The Impact of Tariff and Trade Policy Uncertainty Shocks on Household Asset Portfolios

Siwon Lee*

SEP 2026

Abstract

This paper investigates household portfolio allocation under trade policy uncertainty (TPU) and imperfect information regarding firms' inventories. In a calendar-preserving monthly local projection framework for the post-2015 sample, a one standard deviation shock to log-TPU reduces equity portfolio shares by 0.508 percentage points while increasing cash and bond shares by 0.330 and 0.183 percentage points, respectively. The reallocation out of equity and the contemporaneous decline in the NASDAQ survive within-path simultaneous inference, although stock-share lead predictability limits a strict causal interpretation. To explain these dynamic responses, I develop a continuous-time portfolio model linking undisclosed inventories to ambiguity and regime-transition risk. A conditional calibration matches baseline allocation levels and jointly reproduces the initial drop in equity alongside the increases in cash and bonds, where the cash response operates as a portfolio residual and exceeds the bond response. Overall, the quantitative framework highlights how imperfect inventory information amplifies household portfolio adjustments under policy shocks.

Keywords: trade policy uncertainty; portfolio choice; Knightian uncertainty; robust control; inventories; regime switching

JEL Classification: D81, E21, F13, F41, G11

*Ph.D. Candidate in Economics, Washington University in St.Louis. Email: l.siwon@wustl.edu

1 Introduction

Recent escalations in trade policy tensions and tariff disputes have fundamentally reshaped international trade relations. While an extensive literature examines the impact of tariffs on firm behavior and industry dynamics, how households adjust their financial decisions under trade policy uncertainty (TPU) remains relatively underexplored. Surges in TPU drive macroeconomic adjustments in asset prices, interest rates, and market valuations, which in turn shape household portfolio choices. This paper proposes a structure framework linking tariff uncertainty to household asset allocation, emphasizing hedging motives and firms' inventory decisions. Specifically, I examine how TPU shocks reallocate household wealth across bonds, stocks, and cash, highlighting the channels through which trade policy transmits to household finance.

This research is inspired by the observation that tariff announcements and trade disputes often coincide with asset market volatility. In February 2025, rising trade tensions and tariff announcements were associated with large swings in asset valuations. These sharp market reactions occurred even before governments implemented the proposed policies. Concurrent fluctuations in bond prices and consumption demand indicate that tariffs affect household finance through multiple channels. Data from the American Association of Individual Investors(AAII) Asset Allocation Survey support this view. The asset share allocated to stocks fell from 67.9 percent in February 2025 to 64.3 percent in May 2025. This drop represents a significant relative decline of approximately 5 percent in household equity exposure. The stock share then recovered to 67.5 percent in June as major tariff agreements were finalized. These dynamics suggest that tariff shocks modify household portfolio choices. The policy shocks shift allocations by changing expected returns or reshaping risk perceptions. Understanding these links extends both the trade and household finance literatures. This study provides a more complete view of how trade policy affects household decision.

Although trade policy shocks influence firm profitability and activate household hedging motives, existing literature primarily focuses on aggregate macroeconomic adjustments rather than portfolio transmission mechanisms. For instance, Schmitt-Grohé and Uribe (2025) analyze the macroeconomic impacts of transitory and permanent U.S. tariff

shocks through monetary policy and inflation dynamics, while Boer and Rieth (2024) evaluate trade policy uncertainty (TPU) within a two-country DSGE framework backed by structural VAR estimation. However, these macroeconomic frameworks abstract from the micro-founded mechanism through which TPU shapes household asset allocation via firm profitability and hedging demand.

An empirical literature documents how trade policy uncertainty (TPU) affects financial markets across various asset classes, suggesting corresponding shifts in household financial behavior. At the macroeconomic and aggregate market level, Hoque et al. (2023) find heterogeneous stock market spillovers from trade uncertainty across major and emerging economies, while Rao et al. (2025) investigate market reactions to tariff announcements across nations with different net exporting or net importing positions. Furthermore, Benguria and Saffie (2025) show that higher tariffs depress stock prices without appreciating the local currency, and Yilmazkuday (2025) confirms these negative impacts on major stock indices while controlling for macroeconomic factors. Mechanistically, Carlomagno and Albagli (2022) attribute multi-asset price changes to risk aversion channels rather than growth channels. At the microeconomic and firm level, Wengerek et al. (2025) document negative abnormal stock returns following historical tariff imposition announcements, though Huang et al. (2023) establish that firms with high research intensity experience smaller equity declines. Together, these studies demonstrate how tariffs and TPU affect market-wide prices and aggregate investor expectations. However, a distinct gap remains between these aggregate asset price reactions and actual medium-run household portfolio reallocation.

When TPU increases, firms adjust inventory levels to maximize profits. However, firms do not disclose inventory data continuously. Therefore, households cannot perfectly observe these corporate operational changes. This information asymmetry limits household knowledge regarding the exact magnitude and structure of corporate earnings. Consequently, this imperfect information forces households to initiate hedging strategies. These complex dynamics raise critical questions about economic transmission mechanisms.

This paper structures its inquiry around four central questions. First, how do corporate inventory management and disclosure delays interact during trade uncertainty surges? Second, how does investor hedging demand shape actual household portfolio allocation?

Third, how do these informational frictions affect aggregate household consumption? Finally, do tariff shocks affect market confidence in sovereign bond safety and government solvency? Furthermore, do these shocks increase the volatility of expected sovereign bond returns?

The empirical analysis utilizes the news based US trade policy uncertainty index of Baker, Bloom, and Davis (2016). This index represents the trade policy category of their broader economic policy uncertainty measure. To construct the series, the authors count newspaper articles containing terms about the economy, policy, and uncertainty alongside trade policy keywords. The observed index is log-transformed once for empirical estimation; TPU denotes this log-transformed variable in the estimation section. Trade policy news can shift the likelihood of a regime change and expand household ambiguity, but the statistical TPU shock is not equated numerically with the structural transition intensity. The index measures trade policy uncertainty broadly rather than tariff uncertainty specifically. However, Section 2 documents that its major movements coincide with key tariff episodes. My regime classification effectively isolates these major episodes.

The empirical TPU index also records non-tariff policy news. Upstream restrictions may operate through imported-input costs, but other components need not share the model's transmission mechanism. Their presence is a measurement and interpretation limitation; without further restrictions, it does not imply classical measurement error or attenuation toward zero.

To address these questions, the paper connects the trade, macroeconomic and portfolio literatures. On the production side, Alessandria et al. (2011) show how inventory adjustment magnifies movements in imports relative to sales and domestic production. Khan and Khederlarian (2021, 2025) study how anticipated tariff changes and imported-input costs affect inventory purchases and productivity. Their work motivates an inventory buffer that changes the timing of imports and the firm's exposure to policy shocks. Boehm et al. (2023) distinguish short-run from long-run trade elasticities and document gradual adjustment in the long term following tariff changes. This evidence helps separate the immediate inventory response from the longer-run substitution between domestic and foreign inputs.

Tariff pass-through also links these production choices to household purchasing power

and public revenue. Amiti et al. (2019) document how the 2018 tariffs raised import prices and imposed costs on domestic importers and consumers. Fajgelbaum et al. (2020) quantify the losses from import and retaliatory tariffs and evaluate them together with tariff revenue and gains to domestic producers. These findings motivate the model's connection between imported-input costs, operating profits, consumption expenditure and the fiscal resources backing government bonds.

On the macroeconomic side, Schmitt-Grohé and Uribe (2025) analyze transitory and permanent import tariff shocks in the United States, while Caldara et al. (2020) show how uncertainty about future trade costs reduces investment. Christiano et al. (2005) demonstrate how wage and price rigidities shape the persistence of aggregate responses and corporate profits. These approaches motivate the transmission from policy and input costs to the firm's wage bill and payouts. The present paper adds a household portfolio problem in which imperfect inventory information affects both equity exposure and the fiscal uncertainty carried by sovereign bonds.

The household-income block draws on evidence about factor shares and labor supply. Gollin (2002) emphasizes the treatment of underestimated self-employment and small-firm income when measuring labor's share of production. Chetty et al. (2011) distinguish intensive- and extensive-margin labor-supply responses and assess how microeconomic estimates relate to macroeconomic elasticities. Their contributions guide the interpretation of labor income and labor-disutility parameters when connecting firm production to household consumption and saving. These labor-market foundations complement the inventory and valuation mechanisms without identifying those mechanisms from labor-market evidence alone.

Finally, the foundation of dynamic portfolio optimization originates from continuous-time consumption and investment models. Merton (1969, 1971) characterizes optimal asset allocation in response to stochastic investment opportunities. In this framework, macroeconomic shocks affect investors primarily through modifications in asset volatility and risk-free rates. Later extensions examine how time-varying investment opportunities and shifts in government policy propagate into portfolio decisions. Brennan and Xia (2001) show how learning under unobservable expected dividend growth rates amplifies stock price volatility in general equilibrium. Campbell and Viceira (2002) summarize and

synthesize strategic asset allocation for long-horizon investors under time-varying investment opportunities. Pástor and Veronesi (2012) analyze how government policy changes affect stock prices and generate a positive jump risk premium. These extensions largely preserve rational expectations about the dynamics of returns, with investors responding to fully observable state variables. This paper extends this line of work in two directions. First, it introduces trade policy uncertainty as a regime-switching state variable that governs the tariff environment and, through it, the intertemporal investment opportunity set. Second, it departs from full knowledge of the return dynamics by assigning households Knightian uncertainty about firms' unobserved inventories, and it disciplines the resulting worst-case beliefs through a relative-entropy hedging mechanism in the spirit of Drechsler (2013). In this setting, tariffs act as a fundamental state variable that links trade policy to household wealth dynamics through both a regime-transition risk channel and an ambiguity channel, and the sovereign bond loads on the same inventory-driven fiscal uncertainty as equity.

The empirical analysis utilizes the AAII survey dataset, the news-based TPU index developed by Baker, Bloom, and Davis (2016), and macro-financial series from the FRED database. Together, these data sources provide comprehensive information on household asset allocation, measures of trade policy uncertainty, and macroeconomic indicators including Treasury yields, equity indices, and wage growth. The empirical strategy applies the local projection method developed by Jordà (2005) to estimate the dynamic impact of trade policy uncertainty on household allocations across stocks, bonds, and cash. Furthermore, the framework evaluates how these uncertainty shocks affect major equity market indices and overall wage levels. The primary objective of these empirical results is to capture specific macroeconomic variations and identify a direct empirical association between market dynamics and actual portfolio distribution, tightly linking broad trade uncertainty to observable changes in household financial behavior.

To complement these empirical findings, the study develops a continuous-time theoretical model. The framework features monopolistically competitive firms that earn positive operating profits. Firms utilize inventory management to buffer against incoming tariff shocks. Meanwhile, the government sector operates a simple revenue and expenditure system. This fiscal structure links tariff revenues directly to sovereign bond stability. The

household portfolio problem builds upon the classic Merton framework. Asset dynamics follow standard Brownian motion with a critical informational friction. Households cannot observe corporate inventory levels instantaneously. Investors only learn about these protective buffers through delayed corporate reports. This observational delay generates Knightian uncertainty regarding actual firm inventory. Following Drechsler (2013), the model incorporates an entropy-based hedging mechanism. This mathematical structure demonstrates how information frictions drive household asset allocation.

To operationalize the quantitative framework, the model distinguishes parameters motivated by established literature from those structurally identified by empirical moments. Preferences, production elasticities, factor shares, and fiscal exposure are disciplined by established macroeconomic and trade studies. Meanwhile, macro-financial series inform baseline return dynamics, interest rates, and regime transition intensities. Unobserved friction parameters, specifically hidden inventory risk, entropy bounds, and news dynamics, are structurally linked to total market risk and fitted to match observed household portfolio distributions, with complete details provided in Section 5.

I make three contributions. First, I link tariff risk timing to firm inventory sourcing and household asset exposure. This captures how trade policy uncertainty accumulates and resolves through inventory holdings and disclosures. This framework distinguishes physical inventory levels from net inventory deviations and separates monthly disclosure from physical stock adjustment. The analytical results describe conditional optimality and accounting restrictions.

Second, I translate the inventory-information mechanism into a constrained portfolio calculation that employs a unified state-contingent pricing framework for cash flows, return drift, risk exposure, and regime revaluation. I retain financial-wealth constraints and monthly posterior revelation. The calibration utilizes the complete baseline stock-bond covariance and permits correlated residual risk, making explicit the additional assumptions needed to connect observed total risk to unobserved inventory exposure.

Third, I show that the declared conditional specification fits equity sales together with increases in both bonds and cash, with cash receiving the larger impact reallocation. I obtain this allocation pattern by jointly disciplining stock and bond responses rather than fitting the stock response alone, while cash follows from portfolio adding-up. By analyzing

the joint effects operating through the inventory and hedging channels, this framework enables the identification and evaluation of their significance in driving household portfolio reallocations. The remaining equity-price discrepancy and failed fixed-parameter refinement checks delimit the provisional quantitative contribution.

The remainder of the paper proceeds as follows. Section 2 presents the data and empirical findings. Section 3 introduces the theoretical model. Section 4 provides the mathematical solutions for the economic framework. Section 5 discusses the calibration and quantitative results of simulation. Section 6 concludes the study. The appendix collects the data sources and the continuous-time derivations and their scope conditions, including static production, conditional inventory verification, household optimization, regime-switching valuation, and a sufficient-condition existence argument for the partial-equilibrium closure.

2 Data

2.1 TPU Regimes

The empirical analysis combines household portfolio survey data, macro-financial time series, and a news-based measure of TPU.[1] The empirical analysis documents key co-movements that motivate the theoretical framework. Additionally, it identifies empirical targets to calibrate the continuous-time parameters governing asset return dynamics and volatility.

1. The index is the news-based US TPU index of Baker, Bloom, and Davis (2016). See Table A.1 for details.

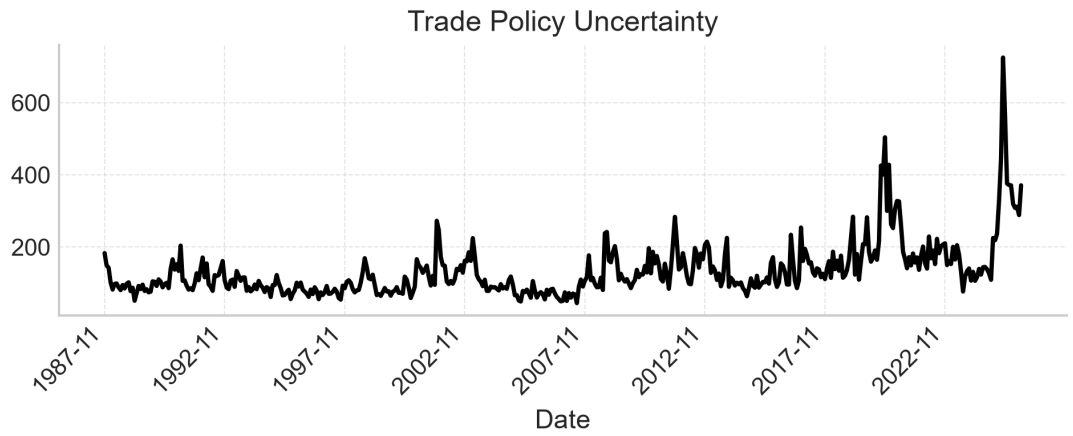


Figure 1: TPU Index

Figure 1 plots the TPU index based on news coverage. The index remains low for extended periods but exhibits distinct spikes around key political events. In the early 1990s, an initial cluster of increases is associated with the North American Free Trade Agreement. Debates over market opening and sectoral adjustment generated temporary uncertainty about the future trade regime.[2]

From the mid-1990s through the mid-2010s, the index remains comparatively stable and rarely departs from its historical mean. This stability reflects a period where multilateral liberalization and regional trade agreements underpinned a free-trade environment. During these years, tariff schedules change only gradually. Major economies largely converge toward lower average tariffs. Consequently, TPU is not a dominant driver of macro-financial uncertainty during this era.

The historical pattern changes markedly with the onset of the United States and China trade conflict in 2018 and 2019. The United States withdrew from the Trans-Pacific Partnership and imposed tariffs on solar panels and metals. Subsequently, the government announced a sequence of Section 301 tariffs on Chinese imports. These actions triggered immediate retaliatory measures from Chinese authorities. The policy shifts generated a sharp and persistent rise in the TPU index. Pronounced spikes occurred in March 2018 when broad steel and aluminum tariffs were announced. Uncertainty rose further in 2018 and 2019 as trade negotiations repeatedly broke down. This initial episode constitutes the first major trade conflict wave. Governments used tariffs on an unprecedented scale

2. Baker, Bloom, and Davis (2016) discuss this agreement-related TPU.

as a bargaining tool between the two largest economies.

A second wave of trade conflict between the United States and China emerges in the mid-2020s. In this phase, traditional tariff restrictions are complemented by export controls targeting advanced technologies. These policies target critical inputs like semiconductors and artificial intelligence hardware. New tariff measures and threats of across-the-board surcharges led to renewed spikes in the index during 2025. These developments reflect a structural shift toward targeted restrictions on technology trade. The escalation reinforces the interpretation of TPU as a key source of macro-financial risk. The structural differences between these historical periods are pronounced. Consequently, the framework characterizes the economic environment using distinct low and high TPU regimes.

A clean time series isolating tariff-specific uncertainty remains highly limited. However, within the sample, the TPU index effectively captures these dynamics and closely aligns with major tariff episodes. These episodes include the Section 232 and 301 actions in 2018 and 2019. They also include the tariff escalation in 2025. The regime classification isolates precisely these historical episodes. Consequently, the tariff regimes in the model map to the empirically dominant component of the index. Non-tariff content limits the mapping from broad TPU to tariff-specific uncertainty. Neither classical measurement error nor attenuation toward zero is assumed.

2.2 Estimation

The empirical strategy uses monthly local projections following Jordà (2005). The observed news-based TPU index is log-transformed once for estimation. In this section, $\text{TPU}_t = \log(\text{TPU}_t^{\text{raw}})$ denotes that analysis variable; the raw index is not assumed to have been logged by its provider. A standardized residual from an AR(6) in this variable defines u_t^{TPU}. The AR is estimated using observed, complete calendar-lag windows, and its residual standard deviation sets the shock unit. This shock is a statistical proxy for unexpected TPU movements, not an identified structural tariff shock or a direct measurement of the model's transition intensity.

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h u_t^{\text{TPU}} + \Gamma_h' W_{t-1} + \varepsilon_{t+h}, \qquad h = 0, \ldots, 10 \tag{1}$$

In this specification, y_{t+h} denotes the outcome variable at horizon $t+h$, with the dependent variable capturing the cumulative change from $t-1$ to $t+h$. The parameter α_h represents the horizon-specific intercept, while β_h measures the dynamic impulse response to the TPU shock. The vector W_{t-1} incorporates macro-financial controls observed at $t-1$ to condition on prevailing economic state variables, and ε_{t+h} represents the residual disturbance term.

To construct impulse response functions (IRFs), I estimate equation (1) separately for each horizon $h = 0, 1, \ldots, 10$. The sequence of estimated coefficients $\{\beta_h\}$ traces out the dynamic response of the outcome variable to a TPU shock. Confidence intervals are constructed using heteroskedasticity and autocorrelation consistent (HAC) standard errors following Newey and West (1987). I plot these coefficients along with their 95% confidence intervals to visualize the IRFs. Negative horizons $h = -3, -2, -1$ are examined separately through the lead regressions described below.

The estimation sample begins in January 2015, incorporating one lag of nominal PCE growth as the baseline control. Impulse responses are evaluated across specific scales: portfolio reallocations are reported in percentage points, whereas economic activity, consumption, income, and equity indices are expressed as log-percentage changes. Because the wage series reflects nominal growth, its impulse response captures shifts in the growth rate rather than fluctuations in real labor income or absolute wage levels. Finally, bond price dynamics are captured through an eight-year duration transformation of Treasury yields rather than observed total returns.

The baseline sample excludes observations from March to June 2020 while maintaining the underlying monthly calendar structure. No adjacent-period averaging or artificial imputation is applied. Growth rates and lags are constructed after masking, ensuring pre- and post-gap observations are never treated as consecutive. For each projection horizon, regressions require fully observed endpoints, complete controls, and an intact outcome path, making any window crossing the excluded block ineligible. Naturally missing values remain unadjusted. Consequently, the resulting estimand reflects eligible non-COVID periods rather than a counterfactual pandemic-free economy. This baseline exclusion is fixed structurally and is not selected based on coefficient significance.

2.3 Empirical Results

The baseline local projections use eligible shock months from January 2015 through January 2026, subject to horizon-specific data availability and the calendar-window exclusions described above. The horizontal axis in Figures 2–4 is the relative response horizon h, measured in calendar months. I define $h = 0$ as the shock month and $h = 1$ as the following month. The vertical axes report the response of each outcome in its stated units, not the TPU shock. These relative horizons are distinct from the calendar dates defining the estimation sample. Separately, the AR(6) used to construct the statistical shock is trained on the longer TPU history and comprises 443 complete windows from May 1988 to January 2026, excluding unavailable COVID observations and their incomplete lag windows. Its residual standard deviation is 0.236291 in log-index units. The portfolio regressions have 122 observations at impact and 102 at horizon ten; the sample changes with the horizon rather than compressing calendar time. All estimates below come from the same completed main run. Table 1 reports pointwise intervals from 1,999 dependent Gaussian multiplier draws.

Table 1: Responses to a one-standard-deviation log-TPU shock at impact

Outcome	Estimate	Pointwise 95%	N
Stock share (pp)	-0.508	[-0.843, -0.173]	122
Cash share (pp)	0.330	[0.045, 0.615]	122
Bond share (pp)	0.183	[0.039, 0.326]	122
NASDAQ	-0.994	[-1.557, -0.431]	122
Industrial production	0.047	[-0.046, 0.141]	122
Employment	-0.010	[-0.027, 0.006]	122
Unemployment (pp)	-0.009	[-0.037, 0.020]	120
Real income	0.006	[-0.029, 0.041]	122
Nominal consumption	-0.057	[-0.141, 0.027]	122
Real consumption	-0.043	[-0.122, 0.036]	122
Implicit deflator	-0.014	[-0.031, 0.003]	122
Wage growth (pp)	-0.021	[-0.106, 0.064]	122
Treasury yield (pp)	-0.025	[-0.054, 0.003]	122
Duration price proxy (%)	0.204	[-0.023, 0.431]	122

Activity, income, consumption and deflator responses are differences in 100 log levels. The wage outcome is nominal wage growth, not a real wage level. Simultaneous coverage is within each outcome's 0–10 month path, not across outcomes or specifications. $N = 122$ at impact except unemployment ($N = 120$). The Treasury price proxy uses duration eight and is not an observed total return.

The central pattern is a reduction in equity exposure accompanied by larger cash and bond shares. The equity share falls by 0.508 percentage points at impact; cash rises by

0.330 and bonds by 0.183 percentage points. Cash is therefore the larger positive allocation response in the point estimates, although a formal test of the cash–bond difference is not provided. Equity remains negative at horizons zero and one under simultaneous inference. Cash is positive at these horizons under pointwise inference only. The bond point estimates are positive at several horizons, but its simultaneous band excludes zero only at horizon ten. The evidence supports reallocation away from equities more strongly than a precisely identified allocation between cash and bonds.

The NASDAQ falls by 0.994 log-percentage points on impact, and that impact remains negative under the simultaneous band. It should not be described as unresponsive. This is an observed market-price response, not a direct observation of corporate operating profit, inventory value, or inventory quantity. Figure 2 displays the baseline paths.

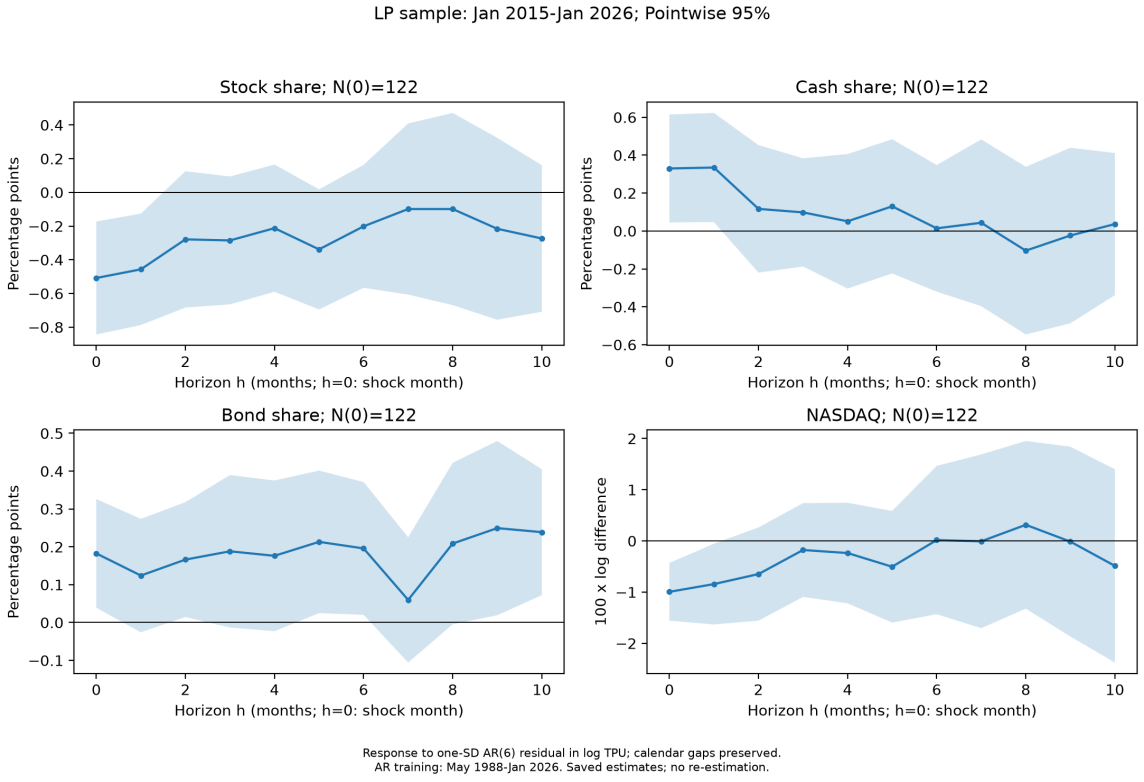


Figure 2: Portfolio and market-price responses. Shading denotes pointwise 95% intervals; horizons are calendar months.

