent histories requires storing objects that grow exponentially over time. In contrast, our method reduces this to a linear recursion in the horizon. For a household problem in a discretized grid with N points in the joint state space and N_θ possible realizations of the idiosyncratic shock, our algorithm reduces the complexity of computing each of the substitution, wealth, and precautionary effects from $O(NN_{\theta}^{T})$ to $O(TNN_{\theta})$. In this way, we can obtain these three components via three single backward passes over the state space. Furthermore, we show how to express the preference-adjusted probability measures as sparse transition operators, so we can exploit the efficiencies of sparse matrix-vector multiplications at each step of the recursion.

We apply the decomposition algorithm to the Huggett (1993) model following a persistent interest rate perturbation. The Huggett model provides a natural environment to test our method. It is an incomplete-markets economy with uninsurable idiosyncratic risk and occasionally binding borrowing constraints that generate non-trivial stopping times, while also remaining tractable enough to compare the decomposition results with a finite-difference benchmark. Under our baseline calibration, the substitution effect explains most of the contraction in aggregate consumption on impact, while the wealth effect partially offsets it and the precautionary effect is negligible. Importantly, the sum of the three effects closely matches the finite-difference response of aggregate consumption. Because we construct the decomposition at the household level before aggregating, we characterize how the substitution, wealth, and precautionary effects vary across the distribution of earnings and assets. Furthermore, the state-level decomposition closely matches the state-level finite-difference responses, highlighting that the method's aggregate-level accuracy does not arise from offsetting approximation errors across households.

Our algorithm also extends to models with discrete choices after accounting for extensive-margin responses. We formalize this result and show that the total consumption response can be decomposed into an extensive-margin component and alternative-specific intensive-margin responses, each of which admits the same substitution, wealth, and precautionary effects as in the single-asset case. To illustrate this extension, we apply the more general version of our method to a standard two-asset model with discrete illiquid-portfolio adjustment. We show that, in this environment, interest rate changes affect aggregate consumption not only via the substitution, wealth, and precautionary channels conditional on households' portfolio choices, but also through changes in the probabilities of households adjusting their portfolios.

Related Literature. Here

The remainder of the paper proceeds as follows. Section 2 presents the price-theoretic decomposition and discusses the computational challenges involved in its computation, and Section 3 develops the backward adjoint algorithm. Section 4 validates the method against finite-difference responses in a Huggett economy. Section 5 extends the decomposition to discrete choices and applies the method to a two-asset model with liquid and illiquid assets. Section 6 concludes.

2 Price-Theoretic Decompositions and the Computational Problem

To describe the components of the decomposition, discuss the intuition behind it, and introduce notation, we start with a standard incomplete-markets model.

2.1 Household Problem

Each household has preferences over consumption, with instantaneous utility function $u(c_t)$, which is strictly increasing, strictly concave, and three times continuously differentiable. They discount the future at the rate $\beta \in (0,1)$.

Households are subject to uninsurable idiosyncratic income shocks θ_t, with Markov transition matrix P. They can save in a one-period asset a_t with gross return R_t, and face an exogenous borrowing limit. Given this, the household's recursive problem is

$$V(a_t, \theta_t) = \max_{c_t, a_{t+1}} \Big\{ u(c_t) + \beta \, E\big[V(a_{t+1}, \theta_{t+1}) \mid \theta_t \big] \Big\} \tag{1}$$

subject to the budget constraint and the borrowing constraint

$$c_t + \frac{a_{t+1}}{R_t} = a_t + y(\theta_t), \tag{2}$$

$$a_{t+1} \;\geq\; \underline{a}, \tag{3}$$

The constraint (3) may bind occasionally, where $\underline{a} \leq 0$ denotes the borrowing limit. Let $\mathcal{B}_s \equiv \{(a,\theta) : a'_s(a,\theta) = \underline{a}\}$ denote the set of constrained states in period s, where a'_s is the date-s savings policy.

Let $c_t(a_t, \theta_t)$ and $a_{t+1}(a_t, \theta_t)$ denote the optimal policy functions solving (1). The marginal propensity to save is

$$MPS_t \;\equiv\; R_t\big(1 - \partial_a c_t\big), \tag{4}$$

where $\partial_a c_t \equiv \partial c(a_t, \theta_t)/\partial a_t$ is the marginal propensity to consume out of wealth.

2.2 Farhi-Olivi-Werning's Decomposition

This subsection revisits the price-theoretic decomposition of Farhi et al. (2022). It is useful to state some of their definitions first. Let $\{dR_s/R_s\}_{s=0}^{T}$ be a sequence of interest-rate perturbations. The stopping time at which the household first hits the borrowing constraint is

$$\tau \;\equiv\; \min\{s \geq 0 : (a_s, \theta_s) \in \mathcal{B}_s\}, \tag{5}$$

and the elasticity of intertemporal substitution at the steady-state is $\varepsilon(c_0) \;\equiv\; -u'(c_0)/[c_0\, u''(c_0)]$.

Farhi et al. (2022) introduce two probability measures that re-weight the objective probability P according to households' preferences:

$$\frac{dQ}{dP}(\theta_{t+1} \mid a_t, \theta_t) \;\equiv\; \frac{u'(c(a_{t+1}, \theta_{t+1}))}{E[u'(c(a_{t+1}, \theta_{t+1})) \mid a_t, \theta_t]}, \tag{6}$$

$$\frac{dQ^I}{dP}(\theta_{t+1} \mid a_t, \theta_t) \;\equiv\; \frac{u''(c(a_{t+1}, \theta_{t+1}))}{E[u''(c(a_{t+1}, \theta_{t+1})) \mid a_t, \theta_t]}. \tag{7}$$

Q and Q^I are the risk- and prudence-adjusted probability measures, respectively. Both are state-dependent through the consumption policy.

They show that the first-order response of consumption to the perturbation $\{dR_s/R_s\}_{s=0}^{T}$ admits the following decomposition

$$\begin{aligned}
\frac{dc_0}{c_0} \;=\; & \underbrace{-\varepsilon(c_0)\; E_0^Q\left[\sum_{t=0}^{\tau}\left(\prod_{s=0}^{t}\frac{MPS_s}{R_s}\right)\frac{dR_t}{R_t}\right]}_{\text{substitution}} + \underbrace{\frac{\partial_a c_0}{c_0}\; E_0^Q\left[\sum_{t=0}^{\tau}\left(\prod_{s=0}^{t-1}\frac{1}{R_s}\right)\frac{a_{t+1}}{R_t}\,\frac{dR_t}{R_t}\right]}_{\text{wealth}} \\
& + \underbrace{\frac{\partial_a c_0}{c_0}\,\mathrm{Cov}_0^Q\left(\frac{dQ^I}{dQ}, \sum_{t=0}^{\tau}\left(\prod_{s=0}^{t-1}\frac{1}{R_s}\right)\frac{a_{t+1}}{R_t}\,\frac{dR_t}{R_t}\right)}_{\text{precautionary}},
\end{aligned} \tag{8}$$

This decomposition has three major features worth emphasizing. First, these three effects involve expectations under the adjusted probability measures Q and Q^I, rather than P. Second, all three sums are truncated at the endogenous stopping time τ. And third, the substitution term includes products of state-dependent MPS_s terms along all households' histories up to that stopping point. Analytically, these features are relatively straightforward, but they create problems in a quantitative setting with discretized state spaces. We discuss these problems in detail in the following section.

The substitution effect term captures responses to changes in interest rates, discounted by the sequence of marginal propensities to save, and truncated when the borrowing constraint binds. The wealth effect measures the present value of changes in return on households' asset positions and converts them into current consumption using the date-0 *MPC*. Finally, the precautionary effect focuses on capturing the divergence between the prudence-adjusted and risk-adjusted valuations of the same wealth gain, as $\mathrm{Cov}_0^Q(dQ^I/dQ, X) = E_0^{Q^I}[X] - E_0^Q[X]$. See Appendix A.1 for the derivation of this change-of-measure identity.

2.3 Computational Challenges

Direct evaluation of the decomposition (8) on the discrete state space confronts three numerical obstacles.

Stopping times. Borrowing constraints and income risk induce a stopping time $\tau = \min\{s \geq 0 : (a_s, \theta_s) \in \mathcal{B}_s\}$ — the first date at which the household's optimal savings choice hits the borrowing constraint. The difficulty is that τ is path-dependent — different income realizations lead the household to hit the constraint at different dates (or never) — so each channel's sum in (8) must be tracked along each possible trajectory and truncated accordingly.

Adjusted probability measures. The decomposition (8) uses the adjusted measures Q and Q^I rather than the objective probability P. The difficulty is that dQ/dP and dQ^I/dP are state-dependent — they vary with (a, θ) through the next-period consumption policy — so the algorithm must construct two distinct sparse transition operators with row-by-row Radon-Nikodym factors, rather than reweighting samples drawn from P.

Multiplicative path products. The substitution channel involves a product of state-dependent terms $MPS_s(a_s, \theta_s)/R_s$ along the trajectory from 0 to t. The difficulty is that estimating $E_0^Q[\prod_{s=0}^t MPS_s/R_s]$ path-by-path requires evaluating the product along all N_θ^T possible income histories over the horizon T, with endogenous asset choices and MPS_s evaluated along each history.

