Too Many Hand-to-Mouth Households? Monetary Policy and the Risk of Inverted Aggregate Demand in Emerging Economies

Badr Elhamzaoui* Faïçal Lakhchen† Radouane Raouf‡

August 2026

Abstract

In emerging economies, limited consumption smoothing may be widespread. When constrained households are sufficiently prevalent, aggregate demand can invert and the Taylor principle can reverse. Yet monetary models often import constrained-household shares from advanced-economy calibrations. Using two nationally representative Moroccan expenditure surveys, we construct a food-share proxy for short-run budget tightness and map it into a tractable two-agent New Keynesian (TANK) model. At a common 40 percent cutoff, 68.8 percent of households in 2007 and 57.2 percent in 2014 are classified as budget-tight, relative to a benchmark critical share of 54.5 percent. Under the benchmark household-share mapping, the 2007 calibration lies inside the inverted aggregate-demand region, while the 2014 calibration lies close to the boundary and moves into the standard region under stricter cutoffs, higher markups, or consumption weighting. In the inverted region, passive inflation responses, sufficiently strong output responses, or commitment targeting can restore determinacy. In the 2007 benchmark, reallocating more than 31.4 percent of cyclical profit income toward constrained households restores the standard demand slope. Where budget tightness is widespread, the sign of monetary transmission should be tested rather than assumed.

Keywords: Hand-to-mouth households; Inverted aggregate demand; Monetary policy; Emerging economies.

JEL classifications: E12; E21; E52; D52.

*Corresponding author. Mohammed V University, Rabat, Morocco. Email: badr.elhamzaoui@um5r.ac.ma.

†Mohammed I University, Oujda, Morocco.

‡Mohammed V University, Rabat, Morocco.

1 Introduction

Modern monetary policy relies on a basic aggregate-demand logic according to which current spending should fall when the real interest rate rises. The Taylor principle, inflation targeting, and forward guidance are usually interpreted through this standard transmission channel. In the representative-agent New Keynesian model, this relationship follows directly from households' consumption-smoothing behavior. A higher real rate increases the return to postponing consumption, so households shift some spending from the present to the future, causing aggregate demand to contract.

This relationship need not hold. Its sign depends on who is doing the spending. A household that consumes mainly out of current income does not respond primarily to the real interest rate through the intertemporal-substitution channel. It responds through current income. When the share of such households is sufficiently low, their behavior does not overturn the aggregate relation. But when their share is sufficiently high, the income channel can dominate the substitution channel. Bilbiie (2008) formalized this mechanism. Once the share of hand-to-mouth (HtM) households crosses a threshold that depends on the labor-supply elasticity and the markup, the elasticity of aggregate demand to the real rate changes sign, and the economy enters the inverted aggregate-demand logic (IADL) region, where a higher expected real interest rate raises rather than lowers aggregate demand. In that region, the Taylor principle is inverted, so an active inflation response can become destabilizing.

This mechanism is more than a theoretical curiosity. Bilbiie and Straub (2013) argue that the pre-Volcker United States may have operated in or near this region, and that the passive monetary stance of the period can be consistent with determinacy once limited asset-market participation is taken into account. Existing empirical applications, however, have mostly focused on data-rich advanced economies. Emerging economies, where the mechanism may be most relevant, are precisely those in which balance-sheet data are often unavailable. This paper studies that case.

The conditions that can make hand-to-mouth behavior macroeconomically important, such as thin financial buffers, limited formal financial access, widespread informality, and high food expenditure shares, are common in developing and emerging economies. Yet many structural models used in such settings either abstract from household heterogeneity or rely on constrained-household

shares imported from other environments. If the constrained share is high enough to cross the threshold, this is not a minor calibration issue. The sign of aggregate-demand transmission itself becomes an empirical question. We take up this question using Morocco as a case study.

The analysis begins with measurement and its limitations. Morocco's household expenditure surveys, conducted by the Haut-Commissariat au Plan (HCP) in 2007 and 2014 and covering 7,062 and 15,970 households, respectively, contain detailed consumption modules but do not observe liquid assets, credit limits, or saving flows (Haut-Commissariat au Plan, 2007, 2014). A more recent HCP living-standards survey was conducted in 2022–2023, but the household-level microdata required to reconstruct our classifier were not available to us; we therefore use the 2007 and 2014 survey rounds for which the classification can be reproduced. We cannot identify HtM status from the asset side in the sense of Kaplan et al. (2014). In the absence of Moroccan household microdata that directly measure liquid assets, saving behavior, or responses to income shocks, expenditure microdata provide the most direct domestically grounded way currently available to discipline the constrained-household share instead of importing HtM calibrations from other economies. We therefore construct an expenditure-based proxy for short-run budget tightness. Following the tradition that interprets the food-expenditure share as an indicator of household budget pressure (Deaton, 1997; Chai and Moneta, 2010), we classify households that devote a large share of total consumption expenditure to food as high-food-share households. The interpretation is that these households have limited room to compress discretionary spending and limited capacity to absorb shocks. We discipline this proxy through cutoff sensitivity, binned expenditure profiles, kink regressions, breakpoint estimation with bootstrap inference, levels-based checks, and socioeconomic checks using poverty and vulnerability indicators. These exercises support a budget-tightness interpretation, but do not identify liquid-asset hand-to-mouth status. Under a common 40 percent food-share benchmark, the survey-weighted prevalence of high-food-share households is 68.8 percent in 2007 and 57.2 percent in 2014.

We then incorporate these mapped shares into a tractable two-agent New Keynesian model in the spirit of Bilbiie (2008). The model is deliberately simple, with a closed economy, no capital, Rotemberg pricing, two household types, zero steady-state profits, and a forward-looking interest-rate rule. The purpose is to isolate how a high constrained-household share makes the distribution of income between wages and profits central to the transmission of output into current aggregate consumption, and where this places the economy relative to the aggregate-demand threshold. At the

benchmark markup and a unit inverse Frisch elasticity, the critical constrained share is about 0.545.

Under the benchmark household-share mapping, the 2007 calibration lies clearly inside the inverted region and remains the more robust of the two cases across the exercises considered. However, it becomes marginal under the consumption weighting and can move into the standard region under sufficiently strict cutoffs. The 2014 calibration is more sensitive. It moves into the standard region under stricter cutoffs, higher markups, or a consumption weighting that gives less weight to large, low-spending households.

The policy analysis characterizes determinacy under alternative policy rules within the tractable model. In the reported IADL calibrations, an active forward-looking Taylor rule generates indeterminacy, whereas the reported passive responses to expected inflation deliver determinacy. An active rule can also deliver determinacy when its response to output is sufficiently strong. A history-dependent commitment targeting criterion delivers determinacy across all admissible constrained-household shares except at the singular regime boundary. The lagged output gap breaks the contemporaneous feedback underlying indeterminacy, and the targeting block remains valid under both standard and inverted aggregate-demand logic.

The model also identifies a structural margin that changes the demand regime itself. Following Bilbiie (2008), we consider a period-by-period balanced-budget redistribution of profit income. A fraction τ_D of aggregate profit income is assigned to constrained households, while Ricardian households retain the remainder. This lowers the elasticity of constrained consumption to aggregate output and, when the redistribution rate is sufficiently high, restores the standard aggregate-demand slope. Because steady-state profits are zero, the rule's first-order effect operates through profit deviations that can be positive or negative. It is therefore a local redistribution mechanism, not a conventional one-sided transfer program. Under the benchmark 2007 calibration, the redistribution rate required to restore the standard slope is about 31 percent.

The paper makes three contributions. First, it provides a microdata-disciplined expenditure proxy for budget tightness and maps it transparently into the model's constrained-household share under alternative correspondence assumptions. Second, it uses these conditional mappings in a two-agent New Keynesian model to characterize which aggregate-demand regimes are compatible with the Moroccan calibration. Third, it identifies which monetary-policy rules deliver determinacy in the inverted region and quantifies, for the Moroccan calibrations, the profit-redistribution rate required

to restore the standard aggregate-demand slope. To our knowledge, this is the first application to Morocco to combine nationally representative household expenditure microdata with a two-agent New Keynesian model to diagnose the aggregate-demand regime and examine its monetary-policy implications. We do not claim to identify liquid-asset hand-to-mouth households, and we do not propose an operational policy rule for Bank Al-Maghrib. The contribution is a regime diagnosis based on a transparent proxy and an explicit mapping.

The rest of the paper proceeds as follows. Section 2 places the paper in the literature. Section 3 describes the surveys and the construction of the proxy. Section 4 presents the empirical evidence and disciplines the cutoff. Section 5 maps the proxy into the model's constrained-household share. Section 6 develops the TANK model and derives the regime threshold. Section 7 locates Morocco relative to that threshold. Section 8 studies determinacy under alternative policy rules and the balanced-budget profit-redistribution rule. Section 9 reports robustness and limitations. Section 10 concludes.

2 Related Literature

Three literatures are directly relevant to this paper. The first studies households that do not fully smooth consumption. The second studies how such households alter monetary transmission in two-agent New Keynesian models. The third studies limited smoothing, informality, and financial inclusion in emerging economies. At their intersection, we discipline the constrained-household share with Moroccan household microdata and use it to diagnose the sign of aggregate demand in a tractable monetary model.

The first strand documents that many households consume current income more closely than permanent-income theory predicts. Campbell and Mankiw (1990) provided seminal time-series evidence that a sizable share of income accrues to households whose consumption tracks current income. Zeldes (1989) showed that low-wealth households display excess sensitivity of consumption to predictable income, consistent with binding liquidity constraints. Johnson et al. (2006) found large short-run consumption responses to the 2001 U.S. tax rebates, especially among low-income and low-liquidity households. Kaplan et al. (2014) sharpened the concept by showing that many households hold little liquid wealth even when their net worth is positive. This distinction is central

for our empirical interpretation. Because the Moroccan surveys used below do not report liquid assets, borrowing limits, or saving flows, we use food-expenditure shares as an expenditure-based proxy for short-run budget tightness, following the broader Engel-curve tradition associated with Deaton (1997) and Chai and Moneta (2010).

The second strand studies how constrained households affect monetary transmission. Galí et al. (2004) showed that rule-of-thumb consumers change the conditions under which interest-rate rules deliver a determinate equilibrium. Bilbiie (2008) provided the analytical mechanism on which this paper builds: once limited asset-market participation is sufficiently high, the elasticity of aggregate demand to the expected real interest rate changes sign, the Taylor principle is inverted, and active interest-rate rules can generate indeterminacy. Bilbiie (2008) also shows that a period-by-period balanced-budget tax on dividend income, redistributed to non-asset holders, can restore standard aggregate-demand logic. Bilbiie (2019) corrects the corresponding derivation when both household types have a common labor-supply elasticity, which is the case considered here. Buffie and Zanna (2018) show that open-economy channels can substantially alter the determinacy implications of limited asset-market participation. Subsequent work clarified the broader relevance of tractable two-agent frameworks. Bilbiie (2020) emphasized the New Keynesian cross and the role of the income elasticity of constrained households. Debortoli and Galí (2025) showed that suitably specified TANK models can approximate important aggregate implications of richer heterogeneous-agent New Keynesian (HANK) economies. Bilbiie (2025) extended the analytical framework to precautionary saving, illiquid assets, and realistic intertemporal marginal propensities to consume. Auclert (2019) made the complementary point that redistribution induced by interest-rate movements is itself a first-order transmission channel when households differ. Closest to our exercise, Bilbiie and Straub (2013) brought the limited-participation framework to U.S. data and argued that the pre-Volcker economy was consistent with the inverted region. Our contribution is to discipline the key participation parameter with local household microdata and to ask whether the resulting emerging-economy calibration lies near or inside the Bilbiie threshold.

The third strand concerns heterogeneity, limited smoothing, and monetary transmission in emerging economies. Aguiar and Gopinath (2007) documented that emerging-market cycles feature consumption that is more volatile than income, a pattern consistent with limited smoothing.

Gertler and Gruber (2002) showed that Indonesian households cannot fully insure consumption against health shocks. Mehrotra and Yetman (2015) argued that weak financial inclusion can weaken the intertemporal channel of monetary policy and increase the importance of income-based transmission. Closest to our setting, Takyi and León-González (2020) estimate a New Keynesian DSGE model with financially included and excluded households for four Sub-Saharan African economies. For Morocco, the Morocco Policy Analysis Model (MOPAM) includes liquidity-constrained households that consume current labor income, remittances, and net transfers (Achour et al., 2021). The model reports only their share of aggregate consumption, equal to 47.82 percent. We construct a microdata-based measure of budget-tight household prevalence and use it to evaluate the IADL regime. Batini et al. (2011) showed that informality can materially affect the design of monetary policy. Cross-country evidence also suggests that constrained-household shares are substantially higher in emerging economies than in advanced economies, with implications for fiscal and monetary transmission (Bracco et al., 2021). Less common is a structural monetary exercise that calibrates the constrained-household share from local microdata and asks, for a specific emerging economy, on which side of the aggregate-demand threshold that share places the economy.

3 Data and Proxy Construction

The central empirical input to the paper is the share of Moroccan households whose expenditure patterns are consistent with limited short-run smoothing capacity. This section describes the surveys from which that share is built, explains the construction of the proxy, and states its limitations.

Surveys.

We use two nationally representative household expenditure surveys conducted by Morocco's Haut-Commissariat au Plan (HCP). The 2007 National Survey on Household Living Standards covers 7,062 households and 36,333 individuals (Haut-Commissariat au Plan, 2007). The 2014 National Survey on Household Consumption and Expenditure covers 15,970 households and 75,691 individuals (Haut-Commissariat au Plan, 2014). Both are repeated cross-sections with detailed expenditure modules and probability sampling designs. We apply the survey weights throughout, so all point estimates are nationally representative.

The two surveys are distinct instruments: the 2007 data come from the Enquête Nationale sur les Niveaux de Vie des Ménages, whereas the 2014 data come from the Enquête Nationale sur la Consommation et les Dépenses des Ménages. Both surveys classify expenditure using the Nomenclature Analytique des Biens et Services, which is based on the Classification of Individual Consumption According to Purpose (COICOP). We use the version applicable to each survey year. To preserve comparability, we aggregate expenditure items to the same broad food, clothing, household-equipment, and total-consumption categories in both surveys. The public files also differ in their survey-provided upper-coded household-size variable—eight members in 2007 and six in 2014. We retain those variables in regressions using the survey size control, while actual household size is used for per-capita quantities and descriptive household-size statistics. Poverty and vulnerability indicators are available in the 2014 public file but not in the 2007 file used here, so they are used only for within-survey validation. We treat the surveys as independent cross-sections and interpret prevalence differences descriptively because detailed survey procedures may not be fully comparable.

Two data limitations shape the empirical strategy. First, the surveys do not contain the balance-sheet information needed to identify liquid-asset hand-to-mouth status: liquid assets, saving flows, borrowing constraints, and credit limits are not consistently observed. Asset-side identification of hand-to-mouth status, in the sense of Kaplan et al. (2014), is therefore not possible in these data. Second, the public microdata do not include the strata and primary-sampling-unit identifiers required for design-based variance estimation. We therefore apply sampling weights to all point estimates but report heteroskedasticity-robust HC1 standard errors for regression-based exercises. These standard errors provide descriptive rather than design-based precision.

Proxy construction.

We construct the food-expenditure share as the ratio of food outlays to total household consumption expenditure. The denominator is restricted to consumption proper: it excludes fiscal obligations, transfers to other households, and loan repayments, so that the share is not mechanically distorted by debt service or taxes that vary across households for reasons unrelated to current consumption needs. Under the benchmark classification, households with a food share at or above 40 percent are classified as budget-tight. The resulting HtM-proxy prevalence is used as the baseline empirical input to discipline the constrained-household share λ in the model.

The economic logic follows Engel's law. Food is among the least compressible categories in the household budget. A household that devotes a large fraction of its spending to food has, by construction, less discretionary margin left: when income falls or prices rise, it cannot easily cut subsistence expenditure, and, if buffers are limited, intertemporal smoothing is also harder. The food share is therefore a proxy for budget tightness, which may be associated with greater sensitivity of consumption to current income. It does not identify the composition of the household's balance sheet. This interpretation has a long pedigree in development economics, from Deaton (1997) to Chai and Moneta (2010), where the food share is used as a workhorse indicator of welfare and budget tightness in settings where balance-sheet data are unavailable.

We use 40 percent as a common cross-survey reference point. Section 4 disciplines it empirically through binned expenditure profiles, kink regressions, a systematic cutoff sweep, breakpoint estimation with bootstrap inference, and levels-based checks. Section 9 then traces how the structural conclusions move under alternative cutoffs.

Section 5 treats the mapping of this empirical classification to the model's constrained type as conditional.

4 Empirical Evidence

This section establishes the empirical foundation of the calibration. We examine prevalence across cutoffs, raw binned profiles, kink regressions around the common benchmark, a systematic cutoff sweep, breakpoint estimation with bootstrap inference, levels-based checks, and validation against independent welfare indicators.

4.1 HtM-proxy prevalence by cutoff

Table 1 reports the survey-weighted share of households classified as budget-tight across food-share cutoffs from 30 to 50 percent.

Table 1: HtM-proxy prevalence by food-share cutoff (survey-weighted, %)

Cutoff (%)	2007	2014	Difference (pp)
30	87.8	81.4	6.4
35	79.8	70.4	9.3
40	68.8	57.2	11.6
45	54.8	43.2	11.6
50	41.2	30.1	11.1
N (households)	7,062	15,970	—

Notes: Prevalence is survey-weighted; differences are calculated using unrounded values.

Prevalence is high across the full cutoff range: even under the strict 50 percent cutoff, more than four households in ten were budget-tight in 2007 and three in ten in 2014. Whatever the exact benchmark, Morocco's budget-tight share is large by the standards of conventional TANK calibrations for advanced economies. Prevalence is also lower in 2014 at every cutoff, by 6 to 12 percentage points overall and by approximately 9 to 12 percentage points over the 35–50 percent cutoff range. The consistency of the difference across cutoffs shows that it is not driven by the choice of a single threshold. Because the data come from distinct survey instruments, however, we interpret this pattern as a descriptive cross-survey difference rather than as a formally estimated trend. The unweighted shares, 70.2 and 57.9 percent at the benchmark, are close to the weighted ones, so the pattern is not driven by a small number of heavily weighted observations.

The qualitative finding of a large budget-tight group is therefore robust across cutoffs. The quantitative mapping into the model parameter λ is not: the regime classification in Section 7 depends on which cutoff is used. Section 9 traces this sensitivity explicitly.

4.2 Binned expenditure profiles

Before imposing any functional form, Figure 1 shows the raw relationship the proxy relies on. We compute the compressible-consumption share, clothing and household equipment as a fraction of total consumption, within narrow food-share bins, using survey weights.

The profiles are similar across the two surveys, seven years apart and with different sample sizes. Below a food share of roughly 40 percent, the compressible share is broadly flat, which is consistent with the idea that food intensity is not yet strongly associated with a shift in total-budget composition. Past the low 40s, the profile turns down, and the decline steepens in the upper tail.

This pattern is consistent with Engel-curve logic: as food absorbs more of the budget, the nonfood share of total expenditure contracts.