Industrial production, employment, unemployment and real income have imprecise impact responses. No individual horizon of production, employment or unemployment excludes zero under pointwise inference. This does not establish that their entire paths are zero: conditional joint post-response tests yield $p = 0.005$, 0.002, and 0.0045, respectively.

Joint Wald tests and maximum-t bands test different features of correlated paths. These supplementary joint results are not adjusted across all outcomes and specifications, and should not be converted into claims of precise horizon-specific effects.

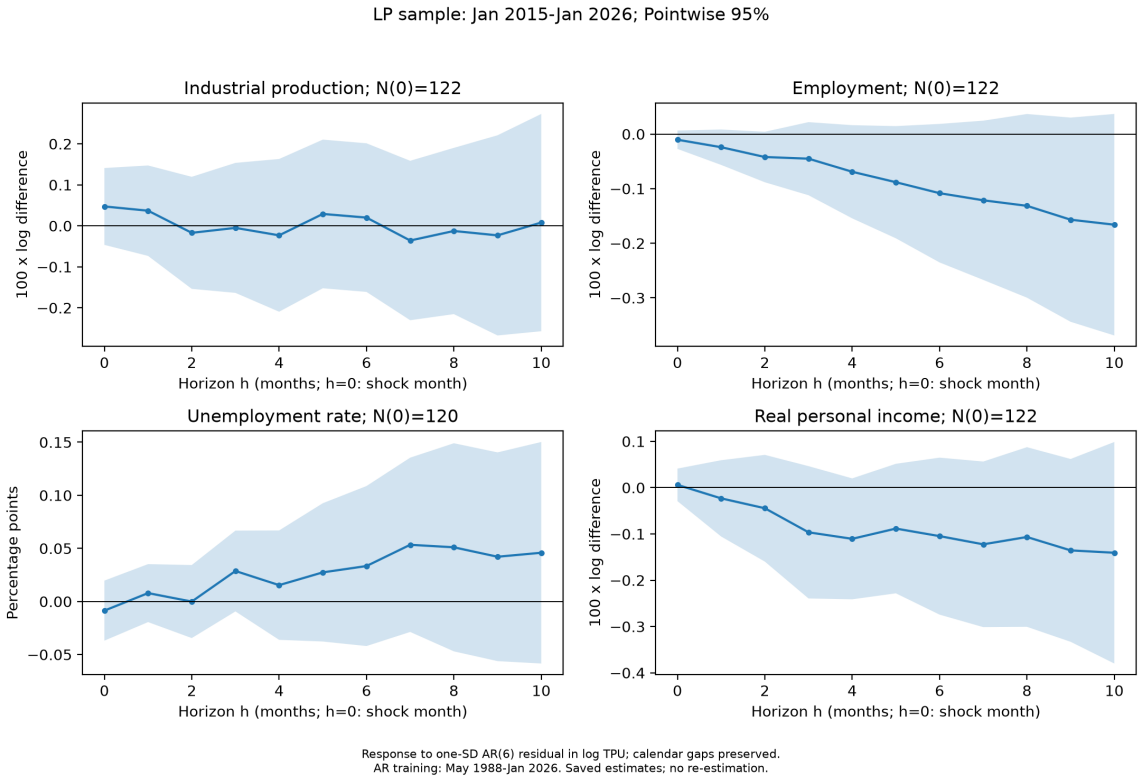


Figure 3: Real activity with pointwise 95% intervals; imprecision is not evidence of equivalence to zero.

Nominal and real consumption fall by 0.057 and 0.043 log-percentage points at impact, respectively, but neither impact nor any individual post horizon is significant at the pointwise level. The nominal-consumption joint post test gives $p = 0.3235$, whereas the real-consumption joint test gives $p = 0.0005$; the latter is qualified by the lead evidence below. The impact deflator response is -0.014 log-percentage points, and its negative response at horizon one survives simultaneous inference. The broad TPU shock therefore does not provide evidence of an immediate increase in the aggregate consumption price level. This finding is not a reason to change the model's maintained price assumption: intensity news and a realized tariff-regime switch are different experiments.

Nominal wage growth is statistically indistinguishable from zero on impact at -0.021 percentage points, with its pointwise confidence interval spanning zero. Although isolated pointwise responses appear at horizons four and nine, no individual horizon remains sig-

nificant under simultaneous inference bands. The conditional joint post-estimation test rejects a zero path at $p = 0.0005$, and the lead test similarly rejects following adjustment. Consequently, the analysis retains a flat short-run wage benchmark without claiming that labor income or wage dynamics are strictly invariant.

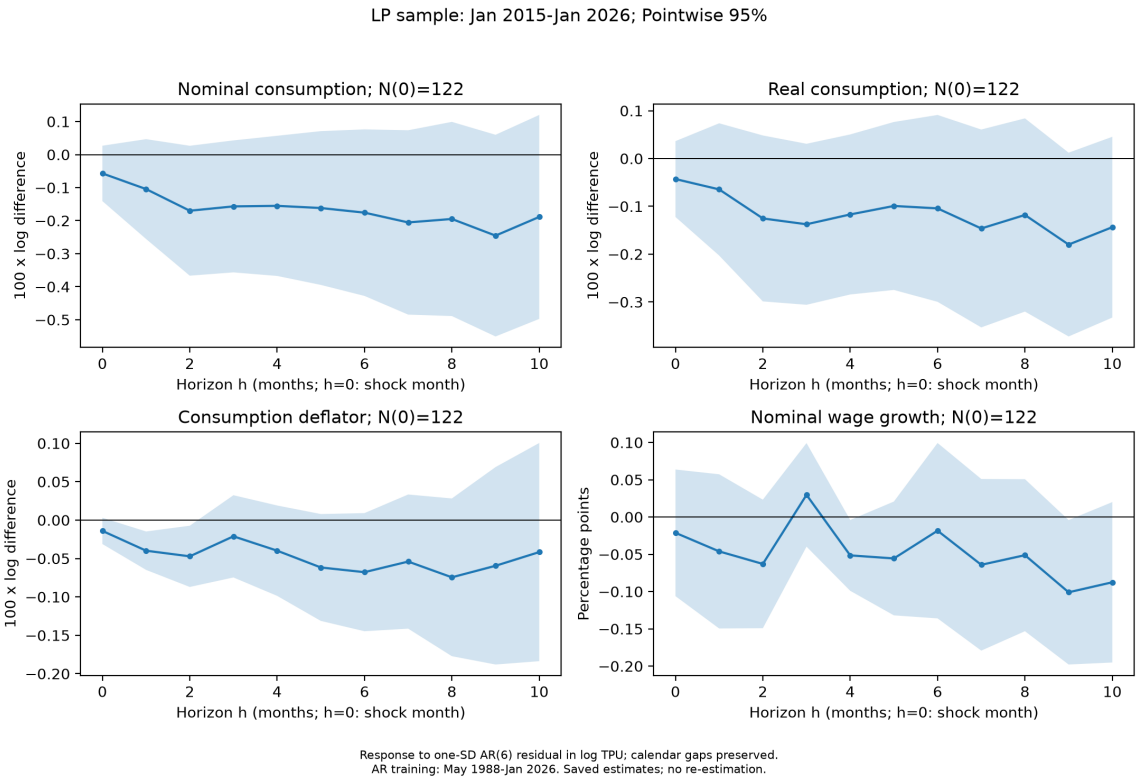


Figure 4: Consumption, prices and nominal wage growth with pointwise 95% intervals.

On common regression support, nominal consumption equals real consumption plus the implicit deflator up to numerical precision. Portfolio adding-up is not exact in the survey data: the sum of the three impact coefficients is 0.00428 percentage points and its largest absolute value over the reported horizons is 0.01496. These residuals are disclosed rather than removed by renormalizing the observed shares.

2.4 Empirical Robustness and Identification Limits

Table 2 keeps the calendar convention fixed while varying the excluded period. The signs of all three impact allocation responses survive no exclusion and longer COVID exclusions. Longer exclusions reduce effective sample size and sometimes remove simultaneous significance. They are sensitivity checks, not a search for the most significant cutoff. Re-

moving 2018–19 weakens the bond evidence; removing 2025 weakens the cash evidence and removes simultaneous impact significance for equities. The 2000-onward sample is a supplementary comparison, not the main specification.

Table 2: Impact allocation sensitivity, percentage points

Specification	Stock	Cash	Bond	N
Main	-0.508	0.330	0.183	122
No COVID exclusion	-0.707	0.568	0.145	132
Exclude calendar 2020	-0.475	0.306	0.174	114
Exclude through June 2021	-0.499	0.328	0.177	110
Exclude through December 2021	-0.529	0.355	0.179	104
Exclude 2018–19	-0.401	0.282	0.122	97
Exclude 2025	-0.457	0.291	0.169	109
Six control lags	-0.497	0.322	0.179	116
AR(3)	-0.467	0.302	0.168	125
AR(12)	-0.479	0.324	0.160	116
State-space AR	-0.498	0.322	0.181	126
Joint financial controls	-0.492	0.310	0.186	122
2000-onward comparison	-0.619	0.532	0.087	302

Point estimates only; this table does not imply uniform statistical significance. The shock is re-estimated under the specified availability and AR assumptions. The full-sample row is reported only as a comparison.

AR orders 3 and 12, pre-2020 AR training, missing-aware state-space filtering, three or six control lags, and lagged financial controls preserve the main impact signs. The state-space AR gives stock, cash and bond impacts of -0.498, 0.322, and 0.181 percentage points. It uses observed forecast errors, not imputed LP outcomes. Standardizing on post-2015 residuals changes the shock unit and hence coefficient magnitudes, rather than constituting a new structural shock. Dependence blocks of six and eighteen months and leverage-adjusted heteroskedasticity-consistent estimators leverage adjustment preserve the baseline pointwise allocation conclusions. They do not establish coverage under arbitrary dependence or nonstationarity. Maximum impact leverage is approximately 0.528 despite a well-conditioned design matrix with a condition number of 2.55, so leverage sensitivity is substantive rather than cosmetic. The generated-shock uncertainty is held fixed, as AR-refit bootstrap and finite-sample coverage analysis have not been performed.

The lead sensitivity regressions specify outcome variables as first differences $y_{t-j} - y_{t-j-1}$ for $j = 1, 2, 3$. Within the portfolio allocation group, Holm-adjusted joint lead tests yield p-values of 0.009 for equities, 0.079 for cash, and 0.4805 for sovereign bonds, leading to a clear rejection of the lead restriction for equity holdings. Among the supple-

mentary outcomes, real consumption growth and nominal wage growth reject the zero-lead restriction with adjusted p-values of 0.010 and 0.0135, respectively, whereas the NASDAQ index exhibits no significant lead response with an adjusted p-value of 0.225.

The joint lead test for nominal consumption is omitted due to a singular covariance matrix. This singularity occurs because the first lead outcome coincides with the lagged nominal PCE growth control variable, generating exact mechanical collinearity. A related structural failure affects the three-lag nominal PCE specification across both lead and post-estimation families, where the growth control operates as a linear combination of the included level lags. Consequently, these two invalidated estimation structures are excluded from model comparisons and are not interpreted as valid robustness tests.

Consequently, the empirical findings reflect conditional associations rather than a strictly identified causal trade policy uncertainty shock. Simultaneous inference bands are evaluated within individual impulse response paths, whereas lead-adjustment corrections apply strictly within designated specification families rather than across the full specification search. Although calendar support evaluations, covariance diagnostics, and explicit documentation of model non-convergence enhance empirical transparency, they do not fully eliminate structural endogeneity. Ultimately, the subsequent calibration framework leverages these empirical estimates to discipline the magnitude and qualitative hierarchy of theoretical responses, avoiding the interpretation of local projection shocks as regime transitions or direct inventory observations.

3 Model Description

This section outlines the theoretical framework that links tariff shocks to household portfolio choices. The model builds on Merton's continuous-time portfolio problem, extended to include tariff-induced disturbances that affect asset returns. The final goods, labor, and fiscal accounts are closed by explicit market clearing conditions in Section 3.6. By contrast, intermediate inputs are supplied with infinite elasticity at exogenous world prices, and the equity, bond, and cash account claims are in exogenous supply with the price processes of Section 3.4, so the model is deliberately a partial equilibrium on the input and asset sides while remaining closed on the goods, labor, and fiscal sides.

3.1 Trade Policy Uncertainty Regimes

The economy transitions between two regimes that differ in the level of trade policy uncertainty. Regime 1 is a low uncertainty regime, a stable trade environment similar to the pre trade war period, in which tariff policy changes only gradually. Regime 2 is a high uncertainty regime, an escalatory environment characterized by the declaration of trade wars, frequent and large tariff revisions, and other trade conflict events. Let $\xi(t) \in \{1, 2\}$ denote the regime in force at time t, modeled as a continuous time Markov chain with infinitesimal generator

$$\mathbf{Q}(t) = \begin{pmatrix} -\lambda_{12}(t) & \lambda_{12}(t) \\ \lambda_{21}(t) & -\lambda_{21}(t) \end{pmatrix} \tag{2}$$

The intensities $\lambda_{12}(t)$ and $\lambda_{21}(t)$ represent the transition rates between the low and high uncertainty regimes. Specifically, $\lambda_{12}(t)$ measures the intensity of moving from the low uncertainty regime 1 to the high uncertainty regime 2. Similarly, $\lambda_{21}(t)$ captures the intensity of the reverse transition. The sojourn time in each regime follows an exponential distribution. Both intensities are strictly positive with $\lambda_{12}(t) > 0$ and $\lambda_{21}(t) > 0$.

The transition intensities are not constant over time because trade policy announcements and political statements actively shift them. Specifically, an escalatory statement raises $\lambda_{12}(t)$ to make a transition to the high uncertainty regime more imminent, whereas a cooperative statement increases $\lambda_{21}(t)$ to make de-escalation more likely. A TPU shock represents such a sudden shift in these transition intensities, and the news-based TPU index in the empirical analysis serves as its observable counterpart. This paper therefore writes these intensities as time-varying functions to emphasize that the underlying transition matrix itself responds directly to the arrival of trade policy news.

The regime governs the intensity of the uncertainty that households face about the firms' inventory driven cash flow. In the high uncertainty regime the ambiguity that households attach to the undisclosed inventory is wider, which I formalize through the relative-entropy bound $\kappa_k(t)$ in Section 3.4. The bound is larger in the high uncertainty regime and rises with the escalation intensity $\lambda_{12}(t)$, so escalatory news widens the set of inventory laws the household entertains even before the regime itself switches. The

regime is publicly observed, so it is a source of risk. What households cannot observe, and what therefore carries Knightian uncertainty, is the firm's inventory.

The regime bundles two distinct components of a trade conflict. The tariff component raises the applied tariff from τ_1 to τ_2, which directly enters both marginal costs and fiscal revenues. Meanwhile, the uncertainty component widens the dispersion parameter $\bar{\nu}_k$ and expands the entropy bound κ_k that households attach to undisclosed inventories. The TPU index shifts the transition intensity $\lambda_{12}(t)$ of this entire bundle. In contrast, the tariff rates τ_k are calibrated directly from applied tariff data rather than being derived from the index itself. Non-tariff instruments recorded by the index, such as upstream export controls, enter the same import cost channel and are consequently absorbed by the uncertainty component of the model.

The tariff rate faced by the firm is higher in the high uncertainty regime.

$$\tau(t) = \tau_{\xi(t)}, \qquad 0 \leq \tau_1 < \tau_2 \tag{3}$$

The parameters τ_1 and τ_2 denote the tariff rates in the low and high uncertainty regimes respectively. These rates enter the import price of newly imported intermediate goods. Economic agents observe both the current regime and the transition intensities. Therefore, the expected future tariff conditional on the current regime is given by a specific conditional expectation.

$$\mathbb{E}\big[\tau(t+s) \mid \xi(t) = k\big] = \big(e^{\mathbf{Q}(t)s}\,\boldsymbol{\tau}\big)_k, \qquad \boldsymbol{\tau} = (\tau_1, \tau_2)^\top \tag{4}$$

The firm uses this expectation when setting inventory policy in anticipation of future tariff changes. The term $e^{\mathbf{Q}(t)s}$ freezes the generator matrix at its value at time t. This adjustment is necessary because the intensities themselves vary with the arrival of new TPU news. The conditional mean summarizes the level channel of the model. Through this channel, a higher expected tariff increases the expected marginal cost. In contrast, the ambiguity about inventory across regimes summarizes the uncertainty channel. The transition intensities govern both channels. Consequently, both channels depend directly on TPU.

3.2 Firm Side Product Economy

The production side of country i consists of a continuum of monopolistically competitive firms indexed by $\omega \in [0, 1]$. These firms operate in continuous time with each producing a differentiated variety of the final good. A CES aggregator with elasticity of substitution $\varepsilon_p > 1$ combines these varieties into the final consumption good. Consequently, each firm faces the downward-sloping demand curve

$$y(\omega, t) = \left(\frac{p(\omega, t)}{P_Y(t)}\right)^{-\varepsilon_p} Y(t), \qquad P_Y(t) = \left(\int_0^1 p(\omega, t)^{1-\varepsilon_p}\, d\omega\right)^{\frac{1}{1-\varepsilon_p}} \tag{5}$$

Here, $Y(t)$ is aggregate final output and $P_Y(t)$ is the producer price index. Because firms are symmetric, this paper drop the index ω and analyze the symmetric equilibrium where $p(\omega, t) = P_Y(t)$ and $y(\omega, t) = Y(t)$. The market power arising from $\varepsilon_p < \infty$ generates the strictly positive operating profits that back the equity claim priced in Section 3.4. This positive markup is structurally necessary because the constant returns technology would otherwise generate zero cash flows. Each firm uses labor and a composite intermediate input to produce its variety.

$$Y(t) = A(t)L(t)^{\alpha} X(t)^{1-\alpha} \tag{6}$$

where $Y(t)$ denotes output, $A(t)$ is total factor productivity, $L(t)$ is labor supplied by households, and $X(t)$ is the composite bundle of intermediate goods. Output is thus determined by technology, labor input, and the quantity of effective intermediates.

The composite intermediate input combines domestically sourced and foreign-sourced inputs through a constant elasticity of substitution(CES) aggregator,

$$X(t) = \left[s^{\frac{1}{\eta}} X_D(t)^{\frac{\eta-1}{\eta}} + (1-s)^{\frac{1}{\eta}} X_F(t)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{7}$$

$$X_F(t) = X_F^N(t) + X_F^I(t) \tag{8}$$

where $X_D(t)$ is the quantity of domestically produced intermediate goods, $X_F(t)$ is the quantity of foreign intermediates effectively used in production, and $s \in (0, 1)$ captures the degree of home bias or trade openness in intermediate input use. The elasticity of

substitution between domestic and foreign intermediates is $\eta > 1$. Foreign intermediates used in production consist of two components $X_F^N(t)$ which are newly imported and used immediately and $X_F^I(t)$ which are drawn from existing inventories.

The unit cost of foreign intermediates actually used in production is a weighted average of the contemporaneous import price and the average cost of inventory. Let

$$P_F^{\text{use}} = \phi(t)P_M(t) + (1 - \phi(t))P_I(t) \tag{9}$$

$$P_M(t) = (1 + \tau(t))\epsilon(t)p_{x,j}(t) \tag{10}$$

Here, $P_F^{\text{use}}(t)$ represents the effective price of foreign inputs used at time t, while $P_M(t)$ is the current domestic currency price of newly imported intermediates. The variable $\epsilon(t)$ is the nominal exchange rate expressed as domestic currency per unit of foreign currency. Meanwhile, $p_{x,j}(t)$ is the foreign currency price of the intermediate good produced in country j. The share $\phi(t) = \frac{X_F^N(t)}{X_F(t)} \in [0, 1]$ represents the fraction of foreign inputs coming from contemporaneous imports. The remaining share $1 - \phi(t)$ is drawn directly from inventories and is thus valued at the average inventory cost $P_I(t)$. When the firm relies more heavily on inventories, the effective cost moves closer to the historical average inventory price. Conversely, when the firm relies more on new imports, the cost stays closer to the current import price and reflects the contemporaneous shock.

The average price of inventory $P_I(t)$ evolves according to an asset accumulation equation that tracks the value of the stock of foreign intermediates held in inventory is

$$\frac{dP_I(t)I(t)}{dt} = P_M(t)Q_X(t) - P_I(t)X_F^I(t) - \delta_I P_I(t)I(t) \tag{11}$$

where $I(t)$ denotes the physical stock of foreign intermediate inventories, $Q_X(t)$ is the flow of newly imported intermediates added to inventory, X_F^I is the flow of inventory drawn down for production, and $\delta_I \in (0, 1)$ is the depreciation or obsolescence rate of inventories. The first term on the right-hand side captures the value of new imports entering inventory at the current import price, the second term reflects the removal of inventory at its average cost when it used in production, and the final term captures the loss of inventory due to depreciation. This structure implies the average inventory price adjusts gradually toward the current import price as new imports are accumulated, with

the speed of adjustment governed by the ratio $\frac{Q_X(t)}{I(t)}$.

Given the effective cost of foreign intermediates, the unit cost of the composite intermediate input used in production is

$$P_X(t) = \left[s p_{x,i}(t)^{1-\eta} + (1-s)\big(P_F^{\text{use}}(t)\big)^{1-\eta} \right]^{\frac{1}{1-\eta}} \tag{12}$$

Here, $p_{x,i}(t)$ is the price of domestically produced intermediates. The final input cost index $P_X(t)$ is determined by the degree of home bias, the domestic intermediate price, and the effective price of foreign intermediates. This effective price itself depends on both the current import price and the average inventory cost. Through this channel, the inventory mechanism directly affects marginal costs and firm profitability.

The law of motion for the inventory is

$$\dot{I}(t) = Q_X(t) - X_F^I(t) - \delta_I I(t) \tag{13}$$

Indicating that the stock of inventories increases with new imports $Q_X(t)$, decreases with the quantity of inventories used in production $X_F^I(t)$, and declines over time due to depreciation. While new imports exceed the amount of inventories drawn down for production and depreciation, the inventory stock expands. When they fall short, the inventory stock is reduced. This formulation captures the idea that inventories allow firms to smooth the effect of shocks to import prices and tariffs over time.

Firms optimally choose the inventory dynamics in (11)–(13), but the resulting state variables remain strictly internal. These hidden states include the physical stock $I(t)$, the average valuation $P_I(t)$, and the sourcing split $\phi(t)$. Outside agents observe these variables only at discrete reporting dates when financial statements are released. While households observe market prices and wages continuously, they cannot directly monitor the inventory position between reporting dates. This undisclosed state enters the effective marginal cost $P_X(t)$ through (9)–(12), thereby shifting firm profits (16) and equity cash flows. As a result, households must form expectations about this unobserved object under incomplete information. The severity of this informational friction depends on both the prevailing tariff $\tau(t)$ and trade policy uncertainty. This structural friction ultimately generates the Knightian uncertainty that households face when pricing equity.

Cost minimization over labor and the composite intermediate implies, by the Cobb–Douglas structure of (6), a constant unit marginal cost

$$MC(t) = \frac{\Theta}{A(t)} w(t)^{\alpha} P_X(t)^{1-\alpha}, \qquad \Theta \equiv \alpha^{-\alpha}(1-\alpha)^{-(1-\alpha)} \tag{14}$$

where $w(t)$ is the wage rate, with conditional factor demands

$$w(t)L(t) = \alpha MC(t)Y(t), \qquad P_X(t)X(t) = (1-\alpha)MC(t)Y(t) \tag{15}$$

Given the demand curve (5) and the sales tax τ_s from Section 3.3, each firm sets its price as a markup over marginal cost. This pricing strategy yields strictly positive operating profits.

Lemma 1 (Markup pricing and positive profit)**.** *In the symmetric equilibrium the profit maximizing price and the resulting operating profit are*

$$\begin{aligned} P_Y(t) &= \frac{\varepsilon_p}{(\varepsilon_p - 1)(1-\tau_s)} MC(t), \\ \Pi(t) &= (1-\tau_s)P_Y(t)Y(t) - MC(t)Y(t) = \frac{(1-\tau_s)P_Y(t)Y(t)}{\varepsilon_p} > 0 \end{aligned} \tag{16}$$

Consequently, profit is strictly positive for every $\varepsilon_p \in (1, \infty)$ and vanishes only in the perfectly competitive limit $\varepsilon_p \to \infty$.

Profits are distributed continuously to shareholders and become the cash flow of the equity claim introduced in Section 3.4. Because $MC(t)$ depends on the effective intermediate price $P_X(t)$ through (9)–(12), the inventory mechanism plays a central role in determining how tariff changes and trade policy uncertainty translate into the firm's marginal costs, prices, and profits.A higher inventory level reduces marginal cost $MC(t)$, thereby lowering the domestic price $P_Y(t)$. With elasticity of substitution $\eta > 1$ between home and foreign final goods, the lower price raises the equilibrium revenue $P_Y(t)Y(t)$ collected by domestic firms, so profit is increasing in the inventory. This monotonicity is what Section 3.4 formalizes as the positive equity loading $\psi_k^S > 0$.