A direct implementation would then require tracking all idiosyncratic shocks, endogenous choices, stopping points, and path-dependent products over a rapidly expanding set of histories. In the next section, we show that we can simplify all of these by running a single backward recursion on the state space.

3 The Backward Adjoint Algorithm

Our paper provides an algorithm that resolves all three obstacles simultaneously: the change of measure is built into the sparse transition operator T, and the operator is iterated backward in time. State-dependent reweighting is encoded into the sparse transition matrices. The path-dependent stopping time is enforced as an absorbing boundary at every backward step. Lastly, the multiplicative path products collapse into linear $O(T \cdot N)$ recursions through auxiliary value functions.

3.1 A General Backward Recursion

Our algorithm is structurally simple. It evaluates the path-dependent objects in Farhi et al. (2022) with a single backward pass. No household trajectory needs to be simulated. Let $\mathcal{X}$ denote the discrete state space, and $\mathcal{B}_s \subset \mathcal{X}$ the set of binding states at date s. The following proposition describes the generic backward representation underlying our algorithm.

Proposition 1 (Backward Representation)**.** *Let $\tau_s \equiv \min\{k \geq s : X_k \in \mathcal{B}_k\}$ denote the first time, starting from date s, that the household hits the borrowing constraint. For arbitrary state-dependent functions $\alpha_s, \beta_s : \mathcal{X} \to \mathbb{R}$ and probability measure $\widetilde{Q}$, define*

$$v_s(x) = \begin{cases} \alpha_s(x), & x \in \mathcal{B}_s, \\ \alpha_s(x) + \beta_s(x) E_s^{\widetilde{Q}}[v_{s+1}(X_{s+1}) \mid X_s = x], & x \notin \mathcal{B}_s, \end{cases} \tag{9}$$

with terminal condition $v_{T+1} \equiv 0$. Then, for every $s \leq T$,

$$v_s(x) = E_s^{\widetilde{Q}} \left[\sum_{k=s}^{T \wedge \tau_s} \alpha_k(X_k) \prod_{j=s}^{k-1} \beta_j(X_j) \,\middle|\, X_s = x \right]. \tag{10}$$

Hence, any functional equation of the form (10) can be evaluated by a single backward pass over the state space, without needing to evaluate all household histories.

The proof of this Proposition is provided in Appendix A.2.

On the discrete state space, the conditional expectation in (9) acts on v_{s+1} as a sparse matrix-vector product:

$$E_s^{\tilde{Q}}[v_{s+1}(X_{s+1}) \mid X_s = x] = \sum_{x' \in \mathcal{X}} T_s^{\tilde{Q}}(x, x')\, v_{s+1}(x') \;\equiv\; (T_s^{\tilde{Q}}\, v_{s+1})(x), \tag{11}$$

where $T_s^{\tilde{Q}}$ is the sparse $N \times N$ transition matrix on $\mathcal{X}$ whose row at x represents the conditional probability of X_{s+1} under $\widetilde{Q}$ given $X_s = x$. Recursion (9) then takes the matrix form

$$v_s = \alpha_s + m_s \odot \beta_s \odot (T_s^{\tilde{Q}}\, v_{s+1}), \qquad s = T, T-1, \ldots, 0, \tag{12}$$

where $\alpha_s, \beta_s, m_s, v_s \in \mathbb{R}^N$ are vectors indexed by states, $\odot$ denotes element-wise multiplication, and $m_s(x) \equiv \mathbf{1}\{x \notin \mathcal{B}_s\}$.

The substitution, wealth, and precautionary effects from Section 2 are special cases of this same representation, with specific α_s, β_s coefficients, and measure $\tilde{Q}$. Below, we derive the corresponding transition operators and show how they map from the general form.

3.2 Risk- and Prudence-Adjusted Transition Operators

Assets are discretized on $\{a_1, \ldots, a_{N_a}\}$ with the borrowing limit $a_1 = \underline{a}$, and income shocks on $\{\theta_1, \ldots, \theta_{N_\theta}\}$ following a Markov chain with transition matrix $P(\theta'|\theta)$. The joint state space $\mathcal{X}$ has $N = N_a \cdot N_\theta$ points indexed by $x = (a, \theta)$. Let $a'_s : \mathcal{X} \to [a_1, a_{N_a}]$ denote the date-s savings policy.

To construct the transition operators associated with the adjusted probability measures in (6) and (7), define $F^{Q}(c) = u'(c)$ and $F^{Q^I}(c) = u''(c)$. For $\tilde{Q} = \{Q, Q^I\}$, the adjusted conditional probability of θ is given by

$$\tilde{Q}_s(\theta'|a, \theta) = \frac{F^{\tilde{Q}}(c_{s+1}(a'_s(a,\theta),\, \theta'))P(\theta'|\theta)}{\sum_{\tilde{\theta}} F^{\tilde{Q}}(c_{s+1}(a'_s(a,\theta),\, \tilde{\theta}))\, P(\tilde{\theta}|\theta)}. \tag{13}$$

Given that, generally speaking, the asset policy lies off the asset grid, we can denote $\phi_s(a, \theta; a_j)$ as the linear interpolation weight for point a_j, so that we can write the transition operator on the full discrete state-space as

$$T_s^{\tilde{Q}}\big((a, \theta), (a_j, \theta')\big) = \phi_s(a, \theta; a_j)\tilde{Q}_s(\theta'|a, \theta) \tag{14}$$

Using this operator, we can then write the discrete state-space counterparts of the risk- and prudence-adjusted conditional expectations in (8), noting that, for any function $f : \mathcal{X} \to \mathbb{R}$, we can compute

$$E_s^{\tilde{Q}}[f(X_{s+1})|X_s = x] = \left(T_s^{\tilde{Q}} f\right)(x). \tag{15}$$

By construction $T_s^{\tilde{Q}}$ is sparse and row-stochastic, with at most $2N_\theta$ non-zero entries per row and $O(N_a N_\theta^2)$ total non-zeros.

Note that the backward recursion has an adjoint interpretation. For a probability measure $\tilde{Q}$, the transition operator $T_s^{\tilde{Q}}$ maps functions of future states backward, while the operator $(T_s^{\tilde{Q}})'$ projects distributions of households forward in time. To see this, let $\mu_s(x)$ denote the distribution of households at time s, then $\mu_{s+1}(x) = (T_s^{\tilde{Q}})'\mu_s(x)$. On the other hand, the continuation term in the backward recursion is given by $T_s^{\tilde{Q}} v_{s+1}$. These two operations are adjoint.

3.3 From the Decomposition Components to the Backward Recursion

As Proposition 1 states, each component of the Farhi et al. (2022) decomposition maps directly into the backward recursion. Each component differs only in their specific flow coefficient α_s, continuation coefficient β_s, and probability measure $\tilde{Q}$. Table 1 summarizes this mapping for each component.

Table 1: Component-Specific Coefficients

Component	**Flow:** $\alpha_s(x)$	**Cont.:** $\beta_s(x)$	**Prob:** $\tilde{Q}$
Substitution	$\frac{MPS_s(x)}{R_s}\frac{dR_s}{R_s}$	$\frac{MPS_s(x)}{R_s}$	Q
Wealth	$\frac{a_s'(x)}{R_s}\frac{dR_s}{R_s}$	$\frac{1}{R_s}$	Q
Precautionary	$\frac{a_s'(x)}{R_s}\frac{dR_s}{R_s}$	$\frac{1}{R_s}$	Q^I

All the auxiliary function that will be defined below share the same terminal condition in Proposition 1, $v_{T+1} = 0$. For the substitution effect, however, there is an additional natural boundary condition. Given that $MPS_s(x) = 0$ whenever $x \in \mathcal{B}_s$, we have that $\alpha_s(x)$ and $\beta_s(x)$ are both zero at the binding states. Hence, the boundary condition in Proposition 1 reduces to $v_s(x) = 0$ at those states.

Substitution effect. Let w_s denote the auxiliary function from Proposition 1. Using the first row of Table 1, we can write

$$w_s(x) = \begin{cases} 0, & x \in \mathcal{B}_s, \\ \dfrac{MPS_s(x)}{R_s} \cdot \left(\dfrac{dR_s}{R_s} + (T_s^Q \, w_{s+1})(x) \right), & x \notin \mathcal{B}_s, \end{cases} \tag{16}$$

At the household-level, the substitution effect is then given by

$$\mathrm{sub}(x) = -\,\varepsilon(c_0(x))\, w_0(x), \tag{17}$$

Wealth Effect. Let H_s denote the corresponding auxiliary function from Proposition 1. Using the second row of Table 1 we obtain

$$H_s(x) = \frac{a_s'(x)}{R_s}\frac{dR_s}{R_s} + m_s(x)\frac{1}{R_s}(T_s^Q \, H_{s+1})(x) \tag{18}$$

Hence, the household-level wealth effect is

$$\mathrm{wealth}(x) = \frac{\partial_a c_0(x)}{c_0(x)}\, H_0(x), \tag{19}$$

Precautionary Effect. Using the fact that $\mathrm{Cov}_0^Q(dQ^I/dQ, X) = E_0^{Q^I}[X] - E_0^Q[X]$, letting H^I denote the corresponding auxiliary function, and using the third row of Table 1 we have that

$$H_s^I(x) = \frac{a_s'(x)}{R_s}\frac{dR_s}{R_s} + m_s(x)\frac{1}{R_s}(T_s^{Q^I} \, H_{s+1}^I)(x) \tag{20}$$

Then, the household-level precautionary effect is given by

$$\mathrm{prec}(x) = \frac{\partial_a c_0(x)}{c_0(x)}\, \big(H_0^I(x) - H_0(x)\big). \tag{21}$$

These results show that the full decomposition requires only three backward passes: one for the substitution functional under Q, and two otherwise identical wealth recursions under Q and Q^I. Algorithm 1 in Appendix B summarizes the complete numerical implementation.

