



Marcus Vinícius Fernandes Gomes de Castro

**Two Essays on Weak Identification in
Macroeconomic Models**

Dissertação de Mestrado

Thesis presented to the Programa de Pós-graduação em Economia da PUC-Rio in partial fulfillment of the requirements for the degree of Mestre em Economia.

Advisor : Prof. Carlos Viana de Carvalho
Co-advisor: Prof. Ruy Monteiro Ribeiro

Rio de Janeiro
August 2019



Marcus Vinícius Fernandes Gomes de Castro

**Two Essays on Weak Identification in
Macroeconomic Models**

Thesis presented to the Programa de Pós-graduação em Economia da PUC-Rio in partial fulfillment of the requirements for the degree of Mestre em Economia. Approved by the Examination Committee.

Prof. Carlos Viana de Carvalho

Advisor

Departamento de Economia – PUC-Rio

Prof. Ruy Monteiro Ribeiro

Co-advisor

Departamento de Economia – PUC-Rio

Prof. João Victor Issler

Escola de Pós-Graduação em Economia – FGV EPGE

Prof. Marco Antonio Cesar Bonomo

Insper Instituto de Ensino e Pesquisa – Insper

Rio de Janeiro, August the 30th, 2019

All rights reserved.

Marcus Vinícius Fernandes Gomes de Castro

B.A., Economics, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), 2015.

Bibliographic data

Castro, Marcus Vinícius Fernandes Gomes de

Two Essays on Weak Identification in Macroeconomic Models / Marcus Vinícius Fernandes Gomes de Castro; advisor: Carlos Viana de Carvalho; co-advisor: Ruy Monteiro Ribeiro. – Rio de Janeiro: PUC-Rio, Departamento de Economia, 2019.

v., 179 f: il. color. ; 30 cm

Dissertação (mestrado) - Pontifícia Universidade Católica do Rio de Janeiro, Departamento de Economia.

Inclui bibliografia

1. Economia – Teses. 2. Identificação Fraca;. 3. Curva de Phillips Novo-Keynesiana;. 4. Elasticidade de Substituição Intertemporal;. 5. Inflação;. 6. Consumo Não Filtrado.. I. Carvalho, Carlos Viana de . II. Ribeiro, Ruy Monteiro. III. Pontifícia Universidade Católica do Rio de Janeiro. Departamento de Economia. IV. Título.

CDD: 330

To my parents, for their essential support. Especially, I would like to thank my mother, who did not refrain from encouraging this work, either directly or indirectly, even under adverse circumstances.

Acknowledgments

First of all, I would like to thank my parents, who demonstrated and have been demonstrating plenty of patience, support and enthusiasm for everything that relates to my professional career. My mother deserves special emphasis. She did not abstain from encouraging any step of this research project, even though suffering from illness during most part of its confection. I have no words that could properly define her understanding with the situation, nor could I be completely clear on how thankful I am for every gesture she did in the last couple of years. Actions that seemed small to me at the time were probably immeasurable to her. I completely recognise her efforts now, and no doubts this paragraph falls short of my obligations as a son. It is still a transparent and humble way to express my feelings, though.

I am grateful to my advisor, Carlos Carvalho, and to my co-advisor, Ruy Ribeiro. This dissertation would not have been possible without their guidance and fundamental contributions. I would like to thank Marcelo Medeiros for relevant collaboration and helpful insights during this work. I also thank João Issler and Marco Bonomo for important comments on this text, as well as Leonardo Rezende for an insightful discussion just after my workshop.

I thank my colleagues and friends. Indeed, a few details of this work are a result of productive conversations.

This study was partially financed by the Coordenação de Aperfeiçoamento Pessoal de Nível Superior – Brasil (CAPES) – Código de Financiamento 001. Financial support from PUC-Rio is also gratefully acknowledged.

Abstract

Castro, Marcus Vinícius Fernandes Gomes de ; Carvalho, Carlos Viana de (Advisor); Ribeiro, Ruy Monteiro (Co-Advisor). **Two Essays on Weak Identification in Macroeconomic Models.** Rio de Janeiro, 2019. 179p. Dissertação de mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

The weak identification problem arises naturally in macroeconomic models. Consequently, instrumental variables methods produce puzzling results more often than what is empirically plausible. We propose novel methods to address puzzles usually featured in two of the main equations in macro models, namely the New-Keynesian Phillips Curve (NKPC) and the Euler Equation (EE). For the former, difficulties to estimate a positive slope without incurring a degree of stickiness incompatible with the micro evidence are widely known. We address the matter in the first chapter, proposing a richer framework of a multi-sector economy with price-setting heterogeneity. The procedure generates positive and roughly unchanging slope coefficients across econometric settings, as well as degrees of stickiness in line with the micro data, both regarding the entire economy and the cross section of sectors. Importantly, all of these estimates move consistently with implications by theory when modifying the model assumptions. The second chapter focuses on the estimation of the elasticity of intertemporal substitution (EIS), central parameter of the EE in models of dynamic choice. There, we argue that the use of officially reported consumption data – which is usually filtered, smoothed, interpolated, etc – distorts estimates of the EIS. A generalised model to “unfilter” available consumption data is proposed, suitable for several types of data – macro and micro – at different frequencies. Estimations based on unfiltered consumption produce considerably more stable estimates of the EIS, regardless of the econometric approach and the type of consumption data used. Results also seem less sensitive to the presence of weak instruments, compared to officially reported data.

Keywords

Weak Identification; New-Keynesian Phillips Curve; Elasticity of Intertemporal Substitution; Inflation; Unfiltered Consumption.

Resumo

Castro, Marcus Vinícius Fernandes Gomes de ; Carvalho, Carlos Viana de ; Ribeiro, Ruy Monteiro. **Dois Ensaios Em Identificação Fraca Em Modelos Macroeconômicos**. Rio de Janeiro, 2019. 179p. Dissertação de Mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

O problema de identificação fraca surge naturalmente em modelos macroeconômicos. Consequentemente, métodos de variáveis instrumentais produzem resultados enigmáticos de forma mais frequente do que seria empiricamente razoável. Neste trabalho, propomos dois novos métodos para tratar destas dificuldades, no que tange a duas das principais equações de modelos macro: a Curva de Phillips Novo-Keynesiana (NKPC) e a Equação de Euler (EE). Sabe-se das dificuldades em se estimar um coeficiente de sensibilidade positivo entre inflação e produto no primeiro caso, e que, mesmo quando se obtém uma estimativa positiva, o nível de rigidez nominal implicado para a economia é incompatível com o que sugerem os micro dados. Nós abordamos essa questão no primeiro capítulo, propondo um modelo de economia multi-setorial com heterogeneidade na fixação de preços entre setores. O método gera coeficientes de sensibilidade positivos e estáveis para diferentes configurações econométricas, assim como níveis de rigidez nominal alinhados com a evidência micro, para a economia como um todo e também para cada setor individualmente. Todas essas estimativas variam em linha com implicações teóricas, quando hipóteses do modelo são alteradas. O foco do segundo capítulo é a estimação da elasticidade de substituição intertemporal (EIS), parâmetro central da EE. Argumentamos como o uso de séries oficiais de consumo – que são estatisticamente tratadas antes de disponibilizadas – distorce estimativas da EIS. Propondo um modelo generalizado para “desfiltrar” diferentes tipos de séries de consumo disponíveis, – micro e macro, com várias frequências –, demonstramos como a utilização de consumo “não filtrado” gera estimativas da EIS que são consideravelmente mais estáveis, independente do arcabouço econométrico e da série de consumo usada. Resultados também parecem menos sensíveis à presença de instrumentos fracos, comparativamente a estimações usando séries oficiais.

Palavras-chave

Identificação Fraca; Curva de Phillips Novo-Keynesiana; Elasticidade de Substituição Intertemporal; Inflação; Consumo Não Filtrado.

Table of contents

1	Inflation Dynamics in a Multi-Sector Framework	14
1.1	Introduction	14
1.2	Data	20
1.3	Heterogeneity, Omitted Terms and the Literature	21
1.4	Inflation Dynamics in a Multi-Sector Framework	30
1.4.1	New-Keynesian Phillips Curves	34
1.4.2	Econometric Approach	37
1.4.3	Main Results	39
1.4.3.1	Reduced Form	39
1.4.3.2	Structural Form	41
1.4.3.3	Behind the Scenes	48
1.4.4	Estimator Uncertainty	49
1.4.4.1	Parametric Stability	50
1.4.4.2	Subsample Stability	52
1.4.5	Additional Robustness Checks	55
1.5	Conclusion	57
2	Elasticity of Intertemporal Substitution with Unfiltered Consumption	58
2.1	Introduction	58
2.2	Model	63
2.2.1	Consumption Volatility	66
2.2.2	Adjusting Asset Returns	67
2.2.3	Calibration	67
2.3	EIS Estimates with Unfiltered Consumption Data	72
2.3.1	Homoscedastic Framework	74
2.3.2	Heteroscedastic Framework	80
2.4	EIS, Limited Participation and Unfiltered Consumption	83
2.4.1	Results	87
2.5	Conclusion	90
	References	93
A	Appendix for Inflation Dynamics in a Multi-Sector Framework	102
A.1	Additional Details on the Data Set	102
A.2	Heterogenous Economy: Completely Specified Model	102
A.2.1	Deriving the NKPCs	104
A.3	Instrument Sets	106
A.4	Additional Results for the Heterogeneous Economy	106
A.4.1	Reduced-Form Estimations	107
A.4.2	Structural Estimations: Implied Slope	110
A.4.3	Structural Estimations: Implied Stickiness	113
A.4.4	Structural Estimations: β and γ Calibrated	118
A.4.5	Alternative Form of Indexation	121
A.4.6	Sectoral Infrequencies: Estimated vs. Implied from Micro Evidence	124

A.4.7	Parametric Stability	129
A.4.8	Rolling-GMM Estimations: Alternative Approach	132
A.4.9	Structural Estimations: Dropping the Aggregate NKPC in (19)	136
A.4.10	Structural Estimations When $\Theta \approx 0.38$	139
A.4.11	Structural Estimations: Sectoral NKPCs to Generate Initial Values	141
B	Appendix for EIS with Unfiltered Consumption	142
B.1	Data	142
B.1.1	Section 3	142
B.1.2	Section 4	142
B.1.3	A Few Notes on Data Sources involving Consumption, Frequencies and Measurement Errors	143
B.1.3.1	Data Sources Comparison: NIPA vs. CEX	143
B.2	Completely-Specified Quasi-Differenced Filter Model	144
B.2.1	Adjusting Unfiltered Consumption for Time-Aggregation Bias	150
B.2.2	Adjusting the Timing of Asset Returns	151
B.2.3	Mapping Unfiltered Quasi-Differenced Consumption Back onto Unfiltered Consumption	152
B.2.4	Consumption Volatility	152
B.2.5	Homoscedastic Counterpart and Proof of Proposition 1	152
B.2.6	Initialising the Model	154
B.2.6.1	Notes on Calibration	154
B.3	Epstein-Zin Preferences Framework	158
B.3.1	Euler Equations	158
B.4	K-Class Estimators and Critical Values	159
B.5	Consumption: Nondurables Only	159
B.6	Implied ψ (EIS) from Estimates of $1/\psi$	165
B.7	Results Using Raw Returns for Both Reported and Unfiltered Consumption	168
B.8	Results with Restricted Sample	172
B.9	Results for Annual Data when Measurement Errors are Persistent	177

List of figures

Figure 1.1	$\hat{\theta}_k$ vs. Micro Benchmarks	50
Figure 1.2	$\hat{\theta}_k$ vs. Micro Benchmarks – Confidence Intervals	51
Figure 1.3	Confidence Sets Constructed from Restricted Estimations	52
Figure 1.4	Subsample Stability: Implied Coefficients	53
Figure 1.5	Subsample Stability: Homogeneous (Left) vs. Heterogeneous (Right)	54
Figure 2.1	Kalman Gain Over Time for Different Values of ρ_ξ	69
Figure A.1	Implied θ_k vs. Micro Benchmarks – \uparrow Real Rigidities	124
Figure A.2	Implied θ_k vs. Micro Benchmarks – \downarrow Real Rigidities	125
Figure A.3	Implied θ_k vs. Micro Benchmarks – Approach II	126
Figure A.4	Implied θ_k vs. Micro Benchmarks – Approach III	127
Figure A.5	Implied θ_k vs. Micro Benchmarks – Approach IV	128
Figure A.6	Confidence Sets for Each λ_k (Estimated Individually)	129
Figure A.7	CIs Constructed Based on Restricted Estimations	130
Figure A.8	CIs Constructed Based on Restricted Estimations	131
Figure A.9	Subsample Stability: Implied Coefficients ($T = 130$)	132
Figure A.10	Subsample Stability: Homogeneous vs. Heterogeneous ($T = 130$)	133
Figure A.11	Subsample Stability: Deep Parameters (Part I)	134
Figure A.12	Subsample Stability: Deep Parameters (Part II)	135
Figure A.13	Implied θ_k vs. Micro Benchmarks ($\Theta \approx 0.38$)	140