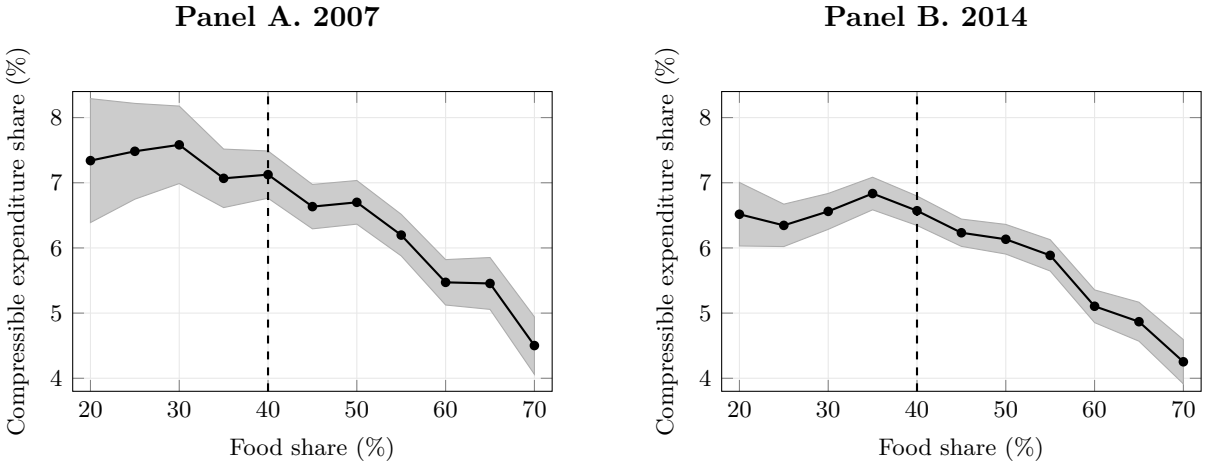


Figure 1: Compressible expenditure share by food-share bin, 2007 (Panel A) and 2014 (Panel B).
Notes: Survey-weighted bin means. Shaded areas show ±2 standard errors of the bin mean under a simple random sampling (SRS) approximation. They provide visual guidance only because design-based uncertainty cannot be computed from the public files. The dashed line marks the 40 percent benchmark.

4.3 Kink regressions around the common benchmark

The fixed-knot kink specification provides a descriptive benchmark around the common cutoff. Because the specification places knots at 40 and 60 percent, its role is to describe how compressible spending behaves around the common benchmark. The cutoff location itself is disciplined by the threshold sweep reported next.

We estimate, by survey-weighted least squares,

$$CS_i = \alpha + \beta_1 FS_i + \beta_2 \max(0, FS_i - 40) + \beta_3 \max(0, FS_i - 60) + X_i'\gamma + u_i, \tag{1}$$

Here, CS_i is household i's compressible-consumption share and $FS_i = 100 \times \mathit{FoodShare}_i$. Both variables are expressed in percentage points. The vector X_i contains controls for household size, rural residence, and household-head characteristics. The specification implies three linear segments: slope β_1 below 40 percent, $\beta_1 + \beta_2$ between 40 and 60, and $\beta_1 + \beta_2 + \beta_3$ above 60. Under the budget-tightness interpretation we expect the slope to be near zero in the first segment and negative afterwards.

Table 2: Kink regression: slopes of compressible share on food share (pp/pp)

	2007		2014	
Segment	Slope	(SE)	Slope	(SE)
$FS < 40\%$	0.002	0.017	0.026	0.010
$40\% \leq FS < 60\%$	−0.072	0.011	−0.076	0.007
$FS \geq 60\%$	−0.117	0.013	−0.120	0.010

Notes: Survey-weighted OLS with HC1 standard errors, which are descriptive rather than design-based (Section 3). Slopes measure the change in the compressible share (pp) per one-percentage-point increase in the food share.

The pattern is the same in both surveys and in the predicted direction. Below the benchmark, the slope is effectively zero in 2007 but small and positive in 2014. Past it, the compressible-expenditure share falls by about 0.07 percentage point per point of food intensity, and the decline steepens beyond 60 percent. Component-level estimates reported in Online Appendix B show that this decline is not driven by a single category: both clothing and household-equipment shares display statistically significant negative slopes above the benchmark in both surveys, with steeper share declines in the upper tail. The adding-up constraint implies that some nonfood share must fall as the food share rises. The separate declines in clothing and household equipment show that the combined result is not driven by a single category, but they do not by themselves establish selective compression relative to all other nonfood expenditure. The estimates locate a change in total-budget composition around the benchmark. Section 4.4 estimates its location. Section 4.6 examines the corresponding evidence in expenditure levels and within the nonfood budget.

4.4 Threshold sweep

We next sweep candidate cutoffs from 30 to 60 percent. For each candidate, we re-estimate the kink specification with the knot placed at that value and record two statistics: the pre-cutoff slope, which should be approximately flat if the candidate provides a plausible empirical onset of the change in total-budget composition, and the survey-weighted sum of squared errors. A candidate is admissible when its pre-cutoff slope is close to zero under a stated tolerance and its post-cutoff slope is negative; among admissible candidates, the selected cutoff minimizes the SSE. The complete cutoff sweep, including alternative near-flat tolerances and weighted SSE profiles, is reported in Online Appendix B.3.

Under the baseline tolerance, $\eta = 0.02$, the procedure selects a cutoff of 41 percent in 2007 and 43 percent in 2014. Under the tighter tolerance, the selected cutoffs are 41 percent in 2007 and 45 percent in 2014, while looser restrictions generate greater sensitivity, particularly in 2007. We retain 40 percent as a transparent common benchmark for cross-survey comparability and treat the year-specific cutoffs as sensitivity diagnostics. The corresponding regime implications are examined in Section 9.

4.5 Breakpoint estimation and bootstrap inference

The tolerance-based sweep in Section 4.4 identifies plausible cutoffs but does not quantify their uncertainty. We therefore estimate a theory-guided breakpoint model:

$$CS_i = \alpha + \beta \max(0, FS_i - c) + X_i'\gamma + u_i,$$

where the compressible-expenditure share is restricted to be flat below an unknown cutoff c and linear above it. The cutoff is selected over the 25–60 percent range, in 0.1-point increments, by minimizing the survey-weighted sum of squared residuals. Uncertainty is assessed using 5,000 HC2-adjusted Rademacher wild-bootstrap replications.

The estimated breakpoint is 41.1 percent in 2007 and 42.5 percent in 2014, close to the 41 and 43 percent cutoffs obtained from the tolerance-based sweep. The post-cutoff slopes are -0.0875 and -0.0869, respectively, and remain negative in every bootstrap replication. The corresponding bootstrap standard deviations, 0.0077 and 0.0057, exceed the conditional HC1 standard errors of 0.0061 and 0.0044 by roughly 27–28 percent, reflecting the additional uncertainty associated with estimating the cutoff. The 95 percent interval for the difference between the two breakpoints is $[-4.1, 9.7]$, so the data do not reject a common low-40s breakpoint in the compressible-share profile.

At the estimated cutoff, the proxy prevalence is 65.9 percent in 2007, above the model threshold $\lambda^* = 54.5$ percent, and 49.9 percent in 2014, below it. The sample prevalence evaluated at each bootstrap cutoff exceeds the model threshold in 95.6 percent of the draws for 2007 and 13.2 percent for 2014. In 83.0 percent of paired replications, 2007 lies above the threshold while 2014 lies below it. The exercise therefore reinforces the central asymmetry: 2007 is the main IADL calibration, whereas 2014 is predominantly SADL under a data-selected cutoff.

These frequencies measure uncertainty about the location of the breakpoint, not the probability that Morocco was in a particular regime. The zero pre-cutoff slope is also a maintained restriction. The estimated breakpoint locates a change in the compressible-share profile; it does not identify a structural hand-to-mouth threshold. We therefore retain the common 40 percent cutoff as a transparent cross-survey reference calibration and use the estimated breakpoints as data-driven sensitivity diagnostics. Full results are reported in Online Appendix B.

Table 3: Breakpoint estimates and bootstrap inference

Survey	$\widehat{c}$	Post-cutoff slope	Bootstrap SD	$\lambda(\widehat{c})$	90% interval	Share of draws with $\lambda(\widehat{c}) > \lambda^*$
2007	41.1	−0.0875 (0.0061)	0.0077	0.659	[35.4, 44.9]	0.956
2014	42.5	−0.0869 (0.0044)	0.0057	0.499	[39.8, 44.7]	0.132

Notes: Parentheses report HC1 standard errors conditional on the estimated cutoff. Bootstrap SD is calculated from 5,000 HC2-adjusted Rademacher wild-bootstrap replications. The pre-cutoff slope is restricted to zero. The 2014 SSE profile is flat near its minimum, but the baseline specification selects 42.5 percent; expanding the grid to 25–60 percent leaves both cutoff estimates unchanged. The final column reports the share of bootstrap cutoffs for which the sample prevalence exceeds $\lambda^* = 0.545$, not the probability that the economy was in the IADL regime.

4.6 Levels-based robustness

Share-on-share regressions invite a mechanical concern: both the dependent and the independent variable are ratios with the same denominator, so some negative association could arise by construction. The standard response is to move to levels. We regress log compressible expenditure per capita on the HtM-proxy indicator, conditional on log total consumption per capita and the same controls, at each cutoff. If the proxy captures genuine compression, classified households should spend less on discretionary items than unclassified households with the same total budget.

Table 4: Levels-based robustness: log compressible expenditure per capita

	2007			2014		
Food-share cutoff	β		(SE)	β		(SE)
35	0.063	*	0.035	0.195	***	0.025
40	0.016		0.029	0.079	***	0.021
45	−0.051	*	0.026	0.000		0.020
50	−0.071	***	0.026	−0.074	***	0.021

Notes: Survey-weighted OLS. The dependent variable is log compressible expenditure per capita. Controls include log total consumption per capita, actual household size, rural residence, and household-head characteristics. Per-capita variables are computed using actual household size. HC1 standard errors are reported in parentheses. *, **, and *** denote significance at the 10, 5, and 1 percent levels. The log specification uses households with positive, non-missing compressible expenditure ($N = 6{,}791$ in 2007 and $N = 15{,}644$ in 2014).

The levels-based evidence is weaker at the 40 percent benchmark. The share regressions and the cutoff sweep support a budget-tightness range around the high-30s to low-40s, but the non-ratio evidence is weakest exactly at the benchmark cutoff. At 40 percent, the coefficient is close to zero in 2007 and positive in 2014: conditional on total consumption, classified households do not yet spend less on discretionary items at that threshold. Compression in levels appears farther into the upper tail, turning negative at 45 percent in 2007 and at 50 percent in both surveys, broadly where the binned profiles steepen. A complementary specification expresses clothing and household equipment as a share of total nonfood expenditure. Their share rises with food intensity in both surveys. The nonfood share of total expenditure contracts, while these categories gain share within it. This result concerns budget composition. The levels estimates above provide evidence on expenditure amounts. Full results are reported in Online Appendix B.5. We therefore retain 40 percent as a transparent common benchmark across surveys and evaluate the macroeconomic implications directly under alternative cutoffs in Section 9. Residual binned profiles underlying this exercise are reported in Online Appendix B.

4.7 Socioeconomic and external validation

We next compare the classification with independent socioeconomic indicators. Within the surveys, budget-tight households display socioeconomic characteristics consistent with short-run budget tightness. They are more rural and larger than unclassified households, and their total consumption is substantially lower. In 2007, total consumption averages 44,687 Moroccan dirhams (MAD) among classified households against 81,565 MAD among others, while food expenditure differs much less. In 2014, their poverty rate is 4.18 percent against 1.43 percent for other households, and their vulnerability rate is 12.48 percent against 4.79 percent. Full socioeconomic profiles and the corresponding logit and probit estimates are reported in Online Appendix B.

External indicators provide further context. The 2022–2023 HCP living-standards survey[1] reports that only 4.2 percent of households were able to save over the preceding twelve months, while 7.6 percent borrowed to cover current expenditure (Haut-Commissariat au Plan, 2024). The same survey reports that the food-expenditure share increased from 37.0 percent in 2014 to 38.2 percent in 2022. Financial inclusion has improved but remains incomplete: Bank Al-Maghrib reports that the account-holding rate rose from 58 percent in 2024 to 62 percent in 2025 (Bank Al-Maghrib, 2026). As an independent modeling benchmark, an IMF staff model for Morocco calibrates the share of liquidity-constrained households at 0.65, defining them as households that consume all current disposable income (Taheri Sanjani et al., 2026). These indicators are consistent with limited short-run smoothing capacity, but none directly identifies liquid-asset hand-to-mouth status.

These results support interpreting the food-share classification as a proxy for budget tightness rather than as direct evidence of structural hand-to-mouth status. Section 5 then makes explicit how this empirical classification is mapped into the model parameter λ.

5 From the HtM Proxy to the Model Parameter λ

Section 4 produces an empirical classification of budget-tight households, whereas the model requires a structural constrained-household share, λ. The mapping between these two objects matters and, especially for 2014, affects the regime classification.

The mapping problem arises because the benchmark TANK model has a symmetric steady state in which the two household types have equal consumption and hours, whereas the Moroccan data do not. Budget-tight households are larger and consume substantially less than other households, a pattern that motivates reporting alternative empirical weightings. We use the household-weighted share λ_H as the benchmark structural mapping. Because λ is a population share in the symmetric model, the individual- and consumption-weighted measures are reported as sensitivity diagnostics rather than alternative structural estimates of λ. The individual-weighted share λ_I measures the fraction of individuals living in budget-tight households, while the consumption-weighted share λ_C measures their aggregate consumption footprint. Because classified households consume less, λ_C provides a conservative diagnostic of the regime classification.

[1]We use this survey only as external context because the household-level microdata required to reconstruct the classifier were not available to us.

Table 5 reports alternative aggregation weights and one- and two-sided proxy-to-structure correspondence exercises.

Table 5: Alternative empirical weightings and proxy-to-structure sensitivity

Panel A. Alternative aggregation weights

Year	Weighting	λ	$\phi^*(\lambda)$	Regime at $\phi = 1$
2007	Household	0.688	0.543	IADL
2007	Individual	0.716	0.477	IADL
2007	Consumption	0.548	0.992	IADL, near boundary
2014	Household	0.572	0.898	IADL, near boundary
2014	Individual	0.597	0.809	IADL
2014	Consumption	0.439	1.532	SADL

Panel B. One-sided proxy-to-structure correspondence

Correspondence rate ρ	2007 λ_ρ	2007 regime at $\phi = 1$	2014 λ_ρ	2014 regime at $\phi = 1$
100%	0.688	IADL	0.572	IADL, near boundary
90%	0.620	IADL	0.515	SADL
80%	0.551	IADL, near boundary	0.458	SADL

Panel C. Two-sided proxy-to-structure sensitivity

p_1	Minimum p_0 for IADL, 2007	Minimum p_0 for IADL, 2014
0.70	0.204	0.339
0.80	—	0.205
0.90	—	0.072
1.00	—	—

Notes: Panel A reports the household-, individual-, and consumption-weighted counterparts of λ. Household and individual shares are survey-weighted. The individual-weighted measure uses the household-level individual expansion weight supplied by the HCP, whereas regressions using household size rely on the survey-provided upper-coded size variable. The consumption-weighted measure uses adjusted total consumption expenditure. In Panel B, ρ is the assumed fraction of households classified by the 40 percent food-share proxy that correspond to the structural constrained type, so that $\lambda_\rho = \rho\lambda_H$; no constrained households are assigned below the cutoff. Regimes are evaluated at $\phi = 1$ and $\mu = 0.2$. IADL applies when $\phi > \phi^*(\lambda)$, where $\phi^*(\lambda) = (1 + \mu)(1 - \lambda)/\lambda$, and SADL when $\phi < \phi^*(\lambda)$. Structural calculations use unrounded shares; displayed values are rounded to three decimals.

In Panel C, $p_1 = \Pr(S = 1 \mid P = 1)$ is the correspondence rate among proxy-classified households and $p_0 = \Pr(S = 1 \mid P = 0)$ is the share of structurally constrained households below the proxy cutoff. Structural prevalence is $\lambda_y(p_1, p_0) = p_1\lambda_{H,y} + p_0(1 - \lambda_{H,y})$. Entries report the minimum p_0 required for $\lambda_y(p_1, p_0) > \lambda^* = 6/11$. A dash indicates that the calibration is already in IADL at $p_0 = 0$. These are sensitivity parameters, not estimated misclassification probabilities.

Panel A shows that the 2007 classification is less sensitive to aggregation than the 2014 classification. Under household and individual weighting, both years lie in IADL at the benchmark $\phi = 1$. Consumption weighting moves 2007 close to the boundary, with $\phi^* = 0.992$, and places 2014 in SADL, with $\phi^* = 1.532$. This difference reflects the fact that proxy-classified households are numerous but account for a smaller share of aggregate consumption. Consumption weighting

therefore highlights the tension between the symmetric model benchmark and the consumption heterogeneity observed in the data. A model with asymmetric steady-state consumption would address this tension structurally; we leave this extension open.

Panel B asks how strong the correspondence between proxy classification and structural constrained status must be. At the benchmark parameters, the critical correspondence rate, $\rho_y^* = \lambda^*/\lambda_{H,y}$, is 0.792 in 2007 and 0.954 in 2014. Thus, under the one-sided mapping, the 2007 calibration remains in IADL with correspondence somewhat below one-to-one, whereas the 2014 calibration requires almost complete correspondence. These are sensitivity parameters, not estimated classification probabilities.

Panel C allows structurally constrained households on both sides of the proxy cutoff. At $p_1 = 0.80$, the 2007 calibration is already in IADL at $p_0 = 0$, whereas 2014 requires $p_0 > 0.205$; at $p_1 = 0.90$, the corresponding 2014 requirement falls to $p_0 > 0.072$. The full two-dimensional region is reported in Online Appendix D. The central result is therefore not a unique estimate of structural λ, but an asymmetry: 2007 is compatible with IADL over a substantially broader set of proxy-to-structure mappings than 2014.

We therefore use the household-share mapping as the transparent benchmark implied by the symmetric model while treating all regime classifications as conditional on the empirical-to-structural mapping. Sections 7 and 9 examine the remaining sensitivity to the cutoff and structural parameters.

6 A Tractable TANK Model and the Regime Threshold

To interpret HtM-proxy prevalence, we need a framework in which the constrained-household share determines an object that matters for monetary policy. This section lays out that framework and derives its central result: a closed-form threshold for λ above which the slope of aggregate demand changes sign. The model is deliberately minimal: a closed economy with no capital, no banks, and no government purchases, public debt, or independent fiscal dynamics. This specification isolates how a large constrained share routes aggregate income through wages into current consumption. Key derivations are reported in Online Appendix A, with the full model system in Online Appendix C.

6.1 Environment

Time is discrete. A unit mass of households consists of a fraction $1-\lambda$ of Ricardian households and a fraction λ of HtM households. In the model, λ denotes the structural constrained-household share; its empirical mapping is discussed in Section 5. Both household types have preferences

$$U(C_t^j, H_t^j) = \frac{(C_t^j)^{1-\sigma}}{1-\sigma} - \kappa_L \frac{(H_t^j)^{1+\phi}}{1+\phi}, \qquad j \in \{o, r\},$$

with the consumption term interpreted as $\log C_t^j$ when $\sigma = 1$.

Ricardian households trade nominal bonds and receive the share $1-\tau_D$ of aggregate profits assigned to them. Their budget constraint is

$$C_t^o + \frac{B_t^o}{P_t} = W_t H_t^o + \frac{R_{t-1} B_{t-1}^o}{P_t} + \frac{1-\tau_D}{1-\lambda} D_t.$$

HtM households neither save nor borrow and consume current disposable income,

$$C_t^r = W_t H_t^r + \frac{\tau_D}{\lambda} D_t.$$

Both types satisfy the intratemporal labor-supply condition

$$\kappa_L (H_t^j)^{\phi} = W_t (C_t^j)^{-\sigma},$$

while only Ricardian households satisfy an Euler equation. The benchmark sets $\tau_D = 0$, so constrained households consume current labor income and monetary policy affects them through equilibrium wages and income.

