Dynamics of Constraint Enforcement and Consumption Responses

Job Boerma
Wisconsin

Dan Mikulincer
University of Washington

Aleh Tsyvinski
Yale*

March 19, 2026

Abstract

We study the dynamic structure of feasibility constraints that bind intermittently. The standard approach characterizes optimality by showing that discounted marginal utility is a supermartingale, which restricts the process but leaves the structure of its dynamics compressed. We formulate the problem in path space and show that the key object is a shadow constraints value process that aggregates expected discounted future constraint pressure. This process admits a semimartingale decomposition into a regular component driven by Euler drift and shocks and a singular component generated by constraint enforcement, whose dynamics we characterize. We develop these results in a continuous-time consumption–savings model with borrowing constraints. We show that consumption responses to income innovations (shock MPCs) can be written as the sensitivity of expected future constraint pressure to current shocks and decompose them into a mean-exposure component and a covariance component. The covariance term captures a curvature-driven amplification mechanism: shocks have large effects when they shift future asset positions toward regions where constraint curvature is high. We then show that impulse responses satisfy an integral equation whose iterated-kernel expansion determines a recursive propagation–feedback mechanism linking shocks, state dynamics, and constraint enforcement interact.

*We are grateful to Joel Flynn for detailed comments.

1 Introduction

Many economic environments are shaped by feasibility constraints that bind intermittently—households face borrowing limits, firms face financing constraints, and intermediaries face liquidity or collateral requirements. The standard approach analyzes these models through dynamic programming, under which the Euler equation becomes a variational inequality and discounted marginal utility becomes a supermartingale. This characterization, however, leaves the behavior of the model at the borrowing constraint implicit. A supermartingale is a restriction on a process, not a description of its structure: it establishes that constraints matter, but it does not describe how constraint enforcement enters the dynamics or how its effects propagate over time. This paper studies the dynamic structure of constraint enforcement and how constraints influence the dynamics of the model. Our main point is that constraint enforcement can be analyzed as a distinct stochastic process that admits a semimartingale decomposition into a regular component driven by Euler drift and shocks, and a singular component generated by constraints. This decomposition allows us to isolate the role of constraints and study how they affect model dynamics beyond the supermartingale or variational inequality characterization, which compresses constraint enforcement into a boundary condition.

We develop these results in a standard continuous-time income-fluctuation problem. We formulate optimality in path space and show that the key object is a shadow value that aggregates the expected future pressure of the borrowing constraints. Marginal utility equals the expected discounted value of all future constraint enforcement—not just whether the constraint binds today, but how likely it is to bind in the future across all time horizons.

Our first result (Proposition 1) gives the semimartingale decomposition of this shadow constraint value. This decomposition separates three forces: a standard Euler drift, a predictable enforcement (compensator) that applies only when assets hit the boundary, and a stochastic innovation that captures revisions in expected future constraint pressure caused by current income shocks. This decomposition holds globally – both in the interior and at the borrowing boundary. The constrained Euler equation is thus not determined only by a supermartingale or a variational inequality; it has a drift–compensator–innovation structure that is valid both in the interior and at the borrowing boundary.

The innovation coefficient governs consumption responses to shocks. We define the shock MPC as the instantaneous response of consumption to the income innovation and develop its structure in three steps.

First, we introduce a smooth penalized version of the hard borrowing constraints, which makes the problem differentiable and allows us to apply Malliavin calculus. Malliavin calcu-

lus provides a way to take derivatives of stochastic processes with respect to the underlying shocks and, in particular, to express the coefficient on the innovation as the conditional expectation of that derivative, so that the innovation coefficient can be interpreted as an impulse response to shocks. The shock MPC can then be written as the sensitivity of discounted future constraint pressure to the current innovation (Proposition 2 and Proposition 3). This sensitivity depends on two objects. The first is the sensitivity of future assets to the current shock—how the innovation propagates forward through the budget constraint. The second is the curvature of the constraint, which is large near the borrowing boundary and small away from it – it is a smooth representation of marginal effects of the hard constraint. Because curvature is concentrated near the borrowing boundary, shocks matter to the extent that they move future asset positions when the household is in regions where constraint curvature is high. Far from the boundary, constraint curvature is negligible and shocks barely affect expected constraint pressure; near the boundary, even small shocks generate large revisions.

Second, we show that consumption responses depend not only on average exposure to shocks, but on whether a current shock shifts future asset holdings toward states where the borrowing constraint is more curved. Specifically, we decompose the shock MPC into two components (Proposition 4). The average, or the mean-exposure, component combines expected constraint curvature with expected asset sensitivity. The covariance component captures something subtler: whether asset sensitivities are systematically different in states where the constraint is tighter. A household can have large average asset sensitivity and still have a modest MPC, if that sensitivity occurs primarily in states far from the boundary. Conversely, moderate asset sensitivity concentrated near the constraint can generate a large MPC. This is constraint curvature-driven amplification: a shock has a large effect not because it is large on impact, but because it moves the household into states where future constraint curvature is large. The standard solutions contain this information implicitly but do not separately identify this amplification channel, as they combine average exposure and state-dependent covariance effects into a single local object.

The covariance component also admits a beta representation (Corollary 5) that separates dispersion in future constraint curvature from the loading of asset exposure on that dispersion. In Corollary 6 we derive a bound on the covariance component of the shock MPCs.

These results also have a natural interpretation in terms of option-like payoffs. For a smooth claim, average exposure provides a good approximation to the response of the shocks. By contrast, for barrier or digital claims, sensitivities are sharply concentrated near the strike or barrier. Responses to shocks depend not only on average exposure but also on how exposure co-moves with the region where sensitivity is largest. The shock MPC

plays the role of the "delta" of the constraint claim with respect to the underlying innovation process and has a similar delta spike or delta concentration behavior present in pricing barrier or digital options. The borrowing constraint generates an analogous boundary layer concentration: responses depend on how shocks move the household into regions where constraint curvature is concentrated.

Third, we study the impulse responses. Specifically, we define the impulse-response MPC as the response of consumption at future dates to the current innovation. We show that impulse responses satisfy an integral equation (Proposition 7) jointly determining shock propagation, asset dynamics, and constraint curvature. Proposition 8 shows that this integral equation admits an iterated-kernel expansion, and determines a recursive propagation–feedback mechanism through constraint enforcement. The iterated kernel captures successive rounds of feedback generated by the shock. The first term captures the direct effect of the shock. Higher-level terms capture further rounds of feedback of the exogenous shocks: initial responses change future states, which changes expected future constraint tightness along the curvature kernel, which in turn generate further responses. That is, constrained dynamics is represented through the recursive interaction between shocks and constraint enforcement determined by the iterated kernels.