3.4 Aggregation

We can aggrete these components over the distribution of households to obtain an economy-wide decomposition and its respective substitution, wealth and precautionary effects. Let $\mu(x)$ denote the stationary distribution of households over $x = (a, \theta)$. We can define aggregate consumption as the stationary equilibrium as $C_0 = \int_{\mathcal{X}} c_0(x)\mu(dx)$. Thus,

$$\frac{dC_0}{C_0} = \int_{\mathcal{X}} \frac{dc_0(x)}{c_0(x)} \frac{c_0(x)\mu(x)}{C_0(x)}. \tag{22}$$

Hence, for $z = \{\text{sub}, \text{wealth}, \text{prec}\}$, we have

$$\widehat{Z} = \int_{\mathcal{X}} z(x) \, \frac{c_0(x)\,\mu(dx)}{C}. \tag{23}$$

This implies that the response of aggregate consumption can be decomposed as

$$\frac{dC_0}{C_0} = \widehat{\text{Sub}} + \widehat{\text{Wealth}} + \widehat{\text{Prec}}. \tag{24}$$

3.5 Complexity

Direct evaluation of the closed form (10) requires tracking up to $O(N_\theta^T)$ income paths for each initial state, which is exponential in the horizon T. The Markov structure of the state process collapses this exponentially expanding path-tracking to a T-step backward recursion of sparse matrix-vector products $T_s^{\tilde{Q}} \, v_{s+1}$. With linear interpolation in the one-asset model, each $N \times X$ transition operator $T_s^{\tilde{Q}}$ has $O(N_\theta N)$ non-zeros, so each backward step requires only $O(N_\theta N)$ operations and storage. Thus, effectively, the full implementation of our method reduces the computational cost from $O(NN_\theta^T)$ to $O(TNN_\theta)$. This highlights how the backward representation converts an exponential problem in T into a linear one.

4 One-Asset Economy: Validation

In this section, we apply the backward adjoint algorithm of Section 3 to a standard Huggett economy and quantify the three channels of (8). We first describe the model's parametrization and the numerical implementation, and then validate the decomposition by checking that the sum of the three components matches the direct finite-difference response.

4.1 Calibration

The benchmark economy is a standard Huggett incomplete-markets model calibrated at quarterly frequency, following Farhi et al. (2022). Table 2 summarizes the parameters.

Preferences. Utility is CRRA. This delivers an elasticity of intertemporal substitution $\varepsilon(c_0) = 1/\sigma$ and ensures $u'''(c) = \sigma(\sigma+1)\,c^{-(\sigma+2)} > 0$ so that the prudence-adjusted measure Q^I differs from Q and the precautionary channel can be active. The discount factor β is chosen to match a ratio of liquid wealth to output of $A/Y = 1.44$ in steady state (McKay, Nakamura and Steinsson, 2015; Farhi et al., 2022).

Earnings shock process. The idiosyncratic earnings shock follows an AR(1) process,

$$\log y_t = \rho_\theta \log y_{t-1} + \epsilon_t, \qquad \epsilon_t \sim \mathcal{N}(0, \sigma_\theta^2),$$

The process is discretized into $N_\theta = 5$ using the Rouwenhorst algorithm.

Numerical implementation. The savings policy is solved by the endogenous grid method on a log-spaced asset grid of 50 points over $[\underline{a}, 40]$; the stationary distribution and the transition operators $T_s^Q, T_s^{Q^I}$ are interpolated on a $N_a = 750$-point distribution grid over the same range. The interest-rate perturbation is propagated over a horizon of $T = 300$ quarters, long enough for the economy to revert to steady state. The state space has $N = N_a N_\theta = 750 \times 5 = 3{,}750$ points, so each sparse T_s carries $2N_\theta N \approx 3.8 \times 10^4$ non-zeros and fits in under 1 MB.

Table 2: Huggett Benchmark Calibration

Parameter	Symbol	Value
Discount factor	β	0.988
CRRA curvature	σ	2.0
Earnings innovation std	σ_θ	$\sqrt{0.017}$
Earnings persistence	ρ_θ	0.966
Borrowing limit	$\underline{a}$	-0.1
Asset-to-income ratio	A/Y	1.44

4.2 Decomposition and Finite-Difference Benchmark

We validate the algorithm by comparing the decomposition against a direct finite-difference benchmark for a single perturbation of the interest-rate path. Specifically, we

Figure 1: Aggregate Consumption Response: Baseline Economy

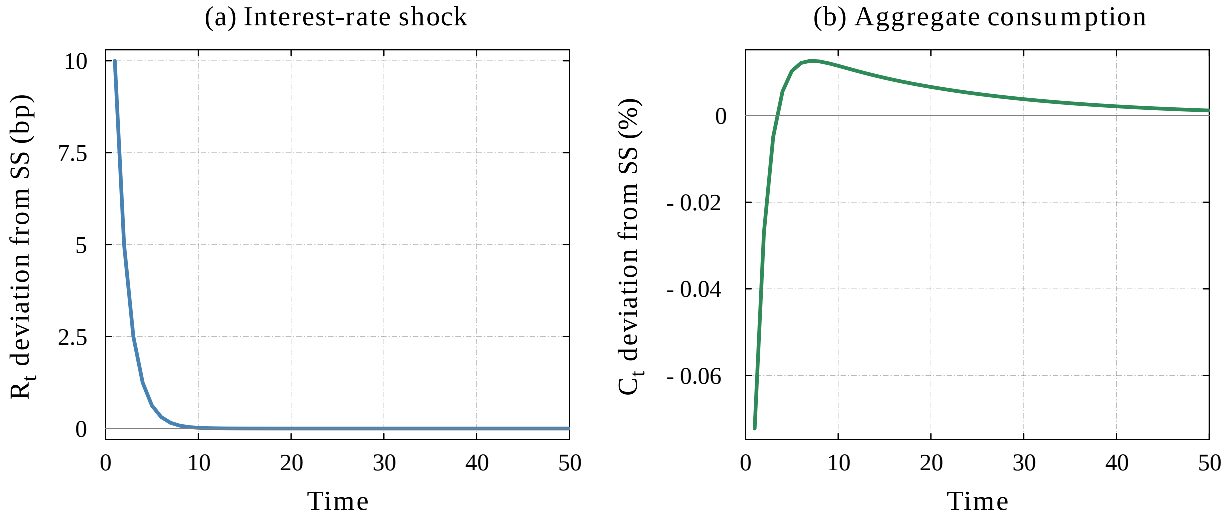

Notes: Panel (a): deviation of the gross rate from steady state, $R_t - R^{ss}$, in basis points ($dR = 10^{-3}$, $\rho_r = 0.5$). Panel (b): aggregate consumption deviation from steady state, $\Delta C_t / C^{ss}$ (%).

perturb the gross rate at $t = 0$ with an AR(1)-decaying impulse,

$$R_t = R^{ss} + dR \cdot \rho_r^t, \qquad t = 0, \ldots, T-1, \qquad R_T = R^{ss}, \tag{25}$$

with $dR = 10^{-3}$ (a 10-basis-point increase at impact) and $\rho_r = 0.5$. The shock reverts to R^{ss} at $t = T$, as shown in Figure 1(a). Panel (b) reports the resulting transition of aggregate consumption C_t relative to its steady-state value. Consumption falls by about 0.07 % on impact and reverts toward steady state.

Table 3 reports the three channels of (8) and compares their sum to the aggregate consumption response $\Delta C_0 / C_0^{ss}$. Specifically, the finite-difference (FD) response is computed by solving the model along the perturbed path $\{R_t\}$ and comparing impact consumption with its no-shock benchmark. The backward decomposition closely reproduces the FD response: summing over the three channels yields a -0.070 % response for the decomposition versus a finite-difference response of -0.072 %.[1] The substitution channel accounts for most of this response (-0.078 %), while the wealth channel partially offsets the substitution effect ($+0.008$ %). Lastly, the precautionary channel is small (-5×10^{-5} %). In this experiment, interest rate changes are deterministic and do not affect idiosyncratic risk, so the precautionary channel plays a relatively small role.

Importantly, the accuracy of the decomposition does not result from household-level errors offsetting each other at the aggregate level. Figure 2 depicts the approximation

[1] The relative gap between the two responses narrows toward 1 % on finer grids.

Table 3: Channel Decomposition at the Baseline calibration.

Decomposition			Total	
Substitution	Wealth	Precautionary	Sum	FD ΔC_0
-0.0777	$+0.0078$	-5×10^{-5}	-0.0700	-0.0722

All entries are percent deviations from steady-state aggregate consumption.