Firms are monopolistically competitive and produce

$$Y_t = A_t H_t - F,$$

with elasticity of substitution ε, net steady-state markup $\mu = 1/(\varepsilon - 1)$, and a fixed cost F chosen to deliver zero steady-state profits. Prices are adjusted à la Rotemberg (1982). Around the benchmark

zero-inflation equilibrium with constant productivity, $\sigma = 1$, and $\tau_D = 0$, inflation satisfies

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa_y \hat{y}_t, \qquad \kappa_y = \frac{\varepsilon - 1}{\kappa_P} \left(1 + \frac{\phi}{1+\mu}\right).$$

Monetary policy follows the forward-looking rule

$$\hat{r}_t = \phi_\pi E_t \hat{\pi}_{t+1} + \phi_y \hat{y}_t + \varepsilon_t^m.$$

The benchmark admits a symmetric steady state in which the two household types have equal consumption and hours; this assumption underlies the household-share mapping of Section 5.

6.2 The income elasticity of constrained consumption

The mechanism runs through one number: how strongly HtM consumption responds to aggregate output. We derive it under $\sigma = 1$ and constant productivity; Online Appendix C reports the general case.

Combining the HtM labor-supply condition with the benchmark budget constraint $C_t^r = W_t H_t^r$ under log utility, the wage cancels out. The substitution effect of a higher wage on hours exactly offsets its income effect. HtM hours are therefore locally invariant, $\hat{h}_t^r = 0$, and HtM consumption moves one for one with the real wage, $\hat{c}_t^r = \hat{w}_t$. In the benchmark, the cycle affects constrained households through the wage.

The wage is pinned down by equilibrium. Around the zero-profit steady state, with constant productivity, labor-market clearing, and goods-market clearing, the real wage is a linear function of output:

$$\hat{w}_t = \left[1 + \frac{\phi}{1+\mu}\right] \hat{y}_t.$$

An expansion raises labor demand. The stiffer the labor-supply schedule, the more the wage must rise to clear the labor market. A higher markup reduces the output elasticity of the real wage in this fixed-cost, zero-profit calibration and therefore dampens the wage response. HtM consumption is therefore proportional to output,

$$\hat{c}_t^r = \chi \hat{y}_t,$$

with

$$\chi = 1 + \frac{\phi}{1+\mu}. \tag{2}$$

Under the benchmark $\phi = 1$ and $\mu = 0.2$, $\chi \approx 1.83$. Constrained consumption is therefore more procyclical than output itself. This amplification—wages reacting strongly to output and constrained households spending the wage-income movement immediately—is the engine of the regime change.

6.3 Aggregate demand and the regime threshold

Aggregate consumption is

$$\hat{c}_t = (1-\lambda)\hat{c}_t^o + \lambda \hat{c}_t^r,$$

and goods-market clearing requires $\hat{y}_t = \hat{c}_t$. Since HtM households contribute $\lambda\chi\hat{y}_t$, Ricardian consumption must satisfy

$$\hat{c}_t^o = \frac{1-\lambda\chi}{1-\lambda}\hat{y}_t. \tag{3}$$

This coefficient contains the regime switch. When $\lambda\chi < 1$, Ricardian consumption is procyclical and the aggregate-demand logic is standard. When $\lambda\chi > 1$, constrained households absorb more than the entire output movement through current-income spending, and equilibrium requires Ricardian consumption to move against output. Substituting (3) into the Ricardian Euler equation yields the aggregate IS equation,

$$\hat{y}_t = E_t\hat{y}_{t+1} - \delta(\hat{r}_t - E_t\hat{\pi}_{t+1}), \qquad \delta = \frac{1-\lambda}{1-\lambda\chi}. \tag{4}$$

In a representative-agent model, δ is positive: a higher expected real rate lowers demand. Here the denominator $1-\lambda\chi$ can turn negative, and the slope of aggregate demand changes sign.

Proposition 1 (Regime threshold)**.** *The economy is in the standard aggregate-demand region (SADL) if* $1-\lambda\chi > 0$ *and in the inverted aggregate-demand region (IADL) if* $1-\lambda\chi < 0$*. The critical constrained share is*

$$\lambda^* = \frac{1}{\chi} = \frac{1}{1+\phi/(1+\mu)}. \tag{5}$$

Equivalently, for a given λ, the critical inverse Frisch elasticity is

$$\phi^*(\lambda) = (1+\mu)\frac{1-\lambda}{\lambda}. \tag{6}$$

A higher ϕ strengthens the wage response to output, raises χ, and lowers λ^*: a less elastic labor supply makes the inversion easier to reach. A higher μ compresses the wage channel, lowers χ, and raises λ^*: market power pushes the inverted region away. Under the benchmark $\mu = 0.2$ and $\phi = 1$, the threshold is $\lambda^* \approx 0.545$, the value against which Section 5 placed the Moroccan mappings. For $\sigma \neq 1$, the income elasticity generalizes to

$$\chi(\sigma) = \frac{\phi+1}{\phi+\sigma} \times \left[\sigma + \frac{\phi}{1+\mu}\right],$$

with the full derivation reported in Online Appendix C and the robustness implications discussed in Section 9.

Below λ^*, monetary transmission works in the standard direction: the Ricardian intertemporal margin dominates, and a higher expected real rate lowers demand. Above λ^*, current-income effects dominate the Ricardian intertemporal margin. An expansion raises wages; a sufficiently large constrained share spends the wage-income movement immediately; and this income loop can overwhelm the interest-rate channel operating through the Ricardian minority. The sign of aggregate-demand transmission is therefore not a structural constant. It depends on the distribution of spending power across household types.

6.4 Calibration

Structural parameters follow standard quarterly values. We set $\beta = 0.99$, $\sigma = 1$, $\varepsilon = 6$, implying $\mu = 0.2$, and a Calvo (1983)-equivalent price rigidity $\theta = 0.75$, which maps into $\kappa_P \approx 58.25$ under Rotemberg pricing. The inverse Frisch elasticity $\phi = 1$ is the benchmark. Chetty et al. (2011) recommend a Frisch elasticity of aggregate hours around 0.75, with roughly 0.5 on the intensive margin, values that correspond to ϕ above one in our notation. The benchmark $\phi = 1$ is therefore conservative relative to calibrations with a stronger wage-income channel. Robustness over $\phi \in \{0.5, 1, 2, 3\}$ is reported in Section 9. The constrained share λ is the only parameter disciplined

by the Moroccan microdata: under the benchmark household-share mapping of Section 5, $\lambda = 0.688$ in 2007 and $\lambda = 0.572$ in 2014.

7 Regime Classification for Morocco

Under the benchmark household-share mapping and parameter values, $\mu = 0.2$ and $\phi = 1$, the regime threshold is $\lambda^* \approx 0.545$. The 2007 calibration, $\lambda = 0.688$, exceeds this threshold by a wide margin. The 2014 calibration, $\lambda = 0.572$, lies only slightly above it. Under the benchmark household-share mapping, this corresponds to $\delta = -1.19$ in 2007 and $\delta = -8.84$ in 2014. The much larger absolute value in 2014 reflects its proximity to the singular boundary $1 - \lambda\chi = 0$, implying substantially greater quantitative sensitivity to small changes in λ or the structural parameters.

The equivalent statement in terms of the labor-supply parameter is useful, because ϕ is calibrated rather than estimated. From Proposition 1, the critical inverse Frisch elasticities are $\phi^*_{2007} \approx 0.543$ and $\phi^*_{2014} \approx 0.898$. The 2007 calibration is in the inverted region for any ϕ above 0.543, equivalent to a Frisch elasticity below 1.84. This covers the benchmark calibration and most micro-oriented estimates of intensive-margin labor-supply elasticities. Moving the 2007 calibration back to the standard region would therefore require a highly elastic labor-supply calibration. By contrast, the 2014 calibration enters the inverted region only once ϕ exceeds 0.898, equivalent to a Frisch elasticity of about 1.1. This is close to the benchmark, which is why we treat 2014 as a near-boundary case rather than as a firm classification.

Figure 2 displays the diagnostic. The figure is a regime diagnostic under the household-share mapping, not an unconditional classification of the Moroccan economy.

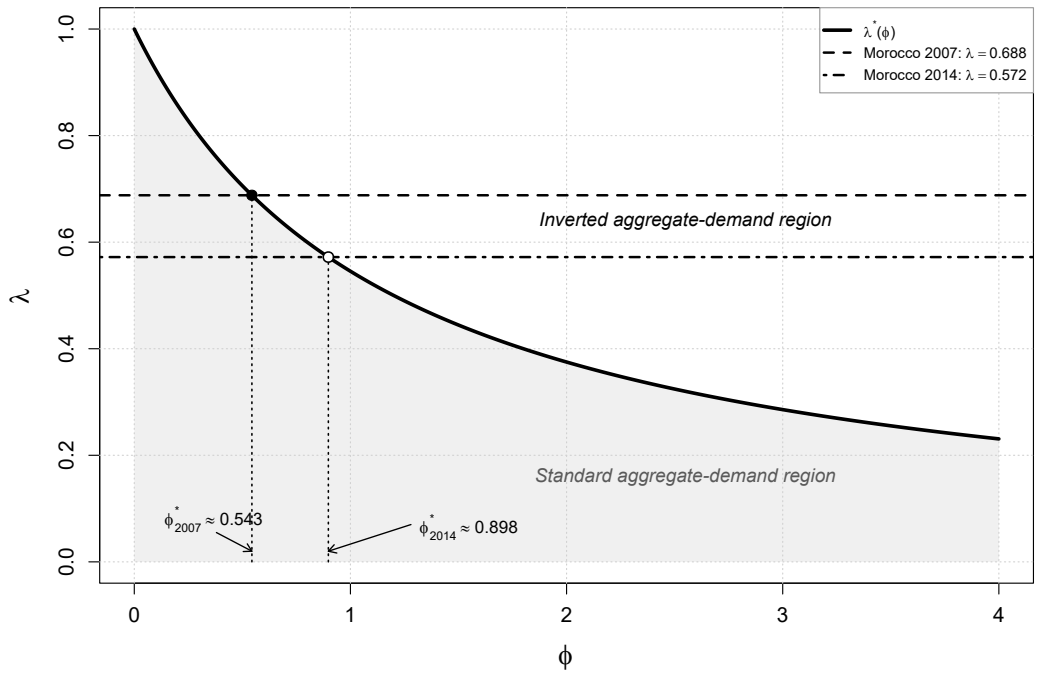


Figure 2: Regime threshold $\lambda^*(\phi)$ and the Moroccan calibrations.

Notes: The curve plots $\lambda^*(\phi) = 1/[1 + \phi/(1 + \mu)]$ under $\mu = 0.2$. Points above the curve are in the inverted aggregate-demand region. Horizontal lines mark $\lambda = 0.688$ for 2007 and $\lambda = 0.572$ for 2014. Vertical dotted lines mark the critical inverse Frisch elasticities $\phi^*_{2007} \approx 0.543$ and $\phi^*_{2014} \approx 0.898$.

The figure summarizes the asymmetry between the two survey years. The 2007 line remains above the threshold curve for most empirically plausible values of ϕ, including the benchmark. The 2014 line intersects the curve just below the benchmark, so modest changes in the labor-supply calibration can move it across the boundary.

8 Monetary Policy Determinacy and Balanced-Budget Profit Redistribution

The benchmark household-share calibration places the Moroccan economy near or inside the inverted region, raising the next question: what does the model imply for the design of interest-rate rules? The exercises below concern determinacy in the calibrated model, not operational policy design. Determinacy is evaluated using the Blanchard–Kahn (1980) condition.

8.1 The policy problem

In the inverted region, $\delta < 0$: a higher expected real rate raises aggregate demand. This turns the standard logic of active rules against itself. Consider a rule with $\phi_\pi > 1$ and suppose agents come to expect higher inflation for no fundamental reason. The rule raises the nominal rate more than one for one, so the expected real rate rises. In the standard region, the resulting fall in demand would invalidate the belief. Here, the higher real rate stimulates demand, output expands, and the Phillips curve delivers the inflation that was expected. The belief validates itself. Nothing in the rule prevents this, so the equilibrium is indeterminate: sunspot fluctuations become possible, and the configuration designed to anchor expectations amplifies them instead.

The standard prescription therefore need not stabilize the IADL regime. Within the calibrated system, determinacy can be restored either by adapting the rule to the inverted mechanism or by changing the distribution of profit income itself.

8.2 The passive forward-looking rule and the inverted Taylor principle

Result 1 (Inverted Taylor-principle pattern)**.** *In the benchmark IADL calibration, under the forward-looking rule with $\phi_y = 0$, the passive inflation responses $\phi_\pi \in \{0.5, 0.8\}$ satisfy the Blanchard–Kahn determinacy condition, whereas the active response $\phi_\pi = 1.5$ does not. Thus, in the reported calibration, determinacy is obtained under passive rather than active inflation responses.*

Under a passive rule, a rise in expected inflation raises the nominal rate less than one for one, so the expected real rate falls. In the inverted region, a falling real rate is contractionary: demand recedes, inflation falls back, and the initial belief is invalidated. The rule works with the inverted demand slope rather than against it.

In the reported calibrations, passive rules satisfy the Blanchard–Kahn count condition, but the critical roots lie close to the unit circle, with modulus around 1.005. The result is therefore a local, small-margin determinacy benchmark. Online Appendix C reports how these roots move under small perturbations of λ, κ_P, β, and ϕ_π.

8.3 Determinacy across rule configurations

Table 6 reports the Blanchard–Kahn verification for the 2007 benchmark across the main rule configurations. Full finite-root calculations for both survey calibrations are reported in Online Appendix C.

Table 6: Blanchard–Kahn determinacy and finite-root moduli

2007 calibration: $\lambda = 0.688$, $\phi = 1$, $\mu = 0.2$

Rule	ϕ_π	ϕ_y	Finite unstable roots	Finite-root moduli	BK status	Interpretation
Active forward-looking	1.5	0.0	1	0.740, 1.364	Indeterminate	Self-fulfilling loop open
Passive forward-looking	0.8	0.0	2	1.005, 1.005	Determinate	Local benchmark, small margin
Passive forward-looking	0.5	0.0	2	1.005, 1.005	Determinate	Local benchmark, small margin
Active + output response	1.5	0.5	1	0.354, 1.156	Indeterminate	Below output threshold
Active + output response	1.5	1.5	1	0.744, 1.065	Indeterminate	Below output threshold
Active + output response	1.5	2.0	2	1.051, 1.326	Determinate	Above output threshold

Notes: Reported counts refer to the two finite generalized eigenvalues associated with output and inflation. The full linearized model contains one additional infinite generalized eigenvalue generated by an algebraic relation. Because it carries no independent dynamics, it is excluded from the finite-root count. When evaluated at the same unrounded calibration, the reduced 2×2 system reproduces the two finite roots of the full model. Determinacy requires both finite roots to lie outside the unit circle. Root moduli are listed in ascending order. In indeterminate configurations, a modulus below one identifies the missing unstable root. Passive-rule determinacy is small-margin because the complex-conjugate pair has modulus $\beta^{-1/2} \simeq 1.005$. Structural calculations use the exact share $\lambda = 0.688339992675652$. Full root sets for both calibrations and perturbation checks are reported in Online Appendix C.

For $\phi_y = 0$, the determinant of the reduced forward-transition matrix is $1/\beta$, so whenever the two finite roots form a complex-conjugate pair their common modulus is $\beta^{-1/2} \simeq 1.005$. Online Appendix C.7 provides the analytical result and the associated perturbation checks. Corresponding impulse responses are reported in Online Appendix C.9.

8.4 The output-response threshold

An alternative to a passive inflation response is to retain $\phi_\pi > 1$ and respond directly to output. In the inverted aggregate-demand region, however, a positive output coefficient restores determinacy only if it exceeds a regime-dependent threshold.

Corollary 1 (Output-response threshold)**.** *Consider the benchmark forward-looking system in the inverted aggregate-demand region, with $\delta < 0$, $\phi_\pi > 1$, and $\phi_y \geq 0$. The determinacy boundary is*

$$\bar{\phi}_y = \frac{\kappa_y(\phi_\pi - 1)}{1+\beta} - \frac{2}{\delta}.$$

For $0 \leq \phi_y < \bar{\phi}_y$, the equilibrium is indeterminate; at $\phi_y = \bar{\phi}_y$, one root lies at -1; and for $\phi_y > \bar{\phi}_y$, the equilibrium is locally determinate. The derivation and root calculations are reported in Online Appendix C.8.

At $\phi_\pi = 1.5$, the threshold is $\bar{\phi}_{y,2007} = 1.72$ under the 2007 calibration and $\bar{\phi}_{y,2014} = 0.27$ under the 2014 calibration. The larger threshold in 2007 reflects its greater distance inside the inverted region, whereas the near-boundary 2014 calibration requires a modest output response to restore determinacy.

The threshold depends on the economy's distance from the regime boundary and therefore on the empirical mapping and structural calibration. This dependence is especially relevant given uncertainty in the constrained-household share and structural calibration.

8.5 Commitment targeting

In the targeting exercise, $\hat{y}_t$ denotes the output gap. Since productivity is fixed and natural output remains at steady state, the output gap coincides with the deviation of output from steady state. Following Bilbiie (2008), consider the history-dependent targeting criterion

$$\hat{\pi}_t + q(\hat{y}_t - \hat{y}_{t-1}) = 0, \qquad q > 0.$$

This is a targeting criterion rather than an instrument rule: the nominal interest rate adjusts endogenously to support the targeted paths of inflation and output. The lagged output gap introduces policy history dependence and breaks the contemporaneous feedback underlying indeterminacy under active forward-looking rules.

Online Appendix C.10 shows analytically that the criterion satisfies the Blanchard–Kahn condition for every $q > 0$, independently of λ, provided the economy is away from the singular boundary $1 - \lambda\chi = 0$. Thus, within the model, commitment targeting delivers determinacy without

requiring the policy authority to condition the rule on whether SADL or IADL prevails. The coefficient q is calibrated rather than derived from a welfare-based optimal-policy problem. The exercise therefore establishes determinacy, not welfare optimality.

8.6 Restoring the standard slope through balanced-budget profit redistribution

A different route changes the regime itself by modifying the income elasticity χ in equation (2): the inversion exists because $\lambda\chi > 1$, so anything that lowers χ raises the threshold λ^* and can pull the economy back into the standard region.

When $\tau_D = 0$, all profit income accrues to Ricardian households, while constrained households consume only labor income. We introduce a period-by-period balanced-budget redistribution of profit income. A fraction τ_D of aggregate profit income D_t is assigned to constrained households, while Ricardian households retain the remaining fraction $1 - \tau_D$. This rule follows the fiscal arrangement studied by Bilbiie (2008).

At the group level, constrained households receive $\tau_D D_t$, while Ricardian households receive $(1 - \tau_D)D_t$. Since the two amounts sum to D_t, the rule changes the distribution of profit income without changing aggregate resources. The budget constraint of each constrained household becomes

$$C_t^r = W_t H_t^r + \frac{\tau_D}{\lambda} D_t.$$

Since steady-state profits are zero, the rule leaves the steady state unchanged. At first order, it operates through the normalized profit deviation $d_t = D_t/\bar{Y}$. In the benchmark model, profit deviations are countercyclical. They fall when output rises and increase when output falls. Assigning part of this profit income to constrained households therefore makes their total income less procyclical. Because both household types have the same labor-supply elasticity in our model, the derivation below corresponds to the corrected case reported in Bilbiie (2019).