In each country, firms choose the price of their variety together with labor, domestic and foreign intermediate inputs, and the inventory policy to maximize the discounted

present value of profits. Formally, the firm solves

$$\max_{P(t),L(t),X_D(t),X_F^N,X_F^I,Q_X} \int_0^\infty e^{-\rho t}\Pi(t)\,dt \tag{17}$$

subject to the demand curve (5), the production technology (6), the definition of the composite intermediate input (7), the decomposition of foreign intermediates into new imports and inventory drawdowns (8), the law of motion for the inventory value (11), the law of motion for the physical inventory stock (13), and the requirement that drawdowns cannot exceed contemporaneous input use, $X_F^I(t) \leq X_F(t)$, together with the nonnegativity of orders and drawdowns, $Q_X(t) \geq 0$ and $X_F^I(t) \geq 0$.

Two observations reduce the firm's problem to a manageable core. First, by Lemma 1, the pricing decision is static and separates from the inventory decision because marginal cost is independent of the firm's own price. Second, define excess inventory by $E(t) = I(t) - K_0$, where K_0 is the constant operational-stock benchmark and $I(t)$ remains physical inventory. Thus $I(t) \geq 0$ is equivalent to $E(t) \geq -K_0$, and $E(t) < 0$ denotes a shortfall relative to the benchmark. The operational stock held for delivery lags and stockout avoidance remains constant outside the model, so $E = 0$ represents the baseline operational level. Thus, the inventory variable captures only the component that responds to trade policy.

Let $\bar{X}_F$ denote the stationary foreign input requirement and let $q(E, k)$ denote the marginal value of one unit of excess inventory in regime k. Under this normalization, the inventory policy solves the cost-minimization problem (71) in Appendix B. The solution is a barrier policy.

Proposition 1 (Firm optimum: the inventory barrier policy)**.** *The optimal inventory policy satisfies the following properties.*

(i) The firm follows a barrier policy. Whenever $q(E, k) > P_{M,k}$, it purchases immediately up to the regime-specific barrier E_k^*. Whenever $q(E, k) < P_{M,k}$, it draws down inventory at the maximal rate $X_F^I = \bar{X}_F$. At the barrier, purchases only replace depreciation.

(ii) The high-uncertainty regime has no excess-inventory target, $E_2^* = 0$. Following a realized tariff escalation, the firm therefore draws down its existing inventory at the maximal rate until it is exhausted, which takes approximately m_1^* months.

(iii) A strictly positive excess-inventory target exists in the low-uncertainty regime if and only if the expected tariff capital gain exceeds the carrying cost,

$$\lambda_{12}\big(P_{M,2} - P_{M,1}\big) > (\rho + \delta_I)P_{M,1} \tag{18}$$

When this condition holds, the optimal inventory coverage, $m_1^* = E_1^*/\bar{X}_F$, is

$$m_1^* = \frac{1}{\rho + \delta_I + \lambda_{21}} \ln \frac{P_{M,2} - \bar{q}}{B - \bar{q}}, \qquad B = \frac{(\rho + \delta_I + \lambda_{12})P_{M,1}}{\lambda_{12}}, \qquad \bar{q} = \frac{\lambda_{21}P_{M,1}}{\rho + \delta_I + \lambda_{21}} \tag{19}$$

(iv) The optimal inventory coverage m_1^* is increasing in λ_{12} and $P_{M,2}/P_{M,1}$, and decreasing in δ_I and λ_{21}.

The proposition establishes front-loading as a threshold phenomenon. In normal times, the escalation intensity remains below the threshold in (18). Interest and depreciation costs therefore dominate the option value of purchasing inputs before a tariff increase, keeping excess inventories at zero. Once the expected escalation intensity exceeds the threshold, the option value becomes sufficiently large to justify immediate stockpiling, and the firm jumps discretely to a strictly positive excess-inventory target. Upon a realized tariff escalation, the firm no longer restores imported inputs and instead runs down the accumulated inventory at the maximal rate until it is exhausted.

The comparative statics in part (iv) follow directly from the option-value logic. A higher escalation intensity λ_{12} or a larger expected price jump $P_{M,2}/P_{M,1}$ raises the expected capital gain from holding low-cost inventory, leading the firm to choose a higher coverage target. By contrast, a faster reversion rate λ_{21} shortens the period over which the inventory provides insurance, while a higher depreciation rate δ_I increases the carrying cost. Both effects reduce the optimal excess-inventory target.

These episodic dynamics closely match the anticipatory import surges documented around announced tariff changes by Khan and Khederlarian (2021, 2025). The calibration in Section 5 maps the comparative statics above into the magnitude of the observed import burst. Furthermore, the effective input price in (9) buffers the subsequent drawdown phase because depreciation enters both the threshold condition and the optimal purchase size. Through this mechanism, the inventory technology disciplines the quantitative strength

of the front-loading channel.

The reduction to cost minimization is exact to first order. Because the effective input price enters marginal cost, a change in the sourcing split alters output and revenue. However, at the optimal price, the envelope theorem renders this induced adjustment a second-order effect for profit, leaving only the direct effect on the discounted import bill. Since the physical requirement is met either through new imports or the accumulated stock, the inventory policy reallocates the timing of imports across regimes to solve the cost minimization problem (71). This timing creates the delayed pass-through recorded by the effective input price P_F^{use} in (9). Through this cost channel, the undisclosed inventory ultimately impacts the profits that households must infer.

3.3 Government

The two production economies are linked to the asset market through a fiscal authority. This authority collects taxes and issues public debt. The government raises revenue from a household consumption tax, a firm sales tax, and tariff receipts on imported intermediates. Total tax revenue is

$$T(t) = \underbrace{\tau_c \, P_C(t) C(t)}_{\text{consumption tax}} + \underbrace{\tau_s \, P_Y(t) Y(t)}_{\text{sales tax}} + \underbrace{\tau(t) \, \epsilon(t) p_{x,j}(t) \left[X_F^N(t) + Q_X(t) \right]}_{\text{tariff revenue}} \tag{20}$$

Here, τ_c and τ_s represent constant consumption and sales tax rates. The equilibrium quantities $C(t)$, $Y(t)$, and the import flows are determined by the market-clearing conditions in Section 3.6. Therefore, the tax bases in (20) match the objects in the household budget constraint (34) and firm revenue (16). Tariff duties accrue to current imports as they cross the border. This mechanism applies whether these inputs are used immediately as $X_F^N(t)$ or added to inventory as $Q_X(t)$. Conversely, inventories drawn from existing stocks $X_F^I(t)$ generate no current duty because they were already taxed at their original import date. Consequently, tariff revenue directly inherits the inventory policy of the firm. When firms buffer against tariff shocks by cutting imports and running down stocks, current tariff receipts fall even as the tariff rate $\tau(t)$ rises.

On the expenditure side, the government services its outstanding debt and, in principle, finances discretionary spending $G(t)$. A fully specified treatment of $G(t)$, in which

public spending feeds back into goods markets and household budgets, would require closing the model on the fiscal side and lies outside the portfolio allocation focus of this paper. This paper therefore adopts a deliberately simple government that abstracts from discretionary expenditure, fixing $G(t) \equiv \bar{G}$, set to zero in the baseline, so that the net fiscal position is governed by the revenue side. Writing the primary balance as $NG(t) \equiv T(t) - G(t)$, the stock of public debt $D(t)$, evolves according to the flow budget constraint

$$dD(t) = \left[r_k^D \, D(t) - NG(t) \right] dt \tag{21}$$

where r_k^D is the rate at which the government services its debt in regime k, so that the government issues new debt whenever debt service exceeds the primary surplus and retires debt otherwise. The bond held by households, denoted $B(t)$ with the dynamics (23), is the market value of the claim on this debt. Fluctuations in $NG(t)$ driven by the firm's unobserved new imports are what expose $B(t)$ to the undisclosed inventory, with loading ψ_k^B.

This structure makes the supply of bonds and their pricing regime-dependent. In the high-uncertainty regime, weaker consumption and output diminish the consumption and sales tax bases. Meanwhile, the net effect on tariff receipts remains ambiguous. The higher tariff rate $\tau(t)$ is partly offset by a contraction in taxable new imports $X_F^N(t) + Q_X(t)$. This contraction occurs because firms rely more heavily on existing inventories. The primary balance $NG(t)$ thus worsens. The associated rise in debt issuance and the debt service rate r_k^D raises future fiscal outlays. This mechanism serves as the channel through which trade policy uncertainty transmits to the bond market. Although the regime $\xi(t)$ is publicly observed, the realized fiscal position remains hidden. Tariff revenue $\tau(t)\epsilon(t)p_{x,j}(t)\left[X_F^N(t) + Q_X(t)\right]$ depends directly on the firm's unobserved new imports. Through these imports, tariff revenue also depends on the unobserved inventory level. The bond therefore inherits the same Knightian uncertainty as equity. Consequently, households form robust beliefs about (μ_k^B, σ_k^B) as outlined in Section 3.4. However, the bond is exposed to the inventory state only through the tariff revenue component of fiscal income. It therefore loads on this uncertainty more lightly than equity does. This lighter loading explains why the bond carries a smaller robustness adjustment and a smaller

estimation variance than the stock.

The inventory policy behind this channel is characterized in Proposition 1, the resulting exposure is priced in Section 4, and Section 5 quantifies it. In particular, the barrier policy implies a fiscal J curve in which tariff receipts collapse during the drawdown phase and overshoot afterward.

3.4 Financial Market Environment and Belief Formation

Households allocate wealth among three assets, a risky equity claim $S(t)$, a government bond $B(t)$, and a cash account $M(t)$. Conditional on regime $\xi(t) = k$, the dynamics are

$$dS(t) = \mu_k^S \, S(t) \, dt + \sigma_k^S \, S(t) \, dZ^S(t), \tag{22}$$

$$dB(t) = \mu_k^B \, B(t) \, dt + \sigma_k^B \, B(t) \, dZ^B(t), \tag{23}$$

$$dM(t) = r_k^M \, M(t) \, dt \tag{24}$$

where r_k^M is the observed interest rate and Z^S and Z^B are standard Brownian motions adapted to the household's filtration. The dynamics in (22) and (23) hold within a regime, and because the coefficients change with the regime the asset values change with it as well. At a transition from regime k to regime l the equity claim and the bond are revalued at the instant of the switch, moving from their regime k values to their regime l values, while the cash account remains continuous. This revaluation is the only jump in the asset prices, and it is the object the household hedges when it chooses its portfolio in Section 4. The coefficients μ_k^i and σ_k^i are not primitives. They are the consequence of the household's belief about the firm's inventory, and the rest of this subsection derives them from that belief. The construction of this belief draws inspiration from Drechsler (2013), specifically from the idea that the household hedges its uncertainty about fundamentals and from the use of relative entropy as the theoretical device that disciplines the beliefs it is willing to hold. The household starts from a reference law for the inventory but does not treat its own expectation as perfectly accurate, and it therefore also entertains alternative laws that depart from the reference.

Everything that matters for pricing is a function of the inventory. Profit is determined by the markup pricing of Lemma 1 and the goods market clearing of Section 3.6. This

process functions through the marginal cost (14). This cost depends on the average inventory price $P_I(t)$ and the sourcing split $\phi(t)$, while the fiscal balance is determined through tariff revenue. This revenue depends on the import flows drawn from the same inventory policy. Equity is a claim on profit and the bond is a claim on public debt, so the asset values are functions of the inventory, $S(t) = S(I(t); k)$ and $B(t) = B(I(t); k)$, with log-value sensitivities $\psi_k^S(I) = \partial \ln S(I; k)/\partial I$ and $\psi_k^B(I) = \partial \ln B(I; k)/\partial I$. The variables $P_I(t)$ and $\phi(t)$ are internal to the firm and are disclosed only at the reporting date, so the household observes asset prices and wages continuously but its information set does not contain the firm's real-time profit. The household's only inference problem is therefore to form a belief about physical inventory $I(t)$. Once that belief is formed, profit, the fiscal balance, the asset values, and the return moments all follow as a consequence.

Under the reference measure $\mathbb{P}_0$, the household's belief about the firm's inventory between reporting dates is

$$dI(t) = b_k\big(I(t)\big)\, dt + \bar{\nu}_k\, dW_t, \qquad \bar{\nu}_2 > \bar{\nu}_1 > 0 \tag{25}$$

where b_k is the believed net accumulation of physical inventory implied by the firm's optimal policy of Proposition 1. Meanwhile, $\bar{\nu}_k$ represents the dispersion the household attaches to an inventory it cannot currently observe. The household knows the firm's policy rule but not the current state. Therefore, b_k serves as the conditional mean drift under that rule. This drift is zero in normal times, when the intensity sits below the threshold (18) and excess stocks are exhausted. It turns negative upon entry into the high uncertainty regime, yielding $b_2 = -(\bar{X}_F + \delta_I \widehat{I})$ during the drawdown phase. Finally, it spikes upward in regime 1 when escalation news pushes the intensity above the threshold. Although the firm-level policy is bang bang, the household faces residual uncertainty about order timing and the remaining stock. This uncertainty makes the believed drift a smooth function of its conditional mean $\widehat{I}_t$. This mean evolves deterministically between reporting dates and resets at disclosures. At each monthly reporting date t_n, the disclosure reveals the physical inventory and associated internal states, so the posterior is reset to that observation and the information gap is zero; between t_n and t_{n+1}, the household propagates the reference law and ambiguity set forward. The randomness in (25) is a

property of the household's belief and not of the inventory itself. The firm knows its inventory at every instant and accumulates it through (13). In contrast, the household's information set contains the inventory only at reporting dates. Between disclosures it therefore describes the evolution by the reference law (25), and because it does not fully trust that law, by the set of laws around it in (27). The diffusion term captures risk within a given law. The Knightian component is the household's inability to single out one law from the set, which is what the entropy bound disciplines. This belief dispersion expands in the high uncertainty regime. In this state, tariffs and the firm's sourcing split move more between disclosures. The entropy apparatus strictly requires a stochastic reference model. A deterministic reference law would be mutually singular with respect to any diffusive alternative. This singularity leaves the relative-entropy neighborhood around it with no nontrivial models.

Because the household does not treat the reference drift as perfectly accurate, it also entertains alternative measures $\mathbb{Q}$ under which the drift of its belief is distorted,

$$dI(t) = \big(b_k(I(t)) + \bar{\nu}_k h_t\big)\, dt + \bar{\nu}_k\, dW_t^{\mathbb{Q}} \tag{26}$$

where $W_t^{\mathbb{Q}}$ is a Brownian motion under $\mathbb{Q}$ and h_t is the drift distortion. The growth rate of relative entropy generated by h_t is $\frac{1}{2}h_t^2$, as in Drechsler (2013), and the set of feasible alternatives is

$$\mathcal{P}_k = \Big\{ \mathbb{Q} : \tfrac{1}{2}h_t^2 \le \kappa_k(t) \text{ for all } t \Big\} \tag{27}$$

The bound widens with trade policy uncertainty,

$$\kappa_k(t) \;=\; \underbrace{\bar{\kappa}\,\mathbf{1}_{\{k=2\}}}_{\text{regime level}} \;+\; \underbrace{\varphi_\kappa\,\frac{\lambda_{12}(t)}{\lambda^{\max}}}_{\text{escalation news}}, \qquad \bar{\kappa},\ \varphi_\kappa > 0 \tag{28}$$

so that it has a level component active only in the high uncertainty regime, where tariff schedules and the firm's sourcing decisions are revised frequently between reporting dates, and a news component that rises with the escalation intensity. A trade policy uncertainty shock, the movement in $\lambda_{12}(t)$ described in Section 3.1, therefore widens the set of inven-

tory laws the household entertains on impact, even before any regime change is realized.

Throughout, this paper maintains $|\psi_k^S| > |\psi_k^B| > 0$. A larger inventory reduces the effective marginal cost of production, thereby increasing firm profits and, in turn, equity value. Bond value also rises with the inventory, although less strongly. The reason is that inventory affects bondholders only through the tariff-revenue component of the fiscal balance. In particular, the loss of tariff revenue due to forgone new imports is second order compared with the increase in the tax base resulting from lower production costs. The perceived volatilities follow from Ito's lemma applied to $S(I;k)$ and $B(I;k)$ under (26),

$$\sigma_k^i = \psi_k^i(I)\,\bar{\nu}_k, \qquad i \in \{S, B\} \tag{29}$$

Consequently, both volatilities rise with $\bar{\nu}_k$ in the high-uncertainty regime. This increase is more pronounced for equity than for the bond. The inventory channel is the dominant component of perceived volatility but not its entirety. Asset-specific disclosure and liquidity noise remain unmodeled. However, these features keep the correlation of the two returns strictly below one. Accordingly, this paper maintains the condition $(\sigma_k^S)^2(\sigma_k^B)^2 > (\sigma_k^{SB})^2$ to ensure that Σ_k is invertible. The robust drift equals the reference return adjusted by the worst-case distortion,

$$\mu_k^i = \mu_k^{i,0}(I) + \psi_k^i(I)\,\bar{\nu}_k\,h_k^*, \qquad i \in \{S, B\} \tag{30}$$

where $\mu_k^{i,0}(I)$ is the reference expected return and h_k^* solves the inner minimization of the household's problem in Section 4. There, this paper show that the worst case equals $h_k^* = -\sqrt{2\kappa_k}$ for a household that is long both risky assets. Consequently, a rise in κ_k lowers both robust drifts. This reduction is more pronounced for equity than for the bond because $\psi_k^S > \psi_k^B$. The expected returns and the volatilities are thus not assumed arbitrarily. Instead, they are inherited directly from the worst-case belief about the inventory. The cash account remains completely free of this channel. Section 5 measures these loadings from return moments. Equation (67) serves as their empirical counterpart.

The drift μ_k^i admits two decompositions that serve different purposes. The belief

decomposition (30) splits it into the reference return and the worst case distortion, and is the object the household optimizes against. The accounting decomposition splits the same total return into payout and capital gain. Equity pays out the profit flow of Lemma 1 continuously and the bond pays the debt service of (21), and I write

$$CF_k^S = \Pi_k, \qquad CF_k^B = r_k^D D \tag{31}$$

for the cash flows of the two claims in regime k. The processes (22)–(23) are cum dividend gain processes, the values of holdings with payouts reinvested, so that

$$\mu_k^S = \frac{CF_k^S}{S(t)} + g_k^S, \qquad \mu_k^B = \frac{CF_k^B}{B(t)} + g_k^B \tag{32}$$

with g_k^i the capital gain component. The reference return $\mu_k^{i,0}$ in (30) is accordingly a total return, payout inclusive. The budget constraint (34) credits payout income through the first components of (32), and the valuation in Section 4 prices the cash flows (31). Only the total returns (μ_k^i, σ_k^i) enter the household problem, so the accounting split affects no portfolio object.

Given these beliefs, the household chooses consumption C and the portfolio shares θ^S and θ^B to maximize expected discounted utility under the worst case,

$$J = \max_{\{C,\theta^S,\theta^B\}} \min_{\mathbb{Q}\in\mathcal{P}_k} \mathbb{E}_0^{\mathbb{Q}} \int_0^\infty e^{-\rho t} \Big(\frac{C^{1-\gamma}}{1-\gamma} - \frac{l^{1+\chi}}{1+\chi} \Big) dt \tag{33}$$

where J is the value function, $\rho > 0$ the discount rate, $\gamma > 0$ the relative risk aversion, l labor supply, and χ the curvature of labor disutility. The optimization is subject to the self financing wealth dynamics and the admissibility conditions in Section 3.5. The reference returns $\mu_k^{i,0}(I)$, the worst-case distortion h_k^*, and the portfolio shares are solved in Section 4 and Appendix B by the HJB and log-linear methods.

3.5 Merton's Portfolio Theory with Tariff Volatility and Three Assets

Households allocate wealth among three assets. These assets include a risky stock, a bond, and cash. Let total wealth be $W_t = S_t + B_t + M_t$. Here, S_t, B_t, and M_t denote dollar holdings in the stock, bond, and cash respectively. The economy is defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The information available to the household is the filtration $\{\mathcal{F}_t^H\}_{t\geq 0}$ introduced in Section 3.4. This filtration is generated by the histories of asset returns and wages together with the publicly observed regime $\xi(t)$. This paper assumes that this information structure satisfies right continuity and completeness. The firm's inventory position is not contained in $\mathcal{F}_t^H$. Therefore, households observe and correctly identify the current regime. However, they form robust beliefs about the inventory-driven returns as described in Section 3.4. All consumption and investment decisions are required to be $\mathcal{F}_t^H$ progressively measurable.

The household maximizes the robust objective (33) under Constant Relative Risk Aversion(CRRA) utility, $U(C) = C^{1-\gamma}/(1-\gamma)$ with $\gamma > 0,\ \gamma \neq 1$, subject to the self financing budget constraint

$$\begin{aligned} dW_t =& \Big[\theta^S(t)W(t)\,\mu_k^S + \theta^B(t)W(t)\,\mu_k^B + \theta^M(t)W(t)\,r_k^M - (1+\tau_c)\,P_C(t)\,C(t) + w(t)l(t)\Big]dt \\ &+ \theta^S(t)W(t)\,\sigma_k^S\,dZ^S(t) + \theta^B(t)W(t)\,\sigma_k^B\,dZ^B(t) \end{aligned} \tag{34}$$

where μ_k^S, μ_k^B are the robust drifts (30) and σ_k^S, σ_k^B are the perceived volatilities (29) formed in Section 3.4, r_k^M is the regime k safe rate, and $Z^S(t), Z^B(t)$ are the standard Brownian motions of (22)–(23). The consumption composite and its price index are the CES aggregators

$$C(t) = \left\{ s_c^{\frac{1}{\eta}} c_{i,i}(t)^{\frac{\eta-1}{\eta}} + (1-s_c)^{\frac{1}{\eta}} c_{j,i}(t)^{\frac{\eta-1}{\eta}} \right\}^{\frac{\eta}{\eta-1}}, \tag{35}$$

$$P_C(t) = \left\{ s_c^{\frac{1}{\eta}} p_{i,i,c}(t)^{\frac{\eta-1}{\eta}} + (1-s_c)^{\frac{1}{\eta}} p_{j,i,c}(t)^{\frac{\eta-1}{\eta}} \right\}^{\frac{\eta}{\eta-1}} \tag{36}$$

where $s_c \in (0,1)$ is the household's home bias in consumption, the consumption analog of the intermediate input home bias s, $c_{i,i}(t)$ and $c_{j,i}(t)$ are consumption of the domestic and foreign goods, $p_{i,i,c}(t)$ and $p_{j,i,c}(t)$ are their prices, and $\eta > 1$ is the elasticity of

substitution between domestic and foreign goods.

Here $\theta^S(t)$ denotes the share of wealth invested in the stock, while $\theta^B(t)$ represents the share in the bond. The remaining share $\theta^M(t) = 1 - \theta^S(t) - \theta^B(t)$ is held in cash. The portfolio shares are adapted to the household's filtration. This filtration contains the publicly observed regime. Therefore, the shares are regime-dependent even though the index k is left implicit. Section 4 writes the optimal shares with the index explicit. By homotheticity, the optimal shares depend on the regime only through specific variables. These variables are the robust excess returns, the perceived covariance, and the hedging weight. This structure reflects both myopic and intertemporal hedging motives.

Consumption expenditure is taxed at the rate τ_c of (20), and I write $\tilde{P}_{C,k} \equiv (1+\tau_c)\, P_{C,k}$ for the effective consumer price in regime k, the object that enters the household's optimality conditions below. The consumption choice is also regime dependent. When tariffs are imposed, the prices of consumption goods adjust differentially across sectors and production technologies, and these relative price changes lead to higher inflation. As a result, the effective expenditure rises even for a given physical quantity of consumption. Consequently, the household's optimal consumption responds to the prevailing regime through the effective consumer price $\tilde{P}_C(t) = (1 + \tau_c)P_C(t)$. This price enters the budget constraint as $\tilde{P}_C(t)C(t)$. This formulation traces how tariff shocks propagate to household consumption through a price channel. In addition, contemporaneous labor income, regime-dependent robust returns, and perceived volatilities affect optimal consumption. This relationship holds because tariff shocks alter the intertemporal trade-off between current and future consumption. These shocks also modify the risk profile of the household's portfolio over the interval $[t, t + dt]$.

By contrast, the baseline assumes wages do not adjust one-for-one with tariff-induced price changes. This friction reflects that nominal wages are typically set in advance through labor contracts. Consequently, they do not adjust instantaneously to trade policy shocks. Because the wage is sticky, employment is entirely demand-determined. The labor market clears at the posted wage through (38). The household supplies exactly the firm's labor demand, so $l(t) = L(t)$. Therefore, labor is not a choice variable in (33). The additively separable disutility term affects neither the consumption nor the portfolio rules. Under market clearing, the wage bill wL_k inherits the regime through demand-determined

output. Corollary 2 makes this relationship precise. However, the calibrated gap across regimes is small because the consumption ratios differ little. The baseline therefore capitalizes labor income at its regime average, $\overline{wl}$. It treats the deviation of the current wage bill from this average as a transitory, uncapitalized component. Appendix B discusses the exact case with regime-dependent labor income alongside the regime-dependent safe rate case.

To ensure the problem is well posed and to rule out Ponzi type financing, I impose standard feasibility conditions. A strategy (C, θ^S, θ^B) is feasible if it is $\mathcal{F}_t^H$ progressively measurable, satisfies $C_t \geq 0$ almost surely for all t, and ensures non negative wealth, $W_t \geq 0$ almost surely. I further impose the transversality condition

$$\lim_{T\to\infty} \mathbb{E}_0\left[e^{-\rho T}\frac{W_T^{1-\gamma}}{1-\gamma}\right] = 0,$$

This condition rules out explosive debt accumulation. It also ensures that the value function remains finite. Let $\mathcal{A}$ denote the set of all feasible strategies. The household's problem is to maximize the robust objective (33) over $\mathcal{A}$ subject to (34). Since H is a positive constant, stating feasibility and transversality in W or in $\widehat{W} = W+H$ is equivalent up to that constant. Therefore, I use whichever formulation is convenient below.