Our analysis builds on the continuous-time incomplete-markets framework of Achdou, Han, Lasry, Lions, and Moll (2022). In that framework, discounted marginal utility is a supermartingale and the constraint is encoded in a boundary condition on the value function. We derive a path-space decomposition of the marginal-value process that remains well defined globally, including at the boundary. We are able to not only state that the discounted marginal utility is a supermartingale but also to characterize the structure of enforcement constraint dynamics. More broadly, our analysis does not rely on the recursive formulation of the problem and instead uses the path-space optimization as in Boerma, Riabov, and Tsyvinski (2023).

Methodologically, the closest connection is to Borovička, Hansen, and Scheinkman (2014) and Borovička and Hansen (2016), who use Malliavin calculus and Clark–Ocone representations to characterize shock elasticities in environments without borrowing constraints. In our model, borrowing constraints generate an enforcement component absent in the unconstrained case. This allows us to move from a representation of shock exposure to a characterization of the effects of constraints and their dynamics.

Our paper is also related to an important paper by Farhi, Olivi, and Werning (2022), who characterize consumption responses using stopped Euler equations and study the interior martingale on the random interval before the constraint first binds. By contrast, we work with the unstopped marginal-value process and its semimartingale decomposition, which

retains the entire future path of constraint enforcement. In this formulation, a shock affects not only whether and when the first constraint binds but also the distribution of all future binding episodes and their intensity. Our framework captures how shocks propagate through this entire path of constraints, which is what gives rise to the covariance decomposition and the recursive propagation–feedback mechanism.

Our covariance decomposition relates to the literature on MPCs and exposure in heterogeneous-agent models. Kaplan and Violante (2022) describe the central role of MPCs for aggregate dynamics and document how borrowing constraints, liquidity, and heterogeneity determine their magnitude and dispersion. Auclert (2019) shows that aggregate responses depend on the covariance between MPCs and balance-sheet exposures, and Patterson (2023) shows that this covariance is quantitatively important. These results operate across households. Our covariance term operates within a household across future states. It isolates a different mechanism: curvature-driven amplification arising when shocks shift the distribution of future asset positions toward states where the constraint curvature is high.

Our impulse-response kernel connects to the literature on intertemporal MPCs and sequence-space methods. Auclert, Rognlie, and Straub (2024) show that intertemporal MPCs summarize dynamic responses to policy, and Auclert, Bardóczy, Rognlie, and Straub (2021) develop the sequence-space Jacobian approach for representing dynamic responses in heterogeneous-agent models. Bilal and Goyal (2025) derive analytic continuous-time sequence-space Jacobians. Our object is not a general-equilibrium Jacobian. Rather, it is an individual-level stochastic forward kernel that determines the micro propagation mechanism—through future constraint enforcement and curvature—that underlies such aggregate objects.

2 Environment

We consider a continuous-time income-fluctuation problem with borrowing constraints. Fix a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ supporting a one-dimensional Brownian motion B_t, where $(\mathcal{F}_t)$ is the completed filtration generated by B_t. All processes are adapted to $(\mathcal{F}_t)$. The interest rate $r > 0$ is exogenous and constant.

2.1 Preferences

A household chooses an adapted consumption process $c_t \geq 0$ to maximize expected discounted utility

$$\mathbb{E}\left[\int_0^\infty e^{-\rho t} u(c_t)\, dt\right],$$

where $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing, strictly concave, and twice continuously differentiable, and $\rho > 0$ is the discount rate. We assume $\rho > r$.

2.2 Income

Income follows a Markov diffusion on $(0, \infty)$

$$dy_t = \mu(y_t)\, dt + \sigma(y_t)\, dB_t,$$

where μ, σ are such that, for every $y_0 > 0$, the SDE admits a unique strong solution remaining in $(0, \infty)$ almost surely for all $t \geq 0$.

2.3 Assets and borrowing constraints

Assets a_t evolve as

$$da_t = (ra_t + y_t - c_t)\, dt,$$

where a_0 is given, subject to borrowing constraints

$$a_t \geq 0 \quad \text{for } t \geq 0.$$

Equivalently,

$$a_t(c) = e^{rt} a_0 + \int_0^t e^{r(t-s)} (y_s - c_s)\, ds. \tag{1}$$

2.4 Feasible allocations

A consumption process c is feasible if

$$c_t \geq 0 \quad \text{and} \quad a_t(c) \geq 0 \qquad \text{for all } t \geq 0.$$

Because the asset path is uniquely determined by consumption from (1), the household's problem reduces to choosing c subject to

$$a_t(c) \geq 0 \qquad \text{for all } t \geq 0.$$

2.5 Household problem

The household's optimization problem is

$$\sup_{c_t \geq 0} \mathbb{E}\left[\int_0^\infty e^{-\rho t} u(c_t)\, dt\right]$$

subject to

$$a_t(c) \geq 0 \qquad \text{for all } t \geq 0.$$

3 Optimality and the Shadow Value of Borrowing Constraints

We derive the first-order conditions in path space and show that optimality is characterized by a shadow state that aggregates the expected future increments of the constraint multiplier and determines the evolution of marginal utility.

3.1 Lagrangian formulation

Let η_t be a nondecreasing adapted finite-variation process enforcing the borrowing constraint:

$$\eta_0 = 0, \qquad d\eta_t \geq 0, \qquad \int_0^\infty e^{-\rho t} a_t(c)\, d\eta_t = 0.$$

Since $e^{-\rho t} > 0$, complementary slackness is equivalent to $\int_0^\infty a_t(c)\, d\eta_t = 0$.