The resulting income elasticity is

$$\chi(\tau_D) = 1 + \frac{\phi}{1+\mu}\left(1 - \frac{\tau_D}{\lambda}\right), \tag{7}$$

which is decreasing in τ_D. The aggregate-demand elasticity is therefore

$$\delta(\tau_D) = \frac{1-\lambda}{1-\lambda\chi(\tau_D)}.$$

Proposition 2 (Profit-redistribution threshold)**.** *The balanced-budget redistribution rule restores the standard aggregate-demand slope, $\delta > 0$, if and only if $\tau_D > \tau_D^*$, where*

$$\tau_D^* = \lambda - \frac{(1+\mu)(1-\lambda)}{\phi}. \tag{8}$$

Under the 2007 benchmark calibration, $\tau_D^* \approx 0.314$; under the 2014 calibration, $\tau_D^* \approx 0.058$.

Table 7: Profit-redistribution threshold τ_D^* across calibrations

λ	ϕ	Regime under $\tau_D = 0$	Required redistribution rate (%)
0.688	0.5	SADL	None
0.688	1.0	IADL	31.4
0.688	2.0	IADL	50.1
0.688	3.0	IADL	56.4
0.572	0.5	SADL	None
0.572	1.0	IADL	5.8
0.572	2.0	IADL	31.5
0.572	3.0	IADL	40.1

Notes: $\tau_D^* = \lambda - [(1+\mu)(1-\lambda)/\phi]$, with $\mu = 0.2$. When $\tau_D^* < 0$, the economy is already in the standard region under $\tau_D = 0$. Because steady-state profits are zero, the rule's first-order allocation can be positive or negative. It should not be interpreted as a conventional one-sided transfer. Structural calculations use the unrounded prevalence shares. Displayed values of λ are rounded to three decimals, so recomputation from the displayed shares may differ in the third decimal.

Because steady-state profits are zero, the first-order allocation operates through $d_t = D_t/\bar{Y}$ and can be positive or negative. The analysis is local: sufficiently small deviations preserve positive consumption, but the exercise does not establish global feasibility. The threshold therefore applies to this cyclical profit-redistribution rule, not to conventional transfer, social-insurance, or debt-financed programs.

The required redistribution rate rises with ϕ because stiffer labor supply strengthens the wage-income loop that the redistribution must offset. It is also much smaller for 2014 than for 2007, 5.8 percent against 31.4 percent at $\phi = 1$, reflecting how close the 2014 calibration sits to the boundary.

9 Robustness and Limitations

The benchmark classification depends on the structural parameters ϕ, μ, and σ, as well as on the empirical cutoff and mapping used to construct λ. This section summarizes the sensitivity of the regime classification along these margins.

9.1 Structural and empirical sensitivity

The inverse Frisch elasticity is the main structural margin governing the strength of the wage-income channel. At the benchmark markup $\mu = 0.2$, the critical inverse Frisch elasticities are $\phi^*_{2007} = 0.543$ and $\phi^*_{2014} = 0.898$. The 2007 household-share calibration therefore remains in IADL for a broad range around the benchmark $\phi = 1$, whereas the 2014 calibration crosses the regime boundary with relatively small changes in the labor-supply calibration. The markup produces the same asymmetry. At $\phi = 1$, the 2007 calibration remains in IADL throughout the grid $\mu \in \{0.10, 0.20, 0.30, 0.50\}$, while the 2014 calibration becomes marginal near $\mu = 0.30$ and moves to SADL at $\mu = 0.50$.

The result is less sensitive to the intertemporal-elasticity parameter. Allowing $\sigma \neq 1$ changes the constrained-consumption elasticity and therefore the regime threshold, but both benchmark household-share calibrations remain in IADL over $\sigma \in \{0.5, 1, 2\}$. At the least favorable value, $\sigma = 0.5$, the threshold rises to 0.5625: the 2007 calibration retains a margin of about 0.126, whereas the 2014 margin falls below one percentage point. Thus, variation in σ reinforces rather than eliminates the distinction between a comparatively robust 2007 calibration and a near-boundary 2014 calibration.

Empirical cutoff sensitivity is more consequential. Under the baseline threshold-selection criterion, the selected cutoffs are 41 percent in 2007 and 43 percent in 2014. At $\phi = 1$, the corresponding household-share mappings place 2007 in IADL and 2014 in SADL. The formal breakpoint estimates of 41.1 and 42.5 percent produce the same qualitative asymmetry. Hence the common 40 percent cutoff is useful as a transparent cross-survey reference, but the data-driven cutoffs materially weaken the 2014 IADL classification while leaving the 2007 classification intact. Complete robustness grids for ϕ, μ, σ, and alternative food-share cutoffs are reported in Online Appendix Tables E.1–E.4. The robustness exercises vary key margins separately rather than jointly, so the regime classifications remain conditional.

9.2 Limitations

The framework is deliberately minimal. The model is a closed economy and therefore omits exchange-rate, net-export, and import-price pass-through channels that may modify the domestic inverted-demand mechanism and its determinacy conditions (Buffie and Zanna, 2018). It also abstracts from capital, banking and credit frictions, government purchases, public debt, and independent fiscal dynamics. Following Bilbiie (2008), the benchmark instrument rule excludes interest-rate smoothing; introducing policy inertia would add a predetermined state and require a separate determinacy analysis. The τ_D exercise should accordingly be interpreted as a period-by-period balanced-budget redistribution of profit income rather than as a general fiscal stabilization policy.

The surveys do not observe liquid assets, borrowing limits, or saving flows, so the empirical measure remains a proxy for budget tightness rather than a direct identification of liquid-asset hand-to-mouth status. Its mapping into the structural share λ is therefore conditional, as discussed in Section 5. The paper also does not identify Moroccan monetary-policy shocks or estimate the actual transmission of Bank Al-Maghrib's policy.

The model assumes a symmetric steady state. In the data, budget-tight households are larger and consume less than other households, whereas the benchmark model abstracts from these level differences. The alternative weightings in Section 5 make this tension explicit, but an asymmetric-steady-state TANK model would address it structurally; we leave that extension open. The results are calibrated regime and determinacy diagnostics for Morocco.

10 Conclusion

This paper asks whether household heterogeneity can change not only the strength of monetary transmission, but its sign. We combine two nationally representative Moroccan expenditure surveys with a tractable two-agent New Keynesian model to examine this possibility. Because the surveys do not observe liquid assets, borrowing limits, or saving flows, food-expenditure shares are used as a proxy for short-run budget tightness rather than as a direct measure of hand-to-mouth status. At the common 40 percent benchmark, 68.8 percent of households in 2007 and 57.2 percent in 2014 are classified as budget-tight.

The empirical-to-structural mapping produces a clear asymmetry across the two survey years.

Under the household-share mapping implied by the symmetric TANK benchmark and at $\phi = 1$, the 2007 calibration lies inside the inverted aggregate-demand region, whereas the 2014 calibration lies close to the regime boundary. This asymmetry persists across the broader sensitivity analysis, although neither classification is unconditional. At the estimated food-share breakpoints of 41.1 percent in 2007 and 42.5 percent in 2014, the corresponding calibrations lie in IADL and SADL, respectively. Alternative aggregation weights and proxy-to-structure mappings reinforce the same ranking: the one-sided correspondence threshold is 79.2 percent in 2007 but 95.4 percent in 2014, while the two-sided exercise shows that 2007 is compatible with IADL over a substantially broader mapping region. The data therefore discipline a set of conditional regime diagnoses rather than identify a unique structural constrained-household share.

The policy results are conditional determinacy results in the tractable benchmark model. In the reported IADL calibration, the passive inflation responses considered satisfy determinacy under the forward-looking instrument rule, although the relevant roots lie close to the unit circle. An active inflation rule can also become determinate when its output response exceeds the model-specific threshold, while the commitment targeting criterion remains determinate for $q > 0$ away from the singular regime boundary. Finally, reallocating a sufficiently large fraction of cyclical profit income toward constrained households lowers their income elasticity and can restore the standard aggregate-demand slope. Because steady-state profits are zero, this mechanism operates locally through profit deviations and should not be interpreted as a conventional one-sided transfer program.

In emerging economies where direct household balance-sheet measures are limited, the aggregate-demand regime should not be imposed through constrained-household calibrations imported mechanically from advanced economies. Household expenditure microdata can provide domestic discipline on budget tightness, provided that the mapping into structural household types is made explicit and its sensitivity is reported. For Morocco, that exercise points to a materially stronger case for inverted-demand logic in the 2007 calibration than in 2014. More generally, when household budget tightness is widespread, the sign of monetary transmission is an empirical and structural question rather than an assumption.

Data availability statement. The household microdata used in this paper are publicly available from Morocco's Haut-Commissariat au Plan. Replication code for the construction of the proxy, the tables and figures, and the model simulations will be made publicly available upon publication.

References

Achour, Aya, Aleš Bulíř, Omar Chafik, and Adam Remo, "The Morocco Policy Analysis Model: Theoretical Framework and Policy Scenarios," IMF Working Paper 2021/122, International Monetary Fund 2021.

Aguiar, Mark and Gita Gopinath, "Emerging market business cycles: The cycle is the trend," *Journal of Political Economy*, 2007, *115* (1), 69–102.

Auclert, Adrien, "Monetary policy and the redistribution channel," *American Economic Review*, 2019, *109* (6), 2333–2367.

Bank Al-Maghrib, "Statistiques sur les comptes bancaires 2025," Technical Report, Bank Al-Maghrib, Rabat 2026. URL: `https://www.bkam.ma/content/download/846906/9158673/TDB%20Stat%20bancaires%202025.pdf`.

Batini, Nicoletta, Paul Levine, Emanuela Lotti, and Bo Yang, "Informality, frictions and monetary policy," Discussion Paper 0711, University of Surrey School of Economics 2011.

Bilbiie, Florin O., "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic," *Journal of Economic Theory*, 2008, *140* (1), 162–196.

———, "Corrigendum to 'Limited asset markets participation, monetary policy and (inverted) aggregate demand logic' [J. Econ. Theory 140 (1) (2008) 162–196]," *Journal of Economic Theory*, 2019, *181*, 421–422. doi: `10.1016/j.jet.2019.03.008`.

———, "The New Keynesian cross," *Journal of Monetary Economics*, 2020, *114*, 90–108.

———, "Monetary policy and heterogeneity: An analytical framework," *The Review of Economic Studies*, 2025, *92* (4), 2398–2436.

Bilbiie, Florin O. and Roland Straub, "Asset market participation, monetary policy rules, and the Great Inflation," *The Review of Economics and Statistics*, 2013, *95* (2), 377–392.

Blanchard, Olivier Jean and Charles M. Kahn, "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 1980, *48* (5), 1305–1311.

Bracco, Jessica Roxana, Luciana Maria Galeano, Pedro Francisco Juarros, Daniel Riera-Crichton, and Guillermo Javier Vuletin, "Social transfer multipliers in developed and emerging countries: The role of hand-to-mouth consumers," Policy Research Working Paper 9627, World Bank 2021.

Buffie, Edward F. and Luis-Felipe Zanna, "Limited asset market participation and determinacy in the open economy," *Macroeconomic Dynamics*, 2018, *22* (8), 1937–1977.

Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

Campbell, John and N. Gregory Mankiw, "Permanent income, current income, and consumption," *Journal of Business & Economic Statistics*, 1990, *8* (3), 265–279.

Chai, Andreas and Alessio Moneta, "Retrospectives: Engel curves," *Journal of Economic Perspectives*, 2010, *24* (1), 225–240.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins," *American Economic Review*, 2011, *101* (3), 471–475.

Deaton, Angus, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, Baltimore: Johns Hopkins University Press, 1997.

Debortoli, Davide and Jordi Galí, "Heterogeneity and aggregate fluctuations: Insights from TANK models," *NBER Macroeconomics Annual*, 2025, *39* (1), 307–357.

Galí, Jordi, J. David López-Salido, and Javier Vallés, "Rule-of-thumb consumers and the design of interest rate rules," *Journal of Money, Credit and Banking*, 2004, *36* (4), 739–763.

Gertler, Paul and Jonathan Gruber, "Insuring consumption against illness," *American Economic Review*, 2002, *92* (1), 51–70.

Haut-Commissariat au Plan, "Enquête nationale sur les niveaux de vie des ménages 2007," Technical Report, Haut-Commissariat au Plan, Rabat 2007.

———, "Enquête nationale sur la consommation et les dépenses des ménages 2014," Technical Report, Haut-Commissariat au Plan, Rabat 2014.

———, "Evolution du niveau de vie de la population à la lumière des résultats de l'enquête nationale sur le niveau de vie des ménages de 2022," Technical Report, Haut-Commissariat au Plan, Rabat 2024. URL: https://www.hcp.ma/file/241059/.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner, "The wealthy hand-to-mouth," *Brookings Papers on Economic Activity*, 2014, pp. 77–138.

Mehrotra, Aaron and James Yetman, "Financial inclusion—issues for central banks," *BIS Quarterly Review*, 2015.

Rotemberg, Julio J., "Sticky Prices in the United States," *Journal of Political Economy*, 1982, *90* (6), 1187–1211.

Taheri Sanjani, Marzie, Kassia Antoine, and Pedro Rodriguez, "Macroeconomic impact of accelerating public infrastructure investments in Morocco," IMF Selected Issues Paper SIP/2026/026, International Monetary Fund, Washington, DC 2026. URL: https://www.imf.org/-/media/files/publications/selected-issues-papers/2026/english/sipea2026026.pdf.

Takyi, P. O. and R. León-González, "Monetary policy and financial exclusion in an estimated DSGE model of Sub-Saharan African economies," *International Economic Journal*, 2020, *34* (2), 317–346. doi: 10.1080/10168737.2020.1729835.

Zeldes, Stephen P., "Consumption and liquidity constraints: An empirical investigation," *Journal of Political Economy*, 1989, *97* (2), 305–346.

Online Appendix to:

"Too Many Hand-to-Mouth Households? Monetary Policy and the Risk of Inverted Aggregate Demand in Emerging Economies"

Badr Elhamzaoui, Faïçal Lakhchen, and Radouane Raouf

August 2026

This Online Appendix provides the supplementary empirical results, full model equations, determinacy calculations, mapping details, and robustness exercises referenced in the main manuscript. It is intended to document the construction of the expenditure-based proxy and to make the model-based regime and determinacy results auditable.

Table of Contents

Appendix A Aggregate-Demand Elasticity and Profit-Redistribution Threshold . 2
Appendix B Empirical Construction and Validation 3
Appendix C Model, Determinacy, and Additional Dynamics 11
Appendix D Empirical-to-Structural Mapping and Calibration 26
Appendix E Additional Robustness 29
References 32

Appendix A Derivation of the Aggregate-Demand Elasticity and the Profit-Redistribution Threshold

We derive the results in Propositions 1 and 2 under $\sigma = 1$ and constant productivity ($A_t = \bar{A}$). All derivations are first-order log-linear approximations around the zero-inflation steady state. Because steady-state profits are zero, profit deviations are represented as normalized level deviations $d_t = D_t/\bar{Y}$ rather than log deviations.

Step 1: Profit deviation. From the production function and aggregate labor-market clearing,

$$d_t = -\frac{1+\phi}{1+\mu}\hat{y}_t.$$

Profits are countercyclical: a rise in output compresses margins because wages rise more than proportionally to output under the fixed-cost structure.

Step 2: HtM income elasticity. The HtM labor-supply condition is $\phi\hat{h}_t^r = \hat{w}_t - \hat{c}_t^r$. The HtM budget constraint is $\hat{c}_t^r = \hat{w}_t + \hat{h}_t^r + (\tau_D/\lambda)d_t$. Combining gives $(1+\phi)\hat{h}_t^r = -(\tau_D/\lambda)d_t$. Using Step 1 and substituting back into the budget constraint yields $\hat{c}_t^r = \chi(\tau_D)\hat{y}_t$ with

$$\chi(\tau_D) = 1 + \frac{\phi}{1+\mu}\left(1 - \frac{\tau_D}{\lambda}\right),$$

which is equation (7). Setting $\tau_D = 0$ recovers (2).

Step 3: Aggregate IS. From consumption aggregation ($\hat{c}_t = (1-\lambda)\hat{c}_t^o + \lambda\hat{c}_t^r$) and goods-market clearing ($\hat{y}_t = \hat{c}_t$), Ricardian consumption satisfies $\hat{c}_t^o = [(1-\lambda\chi(\tau_D))/(1-\lambda)]\hat{y}_t$. Substituting into the Ricardian Euler equation gives equation (4) with $\delta(\tau_D) = (1-\lambda)/(1-\lambda\chi(\tau_D))$.

Step 4: Thresholds. Setting $\lambda\chi(\tau_D) = 1$ and solving for τ_D gives equation (8). Setting $\tau_D = 0$ in the condition $\chi = 1/\lambda$ and solving for λ gives the regime threshold (5).

Appendix B Empirical Construction and Validation

B.1 Variable construction

For household i, the food share is constructed as

$$\text{FoodShare}_i = \frac{\text{FoodExpenditure}_i}{\text{TotalConsumption}_i}. \tag{B.1}$$

Total consumption excludes fiscal obligations, transfers to other households, and loan repayments. For a cutoff c,

$$\text{HtM}_i(c) = \mathbf{1}\{\text{FoodShare}_i \geq c\}. \tag{B.2}$$

The benchmark uses $c = 0.40$. Compressible expenditure includes clothing and household equipment. Expenditure items are aggregated to the same broad categories in both surveys. Both files use the Nomenclature Analytique des Biens et Services, a COICOP-based classification, in the version applicable to each survey year. Survey weights are used throughout. HC1 standard errors are reported because strata and primary-sampling-unit identifiers are unavailable. Except in the levels-based exercise, regressions use the survey-provided household-size variable, upper-coded at eight members in 2007 and six members in 2014. Actual household size is used for per-capita expenditure variables and descriptive household-size statistics. The individual-weighted measure uses the household-level individual expansion weight supplied by the HCP. Household size does not enter the food-share measure or the household-weighted prevalence.

B.2 Exact prevalence

Percentages are rounded to one decimal place and benchmark model shares to three decimals; benchmark structural calculations use the unrounded values reported in Table B.1.

Table B.1. Exact benchmark HtM-proxy prevalence and display convention

Year	N	Exact weighted share λ	Displayed prevalence (%)	Displayed λ
2007	7,062	0.688339992675652	68.8	0.688
2014	15,970	0.571879168720315	57.2	0.572

B.3 Cutoff sweep and SSE profiles

Candidate cutoffs range from 30 to 60 percent. For each cutoff, we estimate the two-segment kink specification and compute the survey-weighted sum of squared residuals (SSE). A cutoff is admissible when

$$|\text{pre-slope}| \leq \eta$$

and the post-cutoff slope is negative. Among admissible candidates, the selected cutoff minimizes the survey-weighted SSE.

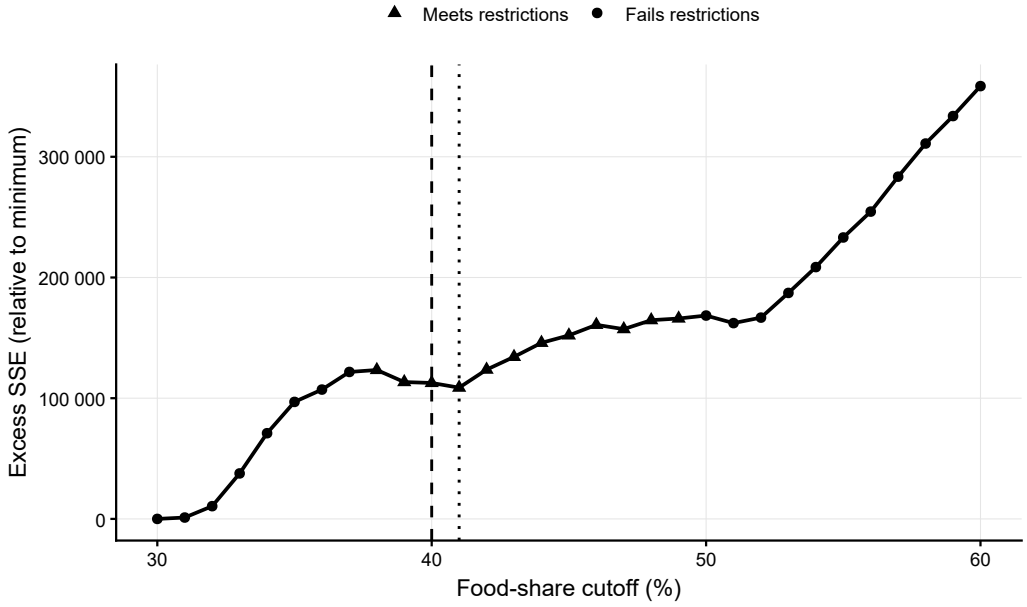


Figure B.1. Weighted cutoff-SSE profile, 2007

Notes: Points show weighted SSE, plotted as excess SSE relative to the minimum. Triangles indicate cutoffs satisfying $|\text{pre-slope}| \leq 0.02$ and post-slope < 0. The dashed line marks the common 40% benchmark; the dotted line marks the best admissible cutoff, 41%.