In addition, the household can neither short an asset nor borrow. The financial shares therefore satisfy $\theta^S \geq 0$, $\theta^B \geq 0$, and $\theta^S + \theta^B \leq 1$, with cash given by the nonnegative residual. The baseline calibration in Section 5 keeps the optimum in the interior of this feasible set, so the analytical results below apply without modification.

In the baseline, this paper take the cash account rate to be regime-invariant, $r_1^M = r_2^M = r^M$. This assumption aligns with the muted response of short rates documented in Section 2. Regime-dependent cash account rates are examined numerically as a robustness exercise. Under this restriction and the regime average capitalization, human wealth represents the present value of the labor income flow. This human wealth is identical in both regimes. It satisfies the expression $H = \overline{wl}/r^M$. Consequently, it carries no revaluation at a regime switch. I define total wealth as

$$\widehat{W}(t) = W(t) + H$$

and define the corresponding shares of total wealth for $i \in \{S, B\}$ as

$$x^i(t) = \theta^i(t) W(t) / \widehat{W}(t)$$

. I substitute $\overline{wl} = r^M H$ for the labor income term in equation (34). This substitution yields the final expression.

$$d\widehat{W}_t = \Big[\widehat{W}_t\big(\boldsymbol{x}^\top \boldsymbol{\pi}_k + r^M\big) - \tilde{P}_{C,k}\, C_t\Big] dt + \widehat{W}_t\big(x^S \sigma_k^S\, dZ_t^S + x^B \sigma_k^B\, dZ_t^B\big) \tag{37}$$

and at a regime switch financial wealth is revalued while H is not, so total wealth moves to $W(1 + \boldsymbol{\theta}^\top \boldsymbol{\Delta}_{kl}) + H = \widehat{W}\big(1 + \boldsymbol{x}^\top \boldsymbol{\Delta}_{kl}\big)$. The problem is therefore homothetic in $\widehat{W}$, all results below are stated in $(\widehat{W}, \boldsymbol{x})$, and the financial shares follow as $\boldsymbol{\theta} = \boldsymbol{x}\, \widehat{W}/W$. Because H behaves like a holding of the safe asset, the scaling $\widehat{W}/W > 1$ shifts the financial portfolio toward the risky assets in the low uncertainty regime and deepens the retrenchment from equity in the high uncertainty regime, as an exact implication of the solution rather than an approximation.

3.6 Aggregation and Market Clearing

The model closes the final goods, labor, and fiscal accounts explicitly. Labor is internationally immobile, so the labor market clears country by country at the posted wage. Aggregating over the firm continuum,

$$\int_0^1 L(\omega, t)\, d\omega \;=\; L(t) \;=\; l(t) \tag{38}$$

where the first equality uses the symmetry of varieties, so employment is determined by the firm's labor demand (15) evaluated at the sticky wage, and the household supplies this quantity passively. The market for each home variety clears internationally, $y(\omega, t) = c_{i,i}(\omega, t) + c_{i,j}(\omega, t)$, which aggregates over the continuum to

$$Y(t) = c_{i,i}(t) + c_{i,j}(t) \tag{39}$$

where $c_{i,j}(t)$ denotes consumption of the country i good by country j households. The mirror conditions hold for country j. I define the final goods trade balance as $NX(t) \equiv p_{i,j,c}(t)\, c_{i,j}(t) - p_{j,i,c}(t)\, c_{j,i}(t)$. Given this definition, the nominal counterpart of (39) is

$$P_Y(t)\, Y(t) = P_C(t)\, C(t) + NX(t) \tag{40}$$

In regime 1, the two countries share a common baseline tariff alongside symmetric preferences and technologies. Consequently, $NX = 0$ holds in regime 1. Therefore, the value of domestic production equals the pre-tax value of domestic consumption expenditure. The escalation is unilateral. In regime 2, the home relative price rises, inducing expenditure switching. However, because intermediate-goods demand need not clear contemporaneously and can be accommodated through inventory adjustment, the resulting reallocation opens a small trade deficit. This deficit remains an order of magnitude below the consumption and revaluation effects. The baseline therefore treats $NX = 0$ as a first-order approximation. Section B carries the correction explicitly. On the fiscal side, the tax bases in (20) are the equilibrium quantities of (38)–(40). Public debt then evolves according to (21). Finally, the profit flow of Lemma 1 is distributed to shareholders. This flow enters the equity return through the payout component of (32).

Two markets are deliberately left open. Domestic and foreign intermediate inputs are supplied with infinite elasticity. They are supplied at the exogenous prices $p_{x,i}(t)$ and $p_{x,j}(t)$. This setup reflects a small open economy assumption on the input side. Consequently, the intermediate market does not clear domestically. Furthermore, equity, bond, and cash account claims are in exogenous supply. These claims follow the price processes described in Section 3.4. Thus, asset markets do not clear domestically either. Their prices are disciplined by the valuation of Lemma 2 rather than by a domestic market-clearing condition. A full general equilibrium with portfolio flow feedback is beyond the scope of this paper. The following result verifies that the closed accounts are mutually consistent.

The accounts that are closed are mutually consistent. Impose the clearing conditions (38)–(40), let dividends equal the profit flow of Lemma 1, and let the government satisfy (20)–(21) with $G(t) = 0$. Then, in drift terms within a regime, every domestic bilateral

flow cancels across the consolidated accounts of the household, the firm, and the government. These flows include wages, dividends, consumption and sales taxes, tariff receipts, and bond debt service. Consequently, the change in consolidated domestic net worth equals the net acquisition of claims on the outside sectors, expressed as

$$\frac{d}{dt}\big[W(t) - D(t)\big] = CG(t) + r^M M(t) - \Big[p_{x,i}(t)X_D(t) + \epsilon(t)p_{x,j}(t)\big(X_F^N(t) + Q_X(t)\big)\Big] + \iota(t) \tag{41}$$

In this expression, $CG(t)$ collects the exogenous capital gain components of the equity and bond returns. The bracket represents the pre-tariff bill paid to outside intermediate suppliers. The variable $\iota(t)$ denotes the accrual cash wedge generated by inventory investment. The complete derivation is provided in Appendix B. Because no domestic flow is left unmatched, Walras' law holds across the goods, labor, and fiscal accounts.

This identity provides the accounting discipline of the framework. The consumption that appears in government revenue is the consumption chosen by the household, and the labor the firm hires is the labor the household supplies. Similarly, the taxes the household and the firm pay equal the income the government receives, and the dividends the household earns match the profits the firm generates. What remains open, by construction, is exactly the set of accounts the paper does not model. These comprise the world market for intermediates, the cash account, and the price formation of the traded assets.

4 Solution

The inventory barrier in Proposition 1 is a maintained analytical benchmark in this version, not a globally verified consequence of the operating-profit problem. Its threshold and coverage expressions are used as candidate rules. The conditional inventory verification and local drawdown calculation in Appendix B do not establish their global optimality. In particular, elapsed drawdown time and linear excess-inventory coverage are distinct at positive depreciation. Claims about front-loading and its comparative statics that use this barrier are conditional on the benchmark, rather than independent proved results. The accounting identities, conditional household optimality conditions, and local results below have their own stated premises.

This section presents the analytical benchmark under its stated interiority, homotheticity and sign restrictions. Return moments and revaluations are followed by portfolio and real-side comparative statics, with derivations in Appendix B. These conditional expressions are not a verified solution of the full state-dependent constrained problem. Section 5 evaluates a numerical implementation and states explicitly where its measured inputs or approximations depart from the benchmark premises.

4.1 Robust return moments and the regime revaluation

Let $X_k^S = S_k$ and $X_k^B = B_k$ denote ex-payout prices, and let G^i denote the corresponding reinvested gain index. Define $\psi_k^i = \partial_I \log X_k^i$ and $\psi_{I,k}^i = \partial_I \psi_k^i$. For a locally stationary price map, Ito's lemma and payout accounting give the within-regime total reference return

$$\mu_k^{i,0} = \frac{CF_k^i}{X_k^i} + \psi_k^i b_k + \tfrac{1}{2}[(\psi_k^i)^2 + \psi_{I,k}^i]\bar{\nu}_k^2, \qquad \mu_k^i = \mu_k^{i,0} + \Psi_k^i h_k^*, \quad \Psi_k^i = \bar{\nu}_k \psi_k^i \tag{42}$$

Payout enters exactly once. Additional explicit time dependence or priced states require their own drift terms. Regime jumps are represented separately in the HJB. Total covariance also includes the mean-zero asset-specific return residual in the calibration. It is not another inventory observation. Log-linearity sets $\psi_{I,k}^i = 0$, not $CF_k^i/X_k^i = 0$. I enforce this payout-inclusive identity in the rebuilt conditional opportunity set rather than treating expected returns as free share-implied inputs.

For any specified admissible receipts and withdrawal policy, the physical drift is $Q_X - X_F^I - \delta_I I$, with $E = I - K_0$ recovered as an indicator. Transition intensities affect regime probabilities and continuation values. This identity does not supply a solved optimal inventory policy. The portfolio exercise conditions on its declared opportunity set.

The same ambiguity premium determines the asset revaluation at a regime switch. The next result characterizes this revaluation in closed form, taking into account that both the risk premium and the cash flow are expected to revert with the economic regime. Define the effective duration by $\tilde{D}^i \equiv \frac{1}{\delta^i - g^i + \lambda_{kl} + \lambda_{lk}}$ and let $\Delta \ln CF_{kl}^i \equiv \ln CF_l^i - \ln CF_k^i$ denote the log cash-flow gap across regimes.

Lemma 2 (Regime revaluation)**.** *Suppose the asset value solves the regime-switching*

valuation

$$\left(\delta_k^i - g^i + \lambda_{kl}\right) X_k^i = CF_k^i + \lambda_{kl} X_l^i \tag{43}$$

where the cash flows CF_k^i *are defined in* (31)*, with* $CF_k^S = \Pi_k > 0$ *by Lemma 1. Then*

$$\Delta_{kl}^i \simeq \tilde{D}^i \Big[(\bar{\delta}^i - g^i) \Delta \ln CF_{kl}^i - \left(p_l^i - p_k^i\right) \Big] \tag{44}$$

Negative revaluation requires the increase in required returns to exceed the cash-flow contribution. Equity falls more than bonds only under the corresponding cross-asset inequality and duration restrictions. Cash flows are first-order terms in (44), even when their coefficients are numerically small. The numerical price map determines whether these analytical restrictions hold at a given state.

The adverse-regime benchmark assumes stronger equity exposure to inventory uncertainty than sovereign-bond exposure. Positive operating profits support the equity cash-flow claim under the maintained payout convention, but positive profits and within-regime loading orderings alone do not establish the ordering of regime jumps.

The effective duration incorporates both transition intensities. Its quantitative importance depends on discount-growth gaps and cash-flow responses. The complete intensity derivative is (48); I solve the exact stationary two-regime system in the numerical specialization, including the declared discount-growth-gap adjustment after news.

Suppress the asset superscript and put $a_k = \delta_k - g$ and $\mathcal{D} = a_1 a_2 + a_1 \lambda_{21} + a_2 \lambda_{12}$. If the discount-minus-generator matrix has eigenvalues with positive real parts, solving (43) gives

$$X_1 = \frac{(a_2 + \lambda_{21}) CF_1 + \lambda_{12} CF_2}{\mathcal{D}}, \qquad X_2 = \frac{\lambda_{21} CF_1 + (a_1 + \lambda_{12}) CF_2}{\mathcal{D}}, \tag{45}$$

$$\Delta_{12} = \frac{a_1 CF_2 - a_2 CF_1}{(a_2 + \lambda_{21}) CF_1 + \lambda_{12} CF_2} \tag{46}$$

These are exact stationary-system solutions, not solutions of a general state-dependent

valuation PDE. Write $X_k = N_k/\mathcal{D}$. At a differentiable interior state,

$$\begin{aligned}\psi_k &= \frac{\partial_I N_k}{N_k} - \frac{\partial_I \mathcal{D}}{\mathcal{D}},\\ \psi_{I,k} &= \frac{\partial_{II} N_k}{N_k} - \left(\frac{\partial_I N_k}{N_k}\right)^2 - \frac{\partial_{II}\mathcal{D}}{\mathcal{D}} + \left(\frac{\partial_I \mathcal{D}}{\mathcal{D}}\right)^2\end{aligned} \tag{47}$$

Thus valuation supplies the loadings in (42). Imposing log-linearity independently is an additional restriction. Writing the system as $\mathcal{A}X = CF$ and differentiating proves

$$\partial_\lambda X = \mathcal{A}^{-1}\left(\partial_\lambda CF - (\partial_\lambda \mathcal{A})X\right) \tag{48}$$

Both intensities remain in $\mathcal{A}$ conditional on the current regime. Cash flows and discount rates can also depend on intensity. The complete TPU effect is therefore not solely the derivative of the entropy penalty.

4.2 The household problem

The following proposition characterizes the homothetic benchmark with average labor-income capitalization and a relaxed financial borrowing constraint. Its rules can be evaluated locally where the financial constraints are slack. It is not a global solution of the fixed-human-wealth constrained problem.

Under the maintained assumption that the money market rate is independent of the regime, $r_1^M = r_2^M = r^M$, labor income can be capitalized at a common rate. Defining human wealth and total wealth as $H \equiv \frac{\overline{wl}}{r^M}$ and $\widehat{W} \equiv W + H$, and expressing portfolio shares relative to total wealth as $x^i \equiv \theta^i \frac{W}{\widehat{W}}$, the following proposition characterizes the household problem in terms of $\widehat{W}$.

Proposition 2 (Household optimum)**.** *(i) Total wealth evolves as*

$$d\widehat{W}_t = \left[\widehat{W}_t(\boldsymbol{x}^\top \boldsymbol{\pi}_k + r^M) - \tilde{P}_{C,k} C_t\right]dt + \widehat{W}_t\left(x^S \sigma_k^S dZ_t^S + x^B \sigma_k^B dZ_t^B\right) \tag{49}$$

and jumps to $\widehat{W}(1 + \boldsymbol{x}^\top \boldsymbol{\Delta}_{kl})$ when the regime switches from k to l. The value function is

$$J_k(\widehat{W}) = A_k \frac{\widehat{W}^{1-\gamma}}{1-\gamma} + \Phi_k,$$

where $A_k > 0$ and Φ_k is a wealth-independent constant.

(ii) Optimal consumption is

$$C_k^* = m_k \widehat{W}, \qquad m_k = (A_k \tilde{P}_{C,k})^{-1/\gamma},$$

and the optimal portfolio satisfies the first-order condition (55)*.*

(iii) Up to terms of order $\|\boldsymbol{\Delta}_{kl}\|^2$, and provided $\boldsymbol{\Delta}_{kl}^\top \Sigma_k^{-1} \boldsymbol{\pi}_k < 1$, the optimal portfolio is

$$\boldsymbol{x}_k^* = \frac{1}{\gamma} \Sigma_k^{-1} \boldsymbol{\pi}_k + w_k \Sigma_k^{-1} \boldsymbol{\Delta}_{kl} \tag{50}$$

where

$$w_k \equiv \frac{a_k}{\gamma} \cdot \frac{1 - \boldsymbol{\Delta}_{kl}^\top \Sigma_k^{-1} \boldsymbol{\pi}_k}{1 + a_k \boldsymbol{\Delta}_{kl}^\top \Sigma_k^{-1} \boldsymbol{\Delta}_{kl}} > 0 \tag{51}$$

The benchmark reduction uses average labor-income capitalization and relaxes the financial borrowing boundary. Capitalization alone does not remove that boundary. Setting $r_1^M = r_2^M = r^M$ gives labor income a common capitalization rate, so human wealth $H = \overline{wl}/r^M$ is the same constant in both regimes. A regime dependent money rate would make H jump at every switch and add a second revaluation term to the wealth dynamics that the portfolio would then also have to hedge. Capitalizing labor income at its regime average $\overline{wl}$ folds the nontraded wage stream into a wealth equivalent, so total wealth $\widehat{W} = W + H$ is the single state and the problem becomes homothetic. Without this step, the regime-dependent wage bill would enter as a separate additive state and destroy homotheticity.

The homothetic ansatz applies to the relaxed benchmark with an exogenous average labor-income stream and wealth-independent labor disutility. Goods and labor clearing subsequently map consumption into output, profits and employment. That mapping does not establish the accuracy of average capitalization or remove the financial borrowing boundary. In particular, household labor is demand-determined at the posted wage rather than an additional household choice.

Part (ii) is the homothetic consumption rule, with m_k the marginal propensity to

consume out of total wealth. The portfolio in part (iii) naturally decomposes into the standard myopic mean–variance demand and a hedging demand against future regime switches. The approximation requires $\boldsymbol{\Delta}_{kl}^{\top}\Sigma_k^{-1}\boldsymbol{\pi}_k < 1$ and truncates at $\|\boldsymbol{\Delta}_{kl}\|^2$. Economically, this condition asks that the priced regime jump remain small relative to the myopic position, so the two portfolio components do not interact at first order and the optimal portfolio remains a clean two-fund decomposition. The hedging weight w_k rises with $a_k = \lambda_{kl}A_l/A_k$ because a regime switch that is both more likely and more painful increases the value of insurance, reflecting the two forces identified in Corollary 1.

The value function solves the robust HJB equation

$$\rho J_k = \max_{C,\boldsymbol{x}} \min_{|h|\le\sqrt{2\kappa_k}} \Big\{ \frac{C^{1-\gamma}}{1-\gamma} - \frac{L_k^{1+\chi}}{1+\chi} + J_k'(\widehat{W})\Big[\widehat{W}\big(\boldsymbol{x}^{\top}\boldsymbol{\pi}_k^0 + r^M + \boldsymbol{x}^{\top}\boldsymbol{\Psi}_k h\big) - \tilde{P}_{C,k}C\Big] + \tfrac{1}{2}J_k''(\widehat{W})\widehat{W}^2\boldsymbol{x}^{\top}\Sigma_k\boldsymbol{x} + \lambda_{kl}\Big[J_l\big(\widehat{W}(1+\boldsymbol{x}^{\top}\boldsymbol{\Delta}_{kl})\big) - J_k(\widehat{W})\Big]\Big\} \quad (52)$$

where $\boldsymbol{\pi}_k^0$ collects the reference excess returns $\mu_k^{i,0} - r^M$.

Lemma 3 (Worst-case distortion)**.** *The inner objective in* (52) *is linear in* h*. The worst case is the corner* $h_k^* = -\sqrt{2\kappa_k}\,\mathrm{sign}(\boldsymbol{x}^{\top}\boldsymbol{\Psi}_k)$*. If the solution satisfies* $\boldsymbol{x}^{*\top}\boldsymbol{\Psi}_k > 0$*, then* $h_k^* = -\sqrt{2\kappa_k}$ *and* (52) *reduces to the outer maximization under the robust excess returns*

$$\boldsymbol{\pi}_k = \boldsymbol{\pi}_k^0 - \sqrt{2\kappa_k}\,\boldsymbol{\Psi}_k \quad (53)$$

The lemma maps robustness into a single return adjustment, with each analytical step serving a distinct purpose. Because the inner objective is linear with respect to the distortion h, every admissible distortion is minimizing at zero exposure. At nonzero exposure the minimum is a corner. Under the entropy constraint, the adversary allocates its entire distortion budget $\sqrt{2\kappa_k}$ in a single direction. Consequently, the worst case occurs at the corner $h_k^* = -\sqrt{2\kappa_k}\,\mathrm{sign}(\boldsymbol{x}^{\top}\boldsymbol{\Psi}_k)$.

The premise $\boldsymbol{x}^{*\top}\boldsymbol{\Psi}_k > 0$ resolves the ambiguity of this sign function. This resolution is necessary because the direction of the worst case depends on the specific position the household holds relative to the ambiguous exposure. A household with a net long position in that exposure fears the exact drift that reduces its return, which leads to the selection of the corner $-\sqrt{2\kappa_k}$. Rather than being exogenously imposed, this premise is

verified consistently at the optimum in Proposition 3(ii), which closes the loop between the assumed sign and the realized sign.

The resulting payoff yields the robust excess return $\boldsymbol{\pi}_k = \boldsymbol{\pi}_k^0 - \sqrt{2\kappa_k}\,\boldsymbol{\Psi}_k$. This shows that ambiguity operates identically to a pessimistic haircut applied to expected returns, scaled by its magnitude $\sqrt{2\kappa_k}$ and by the exposure $\boldsymbol{\Psi}_k$ of each asset. This reduction transforms the robust control problem into a standard Merton portfolio problem with adjusted expected returns, paving the way for the closed-form solution in Proposition 2.

4.3 Optimal rules and the value system

Given Proposition 2 and 3, the first order conditions are standard. The consumption rule is

$$C_k^* = m_k\,\widehat{W}, \qquad m_k = \big(A_k \tilde{P}_{C,k}\big)^{-1/\gamma} \tag{54}$$

where the constraint $C \geq 0$ never binds, since marginal utility diverges as $C \to 0$. And the portfolio satisfies

$$\boldsymbol{\pi}_k - \gamma \Sigma_k \boldsymbol{x} + a_k \big(1 + \boldsymbol{x}^\top \boldsymbol{\Delta}_{kl}\big)^{-\gamma} \boldsymbol{\Delta}_{kl} = \mathbf{0}, \qquad a_k \equiv \lambda_{kl} \frac{A_l}{A_k} = \lambda_{kl} \Big(\frac{m_l}{m_k}\Big)^{-\gamma} \frac{\tilde{P}_{C,k}}{\tilde{P}_{C,l}} \tag{55}$$

Substituting the rules into (52) and dividing by $\widehat{W}^{1-\gamma}$ delivers the coupled value system

$$\frac{\rho A_k}{1-\gamma} = \frac{m_k^{1-\gamma}}{1-\gamma} + A_k\big(\boldsymbol{x}_k^{*\top}\boldsymbol{\pi}_k + r^M - \tilde{P}_{C,k} m_k\big) - \frac{\gamma}{2} A_k \boldsymbol{x}_k^{*\top} \Sigma_k \boldsymbol{x}_k^* + \frac{\lambda_{kl}}{1-\gamma}\Big[A_l\big(1 + \boldsymbol{x}_k^{*\top}\boldsymbol{\Delta}_{kl}\big)^{1-\gamma} - A_k\Big] \tag{56}$$

which is equivalent to the consumption equation

$$\tilde{P}_{C,k} m_k = \frac{\rho}{\gamma} + \frac{\gamma - 1}{\gamma}\big(\boldsymbol{x}_k^{*\top}\boldsymbol{\pi}_k + r^M\big) + \frac{1-\gamma}{2}\boldsymbol{x}_k^{*\top}\Sigma_k \boldsymbol{x}_k^* - \frac{\lambda_{kl}}{\gamma}\Big[\frac{A_l}{A_k}\big(1 + \boldsymbol{x}_k^{*\top}\boldsymbol{\Delta}_{kl}\big)^{1-\gamma} - 1\Big] \tag{57}$$

Setting $\lambda_{kl} = 0$ recovers the Merton benchmark in both equations.

Corollary 1 (Value levels across regimes)**.** *Suppose $\gamma > 1$, let $\tilde{P}_{C,2} \geq \tilde{P}_{C,1}$, and the Sharpe set shrinks in the high uncertainty regime, $\boldsymbol{\pi}_2^\top \Sigma_2^{-1} \boldsymbol{\pi}_2 < \boldsymbol{\pi}_1^\top \Sigma_1^{-1} \boldsymbol{\pi}_1$, a maintained analytical*

restriction rather than a verified property of the calibration. Then, on the local benchmark branch for sufficiently small intensities, $m_2 < m_1$ and a lower wealth-dependent value component when $A_2 > A_1$, and the hedging weight satisfies $a_1 = \lambda_{12}A_2/A_1 > \lambda_{12}$.

The two maintained conditions pinpoint the exact mechanism that causes the high-uncertainty regime to be more severe. The contraction of the Sharpe set, $\boldsymbol{\pi}_2^\top \Sigma_2^{-1} \boldsymbol{\pi}_2 < \boldsymbol{\pi}_1^\top \Sigma_1^{-1} \boldsymbol{\pi}_1$, indicates that the ambiguity discount described in Lemma 3 worsens the risk-return trade-off in regime 2. This deterioration diminishes lifetime utility and, under the condition that $\gamma > 1$, drives up the marginal value of wealth. Meanwhile, the price condition $\tilde{P}_{C,2} \geq \tilde{P}_{C,1}$ ensures that consumption remains weakly more expensive in the adverse regime. As a result, the ranking $m_2 < m_1$ is determined by the narrowing opportunity set rather than being overturned by a cheaper consumption basket. Without the first condition, the welfare ranking could invert. Without the second, the marginal utility ranking $A_2 > A_1$ no longer follows.

The corollary ranks the regimes in welfare and in marginal utility at the same time. The wealth-dependent value component is lower in the high uncertainty regime when $A_2 > A_1$ and $\gamma > 1$. Ranking total lifetime value also requires the wealth-independent labor-disutility terms Φ_k. The hedging weight $a_1 = \lambda_{12}A_2/A_1$ combines the two forces the household cares about. The intensity measures how likely the bad regime is. The value ratio measures how much wealth will matter once it arrives. Escalatory news raises the first factor and the welfare gap raises the second, so insurance demand strengthens on both margins.

Proposition 3 (Admissibility and the conditional value system)**.** *For the homothetic benchmark, the interior equations are evaluated on $1 + \boldsymbol{x}_k^\top \boldsymbol{\Delta}_{kl} > 0$.*

(i) This is the positive-total-wealth domain of the benchmark. The original financial-wealth constraints remain those of the full household problem.