List of tables

Table 1.1	Sectors, Weights and Infrequencies Based on Micro Data	21
Table 1.2	Selected Instruments for (1) – AdaLASSO	26
Table 1.3	Estimations of (1) with $\gamma(L) \in \{1, 1 - \gamma_b L\}$	27
Table 1.4	Estimations of (4) with $\gamma = 0$ and $\gamma \neq 0$	28
Table 1.5	Estimations of Reduced-Form NKPCs – Heterogeneous Economy	40
Table 1.6	Calibration of Parameters Related to Real Rigidities	44
Table 1.7	Implied Slope from Estimations of (19)	46
Table 1.8	Implied Stickiness from the Hybrid Model ($\gamma \neq 0$)	48
Table 1.9	Implied Stickiness from the Forward-Looking Model ($\gamma = 0$)	49
Table 2.1	Calibrated Moments for NIPA Consumption of Nondurables and Services	70
Table 2.2	Calibrated Moments for CEX Consumption: 1982 to 2013	72
Table 2.3	Estimates of the EIS Using K-Class Estimators and Quarterly Data	76
Table 2.4	Estimates of the EIS Using K-Class Estimators and Annual Data	79
Table 2.5	Weak-IV-Robust CIs for the EIS	82
Table 2.6	Heteroscedasticity-Robust Estimates of the EIS	82
Table 2.7	Estimates of the EIS – CEX Data: 1982 to 1996	89
Table 2.8	Estimates of the EIS – CEX Data: 1982 to 2013	91
Table A.1	Data Set for the Heterogeneous Economy	102
Table A.2	Instrument Sets	106
Table A.3	Reduced-Form NKPCs for the Heterogeneous Economy – Approach II	107
Table A.4	Reduced-Form NKPCs for the Heterogeneous Economy – Approach III	108
Table A.5	Reduced-Form NKPCs for the Heterogeneous Economy – Approach IV	109
Table A.6	Implied Slope from Estimations of (19) – Approach II	110
Table A.7	Implied Slope from Estimations of (19) – Approach III	111
Table A.8	Implied Slope from Estimations of (19) – Approach IV	112
Table A.9	Implied Stickiness from the Hybrid Model ($\gamma \neq 0$) – Approach II	113
Table A.10	Implied Stickiness from the Hybrid Model ($\gamma \neq 0$) – Approach III	114
Table A.11	Implied Stickiness from the Hybrid Model ($\gamma \neq 0$) – Approach IV	115
Table A.12	Stickiness from the Forward-Looking Model ($\gamma = 0$) – Approach II	116
Table A.13	Stickiness from the Forward-Looking Model ($\gamma = 0$) – Approach III	117
Table A.14	Stickiness from the Forward-Looking Model ($\gamma = 0$) – Approach IV	118

Table A.15 Implied Slope – $\beta = 0.99$ and $\gamma \in \{0, 0.3, 0.5, 0.7\}$	119
Table A.16 Implied Stickiness – $\beta = 0.99$ and $\gamma \in \{0, 0.3, 0.5, 0.7\}$	120
Table A.17 Estimations of the Model with Indexation à la Gali and Gertler (1999)	123
Table A.18 Implied Slope from (19) <i>Without</i> the Aggregate NKPC	136
Table A.19 Implied Stickiness from (19) <i>Without</i> the Aggregate NKPC	137
Table A.20 Implied Stickiness from (19) <i>Without</i> the Aggregate NKPC ($\gamma = 0$)	138
Table A.21 Estimations of (19) With $\Theta \approx 0.38$	139
Table A.22 Estimations of (19): Initial Values Generated by Sectoral NKPCs	141
Table B.1 Calibrated Moments for NIPA Consumption of Non-durables Only	156
Table B.2 Estimates of the EIS Using K-Class Estimators and Quarterly Data	161
Table B.3 Estimates of the EIS Using K-Class Estimators and Annual Data	162
Table B.4 Weak-IV-Robust CIs for the EIS	163
Table B.5 Heteroscedasticity-Robust Estimates of the EIS	164
Table B.6 Implied EIS from the Estimation of (14) - Nondurables and Services	166
Table B.7 Implied EIS from the Estimation of (14) - Nondurables Only	167
Table B.8 EIS Using K-Class Estimators and Quarterly Data – Raw Returns	169
Table B.9 EIS Using K-Class Estimators and Annual Data – Raw Returns	170
Table B.10 Weak-IV-Robust CIs for the EIS – Raw Returns	171
Table B.11 Heteroscedasticity-Robust Estimates of the EIS – Raw Returns	172
Table B.12 EIS Using K-Class Estimators and Quarterly Data – 1960:2017	173
Table B.13 EIS Using K-Class Estimators and Annual Data – 1940:2017	174
Table B.14 Weak-IV-Robust CIs for the EIS – Restricted Sample	175
Table B.15 Heteroscedasticity-Robust Estimates of the EIS – Restricted Sample	176
Table B.16 Estimates of the EIS – Persistent M.E. and Annual Data	178
Table B.17 Weak-IV-Robust CIs for the EIS – Persistent M.E. and Annual Data	179
Table B.18 Het-Robust Estimates of the EIS – Persistent M.E. and Annual Data	179

List of Abbreviations

AdaLASSO – Adaptive Least Absolute Shrinkage and Selection Operator
CEX – Consumer Expenditures Survey
CUE-GMM – Continuously Updated Estimator
EE – Euler Equation
EIS – Elasticity of Intertemporal Substitution
HAC – Heteroscedasticity and Autocorrelation-Consistent Estimators
HP – Hodrick-Prescott Filter
IV – Instrumental Variables
LAMP – Limited Asset Market Participation Theory
LASSO – Least Absolute Shrinkage and Selection Operator
NIPA – National Income and Product Accounts
NKPC – New-Keynesian Phillips Curve
OLS – Ordinary Least Squares
TSLS – Two-Stage Least Squares
US – United States
WIC – Weighted Irrepresentable Condition
2S-GMM – Two-Step GMM Estimator

Abstract

Estimating the New-Keynesian Phillips Curve (NKPC) can be a perplexing exercise. Its empirical literature is full of uncertainties, which generally arise in the form of econometric puzzles. Perhaps the most crucial one is that even when estimates suggest an economically intuitive positive slope, the implied degree of stickiness in the economy is extremely high to be consistent with the micro evidence. While many have attempted to solve both parts of the problem separately, there seems to be no sufficiently robust method that is immune to criticism. Here we aim to raise the bar by addressing the entire puzzle simultaneously. Specifically, we propose a richer framework of a multi-sector economy with price-setting heterogeneity. Its aggregate NKPC nests many versions in the literature, but also exhibits endogenous terms otherwise omitted. This creates a high dimensional environment, which we exploit using LASSO-based instrument selection. The procedure generates positive and roughly unchanging slope coefficients across econometric settings. Additionally, structural estimations produce degrees of stickiness that substantially approach those exhibited by the micro data, both regarding the entire economy and the cross section of sectors. Importantly, we find that all of these estimates move consistently with implications by theory when model assumptions are modified. Finally, robustness checks indicate that the method performs quite well both in terms of the literature and considering limitations associated with a high dimensional non-linear environment.

1.1**Introduction**

The Phillips Curve has long been a central focus of macroeconomics, and its successes and failures a natural mark of evolution present in the discipline. Laid out back in the 80's and 90's, the New-Keynesian Phillips Curve (henceforth, NKPC) is currently the most widespread theory in that design. However, its early success and widespread adoption contrast with a series of empirical difficulties, which commonly arise in the form of econometric puzzles.

In this paper, we address a central complication in the empirical literature. It is widely known that even when estimations of the NKPC return a positive coefficient for the slope, the implied degree of nominal rigidity (stickiness) is too high, being primarily inconsistent with the micro evidence. Such problem is even more evident when the output gap is used as the forcing (slack) variable in the equation, what has led many authors to propose different models tailored to partially solve this puzzle¹. Nonetheless, few provide sufficiently robust results that are immune to criticism. Here we propose a novel estimation method which simultaneously addresses both parts of that problem, even with the output gap as forcing variable. More precisely, a richer framework of a multi-sector economy with heterogeneity in price setting – à la Calvo (1983) – is proposed. Its aggregate NKPC nests several versions in the literature, but also features sectoral endogenous terms otherwise omitted in the error term. This framework produces a high dimensional environment which we exploit using LASSO-based instrument selection.

Our main finding is that the procedure generates more precise and stable estimates of reduced and structural-form parameters in the NKPC compared to the literature. For instance, structural estimations deliver positive and statistically significant estimates of the slope in *any* of the cases tested. Furthermore, estimated sectoral *and* aggregate Calvo-pricing probabilities materially approach values evidenced in the micro data. We also find that all of our estimations generate theoretically consistent results, in the sense that estimates behave as predicted by theory when we change underlying assumptions of the model. Importantly, our main findings are maintained across different methods and theoretical hypotheses, as regarding the degree of indexation in the economy and that of strategic interactions across sectors.

We conduct a battery of robustness checks with the model, with little to no change in terms of performance. Specifically, we test four approaches to instruments, several calibrations for strategic interactions in price setting, different indexation schemes – including turning this mechanism off altogether – and perturbing which and the number of deep parameters estimated in the

¹Mavroeidis, Plagborg-Møller, and H. Stock (2014) discuss different specifications and methods proposed in the literature. In a seminal paper, Gali and Gertler (1999) attempt to provide more precise estimations of the equation by using the real marginal cost instead of the output gap as the forcing variable. They also introduce lagged inflation terms in the NKPC by assuming that a fraction of firms adopt a backward-looking indexation rule. Estimates of the slope seem to improve substantially with their approach, albeit they still estimate an implied degree of stickiness considerably higher compared to the micro evidence. More recently, Cagliarini, Robinson, and Tran (2011) focus on the latter. Using disaggregated data for Australia, these authors point out that heterogeneity in the frequency of price resetting across firms can help to reconcile micro and macro estimates of price stickiness. They do not evaluate its implications for the slope based on observed data, though, solely conducting a Monte Carlo experiment.

structural form, as well as its sample. Main findings with structural estimations are also not sensitive to the choice for starting values in the algorithm.

The degree of robustness of the method is atypical compared to the empirical literature. Mavroeidis, Plagborg-Møller, and H. Stock (2014) extensively discuss its main puzzles and complications, showing how minor modifications in model specification, instruments and the type of data used can commonly revert the main findings². The model we propose is not tailored to solve specific problems, albeit the high degree of sensitivity evidenced in the literature is nowhere to be found in our estimations.

The NKPC implicitly assumes that inflation is driven by expectations of future real economy activity. Its forward-looking structure typically implies that serial correlation in the error term arises by construction, what complicates limited-information methods³. If the time-series structure of the driving variables and errors were known, one could still reliably estimate the NKPC using heteroscedasticity and autocorrelation-consistent estimators (hereafter, HAC). Nonetheless, as shown by Mavroeidis (2005), a significant caveat of such methods – which also applies to underlying identification tests – is that researchers can not distinguish between weak identification and misspecification of a model⁴.

Specifications are important to the extent that a number of articles in the empirical literature propose different rationales for “intrinsic” persistence in inflation⁵. The motivation for this is essentially empiric, since it implies that lagged inflation terms appear in the NKPC, what often improves the estimations. These terms can be justified through ad-hoc assumptions as: price indexation (Christiano, Eichenbaum, and Evans (2005)); staggered wage contracts (J. Fuhrer and Moore (1995)), or; some backward-looking rule of

²E.g., Rudd and Whelan (2005b) simply replicate Gali and Gertler (1999), the seminal paper in the literature, using the same variables, method and period, but with revised data. Their findings for the slope in the NKPC are substantially different from results originally reported.

³Under a standard approach, expectations are replaced by observed values and expectational errors for future inflation enter the innovation term, which also has the inflation shock (or, the “mark-up shock”). Cross-correlation between the last two implies a serially correlated error term in the NKPC. See Mavroeidis (2004) and Mavroeidis (2005) for a discussion.