Figure 2: Approximation Error by Asset Level

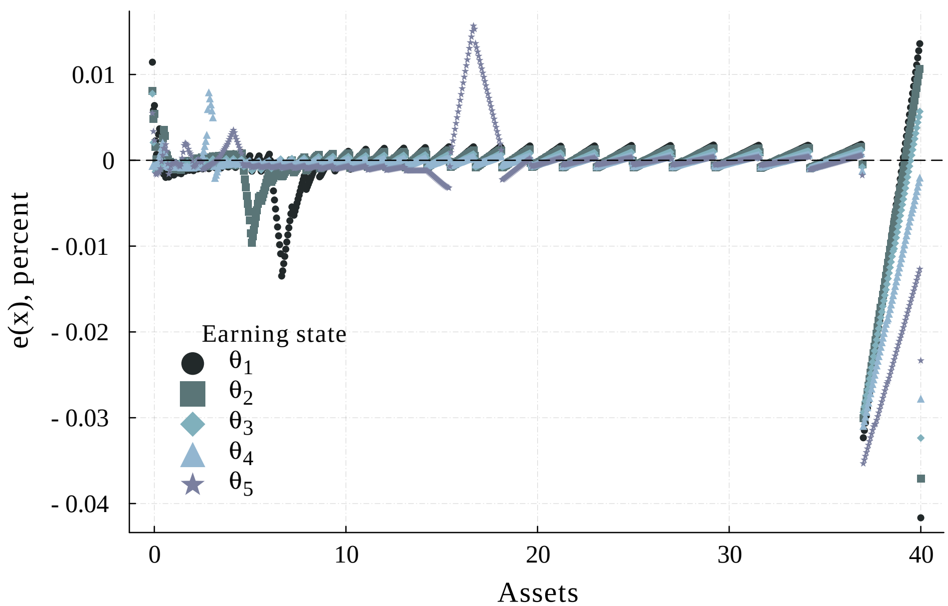

Notes: The figure reports the state-level approximation error $e(x) = [\text{sub}(x) + \text{wealth}(x) + \text{prec}(x)] - \Delta c_0^{FD}(x)/c_0(x)$, for each $x = (a, \theta)$ pair.

error for each income state along the asset grid. The approximation error is very close to zero for most individual-level states, with larger discrepancies near the upper bound of the asset grid.

4.3 Heterogeneity in the Decomposition

Having established that the decomposition closely approximates the finite-difference response, we now use it to characterize how the three channels vary across the model's state space. Figure 3 shows the three decomposition components along percentiles of asset holdings in the stationary distribution. Panels (b)-(d) focus on three specific earning-shock levels, while Panel (a) averages across all N_θ earning-shocks.

Figure 3: Decomposition Along the State Space

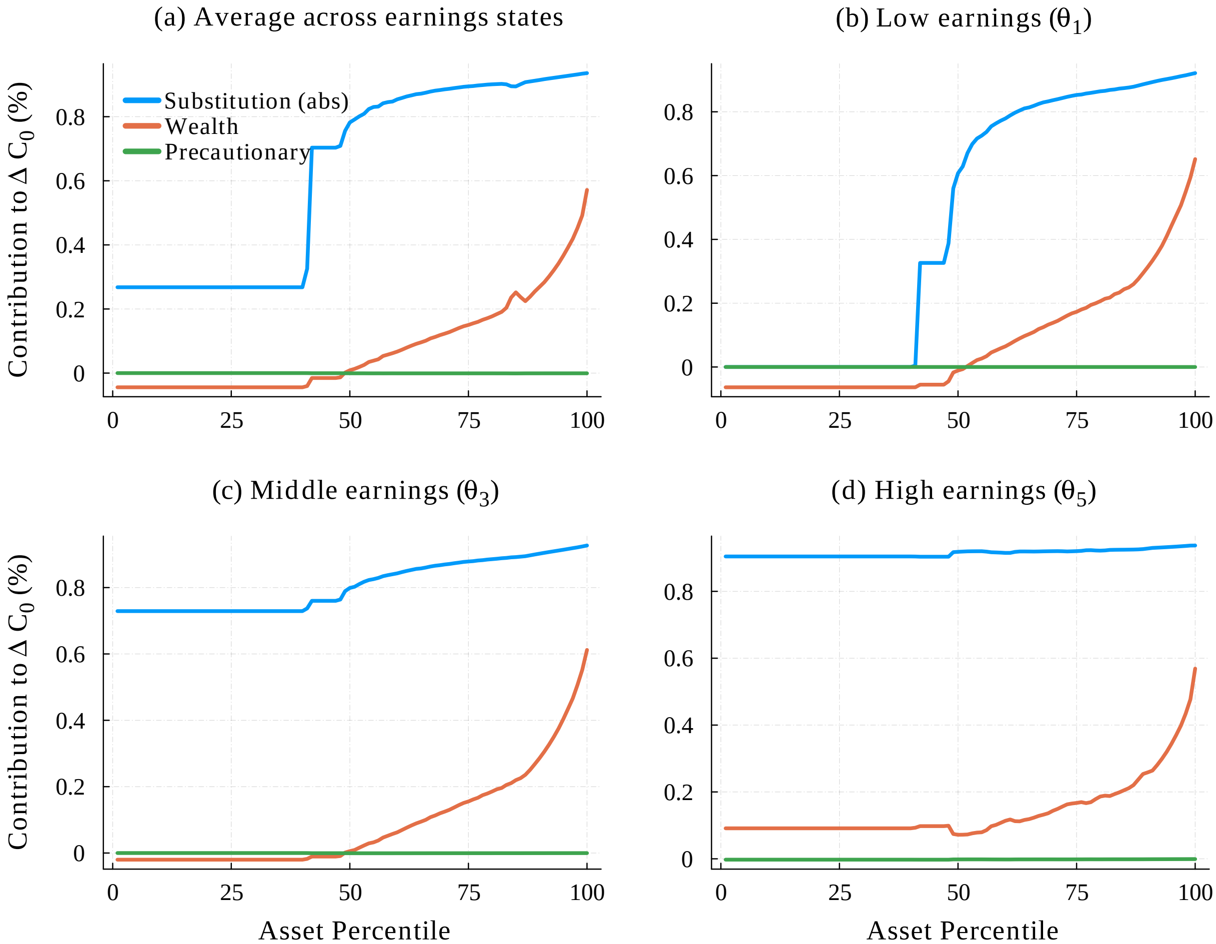

Notes: Contribution of each channel to the total response of consumption at impact. The substitution effect is multiplied by -1.

There are four main results that are worth highlighting. First, the absolute magnitude of the substitution effect raises with wealth, because the effect is attenuated for low-asset households as they approach the borrowing limit and their *MPS* declines. Second, the wealth effect also increases with wealth because richer households are have greater exposure to changes in asset returns. Third, higher earning reinforce both patterns: households with higher earnings save more and are farther from the borrowing constraint, strengthening both the substitution and wealth channels. Finally, the precautionary effect is quantitatively negligible throughout the distribution given that the return shock is deterministic and leaves the idiosyncratic earnings process unchanged.

4.4 Robustness

In this section, we show how the results change with different parameter values (Table 4). Panel A shows that a higher σ (lower EIS) leads to a smaller substitution effect. In Panel B, a higher β puts more weight on the future, so the substitution effect is larger; however, a lower equilibrium interest rate dampens the positive wealth effect. When less borrowing is allowed (Panel C), the precautionary channel becomes larger (more negative), from -4×10^{-5} % to -5×10^{-5} %, as households rely more on precautionary saving. Finally, Panels D and E show that higher income risk, as well as more persistent income risk, raises precautionary saving. Importantly, across all robustness checks, the sum of the three channels closely tracks the aggregate consumption response, confirming the accuracy of the decomposition.

5 Two-Asset Economy: Liquid and Illiquid Assets

We next extend the decomposition to models with discrete choices. In these environments, changes in choice probabilities generate an additional extensive-margin channel in the consumption response. We illustrate this extension using a standard two-asset model with discrete illiquid-portfolio adjustment.

5.1 Model

Consider a two-asset model in which households have access to a liquid asset b and an illiquid asset a. The liquid asset has a gross return R^b and can be adjusted freely each period subject only to the no-short-sale constraint $b' \geq b_{\min} = 0$. The illiquid asset has a higher return $R^a > R^b$, but when changing their illiquid position, households have to pay a portfolio rebalancing cost $\chi(a', a) \geq 0$ out of the liquid budget. As in the one-asset model, households face idiosyncratic risk on their income $y(\theta)$, where θ follows a Markov process. Preferences are CRRA with elasticity of intertemporal substitution ε. Appendix C reports the model calibration and details of the numerical solution.

The household state at the beginning of the period is characterized by the vector $x = (a, b, \theta)$. Given x, households choose whether to adjust their illiquid position. The value associated with adjusting is denoted by $V^a(x)$, while the value of not adjusting is given by $V^n(x)$. Households draw a pair of i.i.d. Type-I Extreme Value (Gumbel) taste shocks ν_a, ν_n with common scale parameter $\sigma_\nu > 0$. This implies that the adjustment

Table 4: Channel decomposition under alternative calibrations.