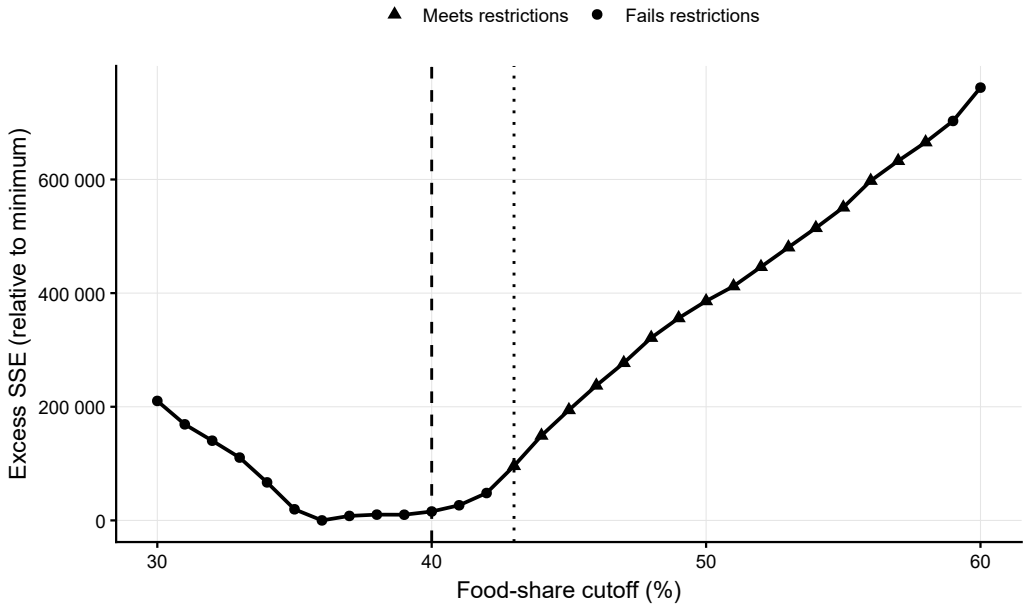


Figure B.2. Weighted cutoff-SSE profile, 2014

Notes: Points show weighted SSE, plotted as excess SSE relative to the minimum. Triangles indicate cutoffs satisfying $|\text{pre-slope}| \leq 0.02$ and post-slope < 0. The dashed line marks the common 40% benchmark; the dotted line marks the best admissible cutoff, 43%.

Table B.2. Preferred admissible cutoff under alternative near-flat tolerances

Year	Tolerance |pre-slope| ≤	Number of admissible cutoffs	Best admissible cutoff (%)	Pre-slope	Post-slope
2007	0.01	6	41	0.004	−0.088
2007	0.02	12	41	0.004	−0.088
2007	0.03	19	36	0.028	−0.081
2014	0.01	8	45	0.010	−0.100
2014	0.02	16	43	0.017	−0.096
2014	0.03	20	41	0.026	−0.091

Notes: η denotes the near-flat tolerance. Candidate cutoffs range from 30% to 60%. The best admissible cutoff minimizes the survey-weighted SSE among candidates satisfying $|\text{pre-slope}| \leq \eta$ and post-slope < 0. The baseline tolerance is $\eta = 0.02$.

Under the baseline tolerance ($\eta = 0.02$), the best admissible cutoff is 41 percent in 2007 and 43 percent in 2014. Under the tighter tolerance ($\eta = 0.01$), the corresponding values are 41 and 45 percent. The common 40 percent cutoff is retained as the benchmark to ensure comparability across the 2007 and 2014 surveys. Candidate-level results are provided with the replication files.

B.4 Share-based regressions

Table B.3 reports the coefficient on $\text{HtM}_i(c)$ from survey-weighted compressible-share regressions controlling for household-head characteristics, rural residence, and household size.

Table B.3. HtM-proxy coefficient in compressible-share regressions across cutoffs

Food-share cutoff (%)	2007 β (SE)	2014 β (SE)
35	-1.048*** (0.213)	-0.413*** (0.111)
40	-1.060*** (0.169)	-0.742*** (0.095)
45	-1.244*** (0.143)	-0.994*** (0.089)
50	-1.302*** (0.133)	-1.180*** (0.089)

Notes: Survey-weighted OLS. The dependent variable is the compressible-consumption share in percent, so coefficients are in percentage points. HtM proxy $= 1\{\text{FoodShare} \geq c\}$. HC1 standard errors are reported in parentheses; these are not design-based standard errors.

To verify that the fixed-knot result reported in Table 2 of the main text is not driven by a single component, we re-estimate the Section 4.3 specification separately for clothing and household equipment.

Table B.4. Component decomposition of fixed-knot slopes

Food-share segment	Clothing 2007	Equipment 2007	Clothing 2014	Equipment 2014
$FS < 40\%$	+0.008 (0.011)	-0.006 (0.012)	+0.020 (0.005)	+0.006 (0.008)
$40\% \leq FS < 60\%$	-0.041 (0.008)	-0.031 (0.007)	-0.056 (0.005)	-0.021 (0.005)
$FS \geq 60\%$	-0.076 (0.009)	-0.040 (0.009)	-0.068 (0.007)	-0.052 (0.006)

Notes: Survey-weighted OLS using the same sample, weights, knots, and household controls as Table 2 of the main text; household size is upper-coded at eight members in 2007 and six members in 2014. Parentheses report HC1 standard errors. Slopes are expressed in percentage points per one-percentage-point increase in the food share. Using unrounded estimates, the two component slopes sum exactly to the corresponding combined slope in Table 2; displayed sums may differ by 0.001 because of rounding.

The post-40 decline appears separately in both components and in both surveys. Clothing accounts for more of the initial adjustment between 40 and 60 percent, while household equipment also contracts and contributes materially in the upper tail. The combined compressible-expenditure result is therefore not driven by a single expenditure category.

B.5 Levels-based robustness

To address common-denominator concerns, we estimate

$$\log(CExp_i) = \alpha + \beta \mathrm{HtM}_i(c) + \gamma \log(TC_i) + X_i'\delta + u_i, \tag{B.3}$$

Here $CExp_i$ and TC_i denote compressible and total consumption expenditure per capita, respectively, both constructed using actual household size.

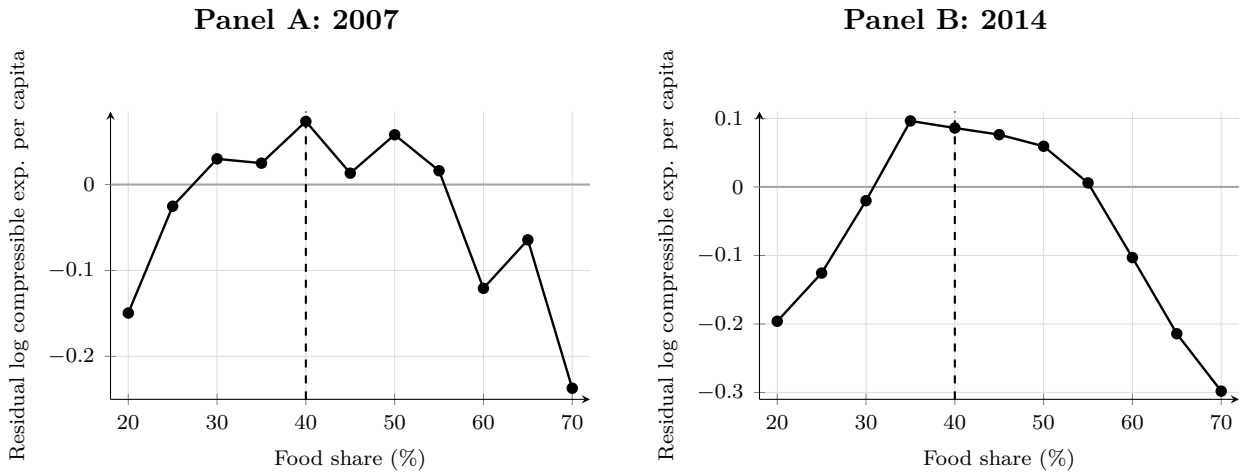


Figure B.3. Residual log compressible expenditure per capita by food-share bin

Notes: Panels A and B report 2007 and 2014. Points are survey-weighted bin means of residual log compressible expenditure per capita after partialling out log total consumption per capita and the controls in equation (B.3), excluding HtM status. The dashed line marks the 40 percent benchmark.

Table B.5. Levels-based robustness across food-share cutoffs

Food-share cutoff (%)	2007 β (SE)	2014 β (SE)
35	+0.063* (0.035)	+0.195*** (0.025)
40	+0.016 (0.029)	+0.079*** (0.021)
45	-0.051* (0.026)	0.000 (0.020)
50	-0.071*** (0.026)	-0.074*** (0.021)

Notes: Survey-weighted OLS. The dependent variable is log compressible expenditure per capita. HtM proxy $= 1\{\text{FoodShare} \geq c\}$. Per-capita expenditure variables and the household-size control use actual household size. HC1 standard errors are reported in parentheses; these are not design-based standard errors. The log specification uses households with positive, non-missing compressible expenditure ($N = 6{,}791$ in 2007 and $N = 15{,}644$ in 2014).

Negative conditional differences emerge only at stricter cutoffs, especially at 50 percent.

Within-nonfood composition check. We re-estimate the fixed-knot specification with clothing and household equipment expressed as a share of total nonfood expenditure, $\left(100C_i/(T_i - F_i)\right)$. This specification separates changes in the nonfood share of total expenditure from changes in the composition of nonfood expenditure. It uses the same knots, survey weights, household controls, and HC1 standard errors as Table B.4.

Table B.5A. Within-nonfood composition check: fixed-knot slopes

Food-share segment	2007	2014
$FS < 40\%$	+0.130 (0.024)***	+0.152 (0.013)***
$40\% \leq FS < 60\%$	+0.135 (0.022)***	+0.109 (0.014)***
$FS \geq 60\%$	+0.087 (0.044)**	+0.084 (0.034)**
Joint test of equal slopes, p-value	0.634	0.023
N	7,062	15,970

Notes: Survey-weighted OLS using the same knots, survey weights, and household controls as Table B.4. The dependent variable is clothing and household equipment as a percentage of total nonfood expenditure. Slopes measure the percentage-point change in this within-nonfood share associated with a one-percentage-point increase in the food share. Parentheses report HC1 standard errors. ** and *** denote significance at the 5 and 1 percent levels. The joint test reports the p-value for H_0: both knot coefficients equal zero. The specification uses the full household samples.

The slopes are positive and statistically significant in every segment of both surveys. Equality of the segment slopes is not rejected in 2007, with a p-value of 0.634, and is rejected in 2014, with a p-value of 0.023. All slopes remain positive. Clothing and household equipment therefore gain share within nonfood expenditure as food intensity rises. Their decline in the total budget reflects

the contraction of the nonfood share, partly offset by reallocation toward these categories. This result concerns composition. The levels-based estimates above provide evidence on expenditure amounts.

B.6 Socioeconomic profiles

Table B.6 reports survey-weighted descriptive statistics for the full sample and by benchmark HtM-proxy status.

Table B.6. Survey-weighted socioeconomic profiles by benchmark HtM-proxy status

Indicator	2007						2014					
	ALL		HtM		N-HtM		ALL		HtM		N-HtM	
	MEAN	SD	MEAN	SD	MEAN	SD	MEAN	SD	MEAN	SD	MEAN	SD
Household characteristics												
Age	50.23	13.85	50.33	13.91	50.01	13.74	52.35	14.00	52.03	14.10	52.78	13.86
Male household head (%)	83.44	37.18	83.75	36.89	82.75	37.78	82.45	38.04	83.83	36.82	80.61	39.54
Married household head (%)	83.64	37.00	84.46	36.23	81.82	38.57	81.46	38.86	82.62	37.89	79.91	40.07
Rural residence (%)	37.62	48.44	45.72	49.82	19.72	39.79	34.55	47.55	47.22	49.92	17.63	38.10
Household size	5.11	2.39	5.29	2.43	4.72	2.24	4.74	2.19	4.91	2.25	4.51	2.08
Consumption												
Total consumption (MAD)	56,180	52,908	44,687	27,876	81,565	79,557	73,614	59,155	56,538	31,290	96,424	77,176
Food expenditure (MAD)	23,347	15,348	23,700	14,151	22,566	17,683	28,071	16,600	29,178	15,669	26,592	17,660
Socioeconomic indicators												
Poverty rate (%)	-	-	-	-	-	-	3.00	-	4.18	-	1.43	-
Vulnerability rate (%)	-	-	-	-	-	-	9.18	-	12.48	-	4.79	-
N	7,062		4,957		2,105		15,970		9,239		6,731	

Notes: Columns report survey-weighted means and standard deviations for the full sample, HtM-proxy households, and non-HtM-proxy households. The HtM proxy is $1\{\text{FoodShare} \geq 40\%\}$. Monetary values are annual household expenditures in MAD. Poverty and vulnerability rates are available only in 2014. N denotes the unweighted number of sampled households. Differences are descriptive and should not be interpreted causally.

B.7 Poverty and vulnerability associations

Using the 2014 survey, we estimate survey-weighted logit and probit models of poverty and vulnerability at food-share cutoffs $c \in \{0.35, 0.40, 0.45, 0.50\}$. All specifications control for rural residence, household size, and household-head age. Sampling weights are normalized to the estimation-sample size, and HC1 standard errors are reported. Table B.6 presents unconditional weighted differences, whereas Tables B.7 and B.8 report covariate-adjusted associations.

Table B.7. Discrete changes in predicted probabilities, logit and probit

Cutoff (%)	Poverty Logit Δp (pp)	Poverty Probit Δp (pp)	Vulnerability Logit Δp (pp)	Vulnerability Probit Δp (pp)
35	+0.580	+0.675	+3.986	+4.424
40	+0.537	+0.657	+3.702	+4.170
45	+0.374	+0.482	+3.300	+3.784
50	+0.389	+0.508	+3.049	+3.563

Notes: Entries report discrete changes in predicted probabilities, in percentage points, when HtM-proxy status changes from zero to one, with rural residence, household size, and household-head age held at their survey-weighted means. These are not average marginal effects. Estimates are obtained from survey-weighted logit and probit models using normalized sampling weights and should be interpreted as descriptive, non-causal associations.

Table B.8. Logit and probit coefficients for poverty and vulnerability

Outcome	Model	35%	40%	45%	50%
Poverty	Logit	0.604***	0.513***	0.333***	0.331***
		(0.155)	(0.135)	(0.114)	(0.110)
Poverty	Probit	0.255***	0.229***	0.156***	0.156***
		(0.066)	(0.059)	(0.052)	(0.052)
Vulnerability	Logit	0.880***	0.727***	0.591***	0.513***
		(0.091)	(0.075)	(0.069)	(0.068)
Vulnerability	Probit	0.435***	0.369***	0.309***	0.274***
		(0.044)	(0.037)	(0.035)	(0.036)

Notes: The table reports coefficients on the HtM-proxy indicator from survey-weighted logit and probit models using normalized sampling weights. All specifications control for rural residence, household size, and household-head age. HC1 standard errors are reported in parentheses; these are not design-based standard errors. The dependent variables are poverty and vulnerability indicators in the 2014 survey. HtM-proxy status is defined separately at each food-share cutoff. The number of observations is 15,970 in all specifications. Significance levels are denoted by * ($p < 0.10$), ** ($p < 0.05$), and *** ($p < 0.01$).

Across cutoffs, HtM-proxy status is positively associated with poverty and vulnerability; these estimates are descriptive and non-causal.

B.8 Breakpoint estimation and bootstrap inference

The tolerance-based cutoff sweep in Section B.3 identifies admissible cutoffs but does not quantify uncertainty about the cutoff location. We therefore estimate the restricted breakpoint model

$$CS_i = \alpha + \beta \max(0, FS_i - c) + X_i'\gamma + u_i, \tag{B.4}$$

where CS_i is the compressible-expenditure share, FS_i is the food share, and the slope below the unknown cutoff c is restricted to zero. The control vector X_i includes upper-coded household size, rural residence, household-head age, a male-head indicator, and a married-head indicator, consistently with the share regressions described in Section B.1. The cutoff is selected over the 25–60 percent range, in 0.1-point increments, by minimizing the survey-weighted sum of squared residuals.

Conditional HC1 standard errors treat the selected cutoff as fixed. To incorporate cutoff-selection uncertainty, we use 5,000 HC2-adjusted Rademacher wild-bootstrap replications. Let $\widehat{CS}_i$ denote the fitted value from the restricted model evaluated at the selected cutoff $\hat{c}$, and let $\hat{u}_i = CS_i - \widehat{CS}_i$. For each replication,

$$u_i^{*(b)} = \frac{\hat{u}_i}{\sqrt{1-h_{ii}}} v_i^{(b)}, \qquad \Pr\Big(v_i^{(b)} = 1\Big) = \Pr\Big(v_i^{(b)} = -1\Big) = \frac{1}{2},$$

where h_{ii} is the corresponding leverage value. Each bootstrap sample is constructed as

$$CS_i^{*(b)} = \widehat{CS}_i + u_i^{*(b)}.$$

The cutoff and the post-cutoff slope are re-estimated in every replication. Food shares, household controls, and survey weights are held fixed.

Table B.9. Breakpoint estimates and wild-bootstrap inference

Survey	$\hat{c}$	Post-cutoff slope	Bootstrap SD	$\lambda(\hat{c})$	90% interval for c	95% interval for c	Share of draws with $\lambda(c) > \lambda^*$
2007	41.1	-0.0875 (0.0061)	0.0077	0.659	[35.4, 44.9]	[33.4, 45.5]	0.956
2014	42.5	-0.0869 (0.0044)	0.0057	0.499	[39.8, 44.7]	[39.3, 45.1]	0.132

Notes: Parentheses report HC1 standard errors conditional on the estimated cutoff. Bootstrap SD is the standard deviation of the post-cutoff slope across 5,000 HC2-adjusted Rademacher wild-bootstrap replications. The cutoff intervals are percentile intervals. For each bootstrap cutoff, $\lambda(c)$ is evaluated using the original survey-weighted sample. The final column reports the fraction of bootstrap cutoffs for which this prevalence exceeds the exact benchmark threshold $\lambda^* = 6/11 \simeq 0.545$; it is not the probability that the economy was in the IADL regime. The pre-cutoff slope is restricted to zero. Because strata and primary-sampling-unit identifiers are unavailable, the exercise does not provide full design-based inference.