⁴Particularly, weak instruments may arise due to the incompleteness of limited-information (usually, single-equation) methods. By neglecting the true dynamics of the model (which necessarily involves the remaining equations), Mavroeidis (2005) argues that weak-identification problems can not be ruled out. This is especially the case when the variables assumed “exogenous” in the equation are actually relevant regarding the entire dynamics of the economy.

⁵“Intrinsic” refers to any form of persistence that is not inherited from other covariates in the model, as the forcing variable and the error term. A typical way to introduce intrinsic persistence for inflation in the model is to rely on some indexation rule. See J. Fuhrer and Moore (1995) for a formal discussion.

thumb (Gali and Gertler (1999) and Galí, Gertler, and López-Salido (2003)). On the one hand, practical results seem to favour persistence in the driving variables⁶. On the other hand, studies using disaggregated price data provide limited support for many of those assumptions – see Bils and P. Klenow (2004), Nakamura and Steinsson (2008) and Nakamura and Steinsson (2013), for example⁷.

A model that does *not* depend on those ad-hoc assumptions to deliver sensible estimates is likely more appropriate. For a multi-sector economy featuring price-setting heterogeneity (as ours), this would signal that the additional structure imposed in the model is capturing relevant information. It can also be a hint that such framework mitigates the misspecification risk in identification.

That is the main motivation of this paper. Fortunately, it is also what we observe in our results. Our proposed framework considerably outperforms typical models in the literature, producing more reliable estimates regardless of *how* and *to which extent* intrinsic persistence is introduced in inflation. Unlike the literature, the method produces substantially stable estimates, maintaining our main findings even when considering a purely forward-looking model that does *not* feature such persistence.

Economies with price-setting heterogeneities across sectors entail different dynamics compared to economies of identical firms. As shown by Carvalho (2006), to replicate the dynamics of the former considering a similar calibration, the latter requires a frequency of price changes that is up to three times lower than the average of the heterogeneous economy. Thus, if one considers an identical-firms economy when the true model features price-setting heterogeneity, estimates of the aggregate nominal rigidity should be biased upwards. With a trivial inverse relationship between this and the slope in that standard model, the latter would be biased towards zero⁸. J. M. Imbs, Jondeau, and Pel-

⁶For example, consider the limit situation where the forcing variable is a white noise process (no persistence). Thus, the slope is zero. In a purely forward-looking NKPC (where inflation is solely determined by its expectations, a slack variable and a cost-push shock), it is straightforward to see that a weak identification setting arises from that. Note that inflation would be solely driven by the cost-push shock, assumed unpredictable. Similarly, good policy and effective anchoring of inflation expectations are also empirically important. If expectations are completely stable, their effect is not statistically discernible in the equation – e.g., see McLeay and Tenreiro (2019).

⁷Typical indexation schemes used in the literature imply that prices would move at a constant rate between more substantial resets, a feature generally not evidenced by disaggregated data.

Most importantly, these articles produce evidence against a constant size of price changes in the data, what directly contradicts typical indexation schemes used in the literature.

⁸As discussed later, such relationship is not so simple for the heterogeneous economy. In that case, the effect of the degree of stickiness on the slope will also depend on how the former is distributed across sectors.

grin (2007), J. Imbs, Jondeau, and Pelgrin (2011) and Cagliarini, Robinson, and Tran (2011) empirically confirm these facts.

Strategic complementarities between price-setting decisions across sectors arise in our model by segmented labour markets. This implies that more sticky sectors exert a disproportionate effect on the degree of adjustment of the aggregate price level⁹. Everything else constant, the heterogeneous economy behaves as if it featured more nominal rigidities. It follows that the *lower* the *actual* nominal rigidity needed to reproduce the same dynamics of an identical-firms model. Hence, the presence of strategic complementarities provides a useful channel through which the model can reconcile macro and micro estimates of stickiness.

This channel should be more evident the higher the degree of strategic complementarities in pricing behaviour and testing it empirically is the second motivation of this paper. If strategic complementarities across sectors are strong enough, the more flexible sectors adjust their prices considerably less than they otherwise would, in response to a higher demand. In the limit, aggregate prices do not change much for a large output effect, what translates into a slope that approaches zero. Similarly, the degree of stickiness to generate such behaviour need not to be as high as estimated in the literature. Evidently, a multi-sector economy provides several forms to empirically test the model for theoretical consistency, which we exploit throughout this paper. In contrast, such tools are unavailable for identical-firms economies.

The model attests to the aforementioned mechanism. To verify that, we estimate deep parameters related to nominal rigidities in the structural form of the model, while calibrating the degree of strategic complementarities across sectors. We systematically estimate a *lower* degree of price stickiness (approaching that implied by the micro evidence) and a *lower* slope as we introduce more strategic complementarities. The converse is also true. Furthermore, we find that our estimates are more sensitive to such parameterisations under the purely forward-looking model (without indexation schemes). In contrast, little sensitivity is found for the hybrid model (with indexation). This is not surprising to the extent that prices of all varieties change every period when indexation is introduced, affecting how strategic interactions across sectors operate in the economy and implications of their channel for both the slope and the degree of stickiness. Nonetheless, when we allow for indexation, our baseline estimations imply an aggregate expected duration of price spells roughly from 5.7 to 7.3 months. These values are broadly in line with the micro evidence,

⁹This is a well known fact when prices of different suppliers are strategic complements. See Woodford (2003, Ch. 3), for example.

which usually delivers estimates in the interval from 4 to 9 months¹⁰.

We adopt a General Method of Moments (GMM) approach with LASSO-based instrument selection as our estimation strategy¹¹. For an NKPC of an identical-firms economy, Berriel, M. Medeiros, and Sena (2016) showed how the application of different regularisation techniques to choose instruments can deliver more disciplined and reliable estimates in face of model uncertainty and potentially weak instruments. Specifically, in light of their findings, we select instruments based on the so-called Adaptive LASSO (AdaLASSO) estimator¹². M. C. Medeiros and Mendes (2016) review properties of such estimator in a flexible econometric environment (which allows for non-Gaussian and conditionally heteroscedastic errors), also conducting an application to forecast US inflation. They show that the AdaLASSO consistently chooses the relevant variables as the number of observations increases (commonly referred to as “model selection consistency”) and estimates coefficients as if the correct set of variables were known ex-ante by the econometrician (“Oracle property”).

Our extension of their approach to a multi-sector economy is natural. First, under rational expectations the set of pre-determined candidates as instruments grows with the number of sectors assumed in the economy. For example, lagged terms related to the output of sectors are manifest candidates. Second, there may be reasons to be sceptic about ad-hoc based instrument selection with a potentially infinite number of candidates. In any case, as robustness check we also present results for an (ad-hoc) instrument set which attempts to mimic standard choices in the literature while adjusting the pattern for a heterogeneous environment. Such setting is also convenient to disentangle contributions of the model from those of the instrument selection routine in producing our results. Our findings suggest that both are important, albeit the former more decisive. In spite of that, main results are maintained with the ad-hoc instrument set.

In the next section, we discuss the data used in this paper. In section 1.3, we address the econometric implications of the presence of price-setting heterogeneity in the economy and the consequences of incorrectly assuming

¹⁰Values obtained with micro data can vary substantially depending on whether sales are included or not. See Bils and P. Klenow (2004), Nakamura and Steinsson (2008) and Nakamura and Steinsson (2013), for useful references. Estimated durations mentioned for our model refer to the baseline instrument set, calibrating for three different degrees of strategic complementarities in price setting. These results are in the main text.

¹¹The Least Absolute Shrinkage and Selection Operator (LASSO) – Tibshirani (1996) – shrinks to zero parameters associated with redundant predictors while estimating coefficients for relevant regressors.

¹²The method we use here, based on the AdaLASSO estimator of Zou (2006), is a direct adaptation of the approach found by Berriel, M. Medeiros, and Sena (2016) to best perform in their simulated environment.

a homogeneous economy. In addition, it shows how an estimation strategy that accounts for such form of heterogeneity when choosing instruments can improve the estimations of the NKPC *even when assuming an identical-firms model*. This not only suggests that heterogeneity is relevant, but also that sectoral terms – which would be included in the model under that hypothesis – can exert an important role in identifying the NKPC. We turn to the complete heterogeneous economy in section 1.4. There, the main findings of this paper are discussed based on reduced and structural-form estimations with the model. Robustness checks are also presented in that section. Finally, section 1.5 concludes.

1.2 Data

The data set consists of aggregate and sectoral quarterly data for the US economy during the interval from 1964:2 to 2017:3. For aggregate data, we use the real gross domestic product (output), the producer price index for all commodities, a 5-year Treasury rate spread (over the Fed Funds rate), the effective Fed Funds rate (interest rate of the economy), the non-farm labour share (as a proxy for the real marginal cost) and changes in average hourly earnings of production and non-supervisory workers (wage inflation). For sectoral data, we use real personal consumption expenditures (output of the sectors), inflation and the price level¹³. Fifteen sectors are assumed in our heterogeneous economy¹⁴.

Table A.1 (appendix) summarises our data set, data sources and applied transformations. Table 1.1 below details the sectors in the economy, providing weights and implied Calvo-pricing probabilities – calculated from micro data in Bils and P. Klenow (2004)¹⁵. These infrequencies are adopted as a benchmark to estimations we later perform in this paper.

¹³There is no available data for income at the sector level. Hence, we proxy it by sectoral consumption.

¹⁴Additionally, cyclical components are extracted by the Hodrick-Prescott filter (HP), setting $\lambda_{HP} = 1600$ (as advised for quarterly data). We are aware of potential problems involving the filter – e.g. James D Hamilton (2018). Main findings of this paper do not change when using the Baxter-King filter, the Beveridge-Nelson decomposition or when quadratically detrending the data. To account for a model with no population growth, we also define main variables in per capita terms, when appropriate.

¹⁵To calculate weights, we average the ratio between the nominal consumption in the sector and that in the entire economy over the sample.

Table 1.1: Sectors, Weights and Infrequencies Based on Micro Data

Sector	Sectoral Weight	Benchmark Infrequency
Motor Vehicles and Parts	5.34%	0.212
Furnishings and Durable Household Equipment	3.61%	0.484
Recreational Goods and Vehicles	2.92%	0.564
Other Durable Goods	1.59%	0.551
Food and Beverages Purchased for Off-Premises Consumption	11.82%	0.327
Clothing and Footwear	5.41%	0.331
Gasoline and Other Energy Goods	3.75%	0.003
Other Nondurable Goods	8.03%	0.541
Housing and Utilities	18.19%	0.212
Health Care	11.79%	0.857
Transportation Services	3.23%	0.375
Recreation Services	3.02%	0.727
Food Services	6.50%	0.590
Financial Services and Insurance	6.44%	0.781
Other Services	8.34%	0.645

Notes: Benchmark infrequencies are implied Calvo-pricing probabilities based on micro data. These come from a mapping between disaggregated PCE price data and early evidence exhibited in Bils and P. Klenow (2004), both expressed monthly. We convert probabilities into quarterly analogues by compounding them geometrically.

1.3

Heterogeneity, Omitted Terms and the Literature

A baseline specification for the NKPC that nests a number of conventional variations proposed in its empirical literature is:

$$\gamma(L)\pi_t = \gamma_f E_t(\pi_{t+1}) + \kappa x_t + \alpha' w_t + u_t, \quad (1)$$

where $\gamma(L) = 1 - \gamma_1 L - \dots - \gamma_p L^p$ is a p -order lag characteristic polynomial, x_t represents the forcing variable (usually either the output gap or the real marginal cost)¹⁶, w_t additional controls and u_t is an unobserved error term (commonly denoted as a mark-up shock). If $\alpha = 0$, the purely forward-looking NKPC is obtained when $\gamma(L) = 1$, while conventional hybrid versions with some indexation scheme – as in Galí and Gertler (1999), Sbordone (2005) and Christiano, Eichenbaum, and Evans (2005) – are obtained with $\gamma(L) = 1 - \gamma_b L$.

If relevant variables in w_t are disregarded in (1), these enter the composite error term, then formed by $e_t \equiv \alpha' w_t + u_t$ ¹⁷. This reason alone is not sufficient to estimate inconsistent coefficients in the NKPC. However, manifest identification issues arise if relevant variables embedded in the error term are also correlated with covariates included in (1)¹⁸.

¹⁶Under a Cobb-Douglas assumption for technology, the output gap and the real marginal cost are proportionately related in terms of log-deviations from the steady state. Thus, both can be used as the forcing variable.

¹⁷A typical instrumental variables approach would also introduce an expectational error relative to inflation next period in e_t . We omit it here for expository purposes.