	R	A/Y	Decomposition: Substitution	Wealth	Precautionary	Total: Sum	FD Δc_0
Panel A: σ (CRRA curvature)							
1.0	0.9998	1.440	-0.1217	$+0.0076$	-4×10^{-5}	-0.1141	-0.1164
2.0	0.9827	1.440	-0.0777	$+0.0078$	-5×10^{-5}	-0.0700	-0.0722
4.0	0.9377	1.440	-0.0428	$+0.0080$	-9×10^{-5}	-0.0349	-0.0367
Panel B: β (discount factor)							
0.970	1.0000	1.440	-0.0632	$+0.0102$	-5×10^{-5}	-0.0531	-0.0555
0.988	0.9827	1.440	-0.0777	$+0.0078$	-5×10^{-5}	-0.0700	-0.0722
0.995	0.9762	1.440	-0.0787	$+0.0070$	-5×10^{-5}	-0.0717	-0.0737
Panel C: B (borrowing limit)							
-0.50	0.9859	1.440	-0.0794	-0.0005	-4×10^{-5}	-0.0799	-0.0853
-0.10	0.9827	1.440	-0.0777	$+0.0078$	-5×10^{-5}	-0.0700	-0.0722
0.00	0.9817	1.440	-0.0642	$+0.0100$	-5×10^{-5}	-0.0542	-0.0539
Panel D: σ_θ^2 (income innovation variance)							
0.0043	1.0034	1.440	-0.0631	$+0.0055$	-7×10^{-6}	-0.0577	-0.0599
0.0170	0.9827	1.440	-0.0777	$+0.0078$	-5×10^{-5}	-0.0700	-0.0722
0.0680	0.9150	1.440	-0.0861	$+0.0194$	-5×10^{-4}	-0.0672	-0.0694
Panel E: ρ_θ (income persistence)							
0.00	1.0119	1.440	-0.0933	$+0.0052$	-4×10^{-6}	-0.0881	-0.0880
0.50	1.0107	1.440	-0.0894	$+0.0078$	-2×10^{-5}	-0.0816	-0.0818
0.97	0.9827	1.440	-0.0777	$+0.0078$	-5×10^{-5}	-0.0700	-0.0722

Notes: Channel and total entries are % deviations of steady-state consumption. Parameters not varied are held at the baseline ($\sigma = 2$, $\beta = 0.988$, $\rho_\theta = 0.966$, $\sigma_\theta^2 = 0.017$, $B = -0.1$). Shock: $dR = 10^{-3}$, $\rho_r = 0.5$, $T = 300$ in all panels. Panel E varies income persistence ρ_θ down from the benchmark 0.966 toward i.i.d.: the precautionary channel is largest at the persistent benchmark and collapses as $\rho_\theta \to 0$, where only the borrowing constraint remains as a source of prudence reweighting.

probability is

$$\pi(x) = \frac{\exp\{V^a(x)/\sigma_\nu\}}{\exp\{V^a(x)/\sigma_\nu\} + \exp\{V^n(x)/\sigma_\nu\}}. \tag{26}$$

Conditional on not adjusting their portfolio, the household's recursive problem is given by

$$V^n(a,b,\theta) = \max_{c,b'} \left\{u(c) + \beta\, E\left[V(R^a a,\, b',\, \theta')|\theta\right]\right\} \tag{27}$$

subject to

$$c + b' = R^b\, b + y(\theta), \qquad b' \geq 0, \qquad a' = R^a\, a.$$

Likewise, conditional on adjusting, the household's recursive problem is

$$V^a(a,b,\theta) = \max_{c,\,a',\,b'} \left\{ u(c) + \beta\, E\big[V(a',\, b',\, \theta')|\theta\big] \right\} \tag{28}$$

subject to

$$c + a' + b' + \chi(a',a) = R^a\, a + R^b\, b + y(\theta), \qquad b' \geq 0 \tag{29}$$

The adjustment cost function is as in Kaplan, Moll and Violante (2018):

$$\chi(a',a) = \frac{\chi_1}{\chi_2}\, \left|a' - R^a\, a\right| \cdot \left(\frac{|a' - R^a\, a|}{R^a\, a + \chi_0} \right)^{\chi_2 - 1}. \tag{30}$$

Let $c^a(x)$ and $c^n(x)$ denote the consumption policies conditional on adjusting and not adjusting that solve (28) and (27). Integrating over the taste shocks, the policy relevant for aggregation is the adjustment-probability mixture

$$c(x) = \pi(x)\, c^a(x) + \big(1 - \pi(x)\big)\, c^n(x). \tag{31}$$

This mixture is key. Changes in the liquid return affect consumption policies and the adjustment probabilities. We will show how the changes in these two margins enter the decomposition.

5.2 Discrete Choices and the Extensive Margin

The two-asset model motivates how the backward adjoint decomposition extends to models with discrete choices. Suppose that a household with state x chooses among $j \in \{1, 2, \ldots, J\}$ discrete alternatives. Let $\pi^j(x)$ denote the probability of choosing alternative j, and $c^j(x)$ the consumption policy associated with that discrete choice. The choice probabilities satisfy that $\sum_{j=1}^{J} \pi^j(x) = 1$. Ex-post consumption is given by

$$c(x) = \sum_{j=1}^{J} \pi^j(x) c^j(x). \tag{32}$$

Differentiating with respect to a perturbation in the price sequence leads to the following proposition.

Proposition 2 (Decomposition with Discrete Choices)**.** *Suppose that probabilities $\pi^j(x)$ are differentiable and that, conditional on the discrete choice j, the backward adjoint decomposition* *applies. Then*

$$dc_0(x) = \underbrace{\sum_{j=1}^{J} \pi^j(x) c_0^j(x) \left[\text{sub}^j(x) + \text{wealth}^j(x) + \text{prec}^j(x)\right]}_{\text{intensive margin}} + \underbrace{\sum_{j=1}^{J} c_0^j(x) d\pi_0^j(x)}_{\text{extensive margin}}. \tag{33}$$

The proof of this Proposition is provided in Appendix A.3.

Proposition 2 shows that discrete choices do not alter the decomposition responses conditional on choosing each specific alternative, but introduces an additional extensive-margin effect that operates via changes in choice probabilities.

Given this, we now apply this result to the two-asset model. Following a perturbation in the liquid-asset return R^b, we can write the response of consumption as

$$dc_0(x) = \underbrace{\pi(x)\, dc_0^a(x) + \big(1 - \pi(x)\big)\, dc_0^n(x)}_{\text{intensive margin}} + \underbrace{\big(c^a(x) - c^n(x)\big)\, d\pi_0(x)}_{\text{extensive margin}}. \tag{34}$$

The intensive margin captures changes in consumption conditional on the adjustment decision, while the extensive margin captures changes in consumption following changes in the probability that a household adjusts their portfolio. The extensive margin is governed by the logit sensitivity of the adjustment probability to the gap between the conditional value functions,

$$d\pi_t(x) = \frac{1}{\sigma_v}\, \pi(x)\big(1 - \pi(x)\big)\big(dV_t^a(x) - dV_t^n(x)\big), \tag{35}$$

derived in Appendix C.3. The extensive margin therefore moves most where the household is near-indifferent and vanishes where adjustment is nearly certain or nearly excluded ($\pi \to 1$ or $\pi \to 0$).

5.3 Decomposition Results

Here

6 Conclusion

Here

References

Farhi, Emmanuel, Alan Olivi, and Iván Werning, "Price Theory for Incomplete Markets," Working Paper 30037, National Bureau of Economic Research 2022.

Huggett, Mark, "The Risk-Free Rate in Heterogeneous-Agent, Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 1993, *17* (5-6), 953–969.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, "Monetary Policy According to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The Power of Forward Guidance Revisited," Working Paper 20882, National Bureau of Economic Research 2015.

Appendices

A Proofs and Additional Results

A.1 Change-of-Measure Identity

When deriving the precautionary component of the decomposition, we used the following change-of-measure identity. For any random variable X measurable with respect to the household's future history:

$$\mathrm{Cov}_0^Q\left(\frac{dQ^I}{dQ}, X\right) \;=\; E_0^{Q^I}[X] - E_0^Q[X]. \tag{36}$$

To get to this result, we start by expanding the covariance using the standard identity $\mathrm{Cov}(Y, Z) = E[Y\,Z] - E[Y]\,E[Z]$ under the measure Q,

$$\mathrm{Cov}_0^Q\left(\frac{dQ^I}{dQ}, X\right) \;=\; E_0^Q\left[\frac{dQ^I}{dQ}\cdot X\right] \;-\; E_0^Q\left[\frac{dQ^I}{dQ}\right]\cdot E_0^Q[X].$$

The first term equals $E_0^{Q^I}[X]$ by the definition of the Radon-Nikodym derivative (change of measure). The second factor satisfies $E_0^Q[dQ^I/dQ] = 1$, since Q^I is a probability measure and the integral of dQ^I/dQ under Q is the total mass of Q^I. Thus, $\mathrm{Cov}_0^Q(dQ^I/dQ,\, X) = E_0^{Q^I}[X] - E_0^Q[X]$.

A.2 Proof of Proposition 1: Backward Representation

We prove Proposition 1 by backward induction. Recall that the definition of the stopping time is:

$$\tau_s \equiv \min\{k \geq s \,:\, X_k \in \mathcal{B}_k\}$$

The stopping time τ_s denotes the first time a household hits the borrowing constraint starting from time s. We now show how to get from Equation (9) to (10).