(ii) The positive-exposure distortion is consistent precisely when $\boldsymbol{x}_k^\top \boldsymbol{\Psi}_k > 0$. For the exact interior root, writing $\zeta_k = a_k(1 + \boldsymbol{x}_k^\top \boldsymbol{\Delta}_{kl})^{-\gamma}$, this is equivalent to

$$\boldsymbol{\Psi}_k^\top \Sigma_k^{-1}(\boldsymbol{\pi}_k + \zeta_k \boldsymbol{\Delta}_{kl}) > 0.$$

(iii) At zero intensity, $d_k = \rho - (1-\gamma)(r^M + CE_k) > 0$ gives the explicit positive bench-

mark solution in (88). *At a strict interior, fixed-sign, jump-solvent solution, the implicit-function argument in Appendix B gives a locally unique continuation for sufficiently small intensities. Verification requires the stated transversality condition.*

(iv) Under the finite-variation belief specification, a changing reference drift enters the local opportunity set through

$$\boldsymbol{\pi}_k(t) = \boldsymbol{\pi}_k^0(0) + \boldsymbol{\psi}_k b_k(t) - \sqrt{2\kappa_k(t)}\boldsymbol{\Psi}_k \tag{58}$$

The value equation retains its time or belief-state derivative. Absence of a belief diffusion eliminates its diffusion covariation term, but does not eliminate its continuation-value term.

The exact exposure test follows by multiplying the interior portfolio equation by $\boldsymbol{\Psi}_k^\top \Sigma_k^{-1}$. The sign test written with the approximate weight w_k applies to that approximation only. The benchmark continuation is local in intensities. A numerical solution at other intensities is assessed by its HJB residual, feasibility and terminal condition. Freezing time-varying coefficients is an additional numerical approximation, for which no error order is asserted here.

4.4 Comparative statics

The analytical expressions describe responses under their maintained sign, interiority and local-approximation conditions. These conditions are not automatically satisfied by the descriptive calibration. Proposition 4 concerns the portfolio response to a realized regime change and to movements in the transition intensities, including the inventory channel of Proposition 1. Proposition 5 isolates the hedging demand by comparison with the no hedging benchmark. Corollary 2 translates the belief channel into consumption, output, and employment through the market clearing of Section 3.6. Suppose the conditions of Propositions 2 and 3 hold, and let $\tilde{D}_\nu^i = 1/(\delta_k^i - g^i + \vartheta)$.

Proposition 4 (Portfolio response to news and regime change)**.** *Within the interior approximation, and under the negative-entry-jump and adverse-regime ranking assumptions, the following expressions apply. The reported calibration is not assumed to satisfy these additional sign restrictions.*

(i) At entry into the high uncertainty regime the portfolio moves by

$$\boldsymbol{x}_2^* - \boldsymbol{x}_1^* = \frac{1}{\gamma}\big(\Sigma_2^{-1}\boldsymbol{\pi}_2 - \Sigma_1^{-1}\boldsymbol{\pi}_1\big) + w_2\Sigma_2^{-1}\boldsymbol{\Delta}_{21} - w_1\Sigma_1^{-1}\boldsymbol{\Delta}_{12} \tag{59}$$

total wealth is revalued by the factor $1 + \boldsymbol{x}_1^{*\top}\boldsymbol{\Delta}_{12} < 1$*, and the price change at the switch equals* Δ_{kl}^i *of Lemma 2. Holding other valuation components fixed, the direct ambiguity-premium contribution to the within-regime log-price change is* $-\tilde{D}_\nu^i\,\Psi_k^i\,d\sqrt{2\kappa_k}$*. This is not the total derivative with respect to* λ_{12}*, which also changes transition probabilities and conditional continuation values.*

(ii) Within regime 1,

$$\frac{d\boldsymbol{x}_1^*}{d\lambda_{12}} = -\frac{1}{\gamma}\Sigma_1^{-1}\boldsymbol{\Psi}_1\frac{\partial\sqrt{2\kappa_1}}{\partial\lambda_{12}} + \frac{\partial w_1}{\partial\lambda_{12}}\Sigma_1^{-1}\boldsymbol{\Delta}_{12} + w_1\Sigma_1^{-1}\frac{\partial\boldsymbol{\Delta}_{12}}{\partial\lambda_{12}} + \frac{1}{\gamma}\Sigma_1^{-1}\boldsymbol{\psi}\,\frac{\partial b_1}{\partial\lambda_{12}} \tag{60}$$

where the displayed decomposition holds with covariance and loadings fixed and with the stated reference-drift mapping. Although $\partial\sqrt{2\kappa_1}/\partial\lambda_{12} > 0$*, the total derivative of* w_1 *must include changes in the value ratio and the priced jump. Its sign is determined by* (100)*.*

(iii) Within regime 2, holding A_1/A_2 *fixed gives* $\partial a_2/\partial\lambda_{21} = A_1/A_2 > 0$*, and* $\boldsymbol{\Delta}_{21} = -\boldsymbol{\Delta}_{12}/(1+\boldsymbol{\Delta}_{12}) > \mathbf{0}$ *entrywise. The hedge moves the portfolio along* $\Sigma_2^{-1}\boldsymbol{\Delta}_{21}$*, where the equity entry has the sign of* $(\sigma_k^B)^2\Delta_{21}^S - \sigma_k^{SB}\Delta_{21}^B$ *and the bond entry carries the sign of* $\sigma_S^2\Delta_{21}^B - \sigma^{SB}\Delta_{21}^S$*.*

Risk aversion $\gamma > 1$ alone does not determine every response sign. The ranking and sign statements additionally require the opportunity-set and revaluation restrictions of the cited results. They should not be transferred to the measured calibration when those restrictions fail.

Part (i) decomposes the portfolio jump at a regime switch into three economically distinct shifts. First, the myopic term $\frac{1}{\gamma}(\Sigma_2^{-1}\boldsymbol{\pi}_2 - \Sigma_1^{-1}\boldsymbol{\pi}_1)$ reprices the investment opportunity set. Second, the term $w_2\Sigma_2^{-1}\boldsymbol{\Delta}_{21}$ establishes the new regime-2 hedge. Third, the term $-w_1\Sigma_1^{-1}\boldsymbol{\Delta}_{12}$ realizes and unwinds the previous regime-1 hedge. The wealth scaling factor $1 + \boldsymbol{x}_1^{*\top}\boldsymbol{\Delta}_{12} < 1$ is strictly less than one because the household holds a long position and the negative jump $\boldsymbol{\Delta}_{12} < \mathbf{0}$ destroys wealth upon entry. Additionally, the within-regime log price movement depends on $\tilde{D}_\nu^i$, which represents the effective duration of a premium

movement decaying at rate ϑ.

Part (ii) isolates the within-regime response to escalation news across four distinct channels, and evaluating all four channels simultaneously is essential because they can drive the portfolio in opposing directions. The ambiguity channel is negative because a higher transition probability λ_{12} widens the uncertainty bound $\sqrt{2\kappa_1}$ and shrinks myopic asset demand. The hedging-weight channel depends jointly on intensity, the destination-to-current value ratio, and repricing of the jump. The third term reflects the revaluation update channel, which operates because the magnitude of the priced jump $\mathbf{\Delta}_{12}$ itself varies with the ambiguity parameter κ through Lemma 2. The fourth term is the optimal-inventory channel. Its derivative must be obtained from the current-profit firm problem rather than assigned a sign from a fixed-input barrier.

Part (iii) establishes the sign of the regime 2 hedge. The negative-entry-jump premise implies an entrywise positive reversion jump. The direct intensity effect on a_2 is positive holding the value ratio fixed. Its total derivative also includes the value-ratio response. Consequently, the hedge aligns with $\Sigma_2^{-1}\mathbf{\Delta}_{21}$. Under that sign configuration, a negative bond entry reduces bond demand. The cash response depends on the sum of the risky-share responses. I compute reversion jumps from the same conditional price map rather than impose this theoretical sign configuration on every numerical state. Together, these components demonstrate that escalation news determines the magnitude of the impact through the ambiguity and hedging channels, while both inventory beliefs and changing intensities can affect transitional dynamics.

The proposition explains the shape of the equity price response and signs the reallocation chain. On impact, the price drops by the news component since a higher escalation intensity widens the ambiguity premium immediately. Over the following months, realized escalations add the revaluation Δ_{12}. The recovery then follows as the intensity path decays and escalated paths revert through λ_{21}. The reallocation operates through the hedging weight w_1, whose strength combines the probability of the loss with the marginal value of wealth in the feared state. It also operates through the episodic inventory term. In announcement episodes, the front loading of Proposition 1 raises the believed accumulation. This front loading partially offsets the ambiguity channel on impact and accelerates the recovery. This is the model counterpart of the hump shaped inventory dynamics

around announced tariff changes.

Let $\boldsymbol{x}_k^{\mathrm{nh}} = \frac{1}{\gamma}\Sigma_k^{-1}\boldsymbol{\pi}_k$ denote the myopic portfolio allocation obtained by deactivating the hedging demand with $w_k = 0$.

Proposition 5 (Hedging demand with and without)**.** *The following hold.*

(i) The hedging demand is given by

$$\boldsymbol{x}_k^* - \boldsymbol{x}_k^{\mathrm{nh}} = w_k \Sigma_k^{-1} \boldsymbol{\Delta}_{kl} \tag{61}$$

(ii) The equity entry of $\Sigma_k^{-1}\boldsymbol{\Delta}_{12}$ is negative when $(\sigma_k^B)^2|\Delta_{12}^S| > \sigma_k^{SB}|\Delta_{12}^B|$, and the bond entry is positive when $\sigma_k^{SB}|\Delta_{12}^S| > (\sigma_k^S)^2|\Delta_{12}^B|$, provided the entry jumps are negative. These inequalities are not asserted to hold in the current calibration.

(iii) Up to terms of order $\|\boldsymbol{\Delta}_{kl}\|$, the inventory channel moves the portfolio by

$$\frac{\partial \boldsymbol{x}_k^*}{\partial b_k} = \frac{1}{\gamma}\Sigma_k^{-1}\boldsymbol{\psi} \tag{62}$$

(iv) Parts (i) and (iii) describe fixed-input interior comparisons. Finite constrained counterfactuals with recomputed consumption need not equal these additive expressions.

The covariance conditions delineate when a single source of distrust divides distinctly across the two risky assets. Because equity experiences the larger revaluation, $|\Delta_{12}^S| > |\Delta_{12}^B|$, and the stock-bond covariance σ_k^{SB} is positive but bounded, the inverse-covariance operator delivers a negative equity entry and a positive bond entry. Consequently, the household reduces its exposure to the asset it distrusts most and acquires the one that co-moves least with it. These inequalities identify the conditions for that sign pattern. Observed total covariance does not by itself establish these inequalities. I evaluate numerical allocations using the complete conditional return law and retain this argument as an interior benchmark.

Monthly disclosure removes the current inventory information gap, while uncertainty about future inventory and regime transitions remains relevant for portfolio choice. Between reports, return signals update the inventory posterior. The two interior comparative-static objects, $\Sigma_k^{-1}\boldsymbol{\Delta}_{kl}$ and $\Sigma_k^{-1}\boldsymbol{\psi}$, distinguish regime revaluation from the inventory-drift channel. Their interpretation does not require disclosure to reset physical stock.

Imposing $w_k = 0$ characterizes a myopic household that optimizes its portfolio under the assumption that the current regime will persist indefinitely. Such a household holds the standard mean-variance portfolio $\frac{1}{\gamma}\Sigma_k^{-1}\boldsymbol{\pi}_k$ and forgoes any insurance against the discrete wealth revaluations triggered by a shift in the tariff regime. Consequently, the wedge $\boldsymbol{x}_k^* - \boldsymbol{x}_k^{\mathrm{nh}} = w_k\Sigma_k^{-1}\boldsymbol{\Delta}_{kl}$ captures the pure hedging demand directed against observed-regime transition risk, within the analytical interior approximation. I do not interpret the fitted cash response as a separately measured value of this wedge.

I implement the quantitative exercise as a conditional household problem under the declared production and valuation inputs. The analytical hedge-off expression does not define a global constrained counterfactual when consumption and continuation values are recomputed. The code instead provides a comparison using a common wealth-dependent continuation value with its information-state variation removed. That diagnostic is not an additive causal decomposition, and it was not completed in the saved run reported below. Equilibrium price feedback remains outside the reported calibration.

Corollary 2 (Consumption, output, and employment)**.** *Suppose $\gamma > 1$ and impose the market clearing conditions of Section 3.6.*

(i) The expenditure ratio satisfies, as a direct effect with the indirect effect through A_l/A_k of order λ_{kl},

$$\frac{\partial\big(\tilde{P}_{C,k}m_k\big)}{\partial\sqrt{2\kappa_k}} = -\frac{\gamma-1}{\gamma}\,\boldsymbol{x}_k^{*\top}\boldsymbol{\Psi}_k < 0 \tag{63}$$

and $m_2 < m_1$ by Corollary 1.

(ii) Given the consumption rule of Proposition 2(ii), equilibrium output, the wage bill, and profit in regime k are

$$\begin{aligned} Y_k &= \frac{P_{C,k}\,m_k\,\widehat{W} + NX_k}{P_{Y,k}}, \qquad wL_k = \alpha\,\frac{(\varepsilon_p-1)(1-\tau_s)}{\varepsilon_p}\,\big(P_{C,k}\,m_k\,\widehat{W} + NX_k\big), \\ \Pi_k &= \frac{1-\tau_s}{\varepsilon_p}\big(P_{C,k}m_k\widehat{W} + NX_k\big) \quad + c_{F,k}X_{F,k} - P_{M,k}(X_{F,k}^N + Q_{X,k}) \end{aligned} \tag{64}$$

with $NX_1 = 0$ in the symmetric baseline. The other regime's trade balance is retained rather than assigned an order without verification.

Consumption falls when the ambiguity bound widens because a larger κ_k compresses

the compensation the household expects from its risky positions. For $\gamma > 1$, the resulting income effect dominates, reducing the expenditure ratio. Because nominal wages are sticky, output is demand-determined, so this decline in expenditure translates one-for-one into lower production and employment. At given production prices, output and the wage bill follow expenditure through clearing. Distributed profit additionally reflects the endogenous sourcing and investment terms in (64). Its sign response cannot be inferred from expenditure proportionality alone.

The proportionality of the wage bill to expenditure also disciplines the human wealth approximation introduced in Proposition 2(i). The wage bill differs across regimes only through $P_{C,k}m_k$ and the wealth revaluation, whose quantitative size depends on the maintained inputs. Capitalizing labor income at its regime average is a benchmark approximation whose accuracy must be checked against the joint solution. The numerical household problem retains the conditional firm wage bill in financial wealth and does not impose a constant total-to-financial-wealth multiplier. This calculation does not establish joint aggregate consistency.

In the joint model, goods demand and trade affect household income and the firm's distributed profit. Holding those inputs fixed defines the conditional portfolio benchmark, not an exact absence of feedback. The size of approximation errors in these objects has not been bounded by the present numerical exercise. The quantitative insignificance of the omitted feedback is therefore not established.

4.5 Numerical implementation of the conditional solution

I retain the preceding expressions as analytical benchmarks under their stated assumptions. For the numerical calculation, I solve directly for consumption and financial shares on financial wealth and the joint information state. This preserves the nonnegative financial-wealth boundary and does not extend the interior homothetic rule mechanically to constrained states. I recover cash as $\theta^M = 1 - \theta^S - \theta^B$ and evaluate adverse drift distortion at the portfolio actually selected.

I use the same conditional price function to construct payout-inclusive returns, covariance exposures and regime revaluations. Explicit calendar dependence contributes its own price drift after news. A change in inventory dispersion therefore changes covariance,

the Itô return term and ambiguity exposure together. The monthly report collapses the posterior information gap, while the physical inventory policy continues to determine receipts, withdrawals and depreciation. Excess inventory remains the derived coordinate $E = I - K_0$.

The numerical solution separates the additive labor-disutility value from the financial optimization, solves the stationary continuation problem and then works backward along the anticipated news path. I propagate the resulting distribution forward under the reference transition law. The completed fitting stage establishes the reported finite-grid allocation results and their recorded numerical residuals. It does not replace the conditional assumptions of the analytical propositions or establish a completed continuous-time convergence result.

5 Calibration and Simulation

The calibration connects the empirical portfolio responses to the inventory-information mechanism in Sections 3 and 4. I solve the constrained household problem using a common conditional valuation map for payouts, return drift, covariance and regime revaluation. The exercise conditions on the production and fiscal reference accounts. Its quantitative objective is to reproduce baseline risky-asset shares and the initial reallocation after TPU news, while checking that the fitted allocation survives changes in numerical resolution.

5.1 Structural parameters and empirical targets

Table 3 reports the maintained structural choices used by the numerical method. The cited literature motivates the economic mechanisms; each numerical choice is distinguished from an estimate taken directly from a source. The cash rate is measured from the main sample and is reported separately in Table 4. Annual discount, interest and inventory-depreciation rates are divided by twelve to express the numerical problem in monthly units.

Table 3: Preset Parameters

Parameter	Meaning	Value	Reference
γ	risk aversion	5	Campbell and Viceira (2002)
ρ	discount rate	0.04/yr	—
χ	labor disutility curvature	1	Chetty et al. (2011)
α	labor share	0.62	Gollin (2002)
η	input substitution elasticity	2	Boehm et al. (2023)
ε_p	variety elasticity	6	Christiano et al. (2005)
$s,\ s_c$	domestic CES weights	0.85	Amiti et al. (2019)
δ_I	inventory depreciation	0.25/yr	Alessandria et al. (2011)
τ_c	consumption tax	0.065	—
τ_s	sales tax	0.05	—
$\tau_1,\ \tau_2$	tariffs by regime	0.03, 0.12	Fajgelbaum et al. (2020)

Note: References provide calibration guidance; selected values may differ from the numerical estimates reported in the cited studies.

The values in this table combine literature evidence with choices made for the present numerical exercise. Campbell and Viceira (2002) provide the dynamic portfolio-choice framework for risk-averse investors. I select $\gamma = 5$ within the $\gamma > 1$ region as a numerical benchmark for substantial risk aversion rather than an estimated coefficient from that source. Chetty et al. (2011) explain that micro estimates of intertemporal substitution are an order of magnitude smaller than the elasticities needed to explain aggregate-hours fluctuations through preferences. In particular, quasi-experimental extensive-margin Frisch elasticities are around 0.25, whereas leading pure-equilibrium macro models imply values around 2. I choose $\chi = 1$ to provide an intermediate unit-elasticity benchmark in the separable labor-choice problem instead of adopting either endpoint. This choice is a modeling compromise rather than an estimate of the extensive-margin elasticity. In our actual numerical method, labor is demand-determined and χ enters the separate additive labor-disutility value. Gollin (2002) report corrected labor shares mostly between 0.65 and 0.80. More recent U.S. labor-compensation shares have moved toward the mid-0.50s, with FRED series LABSHPUSA156NRUG reporting 0.56830 in 2023. I therefore adjust the historical benchmark downward and set $\alpha = 0.62$ to accommodate the more recent

decline while retaining a labor-majority reference production account. The value 0.62 is our calibration adjustment rather than a direct estimate from Gollin or a sample mean of the FRED series. Boehm et al. (2023) estimate a long-run tariff-exclusive trade elasticity of approximately -2. I use its absolute magnitude to choose the CES input elasticity $\eta = 2$, while keeping the distinction between aggregate trade and input substitution. Christiano et al. (2005) report a benchmark markup of 1.20, which gives $\varepsilon_p = 6$ through $\varepsilon_p/(\varepsilon_p - 1) = 1.20$. Amiti et al. (2019) discuss a representative firm's imported inputs accounting for 15 percent of variable costs. I use this exposure to motivate a foreign CES weight of 0.15 and choose the same domestic weight $s = s_c = 0.85$ in production and consumption for parsimony. These weights do not impose realized import expenditure shares. Alessandria et al. (2011) report that inventory adjustments account for 19 to 46 percent of the import wedge over the longer sample beginning in the second quarter of 1972. This range motivates an economically significant inventory-buffer mechanism. I choose annual depreciation $\delta_I = 0.25$ to preserve the buffer while allowing for positive physical degradation of holdings. The reported contribution to the import wedge is not a depreciation-rate estimate, so the annual rate of 0.25 remains our numerical choice. Fajgelbaum et al. (2020) document import tariffs rising from 2.6 to 16.6 percent. I choose regime tariffs of 3 and 12 percent within this observed interval to represent distinct low- and high-tariff environments. Annual discounting $\rho = 0.04$ is our positive time-preference benchmark, and $\tau_c = 0.065$ and $\tau_s = 0.05$ are chosen fiscal wedges that generate consumption- and sales-tax receipts, rather than externally estimated tax averages. The cash rate is measured separately as the main-sample annual mean of 2.071407 percent in Table 4. Annual discount, interest and depreciation rates are divided by twelve in the monthly calculation.

The measurement sample runs from January 2015 through January 2026 and preserves the monthly calendar and the empirical exclusions in Section 2. Regime 2 consists of observations above the sample's 80th percentile of TPU, corresponding to a raw-index threshold of 225.660; the remaining observed months belong to regime 1. I estimate switching probabilities from adjacent observed calendar months and embed the resulting two-state transition matrix in continuous time. If p_{12} and p_{21} are monthly switching

probabilities, the monthly intensities satisfy

$$\begin{pmatrix} \bar{\lambda}_{12} \\ \bar{\lambda}_{21} \end{pmatrix} = \frac{-\log(1 - p_{12} - p_{21})}{p_{12} + p_{21}} \begin{pmatrix} p_{12} \\ p_{21} \end{pmatrix} = \begin{pmatrix} 0.069881 \\ 0.285115 \end{pmatrix} \tag{65}$$

Missing observations do not become consecutive months when either returns or regime transitions are formed.

Within each regime, I target the mean stock and bond shares measured on the same complete-case survey sample. Cash is the accounting residual $1 - \theta^S - \theta^B$; the separately reported survey cash mean remains a measurement diagnostic. Dynamic targets are the stock and bond local-projection responses at horizons zero and one. Their derived cash counterpart is the negative sum of those two responses. Thus the calibration cash target differs slightly from the separately estimated cash coefficient reported in Section 2.

The baseline risk targets are the two standard deviations and the covariance of monthly stock and bond price-return proxies in regime 1. Stock returns use changes in logged NASDAQ prices, while bond returns use an eight-year-duration transformation of Treasury yield changes. The low-regime standard deviations are 0.037354 and 0.015739, and the covariance is 0.000124948. These proxies discipline perceived baseline risk. High-regime risk moments and responses after month one provide comparisons rather than additional independently fitted targets.

I collect eleven moments: four risky-asset allocation levels, four stock and bond responses at horizons zero and one, and three baseline risk moments. Candidate selection minimizes their complete scaled loss,

$$\mathcal{L}(\vartheta) = \sum_{j=1}^{11} \left[\frac{m_j(\vartheta) - \widehat{m}_j}{d_j} \right]^2 \tag{66}$$

where $m_j(\vartheta)$ is computed from the household solution. Allocation-level scales use survey standard errors with a floor of 0.002 in share units; dynamic scales use the empirical response standard errors with a floor of 0.001. The stock and bond standard-deviation scales are 0.005 and 0.003, and the covariance scale uses its sample uncertainty. Because month-one responses enter this loss, their fit is part of calibration.

5.2 Valuation, covariance and inventory information

I normalize initial physical inventory, its operational benchmark and initial financial wealth to $I(0) = K_0 = W(0) = 100$. Excess inventory is the derived coordinate $E = I - K_0$. Inventory withdrawals and depreciation apply to I, and the monthly report reveals physical stock without resetting it. The maintained inventory barrier determines the reference sourcing policy, with immediate replenishment represented by catch-up orders over a numerical time step.

The CES production block retains the structural parameters in Table 3. A low-regime purchase book price of 1.03 and direct foreign-input use of two units per month determine the reference scale. The financial-wealth-to-annual-labor-income ratio is 15. I implement this endowment choice through household mass, dividing firm labor and the wage bill by the same mass. The government's reference primary receipts anchor bond payouts. These production and fiscal accounts supply conditional income and payouts to the portfolio problem.

Reference stock and bond values are log-linear in physical inventory, with positive inventory derivatives and regime-specific valuation losses. Their payout-price ratios and regime losses are fitted jointly with the household allocation. Following news, I solve the two-regime Gordon system in (45). The same price map determines payout-inclusive gains, inventory drift and Itô terms, explicit calendar price changes and the instantaneous revaluation at a regime switch. Expected returns and jump sizes are therefore recomputed from conditional valuation at every candidate.

To connect inventory exposure to observed baseline risk, I assign a common inventory factor the fraction $f_I = 0.55$ of each asset's low-regime variance. In diffusion-exposure units,

$$\Psi_1^S = \sigma_1^S \sqrt{f_I}, \qquad \Psi_1^B = \sigma_1^B \sqrt{f_I}, \qquad R = \widehat{\Sigma}_1 - \mathbf{\Psi}_1 \mathbf{\Psi}_1^\top, \tag{67}$$

$$\Sigma_k = \mathbf{\Psi}_k \mathbf{\Psi}_k^\top + R \tag{68}$$

The residual covariance R must be positive definite. Its implied correlation is approximately -0.749951, while the complete low-regime covariance retains the measured correlation of approximately 0.212522. The factor fraction is a declared assumption: observed

total covariance does not separately identify inventory risk and residual risk.

Households observe the regime and asset-return signals between monthly inventory reports. I update the inventory posterior by Bayes' rule under the reference law and construct a separate adaptive posterior grid for each phase of the reporting month. The grid covers attainable posterior candidates at a common L^1 tolerance of 0.1 during fitting and refinement. Immediately after disclosure, only point-mass beliefs are admissible. The ambiguity distortion shifts risky return drift through inventory exposure; the physical inventory law and the information-state transition retain their reference dynamics.