The Lagrangian is

$$\mathcal{L}(c, \eta) = \mathbb{E}\left[\int_0^\infty e^{-\rho t} u(c_t)\, dt\right] + \mathbb{E}\left[\int_0^\infty e^{-\rho t} a_t(c)\, d\eta_t\right]. \tag{2}$$

Let $c^\varepsilon = c + \varepsilon \delta c$, where δc is a bounded adapted perturbation. The variation in assets is

$$\delta a_t = -\int_0^t e^{r(t-s)} \delta c_s\, ds. \tag{3}$$

Differentiating (2),

$$\delta \mathcal{L} = \mathbb{E}\left[\int_0^\infty e^{-\rho t} u'(c_t) \delta c_t\, dt\right] + \mathbb{E}\left[\int_0^\infty e^{-\rho t} \delta a_t\, d\eta_t\right].$$

Substituting (3),

$$\delta\mathcal{L} = \mathbb{E}\left[\int_0^\infty e^{-\rho t} u'(c_t)\delta c_t \, dt\right] - \mathbb{E}\left[\int_0^\infty e^{-\rho t}\left(\int_0^t e^{r(t-s)}\delta c_s \, ds\right) d\eta_t\right].$$

Since $d\eta_t$ has finite variation, we may interchange integrals:

$$\delta\mathcal{L} = \mathbb{E}\left[\int_0^\infty \delta c_s \, e^{-\rho s}\left(u'(c_s) - \int_s^\infty e^{-(\rho-r)(t-s)} \, d\eta_t\right) ds\right].$$

Because δc_s is $\mathcal{F}_s$-measurable and otherwise arbitrary, the first-order condition is given by taking conditional expectations at time s:

$$u'(c_t) = \mathbb{E}_t\left[\int_t^\infty e^{-(\rho-r)(s-t)} \, d\eta_s\right].$$

Define the **shadow continuation value**

$$\boxed{W_t := \mathbb{E}_t\left[\int_t^\infty e^{-(\rho-r)(s-t)} \, d\eta_s\right].}$$

so the optimality condition becomes

$$\boxed{u'(c_t) = W_t.} \tag{4}$$

Equation (4) is the Euler equation – marginal utility equals the expected discounted value of all future constrains. The shadow value process W_t aggregates the entire expected future path of constraint binding into a single continuation object that depends on the full future evolution of the multiplier measure η, and not only on whether the constraint binds today.

What W_t prices is not simply current constraint binding, but the entire expected future path of enforcement. This is a forward-looking shadow value of constraint feasibility: it is high when the household expects constraint pressure to remain strong in the future, and low when future enforcement is unlikely. This forward representation will be central for everything that follows. It is what allows us to move from the statement that discounted marginal utility is a supermartingale to a more explicit description of the dynamics of constraint enforcement and of its effects on responses to shocks.

Remark. We remark that $W_t > 0$ almost surely. To see this, it would be helpful to consider a variant of the problem with a finite horizon T. In this variant, it is clear that for the optimal plan c_t the assets must satisfy $a_T(c) = 0$. Otherwise, one could increase c_t on a small neighborhood of T while preserving feasibility, contradicting optimality. It follows that, for

every $t < T$, the multiplier measure assigns positive mass to $(t, T]$, and hence $W_t > 0$. In the infinite-horizon problem, the analogous statement is that $\liminf_{t\to\infty} e^{-rt} a_t(c) = 0$. The corresponding terminal multiplier at infinity then contributes to W_t, yielding the same conclusion that $W_t > 0$ almost surely.

3.2 Dynamics of the shadow continuation value: Semimartingale decomposition

We now derive the semimartingale decomposition of the shadow state W_t. We we that that the process has a drift–compensator–innovation structure that is valid both in the interior and at the borrowing boundary.

Because income remains strictly positive and the optimal consumption path is strictly positive on every finite interval, we have $u'(c_t) < \infty$ for every finite t. Hence, by the Euler equation

$$W_t = u'(c_t) < \infty \qquad \text{for every } t < \infty.$$

Assume $W_0 < \infty$.

Proposition 1 (Semimartingale decomposition of the shadow state)**.** *The shadow constraint process W_t is a semimartingale and admits the decomposition*

$$\boxed{dW_t = (\rho - r)W_t\, dt - d\eta_t + Z_t\, dB_t,} \tag{5}$$

for some adapted process Z_t.

Proof. From the definition of W_t,

$$e^{-(\rho-r)t} W_t = \mathbb{E}_t\left[\int_t^\infty e^{-(\rho-r)s}\, d\eta_s\right].$$

Define

$$M_t := \mathbb{E}_t\left[\int_0^\infty e^{-(\rho-r)s}\, d\eta_s\right] = e^{-(\rho-r)t} W_t + \int_0^t e^{-(\rho-r)s}\, d\eta_s.$$

Then M_t is a martingale. By the martingale representation theorem, there exists an adapted process $\widetilde{Z}_t$ such that

$$dM_t = \widetilde{Z}_t\, dB_t.$$

Differentiating

$$e^{-(\rho-r)t} W_t = M_t - \int_0^t e^{-(\rho-r)s}\, d\eta_s$$

gives

$$d\big(e^{-(\rho-r)t}W_t\big) = -e^{-(\rho-r)t}\,d\eta_t + \widetilde{Z}_t\,dB_t.$$

Applying Itô's product rule to

$$W_t = e^{(\rho-r)t}\big(e^{-(\rho-r)t}W_t\big)$$

gives

$$dW_t = (\rho - r)W_t\,dt - d\eta_t + e^{(\rho-r)t}\widetilde{Z}_t\,dB_t.$$

Defining

$$Z_t := e^{(\rho-r)t}\widetilde{Z}_t$$

gives the result. □

Combined with the Euler equation, Proposition 1 gives:

$$d(u'(c_t)) = (\rho - r)u'(c_t)\,dt - d\eta_t + Z_t\,dB_t. \tag{6}$$

When the borrowing constraint is slack, $d\eta_t = 0$, marginal utility follows the standard continuous-time Euler dynamics with drift $(\rho - r)u'(c_t)\,dt$. When the borrowing constraint binds, the process acquires the predictable boundary-supported compensator $-d\eta_t$. The innovation term $Z_t\,dB_t$ captures revisions in expected future constraint binding.

This decomposition separates regular and constraint-driven components of the marginal-value dynamics. The Euler drift and the innovation term describe how marginal utility evolves through standard intertemporal substitution and shocks. The compensator isolates the distinct contribution of constraint feasibility itself: it is the predictable adjustment generated by the borrowing boundary. The constrained Euler condition is therefore not only a restriction on the sign of drift after discounting. It also identifies the specific channel through which constraint enforcement enters the dynamics.

4 Shock MPCs

Section 3 separates the dynamics of marginal utility into three parts: an Euler drift, a constraint compensator, and a stochastic innovation. The drift and compensator govern the predictable evolution of consumption. The innovation $Z_t dB_t$ governs the response to shocks. This section uses the innovation coefficient to define the shock marginal propensity to consume.

4.1 Consumption response to shocks

Start from the Euler equation

$$u'(c_t) = W_t.$$

In order to characterize the response of consumption to shocks, differentiate this equation using the Malliavin derivative[1] with respect to the innovation at time t:

$$u''(c_t)\, D_t c_t = D_t W_t.$$

Solving for the consumption response:

$$D_t c_t = \frac{1}{u''(c_t)} D_t W_t.$$

The response of consumption to the innovation dB_t at time t is determined by the sensitivity of the shadow value of enforcement constraints, W_t, to that innovation.