¹⁸Instrumental variables estimators would need instruments that are little correlated with omitted variables and relevant regarding the endogenous variables in the equation, what easily puzzles macro models.

The assumption of an identical-firms economy is ubiquitous in the empirical literature and under this premise a correlation between covariates and omitted controls becomes central. Compared to that framework, heterogeneity in pricing behaviour across sectors produces a composite endogenous term in (1). With vast empirical support for the latter¹⁹, it must be included in the controls, w_t . The term is proportional to a weighted average of sectoral output gaps when x_t is the aggregate output gap (henceforth, just output gap). If the real marginal cost is the forcing variable, the term encompasses two main components, the first proportional to a weighted average of sectoral output gaps, and a second, proportional to the output gap. In any case, the assumption of an identical-firms economy implies an inflexible correlation between the forcing variable and endogenous components in the composite error.

To illustrate the problem, a typical representation for an economy that features heterogeneity in price stickiness across sectors takes the form²⁰:

$$\gamma(L)\pi_t = \gamma_f E_t(\pi_{t+1}) + \phi y_t + \int_0^1 \phi_k f(k) y_{k,t} dk + \delta' h_t + u_t, \quad (2)$$

where $y_t = \int_{k=1}^K f(k) y_k dk$ is the output gap, y_k denotes the output gap of sector k , ϕ is the slope, ϕ_k is a function of deep parameters specific at the sector level, h_t are additional variables²¹ and $f(k)$ is a density function that measures the distribution of firms across sectors $k \in \{1, \dots, K\}$ – i.e., the weight of such sectors.

Note that when (1) is assumed instead of (2), then $e_t = \int_0^1 \phi_k f(k) y_{k,t} dk + u_t$ ²². In addition, the association between the degree of nominal rigidity in the economy and the slope ϕ is non-trivial with heterogeneity in price setting across sectors, contrasting with the case of an identical-firms economy, κ . Researchers would typically identify a different slope, since ϕ also depends on the cross-sectional distribution of the frequency of price changes as well as on the degree of strategic complementarities across sectors²³.

¹⁹See Blinder et al. (1998), Bils and P. Klenow (2004), Nakamura and Steinsson (2008), P. J. Klenow and Kryvtsov (2008) and Vermeulen et al. (2007), for example.

²⁰Using the output gap as the forcing variable. An analogue of (2) with the real marginal cost is derived in section 4. Examples of (2) in the literature can be seen in Aoki (2001), Carvalho (2006), Carvalho and Lee (2011) and Carvalho and Nechio (2016).

²¹For instance, if capital is also an input in addition to labour, then the aggregate investment must be included in h_t .

²²Again omitting the expectational error – see footnote 17.

²³Reasons for a different mechanism in the slope vary, but frequently relate to the presence of strategic interactions between firms. For example, Sbordone (2007) models an open economy with Kimball (1995) preferences to quantitatively measure the effect of an increase in trade on the slope of the NKPC. Since with Kimball preferences the elasticity of demand depends on the firm's relative sales, another sort of strategic complementarity arises. Everything else constant, this amplifies the effects of nominal disturbances, reducing the slope.

When the true model is (2), it may be obvious how estimates by ordinary least squares (OLS) of the coefficients in (1) can be inconsistent. However, such type of misspecification may also forcibly undermine identification in significant ways, even under linear instrumental variables (IV) estimators.

The first complication of assuming (1) is that it eventually creates a “mirror” between identification strength and serial autocorrelation in the composite error e_t . If the output gap exhibits substantial autocorrelation, – such that its lags are relevant instruments for itself, hence being suggestive candidates –, so would the error term, caused by the presence of the sectoral output gaps in it²⁴. Difficulties to disentangle how persistence in the forcing variable and in the error contribute to the dynamics of inflation (e.g., its own persistence) have been addressed in the literature – e.g., in J. Fuhrer and Moore (1995), Del Negro and Schorfheide (2008), Schorfheide (2008) and Nason and G. W. Smith (2008). Nevertheless, the aforementioned connection between strength of instruments and serial correlation in the error is a singular issue, arising by a very specific form of misspecification in the case the true model exhibits heterogeneity in price setting²⁵. Furthermore, even though HAC methods can be proposed based on autocorrelation of arbitrary form, there would be at least some loss of efficiency by neglecting that particular relationship. Eventually, resulting higher standard errors might ultimately explain why estimates of the slope κ in (1) are frequently not statistically significant in the literature.

Imposing an identical-firms framework also generates a second and more determinant problem regarding the choice of instruments. Since the most widespread procedure in the literature constructs moments by exclusion restrictions (excluding lags of endogenous variables from the model while using them as instruments), then lags of sectoral output gaps in (2) are not taken into

²⁴One can derive a counterpart of (2) where sectoral relative prices ($p_k - p_t$) instead of output gaps (y_k) are embedded in the composite endogenous term. In any case, that “mirror” would still be present, albeit such relationship would be determined by the terms related to aggregate inflation in (2) and the composite error. It happens since, by construction, sectoral prices sum to the aggregate price level exactly as sectoral output gaps sum to the output gap. The model and underlying details are shown in the next section.

²⁵J. M. Imbs, Jondeau, and Pelgrin (2007) and J. Imbs, Jondeau, and Pelgrin (2011) use a similar motivation. However, their Phillips curves (sectoral and aggregate) use the real marginal cost as the forcing variable, what led them to empirically test their predictions based on French data – the sector share of labour income, the straightforward proxy for sectoral marginal costs, is not available for the US economy. They also adjust their model for cross-correlations between industries. We later use a model for the US economy with the output gap as the forcing variable. As mentioned earlier, results in the literature suggest that its use is more prone to generate puzzling results for the NKPC. In contrast, we later demonstrate sensible estimations with it, also showing that our model does *not* need any correction for cross-industry linkages in order to work.

account as potential candidates when estimating (1)²⁶. Yet, some of these may contain valuable information. It is possible that lags of the output gap are not the best instruments for itself in (1), but instead a measure of income in some sector. Since sectoral output gaps enter the error term, e_t , an optimal routine would pick lags of these variables that are most correlated with the output gap, but least correlated with the output gap of other sectors. The literature is oblivious to possible estimation gains by exploiting such relationship²⁷.

This framework also enlightens potential advantages of a data-driven instrument selection routine (e.g., based on LASSO estimators) over main ad-hoc procedures in the literature *even under the possibly misspecified NKPC* (1). First, if the true model exhibits price-setting heterogeneity across sectors, lags of sectoral variables are natural instrument candidates that shall be considered. Clearly, choosing instruments among such an extensive set of candidates is non-trivial. Second, a LASSO-based routine does exactly as depicted above. By construction, it picks those that *best* correlate with endogenous variables in the equation (e.g., the output gap). Additionally, if the so-called Weighted Irrepresentable Condition (WIC)²⁸ holds, it also chooses those that *least* correlate with the other (lagged) sectoral variables (also candidates)²⁹. Hence, autocorrelation in the sectoral (omitted) variables directly corresponds to correlation between them and other candidates in the error. If such autocorrelation is high enough, instruments picked by the AdaLASSO technique will also have *little* correlation with the omitted error³⁰.

To empirically confirm those advantages, we select instruments based on the AdaLASSO estimator – Zou (2006). Berriel, M. Medeiros, and Sena (2016) show how such approach provides more reliable and disciplined estimations for the NKPC of an identical-firms economy, (1). It also outperforms several

²⁶The approach of constructing of moments by exclusion restrictions is central in the critique of Mavroeidis (2005) and has its roots in Gali and Gertler (1999) and Roberts (1995).

²⁷The literature typically considers a very limited set formed by lags of the variables in the equation and some additional instruments outside of the model (e.g., interest rates and some commodity index). We refer to Table 3 in Mavroeidis, Plagborg-Møller, and H. Stock (2014) as an excellent summary of specifications and instruments commonly applied.

²⁸The LASSO estimator requires two simple assumptions in order to work. The first is sparsity (only a small number of variables can be relevant). The second one is the irrepresentable condition, that simply imposes an upper bound for the correlation between relevant and irrelevant variables. The WIC is its less restrictive generalisation for the AdaLASSO.

²⁹The WIC is not formally testable, but a common method – adopted in this paper – is to check whether estimated coefficients increase monotonically when the penalty factor is loosened along the regularisation path.

³⁰In addition to the aforementioned check on the regularisation path, note that the WIC is likely valid in our framework, with a significant portion of the candidates being sectoral data.

other LASSO-based procedures tested in their environment³¹. Here we adapt their estimation strategy to additionally include lags of sectoral variables as candidates. Specifically, for each \mathcal{Y} , endogenous variable in the NKPC, we run:

$$\hat{\rho} = \underset{\rho}{\operatorname{argmin}} \|\mathcal{Y} - Z\rho\|_2^2 + \Lambda \sum_{j=1}^P w_j \|\rho_j\|_1, \quad (3)$$

where $\|\cdot\|_p$ is the ℓ^p norm, ρ is a $P \times 1$ vector of coefficients, Z is a $T \times P$ matrix of instrument candidates, $\{z_1, \dots, z_p, \dots, z_P\}$. Λ controls the shrinkage whereas $w_j = |\tilde{\rho}_j|^{-\tau}$ is a candidate-specific penalty weight formed by a preliminary LASSO estimator $\tilde{\rho}_j$. Finally, $\tau = 1$ is a common choice. Candidate z_p is selected as instrument if $\rho_p \neq 0$, its coefficient in (3). All the results in this paper are maintained when selecting instruments based on the LASSO estimator, instead of (3).

Table 1.2 below shows selected instruments when candidates are lags of inflation and the output gap – relative to the endogenous variables in (1), with no controls³² – as well as of the fifteen sectoral output gaps in the economy and variables predominantly used as instruments in the literature: the non-farm labour share; wage inflation; Fed Funds rate; Treasury spread, and; inflation in commodities. We present results based on lags $t-1$ and $t-2$ of these variables. Lagging a third time produces the same results as lagging twice. Note that the output gap is replaced by the output gap of “other nondurable goods” when we select based on the first lag. In this case, *two* out of only *five* selected variables refer to sectors, what suggests that these type of data are indeed important for the estimation strategy and should not be neglected. In addition, except for wage inflation and the labour share, remaining instruments commonly used in the literature are rejected.

Next, we proceed to conduct GMM estimations of (1) with the instrument sets of Table 1.2. When selecting based on the first lag, we apply the approach found to best perform by Berriel, M. Medeiros, and Sena (2016), setting the first three lags of selected variables as instruments, $\{z_{t-1}, z_{t-2}, z_{t-3}\}$ ³³. This is our baseline approach for instruments. For a selection based on the second lag, we only use that lag as instrument, z_{t-2} . This variation is intended to

³¹Being more specific, the approach that provides the best results in their simulations is called “AdaLASSO observables” in Berriel, M. Medeiros, and Sena (2016). Our baseline instrument selection follows that routine.

³²Following the literature, we consider π_{t-1} as exogenous.

³³One could alternatively construct an instrument candidates set comprised of different lags of each variable – e.g., the first three lags, in that case. Berriel, M. Medeiros, and Sena (2016) showed little differences in performance between this alternative and the routine we apply here. Results of our paper do not seem sensitive to this difference, either, as long as both are constructed based on the same lag structure (e.g., selecting based on the first lag and applying the first three lags of selected variables vs. considering all those three lags as candidates and applying each selected lag separately, for instance).

Table 1.2: Selected Instruments for (1) – AdaLASSO

z_{t-1}	z_{t-2}
Other Nondurable Goods	Output Gap
Gasoline and Other Energy Goods	Gasoline and Other Energy Goods
Inflation	Inflation
Non-Farm Labour Share	Non-Farm Labour Share
Wage Inflation	Wage Inflation

Notes: Selected instruments based on (3) for endogenous covariates in (1). We construct candidate instruments in Z by lagging the following variables once (z_{t-1}) and twice (z_{t-2}): inflation, output gap, output gaps for the fifteen sectors in Table 1.1, non-farm labour share, wage inflation, Fed Funds rate, 5-year Treasury spread and commodities inflation. Lagging a third time gives the same result as lagging twice. For conciseness, for sectoral output gaps we write just the name of the sector.

control for the well known time-aggregation bias in macro data – e.g., R. E. Hall (1988) – as well as to avoid the common pitfall of too many instruments – see Bårdsen, Jansen, and Nymoen (2004), Andrews and Stock (2005) and C. Hansen, Hausman, and W. Newey (2008), for example³⁴.