The news experiment increases escalation intensity and then lets the increment decay:

$$\begin{aligned} \lambda_{12}(t) &= \bar{\lambda}_{12} + \varepsilon_\lambda (0.85)^t, & \lambda_{21}(t) &= \bar{\lambda}_{21}, \\ \bar{\nu}_k(t) &= \bar{\nu}_k(0^-) \exp\{b_\nu[\lambda_{12}(t) - \bar{\lambda}_{12}]\} \end{aligned} \tag{69}$$

Time is measured in months and impact occurs at $t = 0$. The baseline inventory-dispersion scale is 20, with a regime-2 multiplier of 1.5. The entropy bounds use the current escalation intensity,

$$\kappa_1(t) = 0.01\lambda_{12}(t)/0.30, \qquad \kappa_2(t) = 0.02 + 0.01\lambda_{12}(t)/0.30 \tag{70}$$

The intensity scale 0.30, the dispersion response and the entropy coefficients are maintained numerical choices. I calibrate ε_λ and b_ν together; the empirical one-standard-deviation TPU innovation sets the response targets rather than supplying a direct numerical observation of either parameter.

5.3 Household solution and parameter fitting

The household state consists of financial wealth, the public regime, the phase within the reporting month and the inventory posterior. I separate additive labor disutility from the financial value calculation and optimize consumption, stock and bond holdings jointly, subject to nonnegative holdings and the financial-wealth bounds. Cash follows from portfolio adding-up. Conditional probabilities and moments across discrete return states are evaluated deterministically, including correlated residual gains.

Table 4: Numerical inputs for the saved provisional candidate

Object	Value	Basis
Annual cash rate r^M	0.0207141	Main-sample mean
$I(0), K_0, W(0)$	100	Unit normalization
Financial wealth / annual labor income	15	Endowment assumption
Baseline inventory dispersion	20	Maintained scale
Regime-2 / regime-1 dispersion	1.5	Maintained ratio
Inventory fraction of baseline variance	0.55	Covariance decomposition
Residual stock–bond correlation	-0.749951	Covariance decomposition
Stock payout / price	0.01416377	Fitted monthly ratio
Bond payout / price	0.00229207	Fitted monthly ratio
Stock regime log-price loss	0.0442723	Fitted
Bond regime log-price loss	0.0203501	Fitted
News increment ε_λ	4.941761×10^{-6}	Fitted per month
Dispersion coefficient b_ν	2139.7669	Fitted
Monthly news persistence	0.85	Maintained

Payout-price ratios and regime losses are fitted valuation inputs.

The stationary baseline supplies the continuation value for a twelve-month news experiment. I solve the known declining news path backward, then propagate the joint distribution forward under the reference law. At impact, the unexpected price change first revalues the household's existing holdings; households subsequently choose their new allocation. The initial distribution uses stationary regime probabilities at the common financial wealth of 100. Reported portfolio responses are differences between the shocked and baseline expected shares, in percentage points.

Parameter fitting has two nested steps. First, I update the stock and bond payout-price ratios and regime losses using their discrete Euler conditions at the four level targets. Each update is followed by a fresh household value solve, and the computed allocation determines the level-fit error. Second, for each trial dispersion-response coefficient, I bracket an intensity increment that matches the equity impact. I then refine the dispersion coefficient within a verified bracket of bond responses, resolving the full news path at every trial. The selected candidate has the smallest eleven-moment loss in (66) among the evaluated candidates.

The initial numerical resolution uses a half-month step, physical-stock nodes 0, 100 and 200, three return-signal bins and fifteen geometrically spaced wealth nodes between 0.01 and 1000, with an additional node at 100. Portfolio grid points initialize the continuous

constrained optimizer. Consumption remains jointly feasible during allocation updates, and the value-function interpolation uses both one-sided marginal values at a kink. A solved baseline can initialize a larger grid, but its policies are optimized again.

5.4 Refinement and acceptance criteria

The current refinement procedure changes one numerical dimension at a time: return-signal bins, the time step, physical-stock resolution, ambiguity-distortion resolution, and wealth resolution. The twelve-month horizon and the wealth support remain fixed. Each stage rebuilds its phase-specific adaptive posterior grid using the same covering convention as fitting and compares both the eleven calibration moments and the complete stock, bond and cash paths at common monthly reporting dates. Monthly responses are sampled directly at disclosure dates.

A fixed-parameter refinement passes when the maximum moment change is below 0.0025 in the stored moment units, the largest monthly allocation-path change is below 0.25 percentage points, and all continuous-control solves succeed. Selection additionally requires every fitted moment to remain within two calibration scales of its target. If a solved grid loses data fit, I recalibrate on that grid and retain the original failed fixed-parameter comparison. Recalibration changes parameters and is recorded separately from numerical convergence.

Final acceptance combines all five refinement comparisons with the fit, valuation-equation, covariance, posterior-projection, value-shape, budget and monthly-disclosure checks. A low Bellman residual by itself establishes accuracy only for the specified finite-grid problem. The acceptance rule also requires stable re-solves and finite allocation paths. Tables and figures below report a saved provisional candidate while the current refinement sequence is still in progress.

5.5 Provisional allocation results

The model output reflects the optimal candidate under the quarter-month resolution. Its calibrated parameters, targeted moment vector, numerical audit, and simulated allocation paths are cataloged in the accompanying replication package. This candidate follows

recalibration after changes in signal and time resolution; it has not completed the five-stage acceptance procedure. Its scaled loss is 0.4104, and the largest level error across the four independent risky-share targets is 0.0859 percentage points.

Table 5: Provisional allocation levels in percent

Asset	Target	Candidate	Gap (pp)
Stock, regime 1	67.1476	67.0799	-0.0677
Bond, regime 1	15.8126	15.8751	+0.0624
Cash, regime 1	17.0398	17.0450	+0.0052
Stock, regime 2	66.5038	66.4439	-0.0599
Bond, regime 2	16.2692	16.3551	+0.0859
Cash, regime 2	17.2269	17.2009	-0.0260

Both cash columns are residuals. Regime-specific level targets differ from the stationary-weighted initial allocation used in the news experiment.

On impact, the candidate reduces equity by 0.5078 percentage points, increases bonds by 0.1824 and increases cash by 0.3254. Table 6 compares these responses with the empirical calibration targets. Cash receives the larger positive reallocation. This ranking arises from the jointly fitted stock and bond response rather than an independently targeted cash equation. Month-one outcomes also affect candidate selection; later horizons provide comparisons of the dynamic shape.

Table 6: Provisional allocation responses in percentage points

Asset / month	Empirical benchmark	Candidate	Gap
Stock, $h = 0$	-0.5082	-0.5078	+0.0004
Bond, $h = 0$	+0.1826	+0.1824	-0.0002
Cash, $h = 0$	+0.3256	+0.3254	-0.0002
Stock, $h = 1$	-0.4568	-0.4239	+0.0329
Bond, $h = 1$	+0.1242	+0.1543	+0.0302
Cash, $h = 1$	+0.3327	+0.2696	-0.0631
Stock, $h = 2$	-0.2789	-0.3624	-0.0835
Bond, $h = 2$	+0.1663	+0.1270	-0.0393
Cash, $h = 2$	+0.1125	+0.2354	+0.1228
Stock, $h = 3$	-0.2852	-0.3137	-0.0286
Bond, $h = 3$	+0.1882	+0.1022	-0.0860
Cash, $h = 3$	+0.0970	+0.2115	+0.1145

Stock and bond horizons zero and one enter candidate selection. Cash benchmarks are derived residuals. Later horizons compare dynamic shape.

The candidate's equity-price impact is -0.001647 in log-percentage terms, compared with the empirical NASDAQ impact of approximately -0.994. Matching allocation re-

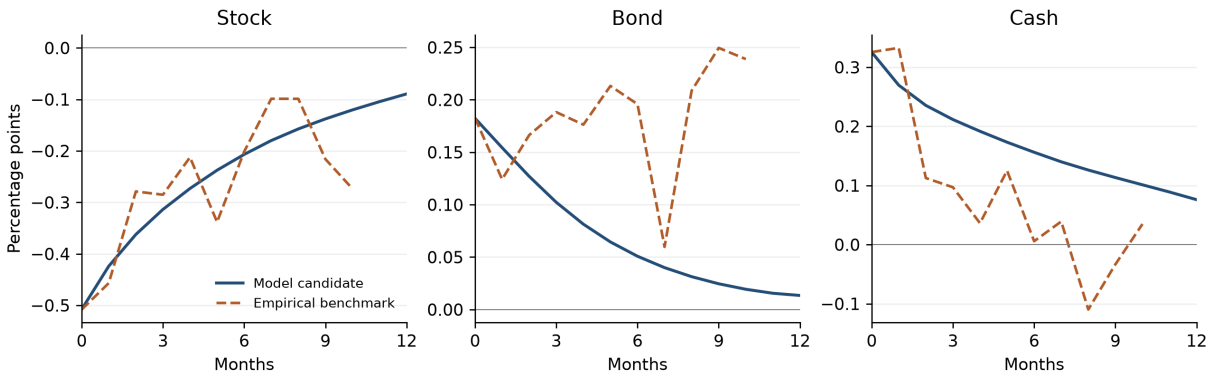


Figure 5: Provisional portfolio responses. Solid lines show model expectations and dashed lines show empirical benchmarks. Model paths are sampled at monthly reporting dates. The cash benchmark is the negative sum of stock and bond responses.

Table 7: Provisional numerical checks and refinement status

Check	Result
Financial Bellman residual	2.87×10^{-9}
Continuous-control failures	0
Maximum share adding-up error	1.11×10^{-16}
Maximum response adding-up error (pp)	6.33×10^{-11}
Repeated-path difference (pp)	5.72×10^{-5}
Posterior variance at disclosure	0.0
Signal fixed-parameter refinement	Failed; recalibrated
Time fixed-parameter refinement	Failed; recalibrated
Five-stage refinement acceptance	Pending
Verified hedging decomposition	Pending

Candidate residuals refer to its finite-grid solution. Recalibration restores fit after a changed grid; it does not reverse a failed fixed-parameter convergence comparison.

sponses therefore leaves a substantial price-response discrepancy. Consumption is stored as a composite-consumption level and requires a measurement mapping before comparison with empirical real-consumption growth. Output and profit come from the conditional reference accounts.

While saved signal and temporal refinements fail initial fixed-parameter comparisons, subsequent recalibration restores empirical fit across those numerical grids. Because the remaining refinement and acceptance stages remain incomplete, the reported cash-bond allocation reflects a provisional finite-grid estimation, leaving a verified numerical hedging decomposition pending. For comparative risk evaluations, a 10% increase in inventory standard deviation scales the inventory covariance component by a factor of 1.21, yielding $1.21\boldsymbol{\Psi}\boldsymbol{\Psi}^{\top} + R$ under a constant residual covariance matrix. Because this adjustment simultaneously alters expected return drift and ambiguity exposure, the counterfactual cannot be interpreted as an isolated variance perturbation or as conclusive evidence that cash commands the dominant intertemporal hedging component.

6 Conclusion

This paper delivers empirical associations alongside a conditional portfolio exercise. Its analytical contribution rests on the formal accounting framework, household optimality conditions conditional on value derivatives, and conditional verification arguments. The closed-form inventory barrier serves as a theoretical benchmark awaiting global verification, as the analysis does not claim to have fully solved and calibrated the joint inventory and portfolio system. Furthermore, the conditional existence result documented in the Appendix relies on explicit strategy-space and aggregate-map regularity rather than a verified construction of the model directly from primitives.

Trade policy uncertainty is associated with an immediate reduction in household equity exposure and positive cash and bond allocation responses. The equity-share and NASDAQ impact declines survive within-path simultaneous inference. The cash–bond split is less precisely estimated, and stock-share lead predictability rules out interpreting the baseline shock as an established causal shock.

The conditional calibration matches allocation levels and jointly disciplines the ini-

tial equity reduction and the increases in bonds and cash. The candidate gives impact responses of approximately -0.508 percentage points for equity, $+0.182$ for bonds and $+0.325$ for cash. Cash is the adding-up residual. Fixed-parameter signal and time refinements fail, and recalibration restores data fit on the changed grids while the remaining acceptance stages are pending. These provisional allocation results leave a substantial NASDAQ price-response discrepancy and do not establish a quantitative hedging or welfare decomposition.

These findings directly address the core research questions regarding firm inventory behavior, informational frictions, and safe-asset demand during uncertainty surges. First, consistent with Proposition 1, sufficiently strong escalation news can raise the regime 1 escalation intensity above the inventory threshold, inducing firms to front-load imports and accumulate excess inventory before a realized tariff escalation. Once the high-tariff regime is entered, the model instead predicts inventory drawdown.

Second, because households cannot directly observe these real-time inventory adjustments prior to formal disclosures, they face a wider set of admissible inventory beliefs as trade-policy uncertainty rises. In the model, the associated expansion of the entropy bound reduces risky-asset demand and contributes to a reallocation away from equity toward bonds and cash. Consequently, households sharply curtail equity exposure while expanding allocations toward sovereign bonds and cash.

Third, while both sovereign bonds and cash absorb the reduction in equity exposure, the substantially larger shift into cash indicates that the portfolio response is dominated by demand for liquidity rather than by a generalized flight into sovereign debt. In the model, this ranking arises because sovereign bonds remain exposed to inventory-linked fiscal risk, whereas cash provides the safer liquid margin of adjustment.

These results therefore leave imperfect inventory information as a plausible economic mechanism, but do not establish it as a quantitatively identified decomposition of the observed portfolio response. They show that the maintained factor mapping, conditional valuation inputs, and numerical opportunity set can reproduce the observed direction of the allocation response only under explicitly declared conditional calibration choices. The high-regime covariance and price-return proxies do not by themselves verify every sign condition invoked by the adverse-regime comparative statics. A theoretical sign result

and a measured calibration satisfying its premises are distinct claims.

The framework retains operating-profit maximization, an unobserved inventory state and inventory valuation between disclosures, and the maintained consumption-price assumption. Observed final-goods prices and market returns do not supply direct observations of inventory or corporate operating returns. Transition intensities affect conditional expectations as well as ambiguity bounds, so a change in TPU cannot be represented solely by a change in the entropy penalty.

Further quantitative work should connect the constrained household problem and finite-horizon inventory beliefs to the complete claim-valuation system, explicitly evaluating counterfactual dynamics by toggling the inventory channel and hedging channel on and off. This extended framework does not require inventory-market clearing, corporate bond issuance, an additional dividend-distribution theory, or security-specific disclosure dates. The current small-news experiment does not test front-loading or a fiscal J-curve. Its quantitative contribution is a documented limitation of the maintained specification, rather than a validated decomposition of the observed portfolio response.

References

Alessandria, G., J. P. Kaboski, and V. Midrigan. 2011. "US Trade and Inventory Dynamics." *American Economic Review* 101 (3): 303–307.

Amiti, M., S. J. Redding, and D. E. Weinstein. 2019. "The Impact of the 2018 Tariffs on Prices and Welfare." *Journal of Economic Perspectives* 33 (4): 187–210.

Baker, S. R., N. Bloom, and S. J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–1636.

Benguria, F., and F. Saffie. 2025. *Rounding Up the Effect of Tariffs on Financial Markets: Evidence from April 2, 2025.* NBER Working Paper 34036. National Bureau of Economic Research.

Boehm, C. E., A. A. Levchenko, and N. Pandalai-Nayar. 2023. "The Long and Short (Run) of Trade Elasticities." *American Economic Review* 113 (4): 861–905.

Boer, L., and M. Rieth. 2024. *The Macroeconomic Consequences of Import Tariffs and Trade Policy Uncertainty.* IMF Working Paper 2024/013. International Monetary Fund.

Brennan, M. J., and Y. Xia. 2001. "Stock Price Volatility and the Equity Premium." *Journal of Monetary Economics* 47 (2): 249–283.

Caldara, D., M. Iacoviello, P. Molligo, A. Prestipino, and A. Raffo. 2020. "The Economic Effects of Trade Policy Uncertainty." *Journal of Monetary Economics* 109:38–59.

Campbell, J. Y., and L. M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors.* Oxford University Press.

Carlomagno, G., and E. Albagli. 2022. "Trade Wars and Asset Prices." *Journal of International Money and Finance* 124:102631.

Chetty, R., A. Guren, D. Manoli, and A. Weber. 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review* 101 (3): 471–475.

Christiano, L. J., M. Eichenbaum, and C. L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

Drechsler, I. 2013. "Uncertainty, Time-Varying Fear, and Asset Prices." *Journal of Finance* 68 (5): 1843–1889.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal. 2020. "The Return to Protectionism." *Quarterly Journal of Economics* 135 (1): 1–55.

Gollin, D. 2002. "Getting Income Shares Right." *Journal of Political Economy* 110 (2): 458–474.

Hoque, M. E., L. Soo-Wah, M. A. Uddin, and A. Rahman. 2023. "International Trade Policy Uncertainty Spillover on Stock Market: Evidence from Fragile Five Economies." *Journal of International Trade and Economic Development* 32 (1): 104–131.

Huang, Y., C. Lin, S. Liu, and H. Tang. 2023. "Trade Networks and Firm Value: Evidence from the US–China Trade War." *Journal of International Economics* 145:103811.

Jordà, Ò. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1): 161–182.

Khan, Shafaat Yar, and Armen Khederlarian. 2021. *How Does Trade Respond to Anticipated Tariff Changes? Evidence from NAFTA.* Policy Research Working Paper 9646. World Bank, April.

Khan, Shafaat Yar, and Armen Khederlarian. 2025. "Inventories, Input Costs, and Productivity Gains from Trade Liberalizations." *International Economic Review* 66 (1): 175–199.

Merton, R. C. 1969. "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case." *Review of Economics and Statistics* 51 (3): 247–257.

Merton, R. C. 1971. "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." *Journal of Economic Theory* 3 (4): 373–413.

Newey, W. K., and K. D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–708.

Pástor, Ľ., and P. Veronesi. 2012. "Uncertainty about Government Policy and Stock Prices." *Journal of Finance* 67 (4): 1219–1264.

Rao, A., B. Lucey, and S. Kumar. 2025. "Financial Market Reactions to US Tariff Announcements: Evidence from Trade-Surplus and Trade-Deficit Countries." *Finance Research Letters,* 107465.

Schmitt-Grohé, S., and M. Uribe. 2025. *Transitory and Permanent Import Tariff Shocks in the United States: An Empirical Investigation.* NBER Working Paper 33997. National Bureau of Economic Research.

Wengerek, S. T., A. Uhde, and B. Hippert. 2025. "Share Price Reactions to Tariff Imposition Announcements during the First Trump Administration." *Finance Research Letters,* 107381.

Yilmazkuday, H. 2025. "US Tariffs and Stock Prices." *Finance Research Letters,* 107708.

A Data Sources

Table A.1: Summary of Data Sources

Title	Source	Period	Frequency
Investor's Asset Allocation Survey	AAII	Nov 1987 – Dec 2025	Monthly
DGS10	FRED	Jan 1962 – Dec 2025	Daily
NASDAQCOM	FRED	Feb 1971 – Dec 2025	Daily
US TPU (news)	Baker et al. (2016)	Jan 1985 – Dec 2025	Monthly
FRBATLWGTUMHWGO	FRED	Jan 1997 – Dec 2025	Monthly

B Continuous-Time Optimization and Equilibrium

B.1 Derivation of the unit cost and markup pricing

The firm takes the posted wage w and effective input price P_X as given. For a target output y, cost minimization gives

$$C(y) = \min_{L,X\geq 0}\{wL + P_X X : AL^{\alpha}X^{1-\alpha} \geq y\} = \frac{\Theta}{A} w^{\alpha} P_X^{1-\alpha} y, \tag{71}$$

$$\Theta = \alpha^{-\alpha}(1-\alpha)^{-(1-\alpha)} \tag{72}$$

Thus $MC = \Theta A^{-1} w^{\alpha} P_X^{1-\alpha}$ and $wL = \alpha MCY$, $P_X X = (1-\alpha)MCY$. Maximizing $(1-\tau_s)py(p) - MCy(p)$ under (5) gives

$$P_Y(t) = \frac{\varepsilon_p}{(\varepsilon_p - 1)(1-\tau_s)} MC(t), \tag{73}$$

$$\Pi(t) = \frac{(1-\tau_s)P_Y(t)Y(t)}{\varepsilon_p} > 0 \tag{74}$$

This proves Lemma 1; the pricing choice is static because MC does not depend on the firm's own price.

B.2 Conditional inventory verification and its scope

Physical inventory I is the stock state. Excess inventory is the signed indicator $E = I - K_0$, where K_0 is the constant operational benchmark in Section 3. It is not an additional stock and receives no separate accumulation or depreciation equation. Hence $E = 0$ means $I = K_0$, whereas stockout means $I = 0$ and $E = -K_0$. The condition $I \geq 0$ does not prohibit $E < 0$. All state dynamics and verification below are written in physical inventory; excess is recovered from its definition after each change in I.

Coordinate identity. For any physical-stock path, $E_t - E_s = I_t - I_s$ for a constant benchmark. For a differentiable value $V_k(I)$ define $C_k(e) = V_k(K_0 + e)$. Then $C_{k,E}(e) = V_{k,I}(K_0 + e)$ and $C_{k,EE}(e) = V_{k,II}(K_0 + e)$. These identities follow immediately by substitution and the chain rule. In particular the generator and the purchase obstacle are invariant under this translation when evaluated at corresponding physical states. No independent depreciation of excess is implied. The physical target corresponding to an excess target E_k^* is $K_0 + E_k^*$, not E_k^*. These statements establish the coordinate interpretation exactly, but do not establish optimality of a target. For a fixed foreign-input requirement $\bar{X}_F$, consider the auxiliary acquisition-cost problem

$$V_k(i) = \inf E_{i,k}\left[\int_0^{\infty} e^{-\rho t} P_{M,\xi_t}(\bar{X}_F + a_t)\, dt + \int_{[0,\infty)} e^{-\rho t} P_{M,\xi_t}\, dA_t\right]. \tag{75}$$

Here $a \in [-\bar{X}_F, 0]$ is the ordinary net flow and A is a nondecreasing cumulative purchase process, including impulses and maintenance. Thus $dI = [a - \delta_I I]dt + dA$; purchases are not counted twice. At $I = 0$, admissibility requires nonnegative physical-stock drift. This auxiliary problem must be distinguished from the operating-profit objective of Section 3:

a first-order envelope calculation does not prove that the two infinite-horizon objectives have identical global optimizers.

Conditional verification. For this auxiliary problem assume positive discounting, finite nonnegative transition intensities, constant regime-specific import prices, adapted controls of finite expected discounted purchase cost, and nonnegative physical stock. Suppose a candidate V_k is continuously differentiable, locally absolutely continuous as needed for the piecewise-deterministic chain rule, and satisfies the transversality condition $\lim_{T\to\infty} E[e^{-\rho T} V_{\xi_T}(I_T)] = 0$ for the policies to which verification is applied. These are sufficient conditions for the following verification result, not assertions that a candidate satisfying them has already been constructed. Define

$$H_k(i) = P_{M,k}\bar{X}_F - \delta_I i V_{k,I} + \lambda_{k\ell}(V_\ell - V_k) - \rho V_k + \inf_{a\in[-\bar{X}_F,0]} (P_{M,k} + V_{k,I})a,$$
$$G_k(i) = P_{M,k} + V_{k,I}.$$

At the physical-stock boundary the ordinary withdrawal control is restricted to $a = 0$. The cost-minimizing QVI is

$$\min\{H_k(i), G_k(i)\} = 0, \qquad i \geq 0.$$

For a finite upward impulse z, verification also uses $V_k(i) \leq P_{M,k}z + V_k(i+z)$, which follows by integrating $G_k \geq 0$. For any admissible policy, the discounted chain rule plus running and purchase costs has nonnegative expected drift, since $H_k \geq 0$ and $G_k \geq 0$. Integration to T and transversality therefore give $V_k(i)$ no greater than that policy's expected total cost. If an admissible policy attains the flow infimum, has $H_k = 0$ whenever time elapses, purchases only where $G_k = 0$, and attains equality across its impulses, all these inequalities are equalities. Its cost equals $V_k(i)$ and it is optimal. This proves the conditional verification result. It does not establish existence or uniqueness of a candidate, or verify the formula in Proposition 1.

Necessary continuation check. In a regime-2 maximal-drawdown interval where the regime-1 marginal value is $P_{M,1}$, set $q_2(i) = -V_{2,I}(i)$ and $r = \rho + \delta_I + \lambda_{21}$. Differentiating the continuation equation yields

$$[\bar{X}_F + \delta_I i]q_2'(i) + rq_2(i) - \lambda_{21}P_{M,1} = 0.$$

Consequently, for any boundary point i_b in that interval,

$$q_2(i) = \bar{q} + [q_2(i_b) - \bar{q}]\left(\frac{\bar{X}_F + \delta_I(i_b)}{\bar{X}_F + \delta_I i}\right)^{r/\delta_I}, \qquad \bar{q} = \frac{\lambda_{21}P_{M,1}}{r}.$$

The exponential in linear coverage is obtained in the zero-depreciation limit, not as an exact solution for positive depreciation. In particular, substituting $q_2 = \bar{q} + c\exp(-ri/\bar{X}_F)$ in the displayed ODE leaves $-r\delta_I i(q_2 - \bar{q})/\bar{X}_F$, generally nonzero. Restricting the state to a compact interval or bounding the transition intensities does not eliminate this interior residual. The boundary data, global objective equivalence, and the stated positive-

depreciation barrier formula therefore remain unverified. No additional economic restriction is imposed here to silently change Section 3.