The post-cutoff slope is negative in every bootstrap replication for both surveys. In 83.0 percent of paired replications, the prevalence evaluated at the 2007 bootstrap cutoff exceeds the model threshold while the corresponding 2014 prevalence falls below it. The 2007 regime margin exceeds the 2014 margin in 97.6 percent of paired draws. The 95 percent interval for the difference $\hat{c}_{2014} - \hat{c}_{2007}$ is $[-4.1, 9.7]$, which includes zero and is consistent with a common low-40s breakpoint in the compressible-share profile.

The 2014 SSE profile is flat around its minimum: candidate cutoffs between approximately 42.4 and 43.0 percent produce nearly identical fit and the same regime classification. Under the baseline control specification, the numerical minimum is attained at 42.5 percent. Expanding the search grid from 30–60 to 25–60 percent leaves both point estimates unchanged. For 2007, only 17 of the 5,000 bootstrap draws, or 0.3 percent, select the 25 percent lower boundary, and 45 draws, or 0.9 percent, select a cutoff below 30 percent. The lower grid limit therefore does not drive the reported interval.

Appendix C Model, Determinacy, and Additional Dynamics

Throughout Appendix C, λ denotes the HtM household share, Λ_t Ricardian marginal utility, and $d_t = D_t/\bar{Y}$ normalized profit deviations.

C.1 Full nonlinear equilibrium conditions

A unit mass of households consists of a fraction $1-\lambda$ of Ricardian households and a fraction λ of hand-to-mouth households. For household type $j \in \{o, r\}$, preferences are

$$U(C_t^j, H_t^j) = \frac{(C_t^j)^{1-\sigma}}{1-\sigma} - \kappa_L \frac{(H_t^j)^{1+\phi}}{1+\phi}.$$

When $\sigma = 1$, the consumption term is interpreted as $\log C_t^j$.

Ricardian households

Ricardian households choose C_t^o, H_t^o, and nominal bond holdings B_t^o subject to

$$C_t^o + \frac{B_t^o}{P_t} = W_t H_t^o + \frac{R_{t-1} B_{t-1}^o}{P_t} + \frac{1-\tau_D}{1-\lambda} D_t.$$

Their marginal utility, Euler equation, and labor-supply condition are

$$\Lambda_t = (C_t^o)^{-\sigma}.$$

$$1 = \beta E_t \left[\frac{\Lambda_{t+1}}{\Lambda_t} \frac{R_t}{\Pi_{t+1}} \right], \qquad \Pi_{t+1} = \frac{P_{t+1}}{P_t}.$$

$$\kappa_L (H_t^o)^{\phi} = W_t (C_t^o)^{-\sigma}.$$

Hand-to-mouth households

Hand-to-mouth households consume current disposable income:

$$C_t^r = W_t H_t^r + \frac{\tau_D}{\lambda} D_t,$$

$$\kappa_L (H_t^r)^{\phi} = W_t (C_t^r)^{-\sigma}.$$

Here D_t denotes aggregate real profits and τ_D denotes the share of aggregate profit income assigned to hand-to-mouth households under the balanced-budget redistribution rule. The benchmark sets $\tau_D = 0$.

The allocation can be implemented through the period-by-period balanced-budget identity

$$\lambda T_t = \tau_D D_t,$$

where T_t is the net lump-sum amount assigned to each constrained household. The fiscal authority makes no purchases and issues no debt. Ricardian households receive $(1-\tau_D)D_t$ in aggregate, while constrained households receive $\tau_D D_t$. Total profit income therefore remains D_t, and the aggregate

resource constraint is unchanged. This is the constant-rate version of the fiscal arrangement studied by Bilbiie (2008). Because both household types in our model have the same labor-supply elasticity, the derivation corresponds to the corrected case reported in Bilbiie (2019).

Firms

Symmetric production and real marginal cost are

$$Y_t = A_t H_t - F,$$

$$MC_t = \frac{W_t}{A_t}.$$

Price setting

Under Rotemberg price adjustment costs, the symmetric pricing condition is

$$(\Pi_t - \bar{\Pi})\Pi_t = \beta E_t \left[\frac{\Lambda_{t+1}}{\Lambda_t} \frac{Y_{t+1}}{Y_t} \Pi_{t+1}(\Pi_{t+1} - \bar{\Pi}) \right] + \frac{\varepsilon}{\kappa_P} \left(MC_t - \frac{\varepsilon - 1}{\varepsilon} \right). \tag{C.1}$$

Aggregation and market clearing

$$D_t = Y_t - W_t H_t - \frac{\kappa_P}{2}(\Pi_t - \bar{\Pi})^2 Y_t.$$

$$C_t = (1-\lambda) C_t^o + \lambda C_t^r,$$

$$H_t = (1-\lambda) H_t^o + \lambda H_t^r,$$

$$Y_t = C_t + \frac{\kappa_P}{2}(\Pi_t - \bar{\Pi})^2 Y_t.$$

$$B_t^o = 0.$$

Monetary policy

Monetary policy follows

$$\frac{R_t}{\bar{R}} = \left(\frac{E_t \Pi_{t+1}}{\bar{\Pi}} \right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_y} \exp(\varepsilon_t^m). \tag{C.2}$$

The baseline holds productivity fixed, $A_t = \bar{A}$, and assumes that ε_t^m is i.i.d.

C.2 Steady state

The deterministic steady state has constant productivity, no shocks, and inflation at target. We normalize

$$\bar{\Pi} = 1.$$

The Ricardian Euler equation and the steady-state pricing condition imply

$$\bar{R} = \frac{\bar{\Pi}}{\beta} = \frac{1}{\beta}, \qquad \overline{MC} = \frac{\varepsilon - 1}{\varepsilon} = \frac{1}{1+\mu}, \qquad \mu = \frac{1}{\varepsilon - 1}.$$

Steady-state profits satisfy

$$\bar{D} = \bar{Y} - \bar{W}\bar{H} = 0.$$

Using

$$\bar{Y} = \bar{A}\bar{H} - F, \qquad \bar{W} = \overline{MC}\bar{A},$$

the zero-profit condition implies

$$F = \mu\bar{Y}.$$

We normalize $\bar{Y} = 1$ and $\bar{H} = 1/3$. It follows that

$$F = \mu, \qquad \bar{A} = \frac{\bar{Y} + F}{\bar{H}} = 3(1+\mu), \qquad \bar{W} = \overline{MC}\bar{A} = 3.$$

Goods-market clearing and the household budget constraints imply the symmetric steady state

$$\bar{C} = \bar{C}^o = \bar{C}^r = 1, \qquad \bar{H}^o = \bar{H}^r = \bar{H} = \frac{1}{3}, \qquad \bar{D} = 0.$$

The labor-disutility parameter is calibrated from the intratemporal condition:

$$\kappa_L \bar{H}^{\phi} = \bar{W}\bar{C}^{-\sigma},$$

so that

$$\kappa_L = \frac{\bar{W}}{\bar{H}^{\phi}\bar{C}^{\sigma}} = \frac{3}{(1/3)^{\phi}}.$$

The Rotemberg parameter is calibrated to reproduce the first-order slope of a Calvo Phillips curve with price-stickiness probability θ:

$$\kappa_P = \frac{(\varepsilon - 1)\theta}{\bar{\Pi}^2(1-\theta)(1-\beta\theta)}.$$

For $\bar{\Pi} = 1$, $\varepsilon = 6$, $\beta = 0.99$, and $\theta = 0.75$,

$$\kappa_P = \frac{5 \times 0.75}{0.25(1 - 0.99 \times 0.75)} \simeq 58.25.$$

The complete steady state is therefore

$$\begin{aligned}
&\bar{\Pi} = 1, && \bar{R} = \frac{1}{\beta}, && \overline{MC} = \frac{\varepsilon - 1}{\varepsilon}, \quad \bar{Y} = 1, \quad \bar{H} = \frac{1}{3}, \\
&F = \frac{1}{\varepsilon - 1} = \mu, && \bar{A} = 3(1+\mu) = \frac{3\varepsilon}{\varepsilon - 1}, && \bar{W} = 3, \\
&\bar{C} = \bar{C}^o = \bar{C}^r = 1, && \bar{H}^o = \bar{H}^r = \frac{1}{3}, && \bar{D} = 0.
\end{aligned} \tag{C.3}$$

C.3 First-order log-linear system

For any strictly positive variable X_t, define $\hat{x}_t = \log(X_t/\bar{X})$. Since $\bar{D} = 0$, profits cannot be log-linearized. The relevant first-order object is the normalized profit deviation

$$d_t = \frac{D_t - \bar{D}}{\bar{Y}} = \frac{D_t}{\bar{Y}}.$$

The following first-order system is taken around the zero-inflation steady state, with $\bar{\Pi} = 1$, constant productivity, and $\sigma = 1$ unless otherwise stated.

Ricardian Euler equation.

$$\hat{c}_t^o = E_t \hat{c}_{t+1}^o - (\hat{r}_t - E_t \hat{\pi}_{t+1}).$$

With general σ, the equation is

$$\hat{c}_t^o = E_t \hat{c}_{t+1}^o - \frac{1}{\sigma}(\hat{r}_t - E_t \hat{\pi}_{t+1}).$$

Labor supply and hand-to-mouth budget.

$$\phi \hat{h}_t^o = \hat{w}_t - \hat{c}_t^o, \qquad \phi \hat{h}_t^r = \hat{w}_t - \hat{c}_t^r.$$

The hand-to-mouth budget constraint is

$$\hat{c}_t^r = \hat{w}_t + \hat{h}_t^r + \frac{\tau_D}{\lambda} d_t. \tag{C.4}$$

The coefficient on d_t follows from $\bar{Y} = \bar{C}^r = 1$.

Production, marginal cost, and profits.

$$\hat{y}_t = (1+\mu)\hat{h}_t, \qquad \widehat{mc}_t = \hat{w}_t - \hat{a}_t.$$

In the baseline exercises $\hat{a}_t = 0$, so $\widehat{mc}_t = \hat{w}_t$. Since the Rotemberg adjustment cost is second order around $\bar{\Pi} = 1$, it drops out at first order, and normalized profits satisfy

$$d_t = \hat{y}_t - \hat{w}_t - \hat{h}_t.$$

New Keynesian Phillips curve. Linearizing the Rotemberg pricing equation gives

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa_{mc} \widehat{mc}_t, \tag{C.5}$$

where

$$\kappa_{mc} = \frac{\varepsilon - 1}{\kappa_P}.$$

Under constant productivity,

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa_{mc} \hat{w}_t.$$

Aggregation, market clearing, and policy.

$$\hat{c}_t = (1-\lambda)\hat{c}_t^o + \lambda \hat{c}_t^r, \qquad \hat{h}_t = (1-\lambda)\hat{h}_t^o + \lambda \hat{h}_t^r.$$

At first order, goods-market clearing is $\hat{y}_t = \hat{c}_t$. The log-linearized forward-looking instrument rule is

$$\hat{r}_t = \phi_\pi E_t \hat{\pi}_{t+1} + \phi_y \hat{y}_t + \varepsilon_t^m. \tag{C.6}$$

C.4 Reduced three-equation system and forward-looking transition matrix

Under the benchmark case $\tau_D = 0$, the hand-to-mouth budget and labor-supply equations imply

$$\hat{h}_t^r = 0, \qquad \hat{c}_t^r = \hat{w}_t.$$

Using production, aggregation, Ricardian labor supply, and goods-market clearing, hand-to-mouth consumption can be written as

$$\hat{c}_t^r = \chi \hat{y}_t, \qquad \chi = 1 + \frac{\phi}{1+\mu}.$$

Aggregate consumption satisfies $\hat{c}_t = \hat{y}_t$, so

$$\hat{c}_t^o = \frac{1-\lambda\chi}{1-\lambda}\hat{y}_t.$$

Substitution into the Ricardian Euler equation gives the aggregate-demand equation

$$\hat{y}_t = E_t\hat{y}_{t+1} - \delta\left(\hat{r}_t - E_t\hat{\pi}_{t+1}\right), \qquad \delta = \frac{1-\lambda}{1-\lambda\chi}.$$

Inflation dynamics reduce to

$$\hat{\pi}_t = \beta E_t\hat{\pi}_{t+1} + \kappa_y \hat{y}_t, \qquad \kappa_y = \frac{\varepsilon - 1}{\kappa_P}\left(1 + \frac{\phi}{1+\mu}\right).$$

The sign of δ determines the aggregate-demand regime. The economy is in the standard aggregate-demand region if $1 - \lambda\chi > 0$ and in the inverted aggregate-demand region if $1 - \lambda\chi < 0$. The boundary $1 - \lambda\chi = 0$ is singular because δ is not finite. The threshold is

$$\lambda^* = \frac{1}{\chi} = \frac{1}{1 + \phi/(1+\mu)}.$$

Under the balanced-budget profit-redistribution rule, the profit block must be treated in normalized level deviations. Let $\Gamma \equiv 1 + \mu$. Production gives $\hat{h}_t = \hat{y}_t/\Gamma$. Averaging labor supply and using $\hat{c}_t = \hat{y}_t$ gives

$$\hat{w}_t = \left(1 + \frac{\phi}{\Gamma}\right)\hat{y}_t.$$

Therefore

$$d_t = \hat{y}_t - \hat{w}_t - \hat{h}_t = -\frac{1+\phi}{\Gamma}\hat{y}_t = -\frac{1+\phi}{1+\mu}\hat{y}_t.$$

The hand-to-mouth budget and labor-supply equations imply

$$\hat{h}_t^r = -\frac{\tau_D}{\lambda(1+\phi)}d_t$$

and

$$\hat{c}_t^r = \hat{w}_t + \frac{\phi}{1+\phi}\frac{\tau_D}{\lambda}d_t.$$

Substituting the expressions for $\hat{w}_t$ and d_t yields

$$\hat{c}_t^r = \left[1 + \frac{\phi}{1+\mu}\left(1 - \frac{\tau_D}{\lambda}\right)\right]\hat{y}_t.$$

Thus

$$\chi(\tau_D) = 1 + \frac{\phi}{1+\mu}\left(1 - \frac{\tau_D}{\lambda}\right), \qquad \delta(\tau_D) = \frac{1-\lambda}{1-\lambda\chi(\tau_D)}.$$

The balanced-budget redistribution rule restores standard aggregate-demand logic when $1 - \lambda\chi(\tau_D) > 0$. The threshold satisfies $\lambda\chi(\tau_D^*) = 1$, so

$$\tau_D^* = \lambda - (1+\mu)\frac{1-\lambda}{\phi}.$$

For $\lambda = 0.688339992675652$, $\mu = 0.2$, and $\phi = 1$, the threshold is

$$\tau_D^* = 0.314347984.$$

For determinacy, the reduced Taylor-rule system is written in expectational form. Let

$$x_t = \begin{pmatrix} \hat{y}_t \\ \hat{\pi}_t \end{pmatrix}.$$

Under the forward-looking rule, after substituting out $\hat{r}_t$,

$$E_t x_{t+1} = A^{FL}(\phi_y) x_t + \begin{pmatrix} \delta \\ 0 \end{pmatrix} \varepsilon_t^m,$$

where

$$A^{FL}(\phi_y) = \begin{pmatrix} 1 + \delta\phi_y - \dfrac{\delta(\phi_\pi - 1)\kappa_y}{\beta} & \dfrac{\delta(\phi_\pi - 1)}{\beta} \\ -\dfrac{\kappa_y}{\beta} & \dfrac{1}{\beta} \end{pmatrix}. \tag{C.7}$$

The reduced 2×2 system reproduces the two finite generalized eigenvalues of the full linearized model. The full representation additionally contains one infinite generalized eigenvalue generated by an algebraic relation. This root carries no independent dynamics and is therefore excluded from the finite-root count. Throughout Appendix C, $z_i(A)$ denotes the eigenvalues of the forward transition matrix A, not those of A^{-1}.

C.5 General case with $\sigma \neq 1$

When $\sigma \neq 1$, the hand-to-mouth budget constraint remains

$$\hat{c}_t^r = \hat{w}_t + \hat{h}_t^r,$$

under $\tau_D = 0$, while labor supply is

$$\hat{w}_t - \sigma\hat{c}_t^r = \phi\hat{h}_t^r.$$

Combining these equations gives

$$\hat{h}_t^r = \frac{1-\sigma}{\phi+\sigma}\hat{w}_t, \qquad \hat{c}_t^r = \frac{\phi+1}{\phi+\sigma}\hat{w}_t.$$

Aggregate equilibrium implies

$$\hat{w}_t = \left(\sigma + \frac{\phi}{1+\mu}\right)\hat{y}_t.$$

Therefore

$$\hat{c}_t^r = \chi(\sigma)\hat{y}_t, \qquad \chi(\sigma) = \frac{\phi+1}{\phi+\sigma}\left(\sigma + \frac{\phi}{1+\mu}\right).$$

The aggregate-demand equation is

$$\hat{y}_t = E_t\hat{y}_{t+1} - \delta(\sigma)\left(\hat{r}_t - E_t\hat{\pi}_{t+1}\right), \qquad \delta(\sigma) = \frac{1-\lambda}{\sigma[1-\lambda\chi(\sigma)]}.$$

The generalized threshold is

$$\lambda^*(\sigma) = \frac{1}{\chi(\sigma)}.$$

A higher σ increases $\chi(\sigma)$ when $\mu > 0$, because

$$\frac{\partial\chi(\sigma)}{\partial\sigma} = \frac{\phi(\phi+1)\mu}{(1+\mu)(\phi+\sigma)^2} > 0.$$

C.6 Blanchard–Kahn determinacy under alternative rules

The determinacy checks use the forward-looking transition matrix $A^{FL}(\phi_y)$ in equation (C.7). The reduced system contains two forward-looking variables, $\hat{y}_t$ and $\hat{\pi}_t$, and no predetermined endogenous variables. Local determinacy therefore requires two unstable roots.

The reported roots are the eigenvalues $z_i(A^{FL})$ of the matrix satisfying

$$E_t x_{t+1} = A^{FL} x_t.$$

A root is classified as unstable when

$$|z_i| > 1 + 10^{-8},$$

stable when

$$|z_i| < 1 - 10^{-8},$$

and on the unit circle when

$$||z_i| - 1| \le 10^{-8}.$$

The benchmark parameters are

$$\beta = 0.99, \qquad \varepsilon = 6, \qquad \theta = 0.75, \qquad \mu = 0.2, \qquad \phi = 1, \qquad \tau_D = 0.$$

These values imply

$$\chi = 1.833333, \qquad \kappa_P = 58.252427, \qquad \kappa_y = 0.157361.$$

Structural calculations use the exact benchmark shares

$$\lambda_{2007} = 0.688339992675652, \qquad \lambda_{2014} = 0.571879168720315.$$

Table C.1. Blanchard–Kahn determinacy: 2007 calibration

Rule	ϕ_π	ϕ_y	Unstable roots	Roots $z_i(A^{FL})$	Moduli $\lvert z_i\rvert$	BK status
Active	1.5	0.0	1	(1.364, 0.740)	(1.364, 0.740)	Indeterminate
Passive	0.8	0.0	2	$0.986 \pm 0.194i$	(1.005, 1.005)	Determinate, small margin
Passive	0.5	0.0	2	$0.958 \pm 0.305i$	(1.005, 1.005)	Determinate, small margin
Active + output	1.5	0.5	1	(1.156, 0.354)	(1.156, 0.354)	Indeterminate
Active + output	1.5	1.0	1	(1.091, −0.176)	(1.091, 0.176)	Indeterminate
Active + output	1.5	1.5	1	(1.065, −0.744)	(1.065, 0.744)	Indeterminate
Active + output	1.5	2.0	2	(−1.326, 1.051)	(1.326, 1.051)	Determinate

Notes: Calculations use $\lambda = 0.688339992675652$, implying $\delta = -1.189739$. Roots are eigenvalues of A^{FL} and are classified using the 10^{-8} tolerance defined above. Under the passive rules, the complex roots have modulus 1.005038, indicating a small determinacy margin.