From Proposition 1, the shadow value satisfies the semimartingale decomposition

$$dW_t = (\rho - r)W_t\, dt - d\eta_t + Z_t\, dB_t.$$

Under the standard regularity conditions, the diagonal Malliavin derivative of W_t coincides with the martingale integrand:

$$D_t W_t = Z_t.$$

Hence,

$$\boxed{D_t c_t = \frac{1}{u''(c_t)}\, Z_t.} \tag{7}$$

This shows that the response of consumption to the innovation dB_t is governed by the innovation coefficient Z_t in the semimartingale decomposition of the shadow continuation value.

4.2 Definition of shock MPC

We define the **shock marginal propensity to consume** as the instantaneous response of consumption to the Brownian innovation:

$$\boxed{\mathrm{MPC}_t^{\mathrm{shock}} := D_t c_t.}$$

[1] In this section, we take Malliavin derivatives formally. The rigorous Mallivin differentiability is given in the penalized formulation of Section 5. The derivative here is understood as the limit of Malliavin derivative in Section 5.

Using (7) gives

$$\boxed{\mathrm{MPC}_t^{\mathrm{shock}} = \frac{1}{u''(c_t)}\, Z_t.}$$

The shock MPC is determined by the innovation coefficient Z_t in the semimartingale decomposition of the shadow continuation value, adjusted by the curvature of the utility function. The innovation term $Z_t dB_t$ captures revisions in expectations about future constraint binding generated by the current shock.[2]

4.3 Relation to the HJB formulation

In a Markovian formulation with policy function $c(a, y)$, Itô's lemma in the interior gives that the diffusion component of consumption satisfies

$$dc_t^{(dB)} = \partial_y c(a_t, y_t)\, \sigma(y_t)\, dB_t.$$

Away from the borrowing boundary, the shock MPC coincides with the usual interior policy derivative. Equivalently, $Z_t = u''(c_t)\, \partial_y c(a_t, y_t)\sigma(y_t)$.

The advantage of the path-space formulation developed here is that it remains well defined even when the borrowing constraint becomes binding. Near the boundary the policy function may fail to be differentiable with respect to assets or income but the shadow continuation value W_t and its martingale component Z_t remain well-defined semimartingale objects. The shock MPC is thus characterized globally, including in the boundary region.

In the HJB formulation, the borrowing constraint enters as a variational inequality, implying that discounted marginal utility is a supermartingale. The decomposition derived here gives the structure of that supermartingale by separating the innovation, which determines the shock MPC, from the compensator, which captures the effect of the borrowing constraint boundary. In the recursive formulation, these effects are summarized by the policy function through its dependence on the current state. The decomposition developed here opens that compression and isolates the constraint enforcement channel. This does not replace the HJB as a way of solving the model, but makes explicit the mechanism through which future constraint pressure enters the response to current shocks.

[2] Alternatively, because income innovations are given by $\sigma(y_t)dB_t$, one may normalize the response by the size of the income innovation and define $\widetilde{\mathrm{MPC}}_t^{\mathrm{shock}} = \frac{1}{u''(c_t)} \frac{Z_t}{\sigma(y_t)}$,which measures the consumption response to a unit income innovation rather than to the Brownian noise itself.

5 Shock MPCs in a Penalized Problem and their Clark–Ocone Representation

The hard borrowing constraint in Section 3 generates a multiplier measure $d\eta_t$ that is singular at the borrowing boundary. In order to obtain a differentiable representation of the innovation coefficient Z_t, we introduce a smooth penalized approximation that converges to the hard-constrained problem. The key payoff of the penalized formulation is that we can expresses the consumption response to shocks as the sensitivity of expected future constraint pressure to the current income innovation.

5.1 Penalized household problem

The household chooses an adapted consumption process $c_t \geq 0$ to maximize

$$J^{\kappa}(c) = \mathbb{E}\left[\int_0^{\infty} e^{-\rho t} u(c_t)\, dt - \kappa \int_0^{\infty} e^{-\rho t} \Phi(a_t)\, dt\right]$$

subject to

$$da_t = (r a_t + y_t - c_t)\, dt, \qquad a_0 \geq 0.$$

The parameter $\kappa > 0$ controls the intensity of the penalty, and the function $\Phi : \mathbb{R} \to \mathbb{R}_+$ is a smooth convex penalty that is close to zero when $a > 0$ and increases rapidly as a becomes negative, so the penalty activates only near the borrowing boundary.

5.2 Penalized Euler equation

Let $c^{\varepsilon} = c + \varepsilon \delta c$. The variation of the asset path is

$$\delta a_t = -\int_0^t e^{r(t-s)} \delta c_s ds.$$

Differentiating the objective and interchanging integrals yields

$$u'(c_t) = -\kappa \mathbb{E}_t \left[\int_t^{\infty} e^{-(\rho - r)(s-t)} \Phi'(a_s)\, ds\right].$$

Define the penalized shadow continuation value

$$W_t^{\kappa} := -\kappa \mathbb{E}_t \left[\int_t^{\infty} e^{-(\rho - r)(s-t)} \Phi'(a_s)\, ds\right].$$

Then the Euler equation has the same structure as before:

$$u'(c_t) = W_t^\kappa.$$

The same argument used in Section 3.2 shows that there exists an adapted process Z_t^κ such that

$$\boxed{dW_t^\kappa = (\rho - r)W_t^\kappa\, dt + \kappa\Phi'(a_t)dt + Z_t^\kappa dB_t.}$$

This is the smooth penalized counterpart of the hard-constraint dynamics

$$dW_t = (\rho - r)W_t\, dt - d\eta_t + Z_t\, dB_t,$$

in which the singular boundary term $-d\eta_t$ is replaced by the absolutely continuous term $\kappa\Phi'(a_t)\, dt$, which penalizes low asset positions continuously rather than enforcing the constraint as a hard barrier. This term can also be interpreted as a shadow cost of borrowing that increases as assets approach the boundary, making borrowing effectively more expensive in a smooth way.

5.3 Sensitivity representation of shock MPCs

Define the tail functional

$$\Xi_t^\kappa = -\kappa \int_t^\infty e^{-(\rho-r)(s-t)}\Phi'(a_s)\, ds,$$

so that the penalized shadow continuation value is

$$W_t^\kappa = \mathbb{E}_t[\Xi_t^\kappa].$$

The tail functional measures the total discounted penalty pressure that the household faces from date t onward. We can now characterize Z_t^κ – and therefore the shock MPC – as a Malliavin derivative of this object.