Base case. Consider $s = T$. If $x \in \mathcal{B}_T$, then $\tau_T = T$. Thus, the recursion in (9) yields $v_T(x) = \alpha_T(x)$. If $x \notin \mathcal{B}_T$, the terminal condition $v_{T+1} = 0$ also implies that $v_T(x) =$

$\alpha_T(x)$. In either case, we have

$$v_T(x) = E_T^{\tilde{Q}}\left[\alpha_T(X_T) \mid X_t = x\right] = \alpha_T(x)$$

This is consistent with the proposed representation given that $T \wedge \tau_T \equiv \min\{T, \tau_T\} = T$.

Inductive step. Suppose that the representation holds at data $s+1$. If $x \in \mathcal{B}_s$, then $\tau_s = s$. Thus,

$$E_s^{\widetilde{Q}}\left[\sum_{k=s}^{T\wedge\tau_s} \alpha_k(X_k) \prod_{j=s}^{k-1} \beta_j(X_j) \,\middle|\, X_s = x\right] = \alpha_s(x),$$

This is exactly the same as the boundary condition of the recursion in the proposition.

Now, suppose that $x \notin \mathcal{B}_s$. Substituting the inductive hypothesis into recursion (9) implies that

$$\begin{aligned} v_s(x) &= \alpha_s(x) + \beta_s(x) \cdot E_s^{\widetilde{Q}}[v_{s+1}(X_{s+1}) \mid X_s = x] \\ &= \alpha_s(x) + \beta_s(x) \cdot E_s^{\widetilde{Q}}\left[E_{s+1}^{\widetilde{Q}}\left[\sum_{k=s+1}^{T\wedge\tau_{s+1}} \alpha_k(X_k) \prod_{j=s+1}^{k-1} \beta_j(X_j) \,\middle|\, X_{s+1}\right] \,\middle|\, X_s = x\right]. \end{aligned}$$

Applying the tower property of conditional expectation, this becomes

$$v_s(x) = \alpha_s(x) + \beta_s(x) \cdot E_s^{\widetilde{Q}}\left[\sum_{k=s+1}^{T\wedge\tau_s} \alpha_k(X_k) \prod_{j=s+1}^{k-1} \beta_j(X_j) \,\middle|\, X_s = x\right].$$

Since $\beta_s(x) = \beta_s(X_s)$ is determined by X_s, it can be moved inside the conditional expectation and absorbed into the product:

$$v_s(x) = \alpha_s(x) + E_s^{\widetilde{Q}}\left[\sum_{k=s+1}^{T\wedge\tau_s} \alpha_k(X_k) \prod_{j=s}^{k-1} \beta_j(X_j) \,\middle|\, X_s = x\right].$$

The first term is the $k = s$ contribution: by the empty-product convention $\prod_{j=s}^{s-1} \beta_j = 1$, so $\alpha_s(x) = \alpha_s(X_s) \prod_{j=s}^{s-1} \beta_j(X_j)$ when $X_s = x$. Folding the $k = s$ term back into the sum yields

$$v_s(x) = E_s^{\widetilde{Q}}\left[\sum_{k=s}^{T\wedge\tau_s} \alpha_k(X_k) \prod_{j=s}^{k-1} \beta_j(X_j) \,\middle|\, X_s = x\right],$$

which matches (10) at s. This completes the induction. $\square$

A.3 Proof of Proposition 2: Decomposition with Discrete Choices

Suppose that a household with state x chooses among J discrete alternatives, and let $\pi^j(x)$ and $c^j(x)$ denote the probability of choosing alternative j, and the optimal consumption policy conditional on j. Ex-post consumption is

$$c_0(x) = \sum_{j=1}^{J} \pi_0^j(x) c_0^j(x).$$

Taking its first-order derivative yields

$$dc_0(x) = \sum_{j=1}^{J} \pi_0^j(x) dc_0^j(x) + \sum_{j=1}^{J} c_0^j(x) d\pi_0^j(x).$$

Assuming that, conditional on j, the backward induction decomposition applies, we obtain

$$\frac{dc_0^j(x)}{c_0^j(x)} = \text{sub}^j(x) + \text{wealth}^j(x) + \text{prec}^j(x)$$

Multiplying by $c_0^j(x)$ and substituting in the derivative of ex-post consumption, we get

$$dc_0(x) = \underbrace{\sum_{j=1}^{J} \pi^j(x) c_0^j(x) \left[\text{sub}^j(x) + \text{wealth}^j(x) + \text{prec}^j(x)\right]}_{\text{intensive margin}} + \underbrace{\sum_{j=1}^{J} c_0^j(x) d\pi_0^j(x)}_{\text{extensive margin}}$$

Which is exactly what Equation (33) states. $\square$

Finally, note that, as $\sum_{j=1}^{J} \pi_0^j(x) = 1$, then $\sum_{j=1}^{J} d\pi_0^j(x) = 0$. Thus, changes in the extensive margin represent reallocation of probability masses across the discrete alternatives.

B Computational Implementation

B.1 General Algorithm

In this appendix, we summarize the implementation of the backward adjoint decomposition from Section 3. After solving the model dynamics under a given sequence of prices, we obtain the resulting policy functions, marginal propensities to save, binding set, and adjusted transition operators. That is, we have

$$\left\{a_s'(x), c_s(x), MPS_s(x), \mathcal{B}_s(x), T_s^{Q}(x), T_s^{Q^I}(x)\right\}_{s=0}^{T}.$$

Algorithm 1 summarizes each step needed to compute the three components of the decomposition at the individual and the aggregate level.

Note that the binding indicator function $m(x)$ takes care of the differential treatment of the continuation value across components. Given that $MPS_s(x) = 0$ when $m_s(x) = 1$, we would have $w_s(x) = 0$. In contrast, for the wealth and prudence-adjusted wealth recursions, even if $m_s(x) = 1$, there is a non-zero contemporary exposure to the price change. That is why $H_s(x) = H_s^I(x) = \gamma_s(x)$ for $x \in \mathcal{B}_s$.

Finally, the computation of the precautionary effect relies on the result presented in Appendix A.1, so that we do not need to compute the covariance explicitly, but the difference between the wealth and prudence-adjusted wealth functions.

B.2 Numerical Implementation Details

We highlight a short list of numerical implementation details relevant for computing the decomposition.

- When numerically identifying the set $\mathcal{B}_s = \{x : a_s'(x) = \underline{a}\}$, it is important to be careful with policy values that are numerically close to $\underline{a}$. We solve the household problem using the Endogenous Grid Method (EGM), which explicitly assigns $a_s'(x) = \underline{a}$ to states at or below the endogenous threshold at which the borrowing constraint begins to bind. Nevertheless, given the further required interpolation of those policy functions onto the distribution grid, it is useful to allow for some tolerance when classifying a state as binding, so $a_s'(x) \leq \underline{a} + \varepsilon_a$, for a small numerical tolerance ε_a, and set $MPS_s(x) = 0$ for these states.

- To obtain MPS_s, we start by computing $MPC_s = \partial_a c_s(x)$, so that $MPS_s = R_s[1 - \partial_a c_s(x)]$. The MPC_s is derived by forward finite difference of the consumption policy with respect to current assets. For small Δa,

$$MPC_s(a, \theta) = \frac{c_s(a + \Delta a, \theta) - c_s(a, \theta)}{\Delta a}$$

  This also yields the date-0 MPC_0, which is also part of the decomposition.

- It is useful to store the transition operators T_s^Q, and $T_s^{Q^I}$ as sparse matrices. Each row of these matrices contains the adjusted transition for each point $x = (a, \theta)$. Thus, after linearly interpolating the optimal saving policy $a_s'(x)$, and taking into account the θ transition probabilities, each row of the operator matrices will have at most $2 \times N_\theta$ non-zero elements. Storing these operators in sparse form therefore reduces both memory requirements and the cost of the matrix-vector products used in the backward recursion.

Algorithm 1: Backward Recursion for the Decomposition

Inputs: Policies $\{a'_s, c_s\}_{s=0}^{T}$, marginal propensity to save $\{MPS_s\}_{s=0}^{T}$, adjusted transition operators $\{T_s^Q, T_s^{Q^I}\}_{s=0}^{T}$, binding indicators $\{m_s\}_{s=0}^{T}$, perturbation path $\{dR_s/R_s\}_{s=0}^{T}$, date-0 marginal propensity to consume, $MPC_0 = \partial_a c_0(x)$, elasticity of intertemporal substitution $\varepsilon(c_0(x))$, stationary distribution μ_0.

Output: $\{w_0(x), H_0(x), H_0^I(x)\}_{x\in\mathcal{X}}$ and aggregate substitution, wealth, and precautionary effects $\widehat{\text{Sub}}, \widehat{\text{Wealth}}, \widehat{\text{Prec}}$.