Exact local solution and the role of coverage. The local algebra can be completed without treating excess as a separate stock. Conditional on a drawdown interval $[i_b, i_u]$, a constant regime-1 marginal value $P_{M,1}$ on that interval, and boundary data $q_b = q_2(i_b) > \bar{q}$, define

$$T(i, i_b) = \frac{1}{\delta_I} \log \frac{\bar{X}_F + \delta_I i}{\bar{X}_F + \delta_I i_b}.$$

This is the time required for physical inventory to fall from i to i_b under maximal drawdown without purchases. Direct integration of the physical inventory equation establishes this identity. Since $T_I = (\bar{X}_F + \delta_I i)^{-1}$, the marginal continuation equation becomes $dq_2/dT = -r(q_2 - \bar{q})$. Its unique solution with the supplied boundary data is

$$q_2(i) = \bar{q} + (q_b - \bar{q}) e^{-rT(i,i_b)}.$$

Existence follows by substitution. For uniqueness, the difference of two solutions satisfies $d\Delta/dT = -r\Delta$ and $\Delta(0) = 0$, hence vanishes. If a proposed marginal matching value satisfies $\bar{q} < B < q_b$, the unique matching point on this branch, provided it lies in the assumed interval, is

$$T^* = \frac{1}{r} \log \frac{q_b - \bar{q}}{B - \bar{q}}, \qquad i^* = i_b + \frac{\bar{X}_F + \delta_I i_b}{\delta_I} (e^{\delta_I T^*} - 1), \qquad E^* = i^* - K_0.$$

These formulas complete the local boundary-value calculation only. A constant marginal branch and its boundary data are conditional premises, not new globally imposed economic assumptions. In particular, when $i_b = K_0$, the excess coverage is

$$\frac{E^*}{\bar{X}_F} = \left(1 + \frac{\delta_I K_0}{\bar{X}_F}\right) \frac{e^{\delta_I T^*} - 1}{\delta_I} > T^* \quad (\delta_I > 0,\ T^* > 0).$$

Thus elapsed drawdown time cannot be identified with linear excess coverage in this branch. Choosing compact bounds or finite transition intensities does not change this strict inequality. The local solution supplies a calibration diagnostic, not a proof of the global policy claimed in Section 3.

B.3 Global solution of the auxiliary acquisition-cost problem

The cost problem (75) can be solved globally under its own stated primitives: constant foreign use $x = \bar{X}_F > 0$, prices $p_2 > p_1 > 0$, $p_k = P_{M,k}$, $\rho > 0$, $\delta = \delta_I \geq 0$, and finite nonnegative switching intensities. Physical stock is nonnegative; no positive operational floor or financing constraint is imposed on this auxiliary problem. Purchase expenditure is charged explicitly in its objective. These conditions define the problem solved here, not an equivalence to the operating-profit objective. The benchmark K_0 affects only the reported indicator $E = I - K_0$.

Put $R = \rho + \delta$, $r = R + \lambda_{21}$, and define

$$L = \begin{pmatrix} -\lambda_{12} & \lambda_{12} \\ \lambda_{21} & -\lambda_{21} \end{pmatrix}, \quad M = R\mathrm{Id} - L, \quad N = \rho\mathrm{Id} - L, \quad \bar{q} = \lambda_{21} p_1 / r.$$

Both M and N have positive eigenvalues. If $\lambda_{12}(p_2 - p_1) > R p_1$, let

$$B = \frac{(R + \lambda_{12}) p_1}{\lambda_{12}}, \quad T^* = \frac{1}{r} \log \frac{p_2 - \bar{q}}{B - \bar{q}}, \quad s = \frac{x}{\delta}(e^{\delta T^*} - 1).$$

At $\delta = 0$ the last expression means $s = xT^*$. Otherwise set $s = 0$. This includes $\lambda_{12} = 0$ without dividing by it. The physical-stock targets are $(s, 0)$.

For the positive-target case define $q(s) = (p_1, B)^\top$; for the zero-target case define $q(0) = (p_1, p_2)^\top$. For $i \geq s$, set

$$\begin{aligned} t_s(i) &= \frac{1}{\delta} \log \frac{x + \delta i}{x + \delta s}, \\ q(i) &= e^{-M t_s(i)} q(s), \qquad V(i) = (x + \delta i) N^{-1} q(i) \end{aligned} \tag{76}$$

At $\delta = 0$, $t_s(i) = (i - s)/x$. When $0 \leq i < s$, set

$$\begin{aligned} &q_1(i) = p_1, \qquad q_2(i) = \bar{q} + (p_2 - \bar{q})(1 + \delta i / x)^{-r/\delta}, \\ &V_1(i) = V_1(s) + p_1(s - i), \\ &V_2(i) = V_2(s) + \int_i^s q_2(u)\, du \end{aligned} \tag{77}$$

The integral is explicit: for $\delta > 0$ it equals

$$\bar{q}(s - i) + \frac{x(p_2 - \bar{q})}{r - \delta}[(1 + \delta i/x)^{1 - r/\delta} - (1 + \delta s/x)^{1 - r/\delta}].$$

At zero depreciation use $\bar{q}(s - i) + x(p_2 - \bar{q})(e^{-ri/x} - e^{-rs/x})/r$.

Global verification. For $i > s$, differentiating (76) gives $(x + \delta i) q' = -Mq$ and $V' = -q$, using $M = N + \delta\mathrm{Id}$. Consequently $(x + \delta i) q + LV - \rho V = 0$: maximal drawdown attains $H_k = 0$. Positivity of $e^{-Mt} = e^{-Rt} e^{Lt}$ and of N^{-1} gives $q > 0$ and $V > 0$. At the positive boundary $-Mq(s) = (0, -r(B - \bar{q}))^\top \leq 0$; at a zero boundary $-Mq(0) \leq 0$ follows from the threshold inequality and $p_2 > p_1$. Thus $q'(i) \leq 0$ throughout the common drawdown interval. In particular $q_k \leq p_k$, so the purchase obstacle is nonnegative there.

Below a positive s, $q_1 = p_1$, $B \leq q_2 \leq p_2$, and the regime-2 continuation equation holds by direct differentiation and matching at s. For regime 1, $G_1 = 0$ and

$$H_1'(i) = \lambda_{12}(B - q_2(i)) \leq 0, \qquad H_1(s) = 0.$$

Hence $H_1(i) \geq 0$ below s. Values and their first derivatives match at s. At stockout, $q_2(0) = p_2$ (and $q_1(0) = p_1$), so forbidding withdrawal at $i = 0$ yields the same boundary equation. Every QVI inequality therefore holds on the entire physical state space.

The attaining policy buys immediately to s in regime 1 if $I < s$, maintains s with

purchase rate δs, and otherwise draws down at rate x. Regime 2 draws down until physical stock is zero and thereafter buys only for current use. At regime switches the same rule is reapplied. Impulses satisfy $V_1(i) = p_1(s-i) + V_1(s)$ and maintenance occurs where $G_1 = 0$. The policy is admissible: its target is finite, intensities are finite, and expected discounted purchase costs are finite. Since $q \geq 0$, $0 \leq V_k(i) \leq \max_j V_j(0)$; thus discounted terminal value vanishes under every admissible policy. The verification argument above now applies without an unverified terminal assumption. These formulas give the unique optimal cost value; the stated policy attains it. No uniqueness of controls at indifference is claimed.

Comparison with the operating-profit benchmark. This global result is for (75). Its excess targets are $(s - K_0, -K_0)$, not $(s, 0)$ when $K_0 > 0$. Its logarithmic quantity is T^*, while $s/x = (e^{\delta T^*} - 1)/\delta$. The discrepancy is not an independent depreciation law for excess inventory. Moreover the Section 3 operating-profit ledger implies, for the same production plan,

$$\Pi^{\text{cash}} = \Pi - \iota, \qquad \iota = \frac{d(P_I I)}{dt} + \delta_I P_I I.$$

Writing $Z = P_I I$, integration by parts on a finite horizon gives

$$\int_0^T e^{-\rho t} \iota_t \, dt = e^{-\rho T} Z_T - Z_0 + (\rho + \delta_I) \int_0^T e^{-\rho t} Z_t \, dt.$$

The last integral generally depends on the inventory policy. Consequently replacing distributed operating profit with cash profit is not merely subtracting a policy-independent constant. This proof supplies an exact auxiliary solution, not a proof of the unchanged Section 3 barrier.

B.4 Link from inventory sourcing to conditional profit maximization

Fix the current inventory policy's sourcing share ϕ and ledger price P_I and put $c = \phi P_M + (1-\phi) P_I > 0$. Let $m(c)$ be the unit cost in (14) and let demand be $y(p) = \mathcal{D} p^{-\varepsilon_p}$, $\mathcal{D} > 0$. Conditional on c, operating profit is

$$\pi(p; c) = [(1 - \tau_s)p - m(c)] \mathcal{D} p^{-\varepsilon_p}.$$

Differentiation gives

$$\pi_p = \mathcal{D} p^{-\varepsilon_p - 1} [\varepsilon_p m(c) - (\varepsilon_p - 1)(1 - \tau_s) p].$$

The bracket is strictly decreasing in p, positive below its unique root and negative above it. Therefore the unique conditional global maximum is

$$p^*(c) = \frac{\varepsilon_p m(c)}{(\varepsilon_p - 1)(1 - \tau_s)}, \quad y^*(c) = \mathcal{D}[p^*(c)]^{-\varepsilon_p}, \quad \Pi^*(c) = \frac{m(c) y^*(c)}{\varepsilon_p - 1}.$$

This completes the price, quantity and operating-profit calculation for a supplied sourcing policy and recovers Lemma 1.

For that calculation to implement a fixed-use inventory solution, the foreign input demand generated by the production plan must also satisfy

$$X_F^*(c) = (1-s)(c/P_X)^{-\eta}\frac{(1-\alpha)m(c)y^*(c)}{P_X} = \bar{X}_F.$$

This is a resource-compatibility condition, not a consequence of positive and well-defined functions. It is checked without changing demand or the inventory requirement to force agreement. Conditional on a fixed sourcing share, Shephard's lemma and the price first-order condition yield

$$\frac{d\Pi^*(c)}{dc} = -y^*(c)m'(c) = -X_F^*(c).$$

The price-envelope term vanishes, but the first-order effect of the input cost remains. A supply policy minimizing the same use-cost functional can therefore be nested in the profit problem; minimizing a different acquisition-cost functional is not justified by the envelope theorem.

More formally, let $\mathcal{R}[Y]$ denote discounted expected after-tax revenue for an admissible sales path Y. Let $\mathcal{A}(Y)$ be its feasible production and inventory plans, and define their discounted use-cost value

$$\mathcal{K}[Y] = \inf_{a\in\mathcal{A}(Y)} E\int_0^\infty e^{-\rho t}[wL + p_{x,i}X_D + P_M X_F^N + P_I X_F^I]dt.$$

Whenever these functionals are finite, the exact nesting identity is

$$\sup_{Y,a\in\mathcal{A}(Y)} J(Y,a) = \sup_Y\{\mathcal{R}[Y] - \mathcal{K}[Y]\}.$$

For each fixed Y, this follows from subtracting the infimum of its use-cost from its revenue, followed by taking the outer supremum. It does not require existence of an attaining policy; if attainment is claimed, that additional property must hold for both stages. The acquisition-cost solution above solves a different inner functional. Its explicit policy-dependent accrual wedge and its fixed-use restriction cannot be removed by this nesting identity. Thus the conditional price calculation is complete, while identification of that auxiliary policy with the Section 3 profit optimizer is not asserted.

B.5 Joint optimization under resource and cost compatibility

The following resource and cost compatibility conditions are maintained for the composition result. They are additional restrictions, rather than consequences of well-defined production functions. They leave the physical inventory law, the meaning of E, and the Section 3 profit notation intact. For each admissible output choice y, let $\mathcal{A}(y)$ denote the jointly feasible production and inventory plans.

(C1) The inner cost minimum is attained by a plan $a_y \in \mathcal{A}(y)$. Its source quantities equal the foreign-input requirement of its production plan. Prices, quantities and other inputs adjust subject to demand and the production technology; quantities are not changed after optimization merely to eliminate an input residual.

(C2) For every admissible plan at the same y, discounted total use-cost and acquisition-cost obey $\mathcal{K}_{\text{use}}(y,a) = \mathcal{K}_{\text{acq}}(y,a) + b(y)$, where $b(y)$ is independent of the inventory policy a. All other input costs and ledger reconciliation terms are included in this identity. Thus the supplied acquisition-cost minimizer also minimizes the exact cost subtracted from revenue in the Section 3 operating-profit objective.

(C3) For the scalar, stationary-demand specialization, the feasible output interval $[\underline{y}, \overline{y}]$ is compact with $0 < \underline{y} < \overline{y}$. The minimized discounted use-cost $\mathcal{C}(y) = \mathcal{K}_{\text{use}}(y, a_y)$ is finite, continuously differentiable and convex. The demand shifter is positive, $\varepsilon_p > 1$, $0 \le \tau_s < 1$, and $\rho > 0$.

Condition (C2) is the precise common-cost condition: equality of accounting labels alone is insufficient. The earlier ledger identity shows that it restricts the admissible policy-dependent inventory-value terms; it is not asserted for unrestricted policies or inferred from the calibration. The correction $b(y)$ must remain in the outer objective and its derivative. Condition (C3) specifies the scalar specialization of the preceding general nesting identity, not a proof that all time-varying policies reduce to it.

Proposition 6 (Cost-minimum and profit-maximum composition)**.** *Under (C1)–(C3), the scalar joint problem has a unique optimal output y^* and an implementing inventory and production plan a_{y^*}. Let $p(y) = (\mathcal{D}/y)^{1/\varepsilon_p}$. The optimal output maximizes*

$$F(y) = \frac{1-\tau_s}{\rho}\mathcal{D}^{1/\varepsilon_p} y^{1-1/\varepsilon_p} - \mathcal{C}(y).$$

For an interior optimum,

$$p^* = \frac{\varepsilon_p}{(\varepsilon_p - 1)(1-\tau_s)}\rho\mathcal{C}'(y^*).$$

At the lower endpoint $F'(\underline{y}) \le 0$ is optimal, and at the upper endpoint $F'(\overline{y}) \ge 0$ is optimal. Inventory controls need not be unique when costs tie.

Proof. For a fixed y, (C2) makes the two inner minimizer sets identical because $b(y)$ is constant across their admissible policies. By (C1) the minimum is attained with all resource equations satisfied. Therefore every feasible pair (y,a) has profit at most $F(y)$, and (y, a_y) attains $F(y)$. Continuity and compactness give an outer maximizer. After-tax revenue is strictly concave in positive y, while $-\mathcal{C}$ is concave; hence F is strictly concave and its maximizing output is unique. Differentiation gives the displayed interior equation and the endpoint conditions. For any feasible (y,a) its profit is at most $F(y) \le F(y^*)$, and the pair (y^*, a_{y^*}) attains equality. This proves joint optimality in the stated class, not merely optimality of two unrelated stages. □

The marginal cost entering this composition is $\rho\mathcal{C}'(y)$. The Section 3 static markup is recovered when its production marginal cost MC is this same marginal cost. This marginal-cost consistency is part of using one common reduced cost, rather than pricing with a different cost after solving inventory. In particular, one must include $\rho b'(y)$. The theorem establishes a complete conditional composition result. It does not identify the original printed excess-coverage expression with the physical-stock target of the auxiliary problem, or assert (C1)–(C3) for parameterizations that have not been checked.

For the exact auxiliary solver, the quantity derivative is available without numerical differentiation. Scaling physical stock and the fixed input rate by the same positive factor scales all costs by that factor. Thus $V_k(I;x) = xV_k(I/x;1)$, and Euler's identity gives

$$\partial_x V_k(I;x) = \frac{V_k(I;x) + Iq_k(I;x)}{x}.$$

When $x = x(y)$ and the common-cost correction is $b(y)$, the implementation therefore uses

$$\mathcal{C}'(y) = \frac{V_k(I;x(y)) + Iq_k(I;x(y))}{x(y)}x'(y) + b'(y).$$

All remaining production expenses belong in the same cost functional and its correction; they cannot be omitted when mapping its derivative to MC.

B.6 Flow-of-funds identity

Using (32), the household drift is $\dot{W} = \Pi + r_k^D D + r_k^M M + CG + wL - (1+\tau_c)P_C C$. The firm identity is $(1-\tau_s)P_Y Y = wL + P_X X + \Pi$, with $P_F^{\text{use}} = \phi P_M + (1-\phi)P_I$. The inventory ledger implies

$$\iota = P_M Q_X - P_I X_F^I = \frac{d(P_I I)}{dt} + \delta_I P_I I \tag{78}$$

Substitution into the household and government accounts, using (20), (21), and (40), gives

$$\frac{d}{dt}[W - D] = CG + r^M M + NX - [p_{x,i}X_D + \epsilon p_{x,j}(X_F^N + Q_X)] + \iota \tag{79}$$

which reduces to (41) under the maintained $NX = 0$ approximation and otherwise retains the final-goods trade correction. Each wage, dividend, tax, tariff, and debt-service term appears once and cancels once.

B.7 Household consumption and portfolio maximization

Let $\mathcal{J}_k(W, z^H)$ be household value under its own information state. The firm's physical inventory is not observed merely because it is a sufficient state for firm optimization. Let $\boldsymbol{z} = (z^S, z^B)^\top$ be dollar holdings, subject to

$$\boldsymbol{z} \geq 0, \qquad \mathbf{1}^\top \boldsymbol{z} \leq W, \qquad W + \boldsymbol{z}^\top \boldsymbol{\Delta}_{kl} \geq 0.$$

The individual takes posted wages and demand-determined employment as given. Current equity payout is the profit in (16); no acquisition expense is subtracted a second time.

With the stated finite-variation information-state specification between updates, the

HJB is

$$\begin{aligned}\rho\mathcal{J}_k = \sup_{C,\boldsymbol{z}} \inf_{|h|\le\sqrt{2\kappa_k}} \Big\{ & \frac{C^{1-\gamma}}{1-\gamma} - \frac{L^{1+\chi}}{1+\chi} \\ & + \mathcal{J}_{W,k}[r^M W + wL - \widetilde{P}_C C + \boldsymbol{z}^\top \boldsymbol{\pi}_k^0 + \boldsymbol{z}^\top \boldsymbol{\Psi}_k h] \\ & + \tfrac{1}{2}\mathcal{J}_{WW,k}\boldsymbol{z}^\top \Sigma_k \boldsymbol{z} + \mathscr{L}_H \mathcal{J}_k \\ & + \lambda_{kl}[\mathcal{J}_l(W + \boldsymbol{z}^\top \boldsymbol{\Delta}_{kl}, z^{H,+}) - \mathcal{J}_k(W, z^H)]\Big\} \end{aligned} \tag{80}$$

An additional information-state diffusion, if specified, requires its own generator and covariation terms; it cannot be omitted from a wealth-slice calculation.

For $\mathcal{J}_W > 0$, minimizing the linear distortion term proves

$$\begin{aligned} h^* &= -\sqrt{2\kappa_k}\,\mathrm{sign}(\boldsymbol{z}^\top \boldsymbol{\Psi}_k), \\ \inf_h \mathcal{J}_W \boldsymbol{z}^\top \boldsymbol{\Psi}_k h &= -\mathcal{J}_W \sqrt{2\kappa_k}|\boldsymbol{z}^\top \boldsymbol{\Psi}_k| \end{aligned} \tag{81}$$

At zero exposure every admissible distortion minimizes that term. The exposure branch is checked at the optimum.

Strict concavity in consumption gives

$$\begin{aligned} C^* &= (\widetilde{P}_C \mathcal{J}_W)^{-1/\gamma}, \\ \sup_C \left\{ \frac{C^{1-\gamma}}{1-\gamma} - \widetilde{P}_C \mathcal{J}_W C \right\} &= \frac{\gamma}{1-\gamma}(\widetilde{P}_C \mathcal{J}_W)^{(\gamma-1)/\gamma} \end{aligned} \tag{82}$$

The actual wage bill, rather than a newly imposed average, remains in the full HJB.

Exact portfolio solution conditional on the value derivatives. Put $v = \mathcal{J}_{W,k}$ and $R = -\mathcal{J}_{WW,k}/v$. At $v > 0, R > 0$, with positive-definite Σ_k and concave destination value, the portfolio objective is strictly concave. This follows from its linear mean term, concave negative absolute-exposure term, negative-definite quadratic, and concave jump continuation. Its supergradient is

$$\begin{aligned} \boldsymbol{G} &= v(\boldsymbol{\pi}_k^0 - \sqrt{2\kappa_k} s \boldsymbol{\Psi}_k) + \mathcal{J}_{WW,k}\Sigma_k \boldsymbol{z} \\ &\quad + \lambda_{kl}\mathcal{J}_{W,l}(W + \boldsymbol{z}^\top \boldsymbol{\Delta}_{kl}, z^{H,+})\boldsymbol{\Delta}_{kl}, \\ s &\in \mathrm{Sign}(\boldsymbol{z}^\top \boldsymbol{\Psi}_k) \end{aligned} \tag{83}$$

Feasibility and the complementary conditions

$$\begin{aligned} \boldsymbol{G} - \eta\mathbf{1} + \boldsymbol{\beta} + \omega\boldsymbol{\Delta}_{kl} &= 0, \\ \eta(W - \mathbf{1}^\top \boldsymbol{z}) &= 0, \\ \beta_i z^i = 0, \qquad \omega(W + \boldsymbol{z}^\top \boldsymbol{\Delta}_{kl}) &= 0 \end{aligned} \tag{84}$$

are sufficient for the within-state maximum; the multipliers are nonnegative.

For an interior fixed-sign branch put $\boldsymbol{b} = \boldsymbol{\pi}^0 - \sqrt{2\kappa} s \boldsymbol{\Psi}$. The exact vector equation

reduces to

$$\begin{aligned} \boldsymbol{z}(h) &= R^{-1}\Sigma^{-1}(\boldsymbol{b} + h\boldsymbol{\Delta}), \\ h &= \frac{\lambda_{kl}}{v}\mathcal{J}_{W,l}(W + \boldsymbol{\Delta}^\top \boldsymbol{z}(h), z^{H,+}) \end{aligned} \tag{85}$$

The right-hand side of the second equation is nonincreasing in h and the left-hand side strictly increasing, giving at most one feasible root on the branch. There is no jump-size approximation.

For independent active constraints $E\boldsymbol{z} = e$, set $H = R\Sigma$ and

$$\begin{aligned} P_E &= H^{-1} - H^{-1}E^\top(EH^{-1}E^\top)^{-1}EH^{-1}, \\ \boldsymbol{z}(h) &= P_E(\boldsymbol{b} + h\boldsymbol{\Delta}) + H^{-1}E^\top(EH^{-1}E^\top)^{-1}e \end{aligned} \tag{86}$$

The same scalar equation applies; P_E is positive semidefinite and preserves its monotonicity. Roots must satisfy all inactive inequalities, multiplier signs and the exposure branch. Zero exposure is treated as the additional face $\boldsymbol{\Psi}^\top \boldsymbol{z} = 0$ with the bounded absolute-value subgradient. Thus the financial boundaries are included rather than removed by capitalization.

On an admissible class closed under convex combinations, individual value is concave in financial wealth. To prove this, combine two feasible strategies driven by the same observations. Affine wealth dynamics preserve the combined initial wealth and the same exogenous labor flow; convex financial constraints preserve feasibility. Consumption utility is concave under every common admissible distortion. Taking the worst case and then near-optimal strategies gives the value concavity inequality. This does not assert strict second derivatives at every boundary.

Robust verification. For a candidate satisfying (80), write the discounted residual as $\ell + \mathscr{G}^{a,h}\mathcal{J} - \rho\mathcal{J}$, where ℓ is flow utility. The maximizing policy makes this residual nonnegative under every admissible distortion. A minimizing response to any other policy makes it nonpositive. Discounted Itô–Dynkin, with the original admissibility, integrability and household terminal condition, gives

$$\inf_h J^{a^*,h} \geq \mathcal{J} \geq \sup_a \inf_h J^{a,h}.$$

Since a^* is feasible, both are equalities. This verifies the max–min objective, not a problem under a fixed distortion chosen independently of the portfolio.

B.8 The conditional homothetic formulas

For the benchmark with exogenous average labor income, fixed $H = \overline{wl}/r^M$ and a relaxed financial boundary, put $\widehat{W} = W + H$. The ansatz

$$\mathcal{J}_k = A_k\widehat{W}^{1-\gamma}/(1-\gamma) + \Phi_k$$

gives $C_k = m_k\widehat{W}$ with $m_k = (A_k\widetilde{P}_{C,k})^{-1/\gamma}$ and yields (55) and (56) by direct substitution. Here A_k is the subscripted household value coefficient, not the firm's TFP $A(t)$. Wealth-

independent labor disutility obeys

$$(\rho + \lambda_{kl})\Phi_k - \lambda_{kl}\Phi_l = -L_k^{1+\chi}/(1+\chi).$$

This system has determinant $\rho(\rho + \lambda_{12} + \lambda_{21}) > 0$.

For a smooth interior exposure branch at zero intensity define

$$\begin{aligned} \boldsymbol{x}_k^0 &= \gamma^{-1}\Sigma_k^{-1}\boldsymbol{\pi}_k, \\ CE_k &= \frac{\boldsymbol{\pi}_k^\top \Sigma_k^{-1}\boldsymbol{\pi}_k}{2\gamma}, \qquad\qquad d_k = \rho - (1-\gamma)(r^M + CE_k) \end{aligned} \tag{87}$$

When $d_k > 0$, direct maximization gives

$$m_k^0 = \frac{d_k}{\gamma \widetilde{P}_{C,k}}, \qquad\qquad A_k^0 = \widetilde{P}_{C,k}^{\gamma-1}(\gamma/d_k)^\gamma \tag{88}$$

The optimized value residual at zero intensity is

$$d_k A_k - \gamma(A_k \widetilde{P}_{C,k})^{(\gamma-1)/\gamma} = 0.$$

Its derivative at A_k^0 is $d_k/\gamma > 0$, while cross-regime derivatives vanish. With strict interiority, exposure and jump-solvency margins, the implicit-function theorem yields a unique local positive branch for sufficiently small intensities. It is a local result for this benchmark.