Proposition 2 (Sensitivity representation of shock MPCs)**.** *The innovation coefficient in the penalized shadow-state dynamics determines the response of consumption to shocks:*

$$\boxed{D_t c_t = \frac{1}{u''(c_t)}\, Z_t^\kappa,}$$

and

$$Z_t^\kappa = \mathbb{E}_t[D_t \Xi_t^\kappa].$$

Consequently the shock marginal propensity to consume satisfies

$$\boxed{MPC_t^{shock} = \frac{1}{u''(c_t)}\, Z_t^{\kappa}.}$$

Proof. From the Euler equation,

$$u'(c_t) = W_t^{\kappa}.$$

Taking the Malliavin derivative with respect to the innovation at time t,

$$u''(c_t)\, D_t c_t = D_t W_t^{\kappa}.$$

Hence,

$$D_t c_t = \frac{1}{u''(c_t)}\, D_t W_t^{\kappa}.$$

From the semimartingale decomposition

$$\boxed{dW_t^{\kappa} = (\rho - r) W_t^{\kappa}\, dt + \kappa \Phi'(a_t) dt + Z_t^{\kappa} dB_t},$$

a standard property of Malliavin calculus for Brownian semimartingales implies that the diagonal Malliavin derivative coincides with the martingale integrand, so that

$$D_t W_t^{\kappa} = Z_t^{\kappa},$$

which is the Clark–Ocone characterization of the martingale integrand. Moreover, since $W_t^{\kappa} = \mathbb{E}_t[\Xi_t^{\kappa}]$ and conditional expectation commutes with the Malliavin derivative at the diagonal, we also have

$$Z_t^{\kappa} = D_t W_t^{\kappa} = \mathbb{E}_t[D_t \Xi_t^{\kappa}].$$

□

Proposition 2 gives a characterization of the innovation coefficient Z_t^{κ} as a Malliavin derivative of the tail functional Ξ_t^{κ}. That is, Z_t^{κ} measures how a current income shock changes the expected future intensity of borrowing constraints. Because the Euler equation links marginal utility to the shadow continuation value W_t^{κ}, this sensitivity determines the consumption response to shocks. The representation says that the response to a shock is governed by how that shock changes the expected future path of constraint feasibility. A positive income innovation matters not only because it raises current resources, but because it changes the expected intensity of future constraint enforcement. This is why the innovation coefficient has the interpretation of an impulse response: it records how current shocks revise

the shadow value of future borrowing pressure.

The next proposition provides a more detailed representation of the shock MPC.

Proposition 3 (Constraint curvature and asset sensitivity representation of shock MPCs.)**.** *The innovation coefficient in the penalized shadow-state dynamics admits representation*

$$\boxed{Z_t^\kappa = -\kappa \mathbb{E}_t \left[\int_t^\infty e^{-(\rho - r)(s-t)} \Phi''(a_s)\, D_t a_s \, ds \right].}$$

Consequently the shock marginal propensity to consume satisfies

$$\boxed{MPC_t^{shock} = -\frac{\kappa}{u''(c_t)}\, \mathbb{E}_t \left[\int_t^\infty e^{-(\rho - r)(s-t)} \Phi''(a_s)\, D_t a_s \, ds \right].}$$

Proof. From the definition of the tail functional,

$$\Xi_t^\kappa = -\kappa \int_t^\infty e^{-(\rho - r)(s-t)} \Phi'(a_s)\, ds.$$

Applying the Malliavin chain rule yields

$$D_t \Xi_t^\kappa = -\kappa \int_t^\infty e^{-(\rho - r)(s-t)} \Phi''(a_s)\, D_t a_s \, ds.$$

Since $W_t^\kappa = \mathbb{E}_t[\Xi_t^\kappa]$, we have

$$Z_t^\kappa = D_t W_t^\kappa = \mathbb{E}_t[D_t \Xi_t^\kappa].$$

□

The expression in this Proposition shows that the shock MPC depends on the interaction of two terms: $\Phi''(a_s)$ and $D_t a_s$. The first term, $D_t a_s$, is the Malliavin derivative of future assets with respect to the current innovation. It measures how a shock at time t affects assets at time s. The second term, $\Phi''(a_s)$, is the curvature of the penalty function evaluated at the future asset position. This curvature determines how sensitive the shadow value of relaxing the constraint is to small changes in assets near the borrowing boundary. When assets are far from the borrowing boundary, the penalty curvature $\Phi''(a_s)$ is close to zero and income shocks have little effect on expected future constraint pressure. As assets approach the borrowing limit, the curvature increases and income shocks generate larger revisions in expected constraint binding, amplifying the consumption response.

5.4 Convergence to the hard-constraint model

Under appropriate regularity conditions, as $\kappa \to \infty$, the penalized model converges to the hard-constraint model of Section 3. In this limit, the enforcement term $\kappa\Phi'(a_t)dt$ converges to the multiplier measure $-d\eta_t$, and the penalized innovation coefficient Z_t^κ converges to the innovation coefficient Z_t in the hard-constraint model.

6 Covariance Decomposition of Shock MPCs

Section 5 showed that the shock MPC is governed by the interaction between the curvature of the penalized borrowing constraint and the asset sensitivity sensitivity. In this section, we decompose that interaction into an average component (mean-curvature×mean exposure) and a covariance (beta) component. That is, we determine how much of the shock MPC comes from the average levels of curvature and exposure, and how much comes from the way they move together across future states

6.1 Covariance decomposition of shock MPCs

For each $s \geq t$, apply the conditional covariance identity

$$\mathbb{E}_t[XY] = \mathbb{E}_t[X]\mathbb{E}_t[Y] + \mathrm{Cov}_t(X, Y)$$

to

$$X = \Phi''(a_s), \qquad Y = D_t a_s.$$

Then

$$\mathbb{E}_t[\Phi''(a_s)D_t a_s] = \mathbb{E}_t[\Phi''(a_s)]\ \mathbb{E}_t[D_t a_s] + \mathrm{Cov}_t(\Phi''(a_s), D_t a_s)\,.$$

Substituting into the representation of the shock MPC yields the following Proposition.

Proposition 4 (Covariance decomposition of shock MPCs)**.** *The penalized shock MPC admits the decomposition*

$$\boxed{MPC_t^{shock} = \mathcal{A}_t^\kappa + \mathcal{C}_t^\kappa,}$$

where

$$\boxed{\mathcal{A}_t^\kappa = -\frac{\kappa}{u''(c_t)} \int_t^\infty e^{-(\rho-r)(s-t)} \mathbb{E}_t[\Phi''(a_s)]\, \mathbb{E}_t[D_t a_s]\, ds.}$$

is the average component, and

$$\boxed{\mathcal{C}_t^{\kappa} = -\frac{\kappa}{u''(c_t)} \int_t^{\infty} e^{-(\rho-r)(s-t)} \operatorname{Cov}_t(\Phi''(a_s), D_t a_s)\, ds.}$$

is the covariance component.