Our results are summarised in Table 1.3³⁵. Expectations in (1) are replaced by actual values, so that that an expectational error enters the disturbance term. We address both the purely forwarding-looking NKPC ($\gamma(L) = 1$) and the typical hybrid NKPC ($\gamma(L) = 1 - \gamma_b L$), considering either the output gap or the real marginal cost (proxied by the non-farm labour share) as the forcing variable in each case. For robustness, the lower end of the table presents estimates when coefficients are calibrated to $\gamma_f = 0.99$ (forward-looking model) and to $\gamma_f = 2/3$ and $\gamma_b = 1/3$ (hybrid version)³⁶. To better compare results, the last two columns consider an instrument set that mimics a typical choice in the literature. It includes as instruments the first three lags of the output gap, inflation, labour share, Treasury spread, wage inflation and commodities inflation. Gali and Gertler (1999) applied very similar instruments.

³⁴A common workaround to control for the presence of time-aggregation bias in macro aggregates is to use instruments lagged at least twice. In addition, the use of too many instruments often biases Two-Stage Least Squares (TSLS) estimators towards the OLS limit distribution. This is more evident the weaker the instruments are. For this reason, other authors – e.g., Mavroeidis, Plagborg-Møller, and H. Stock (2014) – explicitly advice against the use of too many instruments in the estimation of the NKPC.

³⁵We use the standard Two-Step GMM estimator (2S-GMM) in Table 1.3. Main findings are maintained when we use the the continuously updated estimator (CUE-GMM) of L. P. Hansen, Heaton, and Yaron (1996), which is known to be more robust under weak identification – see Stock, Wright, and Yogo (2002), for example.

³⁶For the forward-looking model, it follows that $\gamma_f = \beta$, justifying the former, a typical value used for this parameter in the literature. A commonly used indexation scheme – where firms index based on a fraction γ of past inflation – implies $\gamma_f = \beta/(1 + \beta\gamma)$ and $\gamma_b = \gamma/(1 + \beta\gamma)$ for the hybrid version. Values $2/3$ and $1/3$ are an approximation for $\beta = 0.99$ and $\gamma = 0.5$ in that case. It also produces reduced-form coefficients that are close to estimates obtained when these are free to vary.

Table 1.3: Estimations of (1) with $\gamma(L) \in \{1, 1 - \gamma_b L\}$

Model	Coefficient	AdaLASSO, z_{t-1}		AdaLASSO, z_{t-2}		Literature	
		Output Gap	Marginal Cost	Output Gap	Marginal Cost	Output Gap	Marginal Cost
$\gamma(L) = 1$	κ	-0.027*** (0.004)	0.023*** (0.005)	-0.011 (0.011)	0.015 (0.013)	-0.007 (0.007)	0.010 (0.009)
	γ_f	0.996 (0.004)	0.968 (0.005)	1.002 (0.012)	1.003 (0.014)	1.000 (0.008)	1.003 (0.009)
$\gamma(L) = 1 - \gamma_b L$	κ	0.002** (0.001)	0.004*** (0.000)	0.004 (0.010)	0.009 (0.000)	0.005*** (0.001)	-0.009*** (0.002)
	γ_f	0.688 (0.010)	0.720 (0.001)	0.811 (0.108)	0.769 (0.009)	0.599 (0.005)	0.655 (0.009)
	γ_b	0.308 (0.010)	0.282 (0.001)	0.181 (0.105)	0.226 (0.132)	0.388 (0.004)	0.344 (0.009)
γ_f and γ_b Calibrated							
$\gamma(L) = 1, \gamma_f = 0.99$	κ	-0.024*** (0.004)	-0.013 (0.007)	-0.011 (0.011)	0.014 (0.014)	-0.009 (0.006)	0.011 (0.009)
$\gamma(L) = 1 - \frac{1}{3}L, \gamma_f = \frac{2}{3}$	κ	0.001 (0.001)	0.005*** (0.001)	0.000 (0.003)	0.001 (0.005)	0.002*** (0.001)	-0.007*** (0.001)

Notes: 2S-GMM estimates of reduced-form coefficients in (1) using the output gap and the real marginal cost as the forcing variable. Both the purely forward-looking and the typical hybrid models are tested. We consider two routines for instrument selection. The first selects instruments by (3), considering the first lag of the same candidates – see notes for Table 1.2. Instruments are the first three lags of selected candidates in this case. The second performs the same routine but selecting and applying as instruments only the second lag of those variables. “Literature” mimics a common choice of instruments in the literature – see the main text. The lower end of the table presents restricted estimates when $\gamma_f = 0.99$ (forward-looking model) and $\gamma_f = 2/3$ and $\gamma_b = 1/3$ (hybrid). HAC standard errors are presented in parentheses. The hypothesis of a statistically insignificant coefficient for κ is tested: *p<0.1; **p<0.05; ***p<0.01.

Table 1.3 suggests some advantages of the data-driven instrument selection³⁷. This is especially the case using the marginal cost, for which the slope κ is generally positive but not higher than 0.02. Even when very close to zero, these coefficients are often statistically significant when selecting based on the first lag. The output gap produces similar findings for the hybrid NKPC, albeit providing more mixed results for the forward-looking model, under which we estimate a negative and significant slope in some cases. The instrument set that mimics the literature performs poorly, in comparison. It generates a number of negative estimates of κ , these being statistically significant for the marginal cost. Choosing instruments based on the literature also seems to imply a lower (higher) coefficient of inflation expectations (past inflation).

It is also possible to evaluate how the selection of instruments performs with the structural analogues of (1). We test two NKPCs (forward-looking and hybrid) based on:

$$\pi_t = \frac{\beta}{1 + \beta\gamma} E_t \pi_{t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{t-1} + \frac{(1 - \theta)(1 - \beta\theta)}{\theta(1 + \beta\gamma)} mc_t^r + u_t, \quad (4)$$

where mc_t^r denotes the real marginal cost and θ is the Calvo-pricing probability³⁸. It is assumed that firms that can not reset their prices in t follow

³⁷For expository reasons, we do not present results for significance tests on other coefficients but the slope. These are statistically significant at least at 5%.

³⁸We are not considering the output gap as forcing variable here for simplicity reasons. In this case, the analogue of (4) would involve other deep parameters, – for example, related to the curvature of the production function –, which would have to be calibrated. The model

Table 1.4: Estimations of (4) with $\gamma = 0$ and $\gamma \neq 0$

Model	Coefficient	AdaLASSO, z_{t-1}		AdaLASSO, z_{t-2}		Literature		
		(1)	(2)	(1)	(2)	(1)	(2)	
$\gamma = 0$	κ	0.184***	0.233***	0.062***	0.069***	0.029***	0.099***	
		(0.023)	(0.029)	(0.015)	(0.018)	(0.006)	(0.014)	
	θ	0.653	0.622	0.781	0.772	0.843	0.731	
		(0.018)	(0.017)	(0.023)	(0.025)	(0.014)	(0.016)	
	β	0.999	0.991	0.995	0.994	0.999	0.998	
		(0.024)	(0.026)	(0.015)	(0.015)	(0.013)	(0.016)	
$\gamma \neq 0$	κ	0.081**	0.128***	0.019	0.030*	0.000	0.072***	
		(0.019)	(0.022)	(0.012)	(0.016)	(0.001)	(0.015)	
	θ	0.713	0.658	0.851	0.817	0.999	0.713	
		(0.028)	(0.020)	(0.037)	(0.034)	(0.324)	(0.026)	
	β	0.999	0.998	0.997	0.997	0.999	0.999	
		(0.028)	(0.030)	(0.012)	(0.014)	(0.002)	(0.036)	
	γ	0.425	0.395	0.420	0.369	0.636	0.594	
		(0.094)	(0.080)	(0.237)	(0.234)	(0.008)	(0.094)	
	β and γ Calibrated							
	$\beta = 0.99$	κ	0.000	0.108***	0.003	0.076***	0.000	0.045***
(0.007)			(0.017)	(0.014)	(0.019)	(0.006)	(0.007)	
θ		0.999	0.725	0.948	0.762	0.999	0.813	
		(0.648)	(0.018)	(0.111)	(0.026)	(0.592)	(0.013)	
$\beta = 0.99, \gamma = 0.5$	κ	0.000	0.011***	0.000	0.012**	0.000	0.004***	
		(0.002)	(0.001)	(0.006)	(0.006)	(0.001)	(0.001)	
	θ	0.999	0.880	0.999	0.870	0.999	0.925	
		(0.309)	(0.007)	(0.838)	(0.032)	(0.180)	(0.011)	

Notes: 2S-GMM estimates of deep parameters in (4). Both the textbook NKPC ($\gamma = 0$) and a standard hybrid version ($\gamma \neq 0$) are tested. We construct orthogonality conditions based on equations (5) and (6), here referred as (1) and (2), for expository reasons. Instrument sets are the same as in Table 1.3. Parameters θ , β and γ are free to vary in the range (0, 1). The lower end of the table presents estimates while fixing $\beta = 0.99$ and $\gamma = 0.5$. HAC standard errors are presented in parentheses. The hypothesis of a statistically insignificant coefficient for κ is tested: *p<0.1; **p<0.05; ***p<0.01.

an indexation scheme on past inflation governed by γ ³⁹. The forward-looking model is obtained by fixing $\gamma = 0$. Since non-linear GMM can be sensitive to the normalisation under which moment conditions are derived, we construct orthogonality conditions based on two variations⁴⁰.

we later use for the heterogeneous economy features constant returns to scale, so that such parameters are not present in the remainder of this paper.

³⁹That indexation rule is standard in the literature. Formally, it follows that firm j , which can not re-optimize in t , fixes its prices according to $P_{j,t} = P_{j,t-1} \left(\frac{P_{t-1}}{P_{t-2}} \right)^\gamma$, where P_t denotes the price level of the entire economy in period t .

⁴⁰The first normalisation comes directly from (4). In line with Galí and Gertler (1999), we multiply that by θ to construct a second one. For each instrument z_t , these respectively take the form:

$$E_t[(\pi_t - \theta^{-1}(1 + \beta\gamma)^{-1}(1 - \theta)(1 - \beta\theta)mc_t^r - \beta(1 + \beta\gamma)^{-1}\pi_{t+1} - \gamma(1 + \beta\gamma)^{-1}\pi_{t-1})'z_t] = 0 \quad (5)$$

$$E_t[(\theta\pi_t - (1 + \beta\gamma)^{-1}(1 - \theta)(1 - \beta\theta)mc_t^r - \beta\theta(1 + \beta\gamma)^{-1}\pi_{t+1} - \gamma\theta(1 + \beta\gamma)^{-1}\pi_{t-1})'z_t] = 0 \quad (6)$$

Table 1.4 exhibits our findings⁴¹. When all parameters are estimated (upper half), the baseline instrument selection approach outperforms the others. First, it produces lower estimates of the Calvo-pricing probability. Second, slope coefficients are statistically significant and positive, in line with the intuition. Since under an identical-firms economy both coefficients are directly (inversely) related, those estimates of the slope are also slightly higher. Selecting based on the second lag seems to perform better than mimicking usual choices in the literature, albeit the degree of nominal rigidity is too high, with $\hat{\theta}$ around 0.8. The ad-hoc instrument set based on the literature once more seems to generate less precise results, with estimates generally being more sensitive to the way one normalises the moment conditions. In addition, it often produces unrealistic estimates of the Calvo-pricing parameter, which usually lie around or at the upper bound. The indexation parameter γ is estimated at approximately 0.4 when selecting instruments, but at a higher value of 0.6 when choosing instruments based on the literature. Our findings are weaker for the case when β and γ are calibrated, albeit it is still noticeable that higher coefficients for θ are again produced when using typical instruments in the literature.

The estimations demonstrate how accounting for sectoral terms from the heterogeneous economy to construct orthogonality conditions can provide more reliable results *even when estimating the (misspecified) NKPC of an identical-firms economy*. The resulting presence of many instrument candidates at the sectoral level is exploited using an AdaLASSO-based instrument selection. Nonetheless, researchers could apply alternative regularisation routines. Importantly, the method generates a lower degree of nominal rigidity in the economy, given by θ , which approaches the upper bound implied by micro data. Specifically, we often estimate a Calvo-pricing probability slightly higher than 0.60, while the micro evidence would suggest values not much higher than it. For example, Bils and P. Klenow (2004) report that half of prices in their sample last less than 5.5 months when excluding sales. This would imply $\theta \leq 0.46$ when mapping monthly into quarterly data. Similarly, Nakamura and Steinsson (2008) report a median duration of prices from 7 to 9 months ignoring sales, implying $\theta \in (0.57, 0.66)$.