Table C.2. Blanchard–Kahn determinacy: 2014 calibration

Rule	ϕ_π	ϕ_y	Unstable roots	Roots $z_i(A^{FL})$	Moduli $\lvert z_i\rvert$	BK status
Active	1.5	0.0	1	(2.267, 0.446)	(2.267, 0.446)	Indeterminate
Passive	0.8	0.0	2	$0.865 \pm 0.512i$	(1.005, 1.005)	Determinate, small margin
Passive	0.5	0.0	2	$0.654 \pm 0.763i$	(1.005, 1.005)	Determinate, small margin
Active + output	1.5	0.5	2	(−2.898, 1.192)	(2.898, 1.192)	Determinate
Active + output	1.5	1.0	2	(−7.221, 1.096)	(7.221, 1.096)	Determinate
Active + output	1.5	1.5	2	(−11.610, 1.066)	(11.610, 1.066)	Determinate
Active + output	1.5	2.0	2	(−16.014, 1.052)	(16.014, 1.052)	Determinate

Notes: Calculations use $\lambda = 0.571879168720315$, implying $\delta = -8.837229$. The large absolute value of δ reflects the proximity of the 2014 calibration to the singular regime boundary. Roots are classified using the 10^{-8} tolerance defined above. The active rule with $\phi_y = 0.5$ is determinate because this coefficient exceeds the output-response threshold derived in Section C.8.

C.7 Passive-rule root margins and perturbation checks

Under the passive forward-looking rule, the roots lie only slightly outside the unit circle. The following property establishes the analytical source of this margin.

Property C.1 (Discount-pinned modulus of complex roots). Let z_1 and z_2 denote the eigenvalues of the reduced forward transition matrix $A^{FL}(0)$. Under the forward-looking rule with $\phi_y = 0$,

$$\det A^{FL}(0) = \frac{1}{\beta}.$$

Whenever the roots form a complex-conjugate pair,

$$|z_1| = |z_2| = \beta^{-1/2}.$$

Conditional on the roots remaining complex, λ, κ_P, and ϕ_π affect their trace and angle, but not their modulus.

Table C.3 reports selected one-at-a-time perturbations around the exact 2007 passive-rule benchmark. All retained configurations remain in the inverted region and have two unstable complex roots.

Table C.3. Perturbation checks for passive-rule critical roots

Perturbation	Value	Roots $z_{1,2}$	Modulus $\lvert z\rvert$
Benchmark	–	$0.9861 \pm 0.1940i$	1.0050
λ	0.66834	$0.9817 \pm 0.2156i$	1.0050
λ	0.70834	$0.9895 \pm 0.1759i$	1.0050
λ, 2014	0.571879	$0.8646 \pm 0.5124i$	1.0050
β	0.980	$0.9911 \pm 0.1953i$	1.0102
β	0.995	$0.9837 \pm 0.1933i$	1.0025
κ_P	52.427	$0.9840 \pm 0.2044i$	1.0050
κ_P	64.078	$0.9879 \pm 0.1850i$	1.0050
ϕ_π	0.50	$0.9578 \pm 0.3046i$	1.0050
ϕ_π	0.95	$1.0003 \pm 0.0972i$	1.0050

Notes: Each row changes one parameter relative to the exact 2007 benchmark, $\lambda = 0.688339992675652$, $\beta = 0.99$, $\kappa_P = 58.252427$, $\phi_\pi = 0.8$, and $\phi_y = 0$. The 2014 row uses $\lambda = 0.571879168720315$. In the β perturbations, κ_P is held fixed. All reported configurations have two unstable roots under the 10^{-8} tolerance and therefore satisfy the Blanchard–Kahn condition for local determinacy.

Table C.3 confirms that the small determinacy margin is robust to the selected one-at-a-time perturbations. At $\phi_\pi = 1$, the roots are 1 and $1/\beta$, so strict Blanchard–Kahn determinacy fails.

C.8 Output-response determinacy threshold

For the active forward-looking rule with $\phi_\pi > 1$ and $\phi_y \geq 0$, the output-response boundary can be derived exactly from the two-dimensional transition matrix. Let

$$E_t x_{t+1} = A(\phi_y) x_t, \qquad x_t = (\hat{y}_t, \hat{\pi}_t)'.$$

The matrix is

$$A(\phi_y) = \begin{pmatrix} 1 + \delta\phi_y - \dfrac{\delta(\phi_\pi - 1)\kappa_y}{\beta} & \dfrac{\delta(\phi_\pi - 1)}{\beta} \\ -\dfrac{\kappa_y}{\beta} & \dfrac{1}{\beta} \end{pmatrix}. \tag{C.8}$$

At the determinacy boundary relevant for the output-response exercise, one root crosses the unit circle at -1. The boundary therefore solves

$$\det[A(\bar{\phi}_y) + I] = 0.$$

Expanding the determinant and solving for $\bar{\phi}_y$ gives

$$\bar{\phi}_y = \frac{\kappa_y(\phi_\pi - 1)}{1 + \beta} - \frac{2}{\delta}. \tag{C.9}$$

With $\phi_\pi = 1.5$, $\beta = 0.99$, and $\kappa_y = 0.157361$, the exact benchmark hand-to-mouth shares imply

$$\bar{\phi}_y(2007) = 1.720579, \qquad \bar{\phi}_y(2014) = 0.265853.$$

Table C.4. Roots around the output-response determinacy threshold

Calibration	ϕ_y	Location	Roots $z_i(A)$	Moduli $\lvert z_i\rvert$	Margins	Unst.	Status
2007	1.710579	Below	$(-0.988,\ 1.058)$	$(0.988,\ 1.058)$	$(-0.012,\ +0.058)$	1	Indet.
2007	1.720579	Threshold	$(-1.000,\ 1.058)$	$(1.000,\ 1.058)$	$(0.000,\ +0.058)$	1	Unit-root
2007	1.730579	Above	$(-1.012,\ 1.057)$	$(1.012,\ 1.057)$	$(+0.012,\ +0.057)$	2	Det.
2014	0.255853	Below	$(-0.925,\ 1.377)$	$(0.925,\ 1.377)$	$(-0.075,\ +0.377)$	1	Indet.
2014	0.265853	Threshold	$(-1.000,\ 1.363)$	$(1.000,\ 1.363)$	$(0.000,\ +0.363)$	1	Unit-root
2014	0.275853	Above	$(-1.076,\ 1.350)$	$(1.076,\ 1.350)$	$(+0.076,\ +0.350)$	2	Det.

Notes: The table uses $\phi_y = \bar{\phi}_y - 0.01$, $\phi_y = \bar{\phi}_y$, and $\phi_y = \bar{\phi}_y + 0.01$. Roots are eigenvalues of $A(\phi_y)$, not of $A(\phi_y)^{-1}$. At the threshold, one root lies on the unit circle at -1, so the strict Blanchard–Kahn condition is not satisfied.

C.9 Additional monetary-policy impulse responses

This section collects the impulse responses used to illustrate the determinate policy configurations discussed in the main text. Figure C.1 reports the passive-rule IADL benchmark. Figures C.2 and C.3 provide, respectively, a standard aggregate-demand benchmark under a conventional active rule and the 2007 IADL calibration under an active rule with an output response above the determinacy threshold.

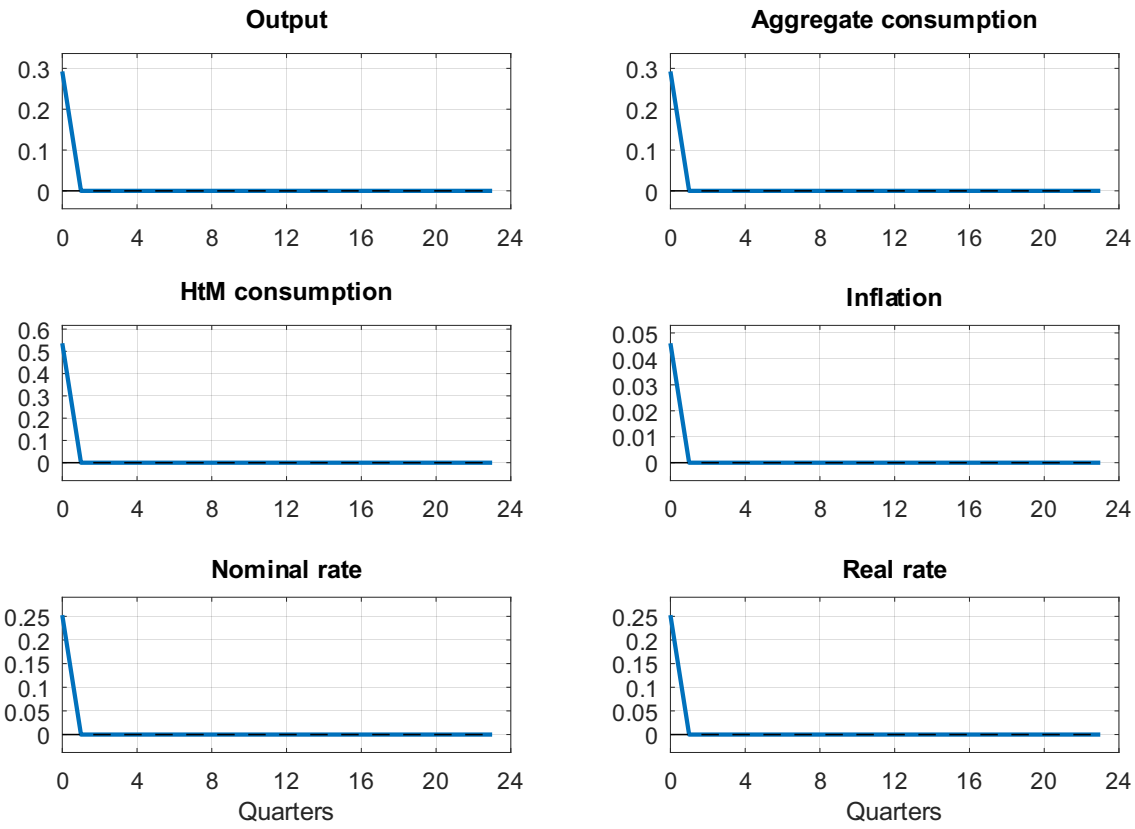


Figure C.1. Impulse responses under the passive forward-looking rule in the IADL calibration

Notes: Responses to a 25-basis-point contractionary monetary-policy shock under the passive forward-looking rule ($\lambda = 0.688$, $\phi = 1$, $\mu = 0.2$, $\phi_\pi = 0.8$, $\phi_y = 0$). Real variables are reported as percent deviations from steady state; nominal and real interest rates are reported in quarterly percentage points.

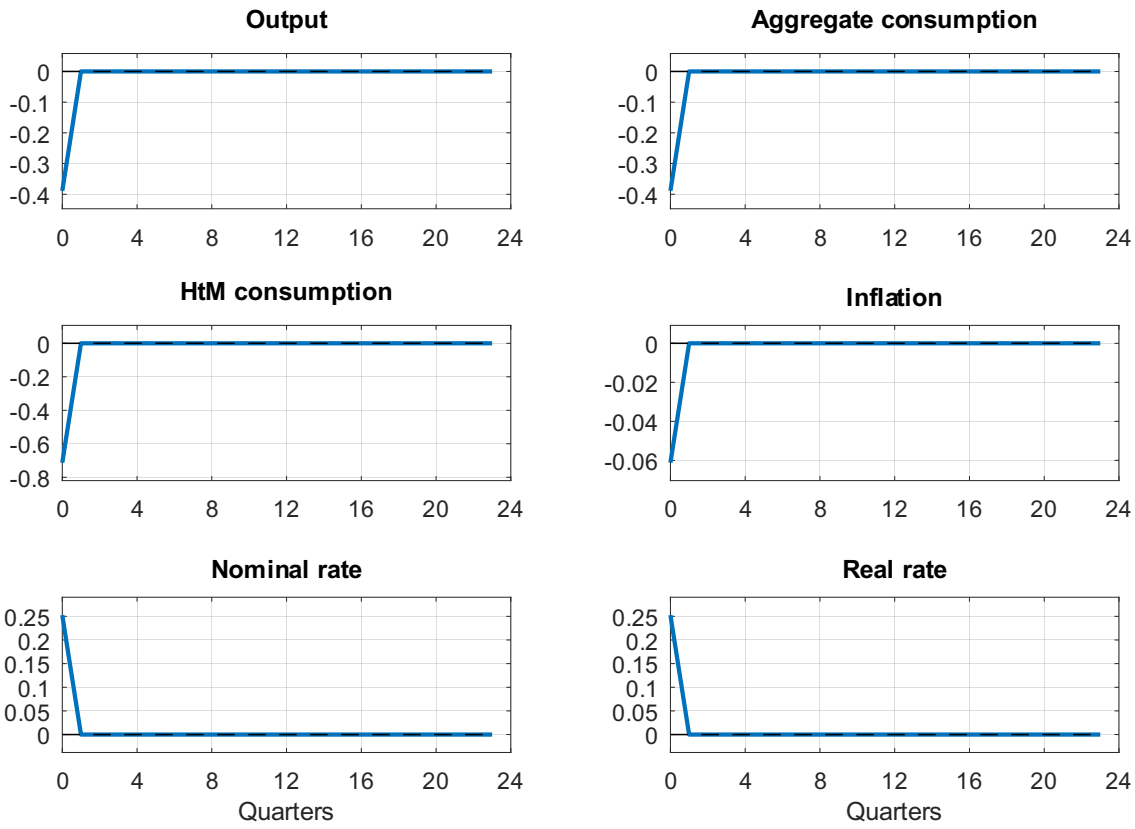


Figure C.2. Monetary-policy shock in the standard aggregate-demand region

Notes: Responses to an i.i.d. 25-basis-point contractionary monetary-policy shock. The calibration uses $\lambda = 0.30 < \lambda^* = 0.545$, $\phi = 1$, $\mu = 0.2$, $\phi_\pi = 1.5$, and $\phi_y = 0$. Real variables are percentage deviations from steady state; interest rates are quarterly percentage-point deviations.

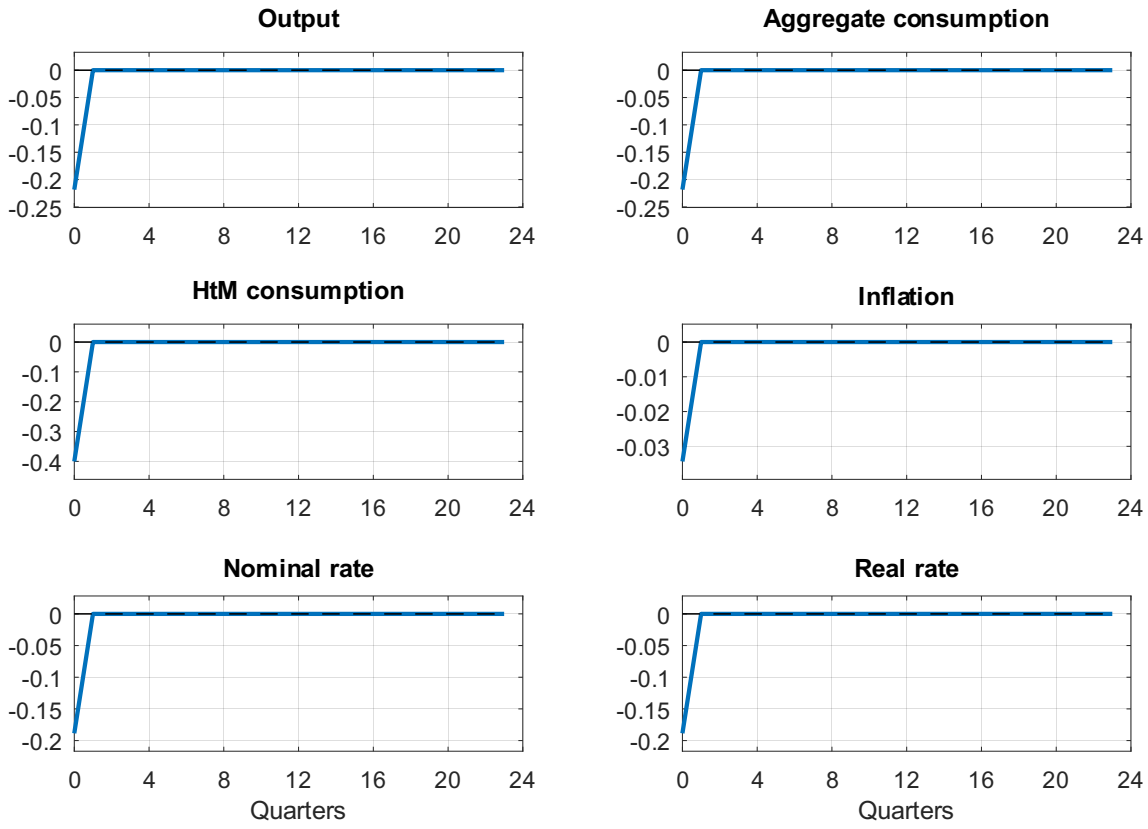


Figure C.3. Active forward-looking rule with an output response above the determinacy threshold

Notes: Responses to an i.i.d. 25-basis-point monetary-policy shock under the 2007 calibration, with $\lambda = 0.688339992675652$, $\phi = 1$, $\mu = 0.2$, and $\tau_D = 0$. The forward-looking rule uses $\phi_\pi = 1.5$ and $\phi_y = 2$, above the determinacy threshold $\bar{\phi}_y = 1.720579$. Real variables are percentage deviations from steady state; interest rates are quarterly percentage-point deviations.

C.10 Commitment targeting: determinacy and impulse responses

The commitment targeting criterion is

$$\hat{\pi}_t + q\left(\hat{y}_t - \hat{y}_{t-1}\right) = 0,$$

where $q > 0$ governs the response of inflation to changes in the output gap. This is a targeting criterion rather than an instrument rule: the nominal interest rate adjusts endogenously through the Ricardian Euler equation to support the targeted paths of output and inflation.

Combining the targeting criterion with the Phillips curve, in the absence of shocks, gives

$$E_t \begin{pmatrix} \hat{y}_{t+1} \\ \hat{y}_t \end{pmatrix} = A^T(q) \begin{pmatrix} \hat{y}_t \\ \hat{y}_{t-1} \end{pmatrix},$$

where

$$A^T(q) = \begin{pmatrix} 1 + \dfrac{q + \kappa_y}{\beta q} & -\dfrac{1}{\beta} \\ 1 & 0 \end{pmatrix}.$$

In this representation, $\hat{y}_t$ is forward-looking and $\hat{y}_{t-1}$ is predetermined, so the Blanchard–Kahn condition requires one unstable root. The characteristic polynomial of $A^T(q)$ is

$$p(z) = z^2 - \left(1 + \frac{q + \kappa_y}{\beta q}\right) z + \frac{1}{\beta}.$$

For $q > 0$, $\kappa_y > 0$, and $0 < \beta < 1$,

$$p(0) = \frac{1}{\beta} > 0, \qquad p(1) = -\frac{\kappa_y}{\beta q} < 0.$$

Hence, one root lies between zero and one. Since the product of the two roots is $1/\beta > 1$, the second root is greater than one. The targeting criterion therefore satisfies the Blanchard–Kahn condition for every $q > 0$. This result is independent of λ, because λ affects the interest-rate path through the aggregate-demand equation but does not enter the reduced targeting block, provided the economy is away from the singular boundary $1 - \lambda\chi = 0$.