The average component $\mathcal{A}_t^{\kappa}$ is a mean-curvature times mean-exposure approximation that combines the expected curvature of the penalized borrowing constraint, $\mathbb{E}_t[\Phi''(a_s)]$, and the expected sensitivity of future assets to the current innovation, $\mathbb{E}_t[D_t a_s]$, at each horizon. This is the component that preserves average curvature and average exposure and shuts down their state-dependent comovement.

The covariance component $\mathcal{C}_t^{\kappa}$ captures the additional effect that arises because curvature and exposure are not independent. It is nonzero whenever shocks to assets are larger in states where the future constraint's curvature is higher. It measures whether shock propagation through assets is concentrated in future states where the borrowing constraint is tighter. In particular, the covariance term isolates a mechanism that is not visible in average responses: shocks have large effects when they move the household into states where constraint curvature is concentrated.

A household may have a large average asset response to shocks without a large shock MPC if those responses occur primarily in states where the borrowing constraint is slack and $\Phi''(a_s) \approx 0$. Conversely, a household with moderate average asset sensitivity can have a large shock MPC if that sensitivity is concentrated in states where the constraint curvature is high. That is, the covariance term gives a source of state-dependent amplification that is not present with only average effects.

This decomposition shows that the response to a shock is not governed only by average exposure. It also depends on where that exposure is realized in the future state space. A shock can have a modest average effect on future assets and still generate a large consumption response if it pushes the household toward states where constraint curvature is concentrated. Conversely, even large average exposure need not imply a large response if that exposure is realized primarily in states far from the boundary. The covariance term isolates exactly this state-dependent amplification mechanism.

Consider two special cases.

First, if $\Phi''(a_s)$ and $D_t a_s$ are conditionally uncorrelated for every $s \geq t$, then

$$\mathcal{C}_t^{\kappa} = 0,$$

and the shock MPC reduces to the average component. In that case all state dependence

works through conditional means.

Second, when the future asset response is strongest in states close to the borrowing boundary, the covariance term becomes important. In such states $\Phi''(a_s)$ is large, so the covariance term captures the amplification created by the interaction between asset sensitivity and local constraint curvature.

Proposition 4 shows that the shock MPC is not determined only by average asset sensitivity or average curvature of the borrowing constraint. What matters is also how these two objects move together across future states. The covariance component can be written in a form that separates the variance of constraint curvature from the loading of exposure on that risk.

Corollary 5. *For each horizon $s \geq t$,*

$$\mathrm{Cov}_t(\Phi''(a_s), D_t a_s) = \beta^{\kappa}_{t,s}\, \mathrm{Var}_t(\Phi''(a_s))\,,$$

where

$$\boxed{\beta^{\kappa}_{t,s} := \frac{\mathrm{Cov}_t(\Phi''(a_s), D_t a_s)}{\mathrm{Var}_t(\Phi''(a_s))}}$$

is the conditional beta of asset exposure with respect to curvature, and

$$\boxed{\mathcal{C}^{\kappa}_t = -\frac{\kappa}{u''(c_t)} \int_t^{\infty} e^{-(\rho-r)(s-t)} \beta^{\kappa}_{t,s}\, \mathrm{Var}_t(\Phi''(a_s))\, ds.}$$

This representation separates two forces: risk arising from dispersion in future constraint curvature, that is, $\mathrm{Var}_t(\Phi''(a_s))$, which measures how dispersed future constraint curvature is across states, and $\beta^{\kappa}_{t,s}$, which measures whether asset exposure is systematically high in states where curvature is high. The covariance term is large when both are present — when constraint tightness risk is substantial and when asset exposure loads positively on it.

6.2 Cauchy–Schwarz bound for Covariance of Shock MPCs

The covariance component of the shock MPC can be bounded using the conditional Cauchy–Schwarz inequality to show when the covariance term may be important.

Corollary 6 (Cauchy–Schwarz bound for Covariance)**.** *For each horizon $s \geq t$, the covariance term satisfies the conditional Cauchy–Schwarz inequality*

$$|\mathrm{Cov}_t(\Phi''(a_s), D_t a_s)| \leq \sqrt{\mathrm{Var}_t(\Phi''(a_s))\, \mathrm{Var}_t(D_t a_s)}.$$

The covariance component of the shock MPC satisfies

$$\boxed{|\mathcal{C}_t^\kappa| \leq \frac{\kappa}{|u''(c_t)|} \int_t^\infty e^{-(\rho-r)(s-t)} \sqrt{\mathrm{Var}_t(\Phi''(a_s))\,\mathrm{Var}_t(D_t a_s)}\, ds.}$$

The bound identifies two conditions under which the covariance component can become large. First, there must be substantial risk or dispersion in the curvature of the penalized borrowing constraint across future states, measured by $\mathrm{Var}_t(\Phi''(a_s))$. Second, the sensitivity of future assets to the current innovation must itself vary across states, captured by $\mathrm{Var}_t(D_t a_s)$. If either of these terms is small, the covariance component is small, and the shock MPC is approximated by the average component $\mathcal{A}_t^\kappa$, that is, by the mean curvature $\times$ mean exposure approximation.

6.3 Shock MPCs as an Option Greek for Digital and Barrier Claims

The covariance decomposition has a useful interpretation through the lens of barrier and digital options. The shadow value W_t represents the expected future value of borrowing-constraint enforcement. In the hard-constraint formulation, this enforcement is the boundary-supported measure $d\eta_t$; in the penalized formulation, it appears through the smooth penalty $-\kappa\Phi'(a_t)$. In both cases, the future enforcement value is concentrated near the borrowing boundary.

This is analogous to barrier and digital claims in option pricing. For a smooth claim – say a standard option – local sensitivities are diffuse across the state space, and average exposure provides a reasonable approximation to the response. For a barrier or digital claim, sensitivities are concentrated near the strike or barrier: $\Phi''(a_s)$ is small when assets are far from the borrowing limit and becomes large only in a thin neighborhood of the boundary.

This concentration of curvature shows that the response of the shadow value to shocks depends not only on average exposure but also on how exposure covaries with the boundary layer. The shock MPC plays the role of the "delta" of the constraint claim with respect to the income innovation process. The covariance term captures that the state-dependent amplification arises because the constraint's sensitivity is concentrated in a thin boundary layer rather than being diffuse across the state space, generating an analogue of the delta spike familiar in digital and barrier options.