The next step is to consider the heterogeneous economy in the estimations – i.e., we shall proceed to estimate (2). The following section briefly discusses the multi-sector model we use for that.

⁴¹Deep parameters in (4) theoretically lie in the range from 0 to 1. We discount 10^{-3} in each bound of this interval in order to properly provide the Jacobian matrix to the algorithm.

1.4

Inflation Dynamics in a Multi-Sector Framework

Our economy features heterogeneity in price setting across firms of different sectors. This framework is similar to that in Carvalho (2006), except for the assumption of an indexation scheme. We later show that our main findings are still valid if such rule is dropped entirely. For conciseness reasons, we only exhibit crucial parts of the model in this section. The completely specified economy is presented in the appendix.

The model features a continuum of firms kj , each producing the consumption variety of the good $j \in [0, 1]$ of sector $k \in [0, 1]$. These firms are monopolistically competitive, hiring labour based on a linear technology function. A representative household, which owns these firms, also supplies firm-specific labour to them⁴². This consumer derives utility from a Dixit-Stiglitz composite of differentiated consumption goods in the economy. It is assumed that each firm kj fixes its price as in Calvo (1983). The probability of a price change is sector-specific, denoted by λ_k ⁴³. A density function $f(\cdot)$ establishes the distribution of firms across sectors.

The representative household solves:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma} - 1}{1-\sigma} - \int_0^1 f(k) \int_0^1 \frac{L_{kj,t}^{1+\varphi^{-1}}}{1+\varphi^{-1}} dj dk \right),$$

subject to:

$$P_t C_t = \int_0^1 f(k) \int_0^1 L_{kj,t} W_{kj,t} dj dk + T_t + I_{t-1} B_{t-1} - B_t,$$

where $f(k)$ denotes the weight of sector k in the aggregate output⁴⁴, C_t is consumption of the composite good, $L_{kj,t}$ labour in firm kj , $W_{kj,t}$ denotes nominal wages related to the latter, P_t is the aggregate price index, T_t are firms' profits distributed by lump sum transfers and B_t denotes bond holdings that collect a gross interest I_t each period⁴⁵. Additionally, β is the discount factor, σ is the inverse of the elasticity of inter-temporal substitution (EIS) in consumption and φ is the Frisch elasticity of labour supply.

⁴²Firm-specific labour is the reason why strategic complementarities in price setting arise in this model. Factor specificity is not always the source, though. See Woodford (2003, chapter 3) for a detailed discussion on this topic.

⁴³Equivalently, the Calvo-pricing probability in sector k is $\theta_k \equiv 1 - \lambda_k$.

⁴⁴In our estimations, we construct $f(k)$ by calculating the weight of (nominal) consumption in sector k as a proportion of the aggregate consumption, measured over the sample. Recall that we use consumption as a proxy for the output in the sectors, since data for the latter are not available.

⁴⁵We assume a cashless economy with one-period maturity for those bonds, which are in zero net supply.

The demand for variety j , produced in sector k , takes the form:

$$Y_{kj,t} = \left(\frac{P_{kj,t}}{P_{k,t}} \right)^{-\epsilon} Y_{k,t}, \quad Y_{kj,t} = C_{kj,t} = N_{kj,t}, \quad (7)$$

where $C_{kj,t}$ and $N_{kj,t}$ denote the consumption of that variety and the specific labour input, respectively. A welcoming advantage of a linear technology is that deep parameters related to the curvature in production are absent in NKPC, so that following results are better compared with those of the previous section. It also implies we have at least one less parameter to calibrate in the estimations⁴⁶. In addition, $\epsilon > 1$ is the elasticity of substitution between varieties, while $Y_{k,t}$ represents the demand in sector k . We also define the latter, as well as the total demand in the economy, Y_t :

$$Y_{k,t} = f(k) \left[\int_0^1 Y_{kj,t}^{\frac{\epsilon-1}{\epsilon}} dj \right]^{\frac{\epsilon}{\epsilon-1}}, \quad Y_t = \left[f(k)^{\epsilon-1} Y_{k,t}^{\frac{\epsilon-1}{\epsilon}} dk \right]^{\frac{\epsilon}{\epsilon-1}}. \quad (8)$$

Usual market clearing conditions in the goods markets imply $Y_t = C_t$, $Y_{k,t} = C_{k,t}$ and $Y_{kj,t} = C_{kj,t}$.

When firm kj can re-optimize, it sets price $X_{kj,t}$ by maximising the following expression for discounted expected future profits:

$$E_t \sum_{s=0}^{\infty} Q_{t,t+s} (1 - \lambda_k)^s [X_{kj,t} Y_{kj,t+s} - W_{kj,t+s} N_{kj,t+s}], \quad (9)$$

subject to:

$$Y_{kj,t+s} = N_{kj,t+s}, \quad Y_{kj,t+s} = \left(\frac{X_{kj,t+s}}{P_{t+s}} \right)^{-\epsilon} Y_{t+s}, \quad (10)$$

where, for instance, $W_{kj,t+s}$ is the nominal wage in the firm kj for the period $t + s$, conditional on time- t information, while $Q_{t,t+s} = \beta(C_{t+s}/C_t)^\sigma (P_t/P_{t+s})$ is the stochastic nominal discount factor between the periods t and $t + s$. In addition:

$$P_{k,t} = \left[\int_0^1 P_{kj,t}^{1-\epsilon} dj \right]^{\frac{1}{1-\epsilon}}, \quad P_t = \left[\int_0^1 f(k) P_{k,t}^{1-\epsilon} dk \right]^{\frac{1}{1-\epsilon}}, \quad (11)$$

representing sectoral and aggregate price indices, respectively.

We will later calibrate our setting for the presence of strategic complementarities across sectors. Strategic interactions in this model arise from factor specificity – in our case, labour at the firm level. Carvalho and Nechio (2016),

⁴⁶Moreover, we need constant returns to scale to derive an aggregate NKPC that features both the marginal cost as the forcing variable and endogenous sectoral terms as controls. We use this version later in this section.

for instance, discuss implications for a similar model where labour and capital are sector-specific. Other sources of real rigidities include diminishing returns to scale in production and sticky intermediate inputs, for instance⁴⁷. One can parameterise our model for either strategic complementarity or strategic substitutability in price setting across sectors. However, as shown in Carvalho (2006), the latter is only achievable in this type of model under unrealistic calibration values⁴⁸.

Compared to the previous section, the combination of strategic complementarities and heterogeneity in price-setting provides the additional benefit of implying a different price dynamics. It follows that more sticky sectors disproportionately influence aggregate prices. The reason is that firms in the sectors with more flexibility avoid setting prices that are too disparate compared to the future aggregate price level. Everything else constant, the price dynamics would be more staggered in the heterogeneous economy, compared to its identical-firms counterpart. Differences should be more evident as one calibrates for a higher degree of strategic complementarities in price setting. Note that, in the limit, a substantial demand shock would generate little effect on aggregate inflation. A first consequence is that the slope – which measures such sensitivity – would tend to zero. A second is that the degree of stickiness in the economy (θ) need not be as high as often estimated in the literature to provide the same dynamics for inflation. Therefore, the implied infrequency of the heterogeneous economy can potentially approach levels observed using micro data.

This setting also modifies the NKPC of the economy in two important ways. First, the simplistic inverse relationship between the degree of nominal rigidities in the economy and the slope becomes non-trivial. The latter will now depend on deep parameters related to strategic interactions in the economy, as well as on specific degrees of nominal rigidity in each sector (and their distribution across sectors). As mentioned in the previous section, a second difference is that price-setting heterogeneity generates a composite endogenous term in the NKPC which is proportional to a weighted average of sectoral

⁴⁷ Another reason for the absence of diminishing returns in our model relates to the P. J. Klenow and Willis (2006) critique. The authors argue that pricing implications of models where strategic complementarities are strong can not be reconciled with the micro evidence, which points to a larger size of price changes. Models in which strategic complementarities arise because of diminishing returns to scale are more subject to the critique. Our model is immune to the critique since these interactions solely arise from segmented labour markets and firms within a sector exhibit the same pricing behaviour. See Nakamura and Steinsson (2013) for a formal discussion on the topic.

⁴⁸ Particularly, for reasonable values of φ and ϵ , one would need to fix a very low value for the elasticity of intertemporal substitution in consumption, σ , to create strategic substitutability in price setting.

output gaps. These would have to be included in the estimation.

The presence of strategic complementarities across sectors also has important consequences regarding the introduction of an indexation scheme in the economy. A rule that introduced noise in the strategic interactions channel might undermine empirical benefits generated by the model in terms of the NKPC. For example, if we were to assume an indexation scheme on the aggregate price level, then sectors with lower frequencies of price changes – which, under no indexation, have a disproportionate effect on aggregate prices – would index based on an aggregate measure that already reflects the price setting of more flexible sectors. Hence, these less sticky sectors would have a higher impact on the aggregate price dynamics compared to the case of no indexation. The otherwise important role of the more sticky sectors in the mechanism would be suppressed. Consequently, differences between our framework and that of an identical-firms economy would be less evident the more firms adjust their prices based on past aggregate inflation. Another problem of an indexation rule based on aggregate prices is that we would eventually lose usual benchmarks for parameter values, provided in the literature. Indexation schemes are usually absent in multi-sector models as ours, and it is not clear how one that affects the mechanism of strategic interactions between sectors may demand a different parameterisation.

Taking those considerations into account, we assume that firms which can *not* readjust their prices in any given period set them following:

$$P_{kj,t} = P_{kj,t-1} \left(\frac{P_{k,t-1}}{P_{k,t-2}} \right)^{\gamma_k} \quad (12)$$

If $\gamma_k = \gamma = 0$, our model becomes that in Carvalho (2006). In the following, we assume $\gamma_k = \gamma$. We do not allow for heterogeneity in the parameter because it implies a more complex NKPC where aggregate inflation no longer appears as a driving variable⁴⁹. Assuming that γ is the same across sectors also lowers the risk of obtaining spurious estimates of λ_k ⁵⁰.

Defining $\Pi_{k,\tau+s,\tau} \equiv P_{k,\tau+s}/P_{k,\tau}$, the aforementioned price setting decisions and first order conditions yield:

$$X_{kj,t} = \frac{\epsilon}{\epsilon - 1} \frac{E_t \sum_{s=0}^{\infty} Q_{t,t+s} (1 - \lambda_k)^s P_{t+s}^{\epsilon} Y_{k,t+s} W_{kj,t+s}}{E_t \sum_{s=0}^{\infty} Q_{t,t+s} (1 - \lambda_k)^s P_{t+s}^{\epsilon} Y_{k,t+s} \Pi_{k,t+s-1,t-1}^{\gamma}}, \quad (13)$$

⁴⁹For instance, we would obtain a forward-looking term $E_t \int_0^1 \frac{\beta}{1+\beta\gamma_k} f(k) \pi_{k,t+1} dk$ instead of $\frac{\beta}{1+\beta\gamma} E_t \pi_{t+1}$, where $\pi_{k,t}$ is the inflation rate in sector k .

⁵⁰The NKPC of an economy where γ_k is sector-specific would also feature non-linear terms where both λ_k and γ_k appear in the same ratio, $\int_0^1 f(k) \left[\frac{\lambda_k}{(1-\lambda_k)(1+\beta\gamma_k)} - \frac{\beta\lambda_k}{(1+\beta\gamma_k)} \right] dk$. In the absence of a reasonable guess for γ_k , we might well be overestimating it at the cost of a higher estimate of λ_k .