Cost-push shock. For the impulse-response exercise, the Phillips curve is augmented with an i.i.d. cost-push innovation:

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa_y \hat{y}_t + u_t,$$

where

$$u_t = \varepsilon_t^u, \qquad E(\varepsilon_t^u) = 0, \qquad \mathrm{sd}(\varepsilon_t^u) = 0.00025.$$

The shock has zero persistence and is used only for the impulse-response exercise; it is set to zero in the determinacy and regime calculations.

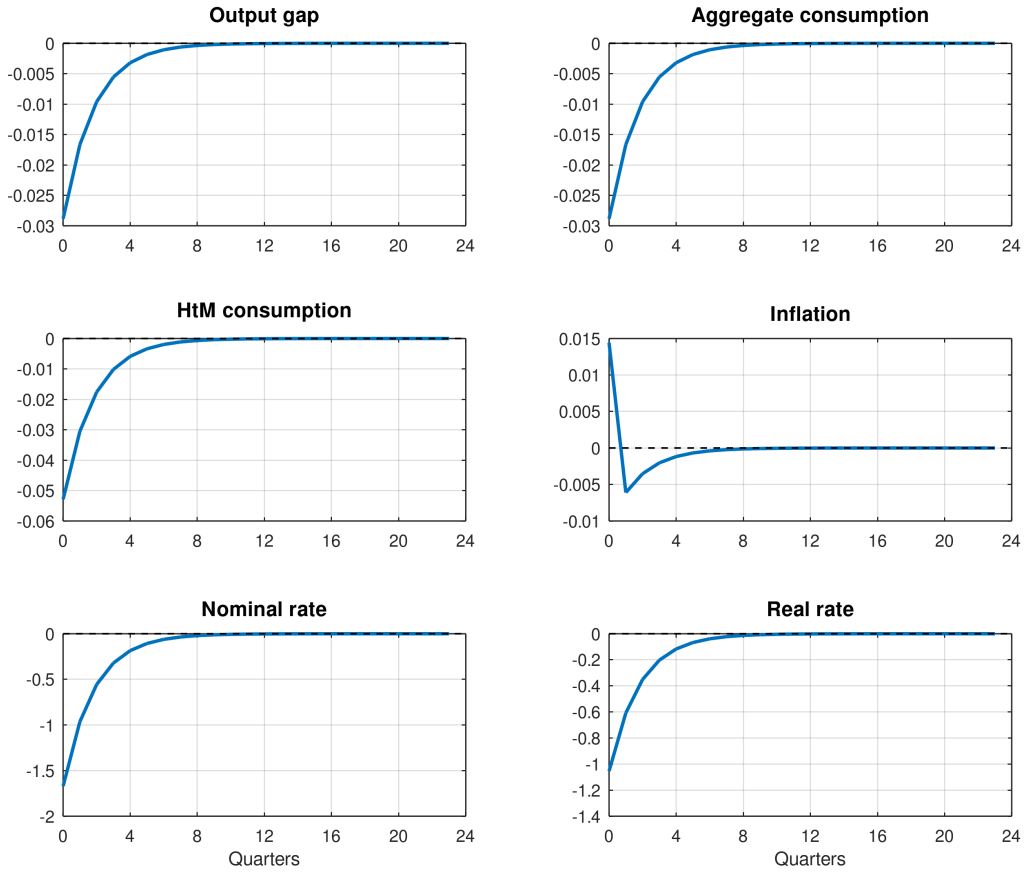


Figure C.4. Impulse responses under commitment targeting

Notes: Responses to a one-standard-deviation positive i.i.d. cost-push shock under the 2007 calibration, with $\lambda = 0.688$, $q = 0.50$, $\phi = 1$, $\sigma = 1$, $\mu = 0.2$, and $\tau_D = 0$. The nominal interest rate adjusts endogenously through the Ricardian Euler equation to satisfy the targeting criterion. Real variables are reported as percent deviations from steady state; nominal and real interest rates are reported in basis points.

Inflation rises on impact following the positive cost-push shock. The targeting criterion requires the output gap to decline relative to its lagged value, so the output gap, aggregate consumption, and hand-to-mouth consumption fall. Under the 2007 IADL calibration, the nominal and real interest rates also decline. Because a lower real interest rate is contractionary in the inverted aggregate-demand region, this endogenous interest-rate adjustment supports the required fall in output.

Appendix D Empirical-to-Structural Mapping and Calibration

This appendix reports the alternative empirical weightings, the exact benchmark shares and associated regime boundaries, and the benchmark calibration and robustness grid.

D.1 Alternative empirical weightings

Let h index households, with household expansion weight ω_h^H, household-level individual expansion weight ω_h^I, total consumption x_h, and classifier HtM_h. Define

$$\lambda_H = \frac{\sum_h \omega_h^H \mathrm{HtM}_h}{\sum_h \omega_h^H},$$

$$\lambda_I = \frac{\sum_h \omega_h^I \mathrm{HtM}_h}{\sum_h \omega_h^I},$$

$$\lambda_C = \frac{\sum_h \omega_h^H x_h \mathrm{HtM}_h}{\sum_h \omega_h^H x_h}.$$

The household-share mapping is the benchmark because λ is the population share of household decision-making units in the symmetric model. The individual- and consumption-weighted measures are reported as alternative diagnostics.

Table D.1. Alternative empirical counterparts of λ

Weighting	2007	2014
Household share λ_H	0.688	0.572
Individual share λ_I	0.716	0.597
Consumption share λ_C	0.548	0.439

Notes: The individual-weighted measure uses the household-level individual expansion weight supplied by the HCP in each public file. The consumption-weighted measure uses the same adjusted total consumption expenditure as the food-share denominator and is reported only as a diagnostic, not as the benchmark structural share.

For the benchmark household-share mapping, the corresponding regime boundaries at $\mu = 0.2$ are

$$\phi^*(\lambda) = (1+\mu)\frac{1-\lambda}{\lambda},$$

giving $\phi^*_{2007} = 0.5433245384$ and $\phi^*_{2014} = 0.8983453597$. Thus, at $\phi = 1$, the 2007 calibration lies clearly in IADL, while the 2014 calibration lies closer to the boundary.

To separate aggregation weights from classification correspondence, let ρ denote the fraction of households classified by the food-share proxy that correspond to the structural constrained type. Under the one-sided exercise reported in Panel B of Table 5,

$$\lambda_{\rho,y} = \rho \lambda_{H,y}.$$

The critical correspondence rate is therefore

$$\rho_y^* = \frac{\lambda^*}{\lambda_{H,y}}.$$

Using the exact household-weighted shares and the exact benchmark threshold $\lambda^* = 6/11 \approx 0.5455$, the critical rates are

$$\rho_{2007}^* = 0.7924, \qquad \rho_{2014}^* = 0.9538.$$

The calibration is in IADL for correspondence rates above these values and in SADL below them; equality corresponds to the singular regime boundary. These quantities are sensitivity parameters, not estimated classification or misclassification probabilities. By assigning no structurally constrained households below the cutoff, the exercise gives the smallest structural constrained share compatible with a given correspondence rate and is therefore conservative with respect to false negatives.

Two-sided proxy-to-structure sensitivity. Let P denote classification by the food-share proxy and S structural constrained status. Define

$$p_1 = \Pr(S = 1 \mid P = 1), \qquad p_0 = \Pr(S = 1 \mid P = 0),$$

and let $q_y = \Pr(P = 1) = \lambda_{H,y}$ denote the survey-weighted proxy prevalence in year y. The structural constrained-household share is

$$\lambda_y(p_1, p_0) = p_1 q_y + p_0(1 - q_y).$$

At the benchmark threshold $\lambda^* = 6/11$, the calibration is in IADL when

$$p_1 q_y + p_0(1 - q_y) > \lambda^*,$$

or, equivalently, when

$$p_0 > \frac{\lambda^* - p_1 q_y}{1 - q_y}.$$

Panel B of Table 5 imposes $p_0 = 0$, which is the most adverse case within this family: raising p_0 increases the structural constrained share for any given p_1. The two-sided region makes the trade-off explicit, since a lower correspondence rate among proxy-classified households can be offset by structurally constrained households located below the cutoff. At $p_1 = 0.70$, IADL requires $p_0 > 0.204$ in 2007 and $p_0 > 0.339$ in 2014. At $p_1 = 0.80$, the 2007 calibration is already in IADL at $p_0 = 0$, whereas 2014 requires $p_0 > 0.205$. At $p_1 = 0.90$, the corresponding 2014 requirement falls to $p_0 > 0.072$. At $p_1 = 1$, both calibrations are in IADL for every $p_0 \geq 0$. The two boundaries intersect at $p_1 = p_0 = \lambda^*$. For $p_1 > \lambda^*$, the requirement on p_0 is stronger in 2014 than in 2007, whereas the ranking is reversed for $p_1 < \lambda^*$. Over the correspondence range emphasized in Table 5, $p_1 \in [0.80, 1]$, the 2014 calibration is uniformly more demanding. p_1 and p_0 are sensitivity parameters, not estimated misclassification probabilities.

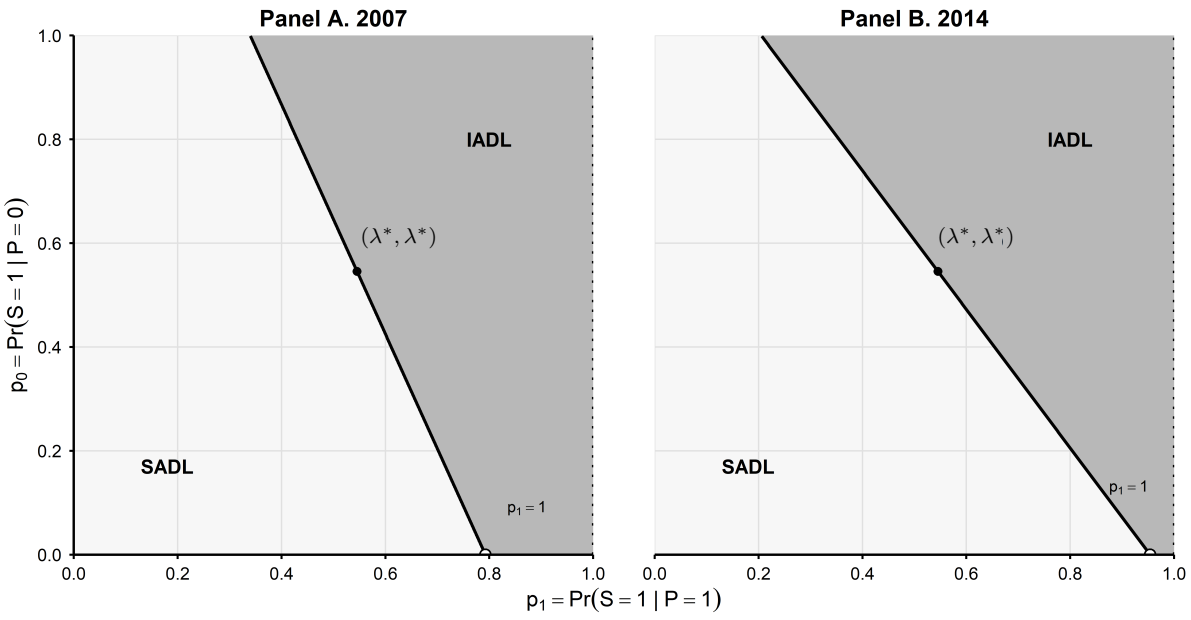


Figure D.1. Two-sided proxy-to-structure sensitivity regions

Notes: Each panel plots the IADL boundary $p_0 = (\lambda^* - p_1 q_y)/(1 - q_y)$, using $\lambda^* = 6/11$ and the survey-weighted household proxy prevalence for the corresponding year. Points above the boundary correspond to IADL and points below it to SADL. The one-sided exercise reported in Panel B of Table 5 corresponds to $p_0 = 0$, marked by the open circle at $p_1 = \rho_y^*$. The solid dot marks $p_1 = p_0 = \lambda^*$, where the proxy carries no information and the two boundaries coincide. The dotted vertical line marks perfect correspondence among proxy-classified households, $p_1 = 1$.

D.2 Calibration

Table D.2 reports the benchmark calibration and robustness grid.

Table D.2. Full benchmark calibration

Parameter	Symbol	Baseline value	Alternative values
Discount factor	β	0.99	—
Inverse intertemporal elasticity of substitution	σ	1	0.5, 2.0
Inverse Frisch elasticity	ϕ	1	0.5, 2, 3
Elasticity of substitution across varieties	ε	6	11.00, 4.33, 3.00
Net markup	μ	0.2	0.10, 0.30, 0.50
Hand-to-mouth share, Morocco 2007	λ	0.688	—
Hand-to-mouth share, Morocco 2014	λ	0.572	—
Calvo-equivalent price rigidity	θ	0.75	—
Rotemberg price-adjustment cost	κ_P	58.25	—
Taylor-rule response to inflation	ϕ_π	1.5	0.5, 0.8
Taylor-rule response to output	ϕ_y	0	0.5, 1.0, 1.5, 2.0
Profit-redistribution rate	τ_D	0	policy-dependent
Monetary policy shock standard deviation	σ_{ε^m}	0.0025	—

$$\kappa_P = \frac{(\varepsilon - 1)\theta}{\bar{\Pi}^2(1-\theta)(1-\beta\theta)}.$$

With $\varepsilon = 6$, $\beta = 0.99$, $\bar{\Pi} = 1$, and $\theta = 0.75$, this gives

$$\kappa_P \approx 58.25.$$

Notes: The model is calibrated at quarterly frequency. Standard parameter values follow Bilbiie (2008) and Galí et al. (2004). The benchmark values of λ are disciplined by the household microdata, while ϕ is calibrated and varied over the reported robustness grid. The parameter θ is the Calvo-equivalent degree of price rigidity used to calibrate κ_P.

Appendix E Additional Robustness

Tables E.1–E.4 report regime classifications under alternative structural parameters and empirical cutoffs.

E.1 Inverse Frisch elasticity

Under the benchmark model,

$$\lambda^*(\phi) = \frac{1+\mu}{1+\mu+\phi}.$$

Table E.1 varies ϕ at $\mu = 0.2$.

Table E.1. Regime classification across inverse Frisch elasticities

ϕ	**Frisch elasticity** $1/\phi$	$\lambda^*(\phi)$	$\lambda = 0.688$ **(2007)**	$\lambda = 0.572$ **(2014)**
0.5	2.00	0.706	SADL	SADL
1.0	1.00	0.545	IADL	IADL, close to threshold
2.0	0.50	0.375	IADL	IADL
3.0	0.33	0.286	IADL	IADL

Notes: IADL applies when $\lambda > \lambda^*(\phi)$; SADL applies when $\lambda < \lambda^*(\phi)$. Calculations use the unrounded benchmark shares.

Both calibrations enter IADL at $\phi = 1$, with a smaller margin for 2014.

E.2 Markup sensitivity

$$\lambda^*(\phi, \mu) = \frac{1 + \mu}{1 + \mu + \phi}.$$

Table E.2 varies μ at $\phi = 1$.

Table E.2. Regime classification across steady-state markups

μ	ε	$\lambda^*(\mu)$	**Morocco 2007,** $\lambda = 0.688$	**Morocco 2014,** $\lambda = 0.572$
0.10	11.00	0.524	IADL	IADL
0.20	6.00	0.545	IADL	IADL
0.30	4.33	0.565	IADL	Near boundary, marginally IADL
0.50	3.00	0.600	IADL	SADL

Notes: $\varepsilon = 1 + 1/\mu$. IADL applies when $\lambda > \lambda^*(\phi, \mu)$; SADL applies otherwise. Calculations use the unrounded benchmark shares.

The 2007 calibration remains in IADL throughout the grid; the 2014 calibration is marginally IADL at $\mu = 0.30$ and moves to SADL at $\mu = 0.50$.

E.3 Inverse intertemporal elasticity

Using the generalized expressions derived in Section C.5, Table E.3 varies σ at $\phi = 1$ and $\mu = 0.2$.

Table E.3. Regime classification across inverse intertemporal elasticities

σ	**IES** $= 1/\sigma$	$\chi(\sigma)$	$\lambda^*(\sigma)$	**2007,** $\lambda = 0.688$	**2014,** $\lambda = 0.572$
0.5	2.000	1.778	0.563	IADL, margin 0.126	IADL, margin 0.009
1.0	1.000	1.833	0.545	IADL, margin 0.143	IADL, margin 0.026
2.0	0.500	1.889	0.529	IADL, margin 0.159	IADL, margin 0.042

Notes: $\lambda^*(\sigma) = 1/\chi(\sigma)$, evaluated at $\phi = 1$ and $\mu = 0.2$. IADL applies when $\lambda > \lambda^*(\sigma)$; margins equal $\lambda - \lambda^*(\sigma)$.

Both calibrations remain in IADL, although the 2014 margin falls below one percentage point at $\sigma = 0.5$.

E.4 Cutoff sensitivity across inverse Frisch elasticities

Table E.4 maps the common benchmark, the year-specific selected cutoffs, and neighboring cutoffs into λ and classifies each calibration across the inverse Frisch elasticity grid.

Table E.4. Regime classification across food-share cutoffs and inverse Frisch elasticities

Cutoff	λ	$\phi^*(\lambda)$	$\phi = 0.5$	$\phi = 1$	$\phi = 2$	$\phi = 3$
Panel A. 2007 survey						
35%	0.798	0.304	IADL	IADL	IADL	IADL
40%	0.688	0.543	SADL	IADL	IADL	IADL
41%†	**0.662**	**0.613**	SADL	IADL	IADL	IADL
43%	0.608	0.773	SADL	IADL	IADL	IADL
45%	0.548	0.989	SADL	IADL, near boundary	IADL	IADL
50%	0.412	1.713	SADL	SADL	IADL	IADL
Panel B. 2014 survey						
35%	0.704	0.504	SADL, near boundary	IADL	IADL	IADL
40%	0.572	0.898	SADL	IADL, near boundary	IADL	IADL
41%	0.541	1.018	SADL	SADL, near boundary	IADL	IADL
43%†	**0.486**	**1.269**	SADL	SADL	IADL	IADL
45%	0.432	1.577	SADL	SADL	IADL	IADL
50%	0.301	2.793	SADL	SADL	SADL	IADL

Notes: λ values are survey-weighted HtM-proxy prevalences under the household-share mapping. The regime boundary is $\phi^*(\lambda) = (1 + \mu)(1 - \lambda)/\lambda$, evaluated at $\mu = 0.2$. IADL applies when $\phi > \phi^*(\lambda)$, while SADL applies when $\phi < \phi^*(\lambda)$. The 40 percent cutoff is the common cross-survey benchmark. † denotes the year-specific cutoff selected under the baseline threshold-selection criterion at tolerance $\eta = 0.02$: 41 percent in 2007 and 43 percent in 2014. Structural calculations use the unrounded prevalence shares; displayed values of λ are rounded to three decimals.

The additional cutoffs sharpen the contrast between the two surveys. At $\phi = 1$, the 2007 calibration remains in IADL at the selected 41 percent cutoff, whereas the 2014 calibration lies in SADL at the selected 43 percent cutoff. In 2014, the regime switch occurs between the common 40 percent benchmark and the 41 percent cutoff.

References

Bilbiie, Florin O. (2008). "Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic." *Journal of Economic Theory*, 140(1), 162–196. https://doi.org/10.1016/j.jet.2007.07.008

Bilbiie, Florin O. (2019). "Corrigendum to 'Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic' [J. Econ. Theory 140 (1) (2008) 162–196]." *Journal of Economic Theory*, 181, 421–422. https://doi.org/10.1016/j.jet.2019.03.008

Galí, Jordi, J. David López-Salido, and Javier Vallés (2004). "Rule-of-Thumb Consumers and the Design of Interest Rate Rules." *Journal of Money, Credit and Banking*, 36(4), 739–763. https://doi.org/10.1353/mcb.2004.0064