6.4 Example: Exponential penalty and Laplace-weighted exposure

A useful example of the penalized problem is that of an exponential penalty

$$\Phi(a) = e^{-\gamma a}, \qquad \gamma > 0.$$

In this case,

$$\Phi'(a) = -\gamma e^{-\gamma a}, \qquad \Phi''(a) = \gamma^2 e^{-\gamma a}.$$

Substituting into the penalized shadow continuation value gives

$$W_t^\kappa = -\kappa \mathbb{E}_t \left[\int_t^\infty e^{-(\rho - r)(s-t)} \Phi'(a_s)\, ds \right] = \kappa\gamma \mathbb{E}_t \left[\int_t^\infty e^{-(\rho - r)(s-t)} e^{-\gamma a_s}\, ds \right].$$

Thus W_t^κ is proportional to the time-integrated conditional Laplace transform of future assets. Intuitively, the exponential weight $e^{-\gamma a_s}$ places greater emphasis on states in which assets are low and the borrowing constraint is close to binding.

The Clark–Ocone representation of the innovation coefficient then becomes

$$Z_t^\kappa = -\kappa\gamma^2 \mathbb{E}_t \left[\int_t^\infty e^{-(\rho - r)(s-t)} e^{-\gamma a_s} D_t a_s\, ds \right].$$

The shock marginal propensity to consume is then

$$\boxed{\mathrm{MPC}_t^{\mathrm{shock}} = -\frac{\kappa\gamma^2}{u''(c_t)} \mathbb{E}_t \left[\int_t^\infty e^{-(\rho - r)(s-t)} e^{-\gamma a_s} D_t a_s\, ds \right].}$$

The exponential kernel $e^{-\gamma a_s}$ tilts expectations toward states with low assets, so the shock MPC depends on how current income innovations shift the future asset distribution under an exponential tilt toward near-boundary states.

Applying the conditional covariance identity yields

$$\mathbb{E}_t[e^{-\gamma a_s} D_t a_s] = \mathbb{E}_t[e^{-\gamma a_s}]\, \mathbb{E}_t[D_t a_s] + \mathrm{Cov}_t(e^{-\gamma a_s}, D_t a_s).$$

Integrating this identity over s with the same discount kernel $e^{-(\rho - r)(s-t)}$ yields the decomposition of the shock MPC into mean-exposure and Laplace-weighted covariance components.

The parameter γ controls the strength of the exponential tilt. When γ is small, the weight $e^{-\gamma a_s}$ varies slowly and the shadow value averages over a broad range of future asset states. When γ is large, the Laplace weight concentrates near the borrowing boundary. Standard Laplace asymptotics (Watson's lemma) imply that expectations of the form $\mathbb{E}[e^{-\gamma a} f(a)]$

become increasingly dominated by the behavior of $f(a)$ near the lower boundary as $\gamma \to \infty$. As γ increases, the Laplace transform concentrates on the boundary layer of the borrowing constraint.

This is the same boundary-concentration mechanism that drives the covariance term in Section 6: shocks matter most when they shift future asset paths toward regions where the borrowing constraint become more binding.

7 Impulse-Response MPCs

The shock MPC captures the response of consumption at time t to an innovation at time t. But income shocks affect not only current consumption but also the future path of consumption. This section discusses the impulse-response marginal propensity to consume (IRF MPC) which is the response of the future consumption to past income innovations.

For $s \geq t$, define the impulse-response MPC

$$\boxed{\mathrm{MPC}^{\mathrm{IRF}}_{t,s} := D_t c_s,}$$

the Malliavin derivative of consumption at horizon s with respect to the income innovation at time t, which gives how a marginal income shock at time t affects consumption at later horizons. Setting $s = t$ in the previous expression gives the instantaneous shock MPC:

$$\mathrm{MPC}^{\mathrm{shock}}_t = \mathrm{MPC}^{\mathrm{IRF}}_{t,t}.$$

From the Euler equation

$$u'(c_s) = W^\kappa_s,$$

we obtain

$$D_t c_s = \frac{1}{u''(c_s)} D_t W^\kappa_s.$$

Thus the IRF MPC is determined by the sensitivity of the shadow continuation value to the innovation at time t.

Recall from Section 5 that

$$W^\kappa_s = \mathbb{E}_s[\Xi^\kappa_s], \qquad \Xi^\kappa_s = -\kappa \int_s^\infty e^{-(\rho-r)(u-s)} \Phi'(a_u)\, du.$$

Taking the Malliavin derivative gives

$$D_t W^\kappa_s = D_t \mathbb{E}_s[\Xi^\kappa_s].$$

Under the Malliavin regularity conditions, conditional expectation commutes with the Malliavin derivative for $t \leq s$:

$$D_t W_s^\kappa = D_t \mathbb{E}_s[\Xi_s^\kappa] = \mathbb{E}_s[D_t \Xi_s^\kappa].$$

Substituting the definition of the tail functional,

$$D_t W_s^\kappa = -\kappa \mathbb{E}_s \left[\int_s^\infty e^{-(\rho - r)(u-s)} \Phi''(a_u) D_t a_u \, du \right].$$

Substituting back

$$\boxed{\mathrm{MPC}_{t,s}^{\mathrm{IRF}} = -\frac{\kappa}{u''(c_s)} \mathbb{E}_s \left[\int_s^\infty e^{-(\rho - r)(u-s)} \Phi''(a_u) D_t a_u \, du \right].}$$

The IRF MPC has the same determinants as the shock MPC. It depends on the interaction between the curvature of the borrowing constraint, $\Phi''(a_u)$, and the propagation of shocks through the household budget constrain, captured by $D_t a_u$. The difference is that the IRF MPC measures this interaction at future horizons s rather than only at impact.

8 Recursive characterization of IRF MPCs and the Propagation-Feedback mechanism

The covariance decomposition in Proposition 4 characterizes shock MPCs through the interaction between constraint curvature $\Phi''(a_s)$ and the sensitivity of future assets to current shocks $D_t a_s$ at each horizon s.

The IRF representation of Section 7 shows that the impulse-response MPC satisfies

$$\mathrm{MPC}_{t,s}^{\mathrm{IRF}} = -\frac{\kappa}{u''(c_s)} \mathbb{E}_s \left[\int_s^\infty e^{-(\rho - r)(u-s)} \Phi''(a_u) D_t a_u \, du \right].$$

We now develop a more explicit characterization of the IRF MPCs.