Finally, the following law of movement for sectoral prices holds:

$$P_{k,t} = [\lambda_k X_{k,t}^{1-\epsilon} + (1 - \lambda_k)(P_{k,t-1} \Pi_{k,t-1,t-2}^\gamma)^{1-\epsilon}]^{\frac{1}{1-\epsilon}}. \quad (14)$$

1.4.1

New-Keynesian Phillips Curves

As shown in the appendix, the log linearised version of the model produces an aggregate NKPC equation that takes the form:

$$\pi_t = \underbrace{\frac{\beta}{1 + \beta\gamma}}_{\equiv \gamma_f} E_t \pi_{t+1} + \underbrace{\frac{\gamma}{1 + \beta\gamma}}_{\equiv \gamma_b} \pi_{t-1} + \underbrace{\psi \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon \varphi^{-1}} - \frac{1}{\epsilon} \right)}_{\text{slope}} y_t + \underbrace{\frac{\psi}{\epsilon} g_t}_{\text{shift-term}} + u_t, \quad (15)$$

where:

$$\psi = \underbrace{\int_0^1 f(k) \left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta \lambda_k}{(1 + \beta\gamma)} \right] dk}_{\text{nominal rigidities}},$$

$$g_t = \int_0^1 \hat{f}(k) y_{k,t} dk,$$

$$\hat{f}(k) = \frac{\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta \lambda_k}{(1 + \beta\gamma)}}{\int_0^1 f(k) \left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta \lambda_k}{(1 + \beta\gamma)} \right] dk} f(k).$$

We refer to (15) as the *generalised NKPC* of the heterogeneous economy⁵¹. Logs of the output gap and the sectoral output gaps are denoted by y_t and $y_{k,t}$, respectively. An identical-firms economy features a direct inverse relationship between the degree of stickiness and the slope of its NKPC, no longer the case in (15). Heterogeneity in price setting produces a slope formed by two components. The first, ψ , summarises the degree of nominal rigidity in the economy, as well as its distribution across sectors. The second term, comprised of $\frac{\varphi^{-1} + \sigma}{1 + \epsilon \varphi^{-1}} - \frac{1}{\epsilon} \equiv \Theta - \frac{1}{\epsilon}$, relates to the degree of real rigidities⁵². As shown by Carvalho (2006), compared to a homogeneous economy calibrated for the average frequency of price changes, the former tends to increase the sensitivity of inflation to the output gap, whereas the latter operates in the opposite

⁵¹“Generalised” in reference to the presence of endogenous sectoral terms in (15). We borrow such name from Carvalho (2006).

⁵²If $\lambda_k = \lambda$ and $\gamma = 0$, the coefficient that multiplies the output gap becomes $\Theta \left(\frac{\lambda}{1 - \lambda} - \beta \lambda \right)$, so that Θ corresponds to the Ball and Romer (1990) coefficient of real rigidities in this model.

direction.

The NKPC in (15) also features an endogenous shift term, $\frac{\psi}{\epsilon}g_t$. This is proportional to a weighted average of sectoral output gaps. Each weight $\hat{f}(k)$ is a transformation of the original weight $f(k)$, but adjusted for the relative degree of flexibility of the sector compared to that of the entire economy. Coefficients that multiply sectoral output gaps in the shift term can have either sign, each of these partially depending on the slope of the NKPC of the corresponding sector. These sectoral equations take the form:

$$\begin{aligned}\pi_{k,t} = & \frac{\beta}{1 + \beta\gamma} E_t \pi_{k,t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{k,t-1} \\ & + \left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta\lambda_k}{(1 + \beta\gamma)} \right] \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon\varphi^{-1}} - \frac{1}{\epsilon} \right) y_t \\ & + \frac{1}{\epsilon} \underbrace{\left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta\lambda_k}{(1 + \beta\gamma)} \right]}_{\equiv \psi_k(\lambda_k, \beta, \gamma)} y_{k,t} + v_{k,t}\end{aligned}\quad (16)$$

The term $\psi_k(\lambda_k, \beta, \gamma)$ is crucial here. It relates to the degree of nominal rigidity present in each sector. It is also relevant for the aggregate NKPC, entering ψ . We later benefit from its presence in both models by conducting estimations with cross-equation restrictions.

Under an identical-firms economy, it is well known that a model with the output gap as forcing variable performs poorly even when it comes to just predicting the sign of a change in inflation. Indeed, several authors propose alternative specifications where a measure of real marginal costs is used as forcing variable, some using sectoral or industry data⁵³. These variations have not been immune to criticism either, with some articles pointing to a limited empirical evidence that proxies for marginal costs can add any information to the dynamics of inflation⁵⁴. To account for these alternatives, as well as the endogenous sectoral terms in (15), we also derive an NKPC that exhibits the real marginal cost as forcing variable. To do so, we need to write $\psi_k(\lambda_k, \beta, \gamma)$ – the function that corresponds to the degree of stickiness in each sector – in terms of its aggregate counterpart (ψ) and a sectoral deviation from this

⁵³Most noticeably: Sbordone (2002) and Gali and Gertler (1999), with aggregate data; J. Imbs, Jondeau, and Pelgrin (2011) and Piazza (2018), with sectoral data, and; Hale Shapiro (2008) and Gwin and VanHoose (2008), who rely on micro-foundations and actual industry costs, respectively.

⁵⁴See Rudd and Whelan (2005a) and Rudd and Whelan (2007), for instance.

weighted average:

$$\psi_k(\lambda_k, \beta, \gamma) = \overbrace{\left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta\lambda_k}{(1 + \beta\gamma)} \right]}^{\text{sectoral stickiness}} = \underbrace{\psi}_{\text{agg. stickiness}} + \underbrace{\zeta_k}_{\text{sectoral deviation}} \quad (17)$$

Under constant returns to scale, the aggregate real marginal cost equals a weighted average of sectoral analogues, $mc_t^r = \int_0^1 f(k)(mc_{k,t} - p_{k,t})dk$. After some manipulations, it is possible to show that⁵⁵:

$$\begin{aligned} \pi_t = & \underbrace{\frac{\beta}{1 + \beta\gamma} E_t \pi_{t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{t-1} + \psi mc_t^r}_{\text{reduced form in the literature (marginal costs)}} \\ & + \underbrace{\left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon\varphi^{-1}} - \frac{1}{\epsilon} \right) \left\{ \int_0^1 f(k)\zeta_k dk \right\} y_t}_{\text{aggregate error}} + \underbrace{\frac{1}{\epsilon} \int_0^1 f(k)\zeta_k y_{k,t} dk + \varepsilon_t}_{\text{idiosyncratic errors}} \end{aligned} \quad (18)$$

We refer to (18) as the *expanded NKPC* for the heterogeneous economy. It nests a number of hybrid NKPCs of identical-firms models seen in the empirical literature. If the true model features heterogeneity in price setting, then the error in the literature would contain two endogenous composite terms: the first a consequence of mismeasurement of nominal rigidities at the aggregate level, the second at the sector level⁵⁶. In that sense, another drawback of a misspecified NKPC is that one would estimate a different slope, ψ , which only captures the distribution of nominal rigidities in the economy. This coefficient is unrelated to the degree of real rigidities, since the function that corresponds to the latter is now embedded in the aggregate error term of (18).

Importantly, if heterogeneities in price setting are not relevant, then $\lambda_k = \lambda$ and thus $\psi_k(\lambda_k, \beta, \gamma) = \psi = \frac{1}{(1 + \beta\gamma)} \left[\frac{\lambda}{(1 - \lambda)} - \beta\lambda \right]$. Since $\lambda = 1 - \theta$, this is the same slope in (4). In addition, the aggregate and idiosyncratic errors in (18) would also disappear, since $\zeta_k = 0$. It follows that reduced-form estimates under (18) would be arbitrarily more similar to those using (4) as heterogeneities across sectors are less important in explaining the dynamics of aggregate inflation⁵⁷. We later show that estimates for ψ in (18) are generally not aligned with those for the slope in (4). Moreover, coefficients for its sectoral

⁵⁵We present the algebra in the appendix.

⁵⁶That is, the resulting composite error would be $\left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon\varphi^{-1}} - \frac{1}{\epsilon} \right) \left\{ \int_0^1 f(k)\zeta_k dk \right\} y_t + \frac{1}{\epsilon} \int_0^1 f(k)\zeta_k y_{k,t} dk + \varepsilon_t$, plus an expectational error term related to inflation in the next period. Approaches that are similar to (18) can be seen in J. M. Imbs, Jondeau, and Pelgrin (2007), Cagliarini, Robinson, and Tran (2011) and Piazza (2018).

⁵⁷Presumably, considering the appropriate econometric setting.

terms are statistically significant.

Apart from differences in the forcing variable, equations (15) and (18) exhibit very similar covariates in the reduced form. Under the assumption of a heterogeneous economy, we can also contrast results under (18) – when the marginal cost is the forcing variable – with those using (15) – which relies on the output gap. A possible complication involving the former is that it features both the real marginal cost and the output gap, though⁵⁸. We once more rely on the selection routine in (3) to properly choose instruments for each of these endogenous variables. Lags of both will be included as candidates in that case. Additionally, if the relationship between inflation and typical proxies for the real marginal cost is indeed spurious, as argued in some articles in the literature, the fact that both the marginal cost and the output gap appear in (18) is less troublesome.

1.4.2

Econometric Approach

The NKPCs (15), (16) and (18) include so many endogenous variables that some ad-hoc choice of instruments would likely face scepticism. A standard approach of exclusion restrictions would assume that all sectors appearing in those equations are in fact relevant in explaining the dynamics of inflation, what is unlikely. However, the data-driven technique in (3) is suitable for a large number of endogenous regressors in the equation, as well as for an even larger set of instrument candidates.

Given the number of endogenous terms in the NKPCs and the potential sensitivity of results, we construct four different approaches to choose instruments. Three of them are based on the AdaLASSO estimator in (3). For each endogenous covariate \mathcal{Y} in the NKPC:

- I. (*Baseline*) Selects instruments by performing (3), when candidates are the first lags of the same endogenous regressors in the equation and variables outside of the model typically considered in the literature (the same as in the previous section: the Fed Funds rate, the Treasury spread, the inflation in commodities and in wages). As before, the first lags of the output gap and the non-farm labour share (proxying marginal costs) are always candidates. For each selected variable, we apply its first three lags as instruments, $\{z_{t-1}, z_{t-2}, z_{t-3}\}$.

⁵⁸In this model, it is not possible to derive an NKPC that features the marginal cost as forcing variable and sectoral terms as controls, but where the output gap does not appear.

- II. Chooses instruments based on (3), considering the exact same candidates set as above, but we only pick the first two lags of each selected variable as instrument, $\{z_{t-1}, z_{t-2}\}$.
- III. Again selects based on (3), but considering variables lagged twice as candidates, z_{t-2} . This accounts for time-aggregation bias in the macro aggregates, as in section 3. It also picks as instruments the same lag of each selected variable. For the same candidates as in I and II, this procedure considerably reduces the number of instruments chosen. We use the fact that the NKPC in (15) also admits an alternative version where sectoral output gaps are replaced by “relative prices” (defined as $p_{k,t} - p_t$, for sector k) to include additional candidates. These are (the second lags of) the inflation rate of each sector and sectoral relative prices.
- IV. Mimics a typical ad-hoc instrument choice in the literature while accounting for the presence of sectoral endogenous terms in the NKPC. Specifically, we choose the first three lags of the output gap, the sectoral output gaps, the non-farm labour share, inflation, the Fed Funds rate, the Treasury spread and the inflation in wages and in commodities.

Table A.2 (appendix) summarises the four approaches. The first one is our *baseline*, shown in the main paper. Results presented here are repeated as robustness checks under the other three procedures in the appendix, reconfirming our findings. By the second approach, we aim to evaluate whether the method is sensitive to the common pitfall of too many instruments in the GMM. The motivation of the third procedure is to test whether it is also robust for different choices of candidates, as well as for their lag structures. The first and the third approaches produce approximately the same number of instruments, what avoids a situation where the number of orthogonality conditions may affect the comparison⁵⁹. Lastly, one could extent the method by imposing specific sectoral variables as candidates for each endogenous variable used in (3) – e.g., including oil drilling measures when \mathcal{Y} is the output gap of the “gasoline and other energy goods” sector. We leave these underlying alternatives for future research.

In the estimations that follow, we replace forward-looking expectations in the NKPCs by their actual values, so that the disturbance term also includes

⁵⁹Given the number of orthogonality conditions, the J-statistic of overidentification is probably unreliable. This is a common pitfall even for the NKPC of an identical-firms model – e.g., Mavroeidis (2005). A more recommended procedure would use robust inference methods, as the S statistic in Stock, Wright, and Yogo (2002). However, the dimension of the GMM in our case also affects the reliability of such procedures. We return to this point later in this section.

an expectational error. For expository purposes, we present the equations in their original form.

1.4.3

Main Results

1.4.3.1

Reduced Form

Table 1.5 shows the estimates of the reduced forms coefficients in the NKPCs (15) and (18)⁶⁰. In addition to the hybrid models, it also presents results under the assumption $\gamma = 0$ (purely forward-looking NKPC) for the former⁶¹. For comparison purposes, the last column provides estimates under the assumption of a homogeneous (identical-firms) economy – the reduced form associated with (4)⁶². The lower end of Table 1.5 once more exhibits estimates when β and γ are calibrated.