1. **Compute auxiliary elements.** For all $x \in \mathcal{X}$, compute

$$\kappa_s(x) = \frac{MPS_s(x)}{R_s}, \quad \gamma_s(x) = \frac{a'_s(x)}{R_s}\frac{dR_s}{R_s}, \quad m_s(x) = \mathbf{1}\{x \notin \mathcal{B}_s\}.$$

2. **Initialize.** For all $x \in \mathcal{X}$, set

$$w_{T+1}(x) = 0, \quad H_{T+1}(x) = 0, \quad H_{T+1}^I(x) = 0.$$

3. **Backward loop.** For $s = T, T-1, \ldots, 0$, update the three auxiliary functions:
   a. Substitution:
   $$w_s = m_s \odot \kappa_s \odot \left[\frac{dR_s}{R_s} \times \mathbf{1} + T_s^Q w_{s+1}\right]$$
   b. Wealth:
   $$H_s = \gamma_s + m_s \odot \left(\frac{1}{R_s}\right) T_s^Q H_{s+1}$$
   c. Prudence-adjusted wealth:
   $$H_s^I = \gamma_s + m_s \odot \left(\frac{1}{R_s}\right) T_s^{Q^I} H_{s+1}^I$$

4. **Recover household-level effects.** At $s = 0$, compute

$$\begin{aligned} \text{sub}(x) &= -\varepsilon(c_0(x)) w_0(x) \\ \text{wealth}(x) &= \partial_a c_0(x)/c_0(x) \cdot H_0(x) \\ \text{prec}(x) &= \partial_a c_0(x)/c_0(x) \cdot [H_0^I(x) - H_0(x)] \end{aligned}$$

5. **Aggregate.** Let $C_0 = \sum_{x\in\mathcal{X}} c_0(x)\mu_0(x)$ denote aggregate consumption. For each component $z(x) \in \{\text{sub}, \text{wealth}, \text{prec}\}$, compute

$$\hat{Z} = \sum_{x\in\mathcal{X}} z(x) \frac{c_0(x)\mu_0(x)}{C_0}$$

C Two-Asset Model

C.1 Calibration and Numerical Solution

The solution to the two-asset model uses log-spaced grids for both assets: $N_b = 40$ liquid points over $b \in [0, 50]$ and $N_a = 50$ illiquid points over $a \in [0, 400]$. The stationary distribution and transition operators are built on twice-finer dense grids ($N_b^d \times N_a^d = 80 \times 100$), giving a state space of $N = N_b^d N_a^d N_\theta = 80 \times 100 \times 5 = 40{,}000$ points. We use the same AR(1) specification for the earning process as in Section 4, and discretize the process over N_θ points using the Rouwenhorst method.

Table 5 presents the model parameterization. We calibrate the model so that ...

Table 5: Two-Asset Model Benchmark Calibration

Parameter	Symbol	Value
Discount factor	β	0.87
CRRA curvature	σ	4.0
Earnings innovation std	σ_θ	$\sqrt{0.037}$
Earnings persistence	ρ_θ	0.87
Borrowing limit	$\underline{b}$	0.0
Liquid asset return	R^b	1.01
Illiquid asset return	R^a	1.03
Adjustment cost parameters	χ_0	0.2
	χ_1	1.5
	χ_2	2.0
Stdev of EV shocks	σ_ν	0.005
Asset-to-income ratio	A/Y	XXX

Two-stage solution. Given the discrete adjustment structure, we solve the household problem by using a nested EGM following Kaplan et al. (2018), Appendix E. Given value function derivatives $V_a(x) \equiv \partial V(x)/\partial a$, and $V_b(x) \equiv \partial V(x)/\partial b$, the inner stage recovers optimal policies (a', b', c) jointly from the first-order conditions

$$
\begin{aligned}
u'(c) &= \beta\, E_{\theta'|\theta}\left[V_b(a', b', \theta')\right], \\
u'(c)\Big(1 + \chi_{a'}(a', a)\Big) &= \beta\, E_{\theta'|\theta}\left[V_a(a', b', \theta')\right].
\end{aligned}
$$

We use a kappa-grid of $N_\kappa = 8$ constrained-multiplier values to handle the liquid borrowing constraint $b' \geq 0$. The outer stage combines the resulting conditional value functions to update (V_a, V_b) for the next backward iteration.

C.2 From the One-Asset Model to the Two-Asset Model

The household problem in the two-asset model has two endogenous asset choices conditional on each discrete choice $j \in \{a, n\}$. These policies are[2]

$$(a')^j_s(a,b,\theta), \quad (b')^j_s(a,b,\theta), \quad \text{for } j \in \{a,n\}.$$

As in the one-asset model, asset choices may lie off-grid and require using bilinear interpolation when constructing the adjusted transition operators. This means that for each row of the operators $T_s^{Q,j}$ and $T_s^{Q^I,j}$, we may have, at most, $4 \times N_\theta$ non-zero elements.

C.3 Computing the Extensive Margin Response

The derivative of the adjustment probabilities in (35), that is, the extensive-margin sensitivity, is derived from the adjustment probability in (26):

$$\pi_s(x) = \frac{\exp(V_s^a(x)/\sigma_\nu)}{\exp(V_s^a(x)/\sigma_\nu) + \exp(V_s^n(x)/\sigma_\nu)}.$$

Dividing numerator and denominator by $\exp(V_s^a(x)/\sigma_\nu)$ and writing the value gap $\Delta_s(x) \equiv V_s^a(x) - V_s^n(x)$ yields,

$$\pi_s(x) = \frac{1}{1+\exp(-\Delta_s(x)/\sigma_\nu)} = L\left(\frac{\Delta_s(x)}{\sigma_\nu}\right), \qquad L(z) \equiv \frac{1}{1+e^{-z}}.$$

Where $L(z)$ is the logistic function, and satisfies $L'(z) = L(z)\big(1 - L(z)\big)$. Differentiating around the stationary equilibrium, only the value gap moves while σ_ν is fixed, so

$$d\left(\frac{\Delta_s(x)}{\sigma_\nu}\right) = \frac{1}{\sigma_\nu}\, d\Delta_s(x) = \frac{1}{\sigma_\nu}\big(dV_s^a(x) - dV_s^n(x)\big).$$

The chain rule implies $d\pi_s(x) = L'(\Delta_s(x)/\sigma_\nu)\, d(\Delta_s(x)/\sigma_\nu)$. Thus,

$$d\pi_s(x) = \frac{1}{\sigma_\nu}\, \pi_s(x)\big(1 - \pi_s(x)\big)\big(dV_s^a(x) - dV_s^n(x)\big),$$

which is (35). The factor $\pi(1-\pi)$ is the variance of the Bernoulli adjust/no-adjust choice, so the extensive margin is most sensitive near indifference.

[2]Conditional on not adjusting, $(a')^j_s(a,b,\theta) = R^a a$, which may also lay off-grid.

To compute the differential of the choice probabilities at date-0, which is what we ultimately need in Equation (34), we use the exact finite change as a way of numerically approximating its value to the first order, so that

$$d\pi_0(x) \approx \Delta\pi_0(x) = \pi_0(x) - \pi(x).$$

D Computational Details and Runtime

Working note: line-by-line derivations of the closed forms

The recursion is

$$v_s(x) = \begin{cases} \bar{v}, & x \in \mathcal{B}_s, \\ \alpha_s(x) + \beta_s(x) \cdot E^{\widetilde{Q}}[v_{s+1}(X_{s+1}) \mid X_s = x], & x \notin \mathcal{B}_s, \end{cases}$$

with terminal condition $v_{T+1}(x) \equiv v_\infty$. The derivations below trace iterations on the unconstrained set; the absorbing boundary truncates the sum or product at $\tau - 1$ because terms with $k \geq \tau$ inherit the absorbing $\bar{v} = 0$ (additive) or interact with the multiplicative chain to terminate it.

Additive case (α_s active, $\bar{v} = v_\infty = 0$)

Step $s = T$**.**

$$v_T(x) = \alpha_T(x) + \beta_T(x) \cdot \underbrace{E^{\widetilde{Q}}[v_{T+1} \mid X_T = x]}_{=0} = \alpha_T(x).$$

Step $s = T - 1$**.**

$$v_{T-1}(x) = \alpha_{T-1}(x) + \beta_{T-1}(x) \cdot E^{\widetilde{Q}}[v_T(X_T) \mid X_{T-1} = x].$$

Substituting $v_T(X_T) = \alpha_T(X_T)$:

$$v_{T-1}(x) = \alpha_{T-1}(x) + \beta_{T-1}(x) \cdot E^{\widetilde{Q}}[\alpha_T(X_T) \mid X_{T-1} = x].$$

Since $\beta_{T-1}(x)$ is deterministic given $X_{T-1} = x$, pull it inside:

$$v_{T-1}(x) = \alpha_{T-1}(x) + E^{\widetilde{Q}}[\beta_{T-1}(X_{T-1})\, \alpha_T(X_T) \mid X_{T-1} = x].$$

Step $s = T - 2$**.**

$$v_{T-2}(x) = \alpha_{T-2}(x) + \beta_{T-2}(x) \cdot E^{\widetilde{Q}}[v_{T-1}(X_{T-1}) \mid X_{T-2} = x].$$

Substituting the previous step:

$$v_{T-2}(x) = \alpha_{T-2}(x) + \beta_{T-2}(x) \cdot E^{\widetilde{Q}}\Big[\alpha_{T-1}(X_{T-1}) + \beta_{T-1}(X_{T-1}) \cdot E^{\widetilde{Q}}[\alpha_T(X_T) \mid X_{T-1}] \;\Big|\; X_{T-2} = x\Big].$$

Apply the tower property ($E[E[\cdot|X_{T-1}]|X_{T-2}] = E[\cdot|X_{T-2}]$):

$$\begin{aligned} v_{T-2}(x) = \alpha_{T-2}(x) &+ \beta_{T-2}(x) \cdot E^{\widetilde{Q}}[\alpha_{T-1}(X_{T-1}) \mid X_{T-2} = x] \\ &+ \beta_{T-2}(x) \cdot E^{\widetilde{Q}}[\beta_{T-1}(X_{T-1})\, \alpha_T(X_T) \mid X_{T-2} = x]. \end{aligned}$$

Pull $\beta_{T-2}(x)$ inside (deterministic given $X_{T-2} = x$):

$$\begin{aligned} v_{T-2}(x) = \alpha_{T-2}(x) &+ E^{\widetilde{Q}}[\beta_{T-2}(X_{T-2})\, \alpha_{T-1}(X_{T-1}) \mid X_{T-2} = x] \\ &+ E^{\widetilde{Q}}[\beta_{T-2}(X_{T-2})\, \beta_{T-1}(X_{T-1})\, \alpha_T(X_T) \mid X_{T-2} = x]. \end{aligned}$$

Pattern. Each term has the form $E^{\widetilde{Q}}\big[\alpha_k(X_k) \prod_{j=s}^{k-1} \beta_j(X_j) \mid X_s = x\big]$ for $k = s, s+1, \ldots, T$. The $k = s$ term uses the empty-product convention $\prod_{j=s}^{s-1} \equiv 1$, giving $\alpha_s(x)$.