From the asset accumulation equation,

$$a_s = e^{rs} a_0 + \int_0^s e^{r(s-k)} (y_k - c_k) \, dk,$$

so that the Malliavin derivative satisfies

$$D_t a_s = \int_t^s e^{r(s-k)} \big(D_t y_k - D_t c_k \big) \, dk.$$

Substituting this expression into the IRF MPC formula yields and integral equation representation of the IRF MPCs.

Proposition 7 (Integral equation for IRF MPCs)**.** *The IRF MPC satisfies the integral equation*

$$\mathrm{MPC}_{t,s}^{\mathrm{IRF}} = -\frac{\kappa}{u''(c_s)}\mathbb{E}_s\left[\int_s^\infty e^{-(\rho-r)(u-s)}\Phi''(a_u)\left(\int_t^u e^{r(u-k)}(D_t y_k - \mathrm{MPC}_{t,k}^{\mathrm{IRF}})\,dk\right)du\right].$$

Separating the exogenous income-propagation and the endogenous feedback components:

$$\begin{aligned}\mathrm{MPC}_{t,s}^{\mathrm{IRF}} = &-\kappa\frac{1}{u''(c_s)}\mathbb{E}_s\left[\int_s^\infty e^{-(\rho-r)(u-s)}\Phi''(a_u)\left(\int_t^u e^{r(u-k)}D_t y_k\,dk\right)du\right]\\ &+\kappa\frac{1}{u''(c_s)}\mathbb{E}_s\left[\int_s^\infty e^{-(\rho-r)(u-s)}\Phi''(a_u)\left(\int_t^u e^{r(u-k)}\mathrm{MPC}_{t,k}^{\mathrm{IRF}}\,dk\right)du\right].\end{aligned}$$

The first term is exogenous and captures the direct propagation of the income innovation through assets and constraint curvature. The second term is the endogenous feedback generated by future consumption responses. The IRF MPC appears on both sides of the equation. In other words, the IRF MPC is determined as the fixed point of a mapping that combines the propagation of income shocks through the asset dynamics with the feedback generated by future consumption responses.

In order to clarify the structure of this mapping, define

$$\Gamma_s(f) = \mathbb{E}_s\left[\int_s^\infty e^{-(\rho-r)(u-s)}\Phi''(a_u)\left(\int_t^u e^{r(u-k)}f_k\,dk\right)du\right],$$

which takes a process f and propagates it forward through the asset dynamics and constraint curvature. The integral equation becomes

$$\mathrm{MPC}_{t,s}^{\mathrm{IRF}} = -\frac{\kappa}{u''(c_s)}\Gamma_s(D_t y) + \frac{\kappa}{u''(c_s)}\Gamma_s(\mathrm{MPC}_{t,\cdot}^{\mathrm{IRF}}).$$

The first term feeds the income innovation through Γ_s. The second term feeds the endogenous consumption response itself through the same map. The iterated-kernel expansion below makes these successive rounds explicit.

Proposition 8 (Iterated kernel expansion)**.** *The integral equation admits the recursive expansion*

$$\mathrm{MPC}_{t,s}^{\mathrm{IRF}} = \mathrm{MPC}_{t,s}^{(1)} + \mathrm{MPC}_{t,s}^{(2)} + \mathrm{MPC}_{t,s}^{(3)} + \cdots.$$

The first term gives the direct propagation of the income innovation

$$\mathrm{MPC}^{(1)}_{t,s} = -\frac{\kappa}{u''(c_s)}\Gamma_s(D_t y),$$

The second term gives the first endogenous feedback generated by future consumption responses

$$\mathrm{MPC}^{(2)}_{t,s} = -\frac{\kappa^2}{u''(c_s)}\Gamma_s\left(\frac{1}{u''(c_\cdot)}\Gamma_\cdot(D_t y)\right).$$

Written explicitly:

$$\mathrm{MPC}^{(2)}_{t,s} = -\frac{\kappa^2}{u''(c_s)}\mathbb{E}_s\Big[\int_s^\infty e^{-(\rho-r)(u-s)}\Phi''(a_u) \times\Big(\int_t^u e^{r(u-k)}\frac{1}{u''(c_k)}\mathbb{E}_k\Big[\int_k^\infty e^{-(\rho-r)(v-k)}\Phi''(a_v)\left(\int_t^v e^{r(v-\ell)}D_t y_\ell\, d\ell\right)dv\Big]\,dk\Big)du\Big].$$

Higher-order terms are generated by repeated application of the same propagation map.

The first term $\mathrm{MPC}^{(1)}_{t,s}$ captures the direct propagation of the income innovation through the asset dynamics when consumption initially does not respond. The higher-order terms capture successive rounds of adjustment. The initial consumption response changes future asset positions, which changes the expected tightness of the borrowing constraint along the curvature kernel, which in turn feeds back into additional consumption responses. Each round of consumption adjustment feeds into future constraint pressure and generates further rounds of response.

9 Conclusion

We develop a path-space approach to the study of intermittently binding feasibility constraints and show that the shadow value of constraints process can be decomposed into regular and constraint-enforcement components. In a continuous-time consumption–savings model, this decomposition delivers a characterization of shock MPCs, a covariance decomposition of state-dependent amplification, and an integral-equation representation of impulse responses, with an iterated-kernel expansion of the propagation–feedback mechanism of shock transmission.

References

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," Review of Economic Studies, 89, 45–86.

Auclert, A. (2019): "Monetary Policy and the Redistribution Channel," American Economic Review, 109, 2333–2367.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," Econometrica, 89, 2375–2408.

Auclert, A., Rognlie, M. and Straub, L., 2024. The intertemporal keynesian cross. Journal of Political Economy, 132(12), pp.4068-4121.

Bilal, A. and Goyal, S., 2025. "Some Pleasant Sequence-Space Arithmetic in Continuous Time" (No. w33525). National Bureau of Economic Research.

Boerma, Job, Georgii Riabov, and Aleh Tsyvinski. "Policy with Stochastic Hysteresis." (2023).

Borovička, J., and L. P. Hansen (2016): "Examining Macroeconomic Models through the Lens of Asset Pricing," Journal of Econometrics, 183, 67–90.

Borovička, J., L. P. Hansen, and J. A. Scheinkman (2014): "Shock Elasticities and Impulse Responses," Mathematics and Financial Economics, 8, 333–354.

Farhi, E., Olivi, A. and Werning, I., (2022) "Price theory for incomplete markets,(w30037)." National Bureau of Economic Research.

Kaplan, G., and G. L. Violante (2022): "The Marginal Propensity to Consume in Heterogeneous Agent Models," Annual Review of Economics, 14, 747–775.

Patterson, C. (2023): "The Matching Multiplier and the Amplification of Recessions," American Economic Review, 113, 982–1012.