Under a fixed benchmark policy the discounted $(1-\gamma)$ wealth moment has diagonal regime coefficients

$$(1-\gamma)[r^M + \boldsymbol{x}_k^\top \boldsymbol{\pi}_k - \widetilde{P}_{C,k} m_k - \tfrac{\gamma}{2}\boldsymbol{x}_k^\top \Sigma_k \boldsymbol{x}_k] - \rho - \lambda_{kl}$$

and off-diagonal coefficients $\lambda_{kl}(1 + \boldsymbol{x}_k^\top \boldsymbol{\Delta}_{kl})^{1-\gamma}$. At zero intensity its eigenvalues are $-d_k/\gamma < 0$. Their continuity proves moment decay for sufficiently small intensities on the same branch. For $\gamma > 1$ and the maintained exposure sign, the worst-case drift bounds these negative wealth moments. Numerical convergence is not used as a transversality proof.

Linearizing the exact jump factor gives

$$\gamma(\Sigma + a\boldsymbol{\Delta}\boldsymbol{\Delta}^\top)\boldsymbol{x} = \boldsymbol{\pi} + a\boldsymbol{\Delta}.$$

The rank-one inverse yields the approximate portfolio formula in Proposition 2. This approximation is distinct from the exact constrained root (85). In the original problem $\mathbf{1}^\top \boldsymbol{x} \le W/(W+H)$ remains state dependent, and equilibrium labor income must be retained.

B.9 Valuation, beliefs, and the observation gap

Reporting dates are one month apart: $t_{n+1} = t_n + 1$ in monthly units. At each report the current inventory information gap is resolved, so $\widehat{I}_{t_n} = I_{t_n}$ and $\mathrm{Var}(I_{t_n} \mid \mathcal{H}_{t_n}) = 0$; consequently $\widehat{E}_{t_n} = I_{t_n} - K_0$ with the same zero conditional variance. This is an update

of information, not a reset of physical inventory or of future regime risk. The moment formulas below apply between reports and are restarted from the disclosed state at the next report. The monthly reporting convention is not a newly introduced uncertainty assumption.

Let X_k^i denote the value of claim $i \in \{S, B\}$ in regime k, and let $a_k^i = \delta_k^i - g_k^i$. With current-profit payout $CF_k^S = \Pi_k$ and bond service $CF_k^B = r_k^D D_k$, the stationary switching valuation satisfies

$$(a_k^i + \lambda_{kl}) X_k^i = CF_k^i + \lambda_{kl} X_l^i \tag{89}$$

For two regimes, this is a two-by-two linear system; its determinant is $a_1^i a_2^i + a_1^i \lambda_{21} + a_2^i \lambda_{12} > 0$ under positive effective discount rates. Hence values are unique conditional on the payout and growth objects. If a primitive θ changes,

$$\partial_\theta X = \mathcal{A}^{-1}\{\partial_\theta CF - (\partial_\theta \mathcal{A}) X\},$$

which separates payout, discounting, and transition-intensity channels. The return process follows from the total differential of $X(I_t, k_t)$:

$$\frac{dX}{X} = \frac{X_I}{X} dI + \frac{1}{2}\frac{X_{II}}{X} d\langle I \rangle + \frac{X_l - X_k}{X} dN_{kl} + \text{drift terms.}$$

Thus the inventory shock enters both expected returns and quadratic variation; an observed final price does not identify I, P_I, or the firm's latent profit separately.

During the assumed non-observation interval, beliefs propagate under the stated reference law. Let $\mathcal{H}_t$ denote household information at the start of the interval and let E_t and Cov_t condition on $\mathcal{H}_t$. The latent initial stock is measurable in the full model filtration $\mathcal{F}_t \supseteq \mathcal{H}_t$, with $\widehat{I}_t = E_t I_t$ and $V_t = \mathrm{Var}_t(I_t)$. Conditioning on $\mathcal{H}_t$ does not make inventory observable.

Proposition 7 (Conditional inventory moments during non-observation)**.** *For the reference SDE* (25)*, on an interval where its integrated drift and diffusion have finite second moments, define*

$$B_{t,u} = \int_t^{t+u} b_{\xi_s}(I_s)\, ds, \qquad M_{t,u} = \int_t^{t+u} \bar{\nu}_{\xi_{s-}}\, dW_s \tag{90}$$

The predictable regime version is used in the stochastic integral. Then

$$E_t I_{t+u} = \widehat{I}_t + \int_t^{t+u} E_t b_{\xi_s}(I_s)\, ds, \tag{91}$$

$$\begin{aligned}\mathrm{Var}_t(I_{t+u}) = V_t + \mathrm{Var}_t(B_{t,u}) + E_t \int_t^{t+u} \bar{\nu}_{\xi_s}^2\, ds \\ + 2\,\mathrm{Cov}_t(I_t, B_{t,u}) + 2\,\mathrm{Cov}_t(B_{t,u}, M_{t,u})\end{aligned} \tag{92}$$

Neither covariance in the last line is assumed zero.

Proof. Integrating the SDE yields $I_{t+u} = I_t + B_{t,u} + M_{t,u}$. The martingale-increment property gives $E[M_{t,u} \mid \mathcal{F}_t] = 0$. By iterated conditioning, $E_t M_{t,u} = 0$ and $E_t[(I_t -$

$\widehat{I}_t)M_{t,u}] = 0$. Thus the covariance of the initial stock with the future martingale increment vanishes, even though the household does not observe that stock. Conditional Itô isometry gives

$$E_t M_{t,u}^2 = E_t \int_t^{t+u} \bar{\nu}_{\xi_s}^2 \, ds.$$

Taking the conditional mean and expanding the variance of $I_t + B_{t,u} + M_{t,u}$ proves both identities. The future integrated drift depends on the evolving stock, so its covariance with the martingale increment cannot be removed by the martingale property. □

Moment dynamics and regime transitions. An equivalent differential form follows directly from Itô's formula. Write $m(u) = E_t I_{t+u}$, $v(u) = \mathrm{Var}_t(I_{t+u})$, and $b_u = b_{\xi_{t+u}}(I_{t+u})$. At times where the moment derivatives exist,

$$\begin{aligned} m'(u) &= E_t b_u, \\ \frac{d}{du} E_t I_{t+u}^2 &= 2E_t[I_{t+u} b_u] + E_t \bar{\nu}_{\xi_{t+u}}^2, \\ v'(u) &= 2\,\mathrm{Cov}_t(I_{t+u}, b_u) + E_t \bar{\nu}_{\xi_{t+u}}^2 \end{aligned} \tag{93}$$

Indeed, $d(I^2) = (2Ib + \bar{\nu}^2)du + 2I\bar{\nu}\, dW$; subtracting $d(m^2)/du = 2mE_t b$ gives the last equation. These equations retain state–drift dependence for nonlinear inventory policies.

For completeness, let $p_k = E_t \mathbf{1}_{\{\xi_{t+u}=k\}}$, $m_k = E_t[I_{t+u}\mathbf{1}_{\{\xi_{t+u}=k\}}]$, and $s_k = E_t[I_{t+u}^2 \mathbf{1}_{\{\xi_{t+u}=k\}}]$. In this paragraph m_k denotes an unnormalized conditional first moment, not the consumption ratio. For deterministic intensities over the interval, the joint generator gives

$$\begin{aligned} p_k' &= \sum_j Q_{jk} p_j, \\ m_k' &= E_t[b_k(I_{t+u})\mathbf{1}_{\{\xi_{t+u}=k\}}] + \sum_j Q_{jk} m_j, \\ s_k' &= E_t[(2I_{t+u} b_k(I_{t+u}) + \bar{\nu}_k^2)\mathbf{1}_{\{\xi_{t+u}=k\}}] + \sum_j Q_{jk} s_j \end{aligned} \tag{94}$$

Apply the generator to $\mathbf{1}_k$, $I\mathbf{1}_k$, and $I^2\mathbf{1}_k$ to obtain these expressions: a regime switch changes the indicator but does not reset inventory. Consequently $m = \sum_k m_k$ and $v = \sum_k s_k - m^2$. Transition intensities remain in the conditional moments even before a realized regime change. For stochastic intensities their products with the state remain inside the conditional expectations. Affine regime drifts close this moment system; nonlinear drifts require the conditional distribution or an explicit numerical approximation of the displayed expectations.

Affine-drift verification. Within a fixed regime, take the drift already specified in the main text, $b(I) = c - \delta_I I$, away from a boundary. The integrating-factor solution is

$$I_{t+u} = e^{-\delta_I u} I_t + \frac{c}{\delta_I}(1 - e^{-\delta_I u}) + \bar{\nu} \int_t^{t+u} e^{-\delta_I (t+u-s)} \, dW_s \tag{95}$$

Conditional expectation and Itô isometry therefore yield

$$m(u) = e^{-\delta_I u}\widehat{I}_t + \frac{c}{\delta_I}(1 - e^{-\delta_I u}),$$
$$v(u) = e^{-2\delta_I u}V_t + \frac{\bar{\nu}^2}{2\delta_I}(1 - e^{-2\delta_I u}) \tag{96}$$

The same result follows from $\mathrm{Cov}_t(I, c - \delta_I I) = -\delta_I v$ in (93), giving $v' = -2\delta_I v + \bar{\nu}^2$ and $v(0) = V_t$. This independent calculation agrees with the main-text formula and verifies that the covariance terms in (92) are essential. The limit $\delta_I \to 0$ gives $v(u) = V_t + \bar{\nu}^2 u$, precisely the constant-drift case.

Ambiguity and boundaries. Under an admissible distorted law $\mathbb{Q}$, apply the same proof with $b_{\xi_s}(I_s) + \bar{\nu}_{\xi_s} h_s$ and $W^{\mathbb{Q}}$. The distortion bound controls its direct integrated contribution:

$$\left| \int_t^{t+u} \bar{\nu}_{\xi_s} h_s \, ds \right| \leq \int_t^{t+u} \bar{\nu}_{\xi_s} \sqrt{2\kappa_{\xi_s}} \, ds.$$

With state-dependent drift this is not by itself a bound on the difference of terminal means, since distortion also changes subsequent reference drift. For the fixed-regime affine case and the same initial belief, the integrating-factor solution instead proves

$$\left| E_t^{\mathbb{Q}} I_{t+u} - E_t I_{t+u} \right| \leq \bar{\nu}\sqrt{2\kappa}\frac{1 - e^{-\delta_I u}}{\delta_I} \tag{97}$$

For the zero-depreciation disclosure benchmark, the corresponding mean interval is

$$E_t^{\mathbb{Q}}[I_{t+u}] \in \widehat{I}_t + b_k u \pm \bar{\nu}_k \sqrt{2\kappa_k}\, u \tag{98}$$

Constant extremal distortions attain the two bounds while this unrestricted affine law applies. Its zero-depreciation limit is the constant-drift mean interval in (98). A general adapted distortion need not leave the variance equal to the reference variance.

The results concern the stated subjective inventory law and introduce no observation likelihood. They apply across policy switches whenever those switches change drift without adding a state jump or another finite-variation term. If a boundary specification introduces reflection, absorption, or a reset, its actual stopped or boundary generator must be used. For example, an added regulator dL contributes $L_{t+u} - L_t$ to the integrated state and $2I\, dL$ to the second-moment equation. Such a boundary mechanism is not silently imposed on the reference diffusion. Physical stock feasibility in the firm's problem remains separate from the household's subjective law.

B.10 Exact portfolio differentials and the rank-one approximation

Write $a = \lambda_{kl} A_l / A_k$ and $z = 1 + \boldsymbol{x}^\top \boldsymbol{\Delta} > 0$ on a smooth interior, fixed-exposure branch of the conditional homothetic problem. Its exact first-order condition is $\boldsymbol{\pi} - \gamma\Sigma\boldsymbol{x} + az^{-\gamma}\boldsymbol{\Delta} =$

0. Differentiation gives

$$\begin{aligned}\mathcal{M} &= \gamma\Sigma + \gamma a z^{-\gamma-1}\boldsymbol{\Delta}\boldsymbol{\Delta}^\top,\\ \mathcal{M}\, d\boldsymbol{x} &= d\boldsymbol{\pi} - \gamma(d\Sigma)\boldsymbol{x} + z^{-\gamma}\boldsymbol{\Delta}\, da + a z^{-\gamma} d\boldsymbol{\Delta}\\ &\quad - \gamma a z^{-\gamma-1}\boldsymbol{\Delta}(\boldsymbol{x}^\top d\boldsymbol{\Delta})\end{aligned} \tag{99}$$

For Σ positive definite and $a \geq 0$, $\mathcal{M}$ is positive definite, so this derivative is uniquely determined conditional on the value-coefficient differentials. Those differentials come from the coupled value equations; holding them fixed is a partial derivative, not a total equilibrium response.

For the stated first-order jump expansion, put $D = \boldsymbol{\Delta}^\top\Sigma^{-1}\boldsymbol{\pi}$ and $E = \boldsymbol{\Delta}^\top\Sigma^{-1}\boldsymbol{\Delta}$. The identity

$$(\Sigma + a\boldsymbol{\Delta}\boldsymbol{\Delta}^\top)^{-1} = \Sigma^{-1} - \frac{a\Sigma^{-1}\boldsymbol{\Delta}\boldsymbol{\Delta}^\top\Sigma^{-1}}{1+aE}$$

is verified by multiplication. Multiplication by $(\boldsymbol{\pi} + a\boldsymbol{\Delta})/\gamma$ yields $\boldsymbol{x} = \Sigma^{-1}\boldsymbol{\pi}/\gamma + w\Sigma^{-1}\boldsymbol{\Delta}$, where $w = a(1-D)/[\gamma(1+aE)]$. Therefore

$$\begin{aligned}dw &= \frac{1-D}{\gamma(1+aE)^2}\, da - \frac{a}{\gamma(1+aE)}\, dD - \frac{a^2(1-D)}{\gamma(1+aE)^2}\, dE,\\ da &= \frac{A_l}{A_k}\, d\lambda_{kl} + \lambda_{kl}\, d(A_l/A_k)\end{aligned} \tag{100}$$

In particular, $D < 1$ makes the partial derivative with respect to a positive, not every total derivative with respect to intensity positive. Intensity remains in the transition probabilities and value ratios even without a realized regime change. Changes in TPU can therefore affect beliefs and portfolios through this channel. At active constraints these differentials are replaced by the active-face KKT system.

B.11 Comparative statics and identification

For an interior, fixed-sign portfolio branch, the exact first-order condition is the scalar root in (85); its implicit derivative with respect to any primitive θ is

$$\frac{dh}{d\theta} = \frac{\partial_\theta R(h,\theta)}{1 - \partial_h R(h,\theta)}$$

whenever the denominator is nonzero. Portfolio, consumption, and output responses are then obtained by differentiating the exact root and the firm policy map. At a binding financial or zero-exposure face, the KKT system (84) replaces an unconstrained derivative. Therefore sign claims require the maintained interiority, sign, and solvency conditions; the formulas do not support unconditional monotonicity claims.

At zero regime intensity, the Merton benchmark gives $\boldsymbol{x}_k^0 = \gamma^{-1}\Sigma_k^{-1}\boldsymbol{\pi}_k$ and $CE_k = \boldsymbol{\pi}_k^\top\Sigma_k^{-1}\boldsymbol{\pi}_k/(2\gamma)$. For sufficiently small intensities, the local branch in (88) inherits these objects by the implicit-function argument. For moderate intensities or active constraints, the exact root and the coupled value equations are evaluated by the corresponding numerical HJB solution; the calibration tables are descriptive targets and do not by themselves establish those global claims.

The common-factor covariance restriction is

$$\Psi^S = \sigma_S\sqrt{c_S}, \qquad \Psi^B = \frac{\mathrm{Cov}(dG^S/G^S, dG^B/G^B)}{\Psi^S}, \qquad c_S \geq \mathrm{Corr}(dG^S/G^S, dG^B/G^B)^2.$$

It identifies the ambiguity exposure only up to the admissible covariance region. Identification of inventory effects therefore requires variation in disclosure timing, TPU, and regime transitions jointly; a price-only inversion is not available. The empirical confidence intervals are interpreted as uncertainty about reduced-form responses, not as direct confidence intervals for the latent stock or its carrying price.

B.12 Consolidation of goods, labor and fiscal flows

The nominal goods identity is $P_Y Y = P_C C + NX$. Under the maintained normalization of domestic dividend and debt-service receipts, the household and government drift accounts are

$$\begin{aligned} \dot{W} &= \Pi + wL + r^D D + r^M M + CG - (1+\tau_c)P_C C, \\ \dot{D} &= r^D D - T, \\ T &= \tau_c P_C C + \tau_s P_Y Y + \tau \epsilon p_{x,j}(n+Q) \end{aligned} \tag{101}$$

Here these are drift identities; asset diffusion and jump gains remain in their corresponding wealth processes. Substitute $\Pi = (1-\tau_s)P_Y Y - wL - p_{x,i}X_D - (1+\tau)\epsilon p_{x,j}(n+Q) + \iota$ and subtract the debt account:

$$\begin{aligned} \dot{W} - \dot{D} &= P_Y Y - P_C C - p_{x,i}X_D - \epsilon p_{x,j}(n+Q) + r^M M + CG + \iota \\ &= NX - p_{x,i}X_D - \epsilon p_{x,j}(n+Q) + r^M M + CG + \iota \end{aligned} \tag{102}$$

This proves (41): wages, domestic taxes, tariff receipts and domestic debt service cancel exactly. Inventory withdrawal does not appear as a second expense. Employment clearing identifies supplied labor with firm labor demand at the posted wage; it is not an additional household labor first-order condition. These identities establish accounting consistency conditional on the stated holdings and goods clearing, not existence of an asset-market or inventory-market equilibrium.

B.13 Optimal production and goods-demand consistency

The model's demand-determined goods closure is distinct from competitive asset-market equilibrium. In particular, existence of a solution does not require a closed-form HJB solution, uniqueness, global smoothness, or a stationary physical inventory.

Lemma 4 (Optimal supply accommodates final-goods demand)**.** *Fix an admissible aggregate state and a positive effective unit cost m on the interior production branch. For any positive aggregate demand Y that is feasible on this branch, the symmetric markup price and the firm's optimal production satisfy $P = P_Y$ and $y^* = Y$. Consequently an admissible household consumption choice and its CES goods allocation can be implemented by optimal firm production at these prices, conditional on their consistency with the same aggregate state and continuation values.*

Proof. Put $\upsilon = (\varepsilon_p - 1)/\varepsilon_p$. With the individual demand shifter $\mathcal{D} = P_Y^{\varepsilon_p} Y$, the production solution is

$$y^* = \left[\frac{\upsilon(1-\tau_s)P_Y}{m}\right]^{\varepsilon_p} Y,$$
$$P^* = \frac{m}{(1-\tau_s)\upsilon} \tag{103}$$

Symmetry gives $P_Y = P^*$, so the bracket equals one and $y^* = Y$. The strictly concave revenue-minus-effective-cost objective established above makes this the global production maximum on the indicated branch. At empty inventory the same static calculation uses the fresh-input cost. This conditional production calculation does not solve the acquisition decision or establish the inventory barrier.

CES household demands allocate spending across the two countries' goods. Setting each country's Y equal to domestic and foreign demand for its goods therefore implements variety clearing and its nominal counterpart $P_Y Y = P_C C + NX$. The labor demand associated with that production is $L = \alpha m Y/w$, supplied at the posted wage. Current acquisitions, withdrawals and profit remain those of the firm optimum; there is no extra inventory-clearing equation. Binding production bounds require the constrained production correspondence instead of the interior formula. □

The relevant curvature can be stated precisely. At a fixed positive demand shifter, inverse demand is $P(y) = \mathcal{D}^{1/\varepsilon_p} y^{-1/\varepsilon_p}$: $P'(y) < 0$ and $P''(y) > 0$. Revenue, however, is strictly concave in y, as is revenue minus a positive linear effective cost. Household consumption utility satisfies $u'(C) = C^{-\gamma} > 0$ and $u''(C) = -\gamma C^{-\gamma-1} < 0$. Households maximize discounted utility, not the physical quantity of current consumption without regard to saving. These curvature properties justify the individual maximizations. Price convexity alone is not the fixed-point argument for endogenous income and continuation values.

B.14 Existence under an explicit consistency-map formulation

The following are sufficient regularity conditions for a conditional fixed-point result. They are not derived from the Section 3 primitives or verified by the calibration. Let b collect trial aggregate prices, income and continuation objects; $\mathcal{B}$ is their trial domain. Let $\Gamma(b)$ consist of the aggregate objects generated by optimizing policies at b that satisfy the goods, labor and fiscal equations.

(i) At every $b \in \mathcal{B}$, each agent's admissible strategy set is nonempty and compact in a specified topology, and the discounted objective is finite and upper semicontinuous on that set. For households this condition applies to the worst-case objective itself. It includes the needed infinite-horizon integrability conditions.

(ii) $\mathcal{B}$ is a nonempty compact convex subset of a Hausdorff locally convex space. The correspondence Γ takes nonempty compact convex values in $\mathcal{B}$; every member has implementing optimal policies satisfying the stated goods, labor and fiscal equations.

(iii) Γ is upper hemicontinuous and is a self-map of $\mathcal{B}$.

Compactness here is a sufficient property of strategy and aggregate spaces, not an asserted bound on physical stock. Convex-valuedness and nonemptiness are additional premises; neither follows merely from concavity of static production or from the inventory barrier formula. No claim is made that imposing numerical grid bounds verifies these premises.

Proposition 8 (Joint consistency under fixed-point regularity)**.** *Under the maintained regularity conditions in Section B.14, there exist optimal admissible firm and household policies and a jointly consistent aggregate environment $b^* \in \mathcal{B}$ satisfying $b^* \in \Gamma(b^*)$. The policies maximize firm discounted profit and household worst-case discounted utility, respectively, and satisfy the stated goods, labor and fiscal closure.*

Proof. First fix $b \in \mathcal{B}$. Each objective is finite and upper semicontinuous on its nonempty compact admissible strategy set by condition (i). An upper semicontinuous real-valued function on a compact set attains its maximum: its upper level sets along a maximizing sequence are nonempty nested closed sets, and compactness gives a point in their intersection attaining the supremum. This applies directly to discounted firm profit and to the household's worst-case discounted utility. It does not require interchange of the household supremum and distortion infimum. Hence both conditional optimization problems have optimal admissible strategies at every trial environment.

Conditions (ii)–(iii) place their aggregate update correspondence on a nonempty compact convex subset of a Hausdorff locally convex space, with nonempty compact convex values, upper hemicontinuity and the self-map property. The Kakutani–Fan–Glicksberg fixed-point theorem therefore gives $b^* \in \Gamma(b^*)$. By the definition of Γ, optimal admissible policies generate this member. Because it is the same object as their input b^*, their income, prices and continuation objects are mutually consistent. Lemma 4 implements interior goods production at the resulting demand; where bounds bind, the definition of Γ uses constrained clearing instead. Labor and fiscal closure follow from the stated accounts. The firm financing and stock restrictions and household financial restrictions are preserved by admissibility. This proves existence, not uniqueness. The map concerns complete continuous-time objects; a fixed point of this map is not a stationary inventory point or a discretization of decisions. □

When a single scalar output coordinate suffices after the other objects have been solved, the same existence step is elementary. Let

$$\mathcal{T}(Y) = \frac{P_C(Y)C^*(Y) + NX(Y)}{P_Y(Y)}, \qquad Z(Y) = \mathcal{T}(Y) - Y \tag{104}$$

If $\mathcal{T}$ is continuous and maps a feasible interval $[\underline{Y}, \overline{Y}]$ into itself, then $Z(\underline{Y}) \geq 0$ and $Z(\overline{Y}) \leq 0$. The intermediate value theorem gives $Z(Y^*) = 0$ and thus goods clearing. This proves existence over that entire interval, without a local small-intensity argument. It does not assume that the full stochastic problem always reduces to this scalar map.

The existence conclusion is global over the maintained admissible domain: it does not rely on a small transition intensity or an interior portfolio. It is an existence result under the explicit aggregate regularity restrictions, not a claim that every primitive parameterization satisfies those restrictions. In a finite-dimensional single-valued specialization, Brouwer's theorem gives the same conclusion. Transition intensities remain the model's observed TPU-driven inputs, not additional markets to clear.

B.15 Established results and solution scope

The physical/excess coordinate identities and the accounting identities are algebraic results; the consolidated account retains NX unless the main-text approximation sets it to zero. Static markup production is conditional on its effective cost and demand environment. The auxiliary inventory-cost problem has an explicit global value function and a verified attaining policy under its stated constant-input primitives. Neither result proves the global barrier or its equivalence to maximizing the operating profit distributed in Section 3.

Proposition 6 completes the joint scalar optimization argument under the explicitly maintained resource, common-cost and convexity conditions (C1)–(C3). It establishes that the inner cost optimizer and outer revenue choice jointly maximize the same profit objective in that class. This is a proved conditional result, distinct from claiming that arbitrary primitives satisfy those restrictions.

Household consumption and portfolio conditions are conditional on the value derivatives, admissibility and concavity premises stated above. The robust verification argument additionally requires an HJB candidate and terminal conditions. The homothetic calculation concerns its stated local benchmark; it is not a construction of a globally smooth constrained value function. Conditional belief moments retain drift covariance terms.

Proposition 8 is a sufficient-condition theorem under the explicit regularity premises in this Appendix. Its proof is complete conditional on those premises, but their validity for the full Section 3 economy has not been established. The scalar alternative has the same conditional status. No unconditional joint-existence, uniqueness, globally optimal closed-form inventory policy, or calibrated joint equilibrium is claimed in this version.