Closed form (with absorbing boundary).

$$\boxed{v_s(x) = E^{\widetilde{Q}}\left[\sum_{k=s}^{T\wedge(\tau-1)} \alpha_k(X_k) \prod_{j=s}^{k-1} \beta_j(X_j) \;\middle|\; X_s = x\right].}$$

Multiplicative case ($\alpha_s = 0$, $\bar{v} = v_\infty \neq 0$)

Step $s = T$.

$$v_T(x) = \beta_T(x) \cdot \underbrace{E^{\widetilde{Q}}[v_{T+1} \mid X_T = x]}_{=v_\infty} = \beta_T(x) \cdot v_\infty.$$

Step $s = T - 1$.

$$v_{T-1}(x) = \beta_{T-1}(x) \cdot E^{\widetilde{Q}}[v_T(X_T) \mid X_{T-1} = x].$$

Substituting $v_T(X_T) = \beta_T(X_T) \cdot v_\infty$:

$$v_{T-1}(x) = \beta_{T-1}(x) \cdot E^{\widetilde{Q}}[\beta_T(X_T) \mid X_{T-1} = x] \cdot v_\infty = E^{\widetilde{Q}}[\beta_{T-1}(X_{T-1})\, \beta_T(X_T) \mid X_{T-1} = x] \cdot v_\infty.$$

Step $s = T - 2$**.**

$$v_{T-2}(x) = \beta_{T-2}(x) \cdot E^{\widetilde{Q}}[v_{T-1}(X_{T-1}) \mid X_{T-2} = x].$$

Substituting the previous step:

$$v_{T-2}(x) = \beta_{T-2}(x) \cdot E^{\widetilde{Q}}\Big[E^{\widetilde{Q}}[\beta_{T-1}(X_{T-1})\, \beta_T(X_T) \mid X_{T-1}] \cdot v_\infty \;\Big|\; X_{T-2} = x\Big].$$

Apply the tower property and pull out the constant v_∞:

$$v_{T-2}(x) = \beta_{T-2}(x) \cdot E^{\widetilde{Q}}[\beta_{T-1}(X_{T-1})\, \beta_T(X_T) \mid X_{T-2} = x] \cdot v_\infty.$$

Pull $\beta_{T-2}(x)$ inside:

$$v_{T-2}(x) = E^{\widetilde{Q}}[\beta_{T-2}(X_{T-2})\, \beta_{T-1}(X_{T-1})\, \beta_T(X_T) \mid X_{T-2} = x] \cdot v_\infty.$$

Pattern. Each step adds one more β factor inside the expectation; the constant v_∞ rides through unchanged.

Closed form (with absorbing boundary).

$$\boxed{v_s(x) = E^{\widetilde{Q}}\left[\prod_{j=s}^{T\wedge(\tau-1)} \beta_j(X_j) \cdot v_\infty \;\middle|\; X_s = x\right].}$$

Wealth channel: instantiation of the additive case

For the wealth channel, $H_s(x) \equiv v_s(x)$ with parameters $\alpha_s(x) = (a_s'(x)/R_s) \cdot (dR_s/R_s)$, $\beta_s(x) = 1/R_s$, $\bar{v} = v_\infty = 0$, $T_s = T_s^Q$. The recursion is

$$H_s(x) = \frac{a_s'(x)}{R_s} \cdot \frac{dR_s}{R_s} + \frac{1}{R_s} \cdot E^Q[H_{s+1}(X_{s+1}) \mid X_s = x] \qquad \text{on } x \notin \mathcal{B}_s,$$

with $H_s(x) = 0$ on $x \in \mathcal{B}_s$ and terminal $H_{T+1}(x) \equiv 0$.

Step $s = T$**.**

$$H_T(x) = \frac{a_T'(x)}{R_T} \cdot \frac{dR_T}{R_T} + \underbrace{\frac{1}{R_T} \cdot E^Q[H_{T+1} \mid X_T = x]}_{=0} = \frac{a_T'(x)}{R_T} \cdot \frac{dR_T}{R_T}.$$

Step $s = T - 1$**.**

$$H_{T-1}(x) = \frac{a'_{T-1}(x)}{R_{T-1}} \cdot \frac{dR_{T-1}}{R_{T-1}} + \frac{1}{R_{T-1}} \cdot E^Q[H_T(X_T) \mid X_{T-1} = x].$$

Substituting $H_T(X_T) = (a'_T(X_T)/R_T)(dR_T/R_T)$ and pulling $1/R_{T-1}$ inside (deterministic given $X_{T-1} = x$):

$$H_{T-1}(x) = \frac{a'_{T-1}(x)}{R_{T-1}} \cdot \frac{dR_{T-1}}{R_{T-1}} + E^Q\left[\frac{1}{R_{T-1}} \cdot \frac{a'_T(X_T)}{R_T} \cdot \frac{dR_T}{R_T} \,\middle|\, X_{T-1} = x\right].$$

Step $s = T - 2$**.**

$$H_{T-2}(x) = \frac{a'_{T-2}(x)}{R_{T-2}} \cdot \frac{dR_{T-2}}{R_{T-2}} + \frac{1}{R_{T-2}} \cdot E^Q[H_{T-1}(X_{T-1}) \mid X_{T-2} = x].$$

Substituting H_{T-1} and applying the tower property:

$$\begin{aligned} H_{T-2}(x) = {} & \frac{a'_{T-2}(x)}{R_{T-2}} \cdot \frac{dR_{T-2}}{R_{T-2}} \\ & + E^Q\left[\frac{1}{R_{T-2}} \cdot \frac{a'_{T-1}(X_{T-1})}{R_{T-1}} \cdot \frac{dR_{T-1}}{R_{T-1}} \,\middle|\, X_{T-2} = x\right] \\ & + E^Q\left[\frac{1}{R_{T-2}} \cdot \frac{1}{R_{T-1}} \cdot \frac{a'_T(X_T)}{R_T} \cdot \frac{dR_T}{R_T} \,\middle|\, X_{T-2} = x\right]. \end{aligned}$$

Pattern. Each term has the form $E^Q\big[\big(\prod_{j=s}^{k-1} \frac{1}{R_j}\big) \cdot \frac{a'_k(X_k)}{R_k} \cdot \frac{dR_k}{R_k} \mid X_s = x\big]$ for $k = s, s+1, \ldots, T$. Combining the discount factors,

$$\prod_{j=s}^{k-1} \frac{1}{R_j} \cdot \frac{1}{R_k} = \prod_{j=s}^{k} \frac{1}{R_j},$$

so each term can equivalently be written $E^Q\big[a'_k(X_k) \cdot \prod_{j=s}^{k} \frac{1}{R_j} \cdot \frac{dR_k}{R_k} \mid X_s = x\big]$.

Closed form (sequence of shocks).

$$\boxed{H_s(x) = E^Q\left[\sum_{k=s}^{T\wedge(\tau-1)} a'_k(X_k) \cdot \prod_{j=s}^{k} \frac{1}{R_j} \cdot \frac{dR_k}{R_k} \,\middle|\, X_s = x\right].}$$

Each date the household is unconstrained contributes its asset position a'_k, discounted from s to k at the gross rates, weighted by the shock at that date.

One-time shock at date t. Setting $dR_s/R_s = 0$ for $s \neq t$ collapses the sum to its $k = t$ term:

$$\boxed{H_s(x) = \frac{dR_t}{R_t} \cdot E^Q\left[a'_t(X_t) \cdot \prod_{j=s}^{t} \frac{1}{R_j} \,\middle|\, X_s = x\right] \qquad \text{for } s \leq t.}$$

This is the discounted expected asset position at date t, multiplied by the date-t shock magnitude. The additive form handles both scenarios with the same $\beta_s = 1/R_s$ structure — no multiplicative recursion is needed.

Summary

Case	Driver	Boundary	What accumulates
Additive	$\alpha_s \neq 0$	$\bar{v} = v_\infty = 0$	A sum of dated payoffs α_k, each compounded by past β_j's
Multiplicative	$\alpha_s = 0$	$\bar{v} = v_\infty \neq 0$	A product of β_j's, scaled by the terminal v_∞