First, note that one estimates a positive coefficient for the marginal cost under the assumption of an identical-firms economy. When β and γ are not fixed, this estimate is also statistically significant. In contrast, this is not the case when we use the marginal cost as the forcing variable for the heterogeneous economy (expanded NKPC)⁶³. Indeed, no positive coefficient is estimated for the slope and we even obtain a negative and statistically significant estimate for the calibrated model. Differences between the expanded and the homogeneous models for the slope are also observed under approaches II to IV (appendix) and favour the view that heterogeneity in price setting is indeed important⁶⁴. The coefficient associated with the output gap in such equation is positive and statistically significant, but recall that this term relates to an aggregate error rather than the slope. In fact, it is shown in the appendix that the other approaches also produce no clear pattern for that coefficient⁶⁵.

⁶⁰For expository purposes, we exhibit results for CUE-GMM. As already mentioned, this estimator is known to be more robust than the more standard 2S-GMM in a potentially weak-IV setting. Our main findings are still valid under the latter, but results are slightly weaker and the differences between models less clear.

⁶¹Imposing $\gamma = 0$ in the expanded version (omitted) roughly reconfirms the findings we discuss for its hybrid analogue.

⁶²Note that results exhibited in the last column do not precisely repeat those presented for the second column in Table 1.3 (since we used 2S-GMM in the previous section).

⁶³Recall that the term that multiplies the output gap in (18) relates to an aggregate error rather than the slope of that equation. Therefore, we refer to the coefficient that multiplies the marginal cost as its slope and to the marginal cost itself as the forcing variable.

⁶⁴Additionally, estimated coefficients related to sectoral terms in the expanded NKPC, (18), are generally positive (between 0.01 and 0.05) and statistically significant (all but one at 1%). We omit estimates of those coefficients for expository purposes.

⁶⁵Being more specific, although approaches I and IV tend to generate a positive coefficient for the output gap in the expanded NKPC, methods II and III indicate the opposite.

Table 1.5: Estimations of Reduced-Form NKPCs – Heterogeneous Economy

Coefficient	Generalised		Expanded	Homogeneous
	$\gamma = 0$	$\gamma \neq 0$	$\gamma \neq 0$	$\gamma \neq 0$
γ_f	1.030 (0.005)	0.703 (0.002)	0.662 (0.002)	0.554 (0.011)
γ_b	-	0.323 (0.002)	0.376 (0.002)	0.478 (0.001)
Output Gap	0.066*** (0.010)	0.010*** (0.002)	0.012*** (0.001)	-
Marginal Cost	-	-	-0.001 (0.001)	0.028*** (0.001)
$\beta = 0.99$ and $\gamma = 0.5$				
Output Gap	0.091*** (0.006)	0.004*** (0.002)	0.028*** (0.001)	-
Marginal Cost	-	-	-0.015*** (0.001)	0.002 (0.001)

Notes: Estimates of the NKPCs in (15) and (18) using CUE-GMM. For the former, we present results for both the forward-looking ($\gamma = 0$) and the hybrid ($\gamma \neq 0$) versions of the model. The instrument set is constructed based on approach I (see the main text). The lower end of the table presents estimates while fixing $\beta = 0.99$ and $\gamma = 0.5$. HAC standard errors are presented in parentheses. The hypotheses of statistically insignificant coefficients for the output gap and the labour share are tested: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Generally, our results indicate that the marginal cost as forcing variable – a common workaround to estimate a positive slope in the literature that assumes an identical-firms model – does not quite work as expected regarding the heterogeneous economy. This favours the views in Rudd and Whelan (2005a), pointing to a potentially spurious relationship between that variable and inflation. In contrast, the coefficient for the output gap in (15) is positive and statistically significant. This finding is reconfirmed for estimations based on the other approaches (see the appendix)⁶⁶.

If the true model features heterogeneity in price setting, our results suggest that a model with the output gap as forcing variable outperforms one that relies on the marginal cost – in that the former more likely establishes a positive dynamic relationship with inflation. This finding seems to oppose a common view that a negative slope is often obtained because the output gap leads inflation in the data – the canonical New-Keynesian model implies the

⁶⁶Specifically, relying on approaches II, III and IV we obtain twelve estimates of the slope in the forward-looking and the hybrid models based on (15). Only three out of those twelve are negative, none of which also statistically significant.

opposite, with inflation depending on a discounted stream of expected output gaps⁶⁷. Too many variables drive inflation for such a simplistic explanation to be essential. Instead, the true nature of the problem may be a consequence of two important factors. The first is a wrong assumption of an oversimplified identical-firms model. Indeed, introducing price-setting heterogeneity across sectors not only improves estimates based on the reduced form but also produce results that are substantially more reliable when estimating the structural form – as to be seen in the following. Second, sectoral terms otherwise omitted are likely econometrically important.

1.4.3.2

Structural Form

Estimating the structural form in (15) is non-trivial. It combines several endogenous variables with a highly non-linear setting. An alternative is to calibrate a number of its deep parameters – eventually, λ_k for some sectors – while estimating others. The main disadvantage of this procedure is that it biases the comparison between the estimated degree of stickiness of the heterogeneous economy and actual micro data⁶⁸.

Another option – the one we adopt here – is to jointly exploit the structures of the aggregate NKPC in (15) and its fifteen sectoral analogues in (16), gaining efficiency from cross-equation restrictions. For convenience, we summarise the system below:

$$\pi_t = \underbrace{\frac{\beta}{1 + \beta\gamma} E_t \pi_{t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{t-1} + \int_0^1 f(k) \psi_k(\beta, \gamma, \boldsymbol{\lambda}_k) dk \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon \varphi^{-1}} - \frac{1}{\epsilon} \right) y_t}_{\text{aggregate slope}} + \frac{1}{\epsilon} \int_0^1 f(k) \psi_k(\beta, \gamma, \boldsymbol{\lambda}_k) y_{k,t} dk + u_t, \quad (19)$$

$$\pi_{k,t} = \frac{\beta}{1 + \beta\gamma} E_t \pi_{k,t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{k,t-1} + \psi_k(\beta, \gamma, \boldsymbol{\lambda}_k) \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon \varphi^{-1}} - \frac{1}{\epsilon} \right) y_t + \frac{1}{\epsilon} \psi_k(\beta, \gamma, \boldsymbol{\lambda}_k) y_{k,t} + v_{k,t}.$$

⁶⁷Figure 1 in Galí and Gertler (1999) exemplifies such argument in the literature.

⁶⁸For example, a method that calibrates some of the λ_k based on the micro evidence would bias the implied degree of nominal rigidities in the economy, $\int_0^1 f(k)(1 - \lambda_k)dk$, towards that of actual disaggregated data. Additionally, in such non-linear setting, final estimates for the sectors that are not parameterised could be highly sensitive to the adopted values for the ones that are calibrated, depending on the surface of the underlying likelihood function. In this sense, a technique that estimates *all* the parameters related to the degree of stickiness in the sectors is preferred.

Our strategy relies on estimating parameters λ_k , associated with sectoral degrees of stickiness, while calibrating the remaining parameters, σ , φ and ϵ , which directly govern the degree of strategic complementarities in price setting across sectors. We later verify how obtained estimates vary based on different parameterisations for the latter, comparing these results to predictions implied by theory. In our baseline framework, we once more estimate β and γ ⁶⁹. Unfortunately, it is not possible to conduct joint estimations based on (18) because the labour share – the proxy we use for the real marginal cost – is not available at the sector level. Therefore, we subsequently focus on the identification of (15), through the system in (19).

To be able to estimate the system in (19), we slightly modify the aforementioned approaches I to IV to reduce the resulting number of orthogonality conditions in the GMM. Specifically, for each approach, we remove the last lag of selected variables (e.g., approach I would apply the first two lags of selected variables rather than the first three)⁷⁰. Additionally, we do not apply instrument selection routines for the sectoral NKPCs in (19). It is quite difficult to improve the first-stage predictability for the NKPCs of some sectors (e.g., “gasoline and other energy goods”) when specific candidates (outside of the model) are absent. As a consequence, underidentification is an issue for a few sectors, when selecting based on (3). Thus, we opt to instrument each endogenous variable in the sectoral NKPCs by its own lags. In spite of that, the number of lags applied still follows the approach we use for the aggregate equation. For example, if we apply the first two lags of selected variables for (15), we instrument each endogenous variables in the sectoral NKPCs with their first two lags⁷¹. Table A.2 (appendix) summarises the approaches used in the structural estimations.

For each of these estimations, we are interested in two implied coefficients. Naturally, the first is the aggregate slope in (19). The second is the implied degree of nominal rigidities (or implied infrequency) in the economy:

$$\theta \equiv 1 - \lambda = \int_0^1 f(k)(1 - \lambda_k)dk \quad (20)$$

We later show that the model delivers a positive slope and an implied

⁶⁹It is possible to extend the method by allowing one of the remaining parameters (σ , φ and ϵ) to be free in the estimation. However, the resulting non-linear structure in the term related to real rigidities (Θ) further complicates the method and our estimations when allowing more than one did not improve our results. In addition, choosing which of those parameters should be estimated is clearly arbitrary. We leave the challenging extension of the method where all structural parameters are estimated for future research.

⁷⁰Approach III is not modified since it only applies the second lag of variables.

⁷¹In line with the established for the aggregate equation, $\pi_{k,t-1}$ is considered exogenous in our estimations.

infrequency that approaches the lower values suggested by the micro evidence. These findings are *not* sensitive to the remaining assumptions for the model – e.g., calibrated values and ad-hoc indexation schemes. Additionally, *both* the implied stickiness and the aggregate slope are *lower* as we introduce more strategic complementarities across sectors, exactly as predicted by theory, as already mentioned. To compare estimates of (20) with the micro evidence, implied data from Bils and P. Klenow (2004) are used, the microeconomic benchmark we assume hereafter. The data are also available at the sector level, – see Table 1.1 –, so that we are also able to evaluate the model based on estimates for each sectoral infrequency, $1 - \lambda_k$.

A number of authors has addressed empirical benefits of the use of sectoral, regional or stage-of-processing NKPCs as an extension of the basic model when aiming to identify the aggregate equation. See J. Imbs, Jondeau, and Pelgrin (2011), Cagliarini, Robinson, and Tran (2011), Kiley (2015), Hooper, Mishkin, and Sufi (2019) and McLeay and Tenreyro (2019), for instance. Furthermore, as discussed in the last two articles, a powerful advantage of the disaggregated Phillips curves is that monetary policy does not respond to their shocks as it does to that of the aggregate equation. In other words, these articles point out that the identification of the aggregate NKPC may be challenging because monetary policy aims to offset aggregate demand shocks that are essential to identify the equation. Hence, under good policy and in the absence of exogenous variation, the econometrician does not purely observe the NKPC in the data, but instead an intersection of that equation and a monetary policy targeting rule, a classic situation of simultaneity bias⁷².

If the target rule only attempts to offset aggregate shocks, then the sectoral terms appearing in the aggregate NKPC in (19) may – even under a good policy – provide the additional variation needed to identify the equation⁷³. In addition, unresponsiveness of monetary policy at the sectoral level is possibly more important to the identification of the aggregate NKPC as one increases the number of sectors in the economy – since, in the limit, idiosyncratic shocks would be minor, compared to the aggregate analogues⁷⁴. With fifteen sectors, almost all the sectoral weights in our economy are not higher than 10% – see Table 1.1.

A joint estimation that includes sectoral NKPCs is also more likely to succeed with substantial cross-sectional variation in inflation expectations.

⁷²A simple example is the optimal policy under discretion in the canonical New-Keynesian model, which fixes a zero output gap and an inflation that equals the target.

⁷³Note that the aggregate NKPC in (19) considers idiosyncratic shocks since sectoral output gaps are present.

⁷⁴Not to mention the advantages of more cross-sectional variation due to the presence of more sectoral NKPCs in the system (19).

Note that disturbances in each equation include corresponding expectational errors, once we substitute the expected sectoral inflation rates by their actual values. Thus, we also gain efficiency from exploiting cross-equation correlations in the error terms, due to the presence of correlated expectational errors^{75,76}.

We exhibit results for our structural estimations for three calibrations of the model, shown in Table 1.6. In the baseline, we set $\sigma = 0.5$, $\epsilon = 9$ and $\varphi = 0.5$. The remaining two increase (reduce) the degree of strategic complementarities in price setting across sectors (shown as “real rigidities” for conciseness reasons)⁷⁷. In the appendix, we also show that our results are robust to a much more conservative⁷⁸ calibration, used in Carvalho (2006): $\sigma = 1$, $\varphi = 1.5$ and $\epsilon = 5$, implying $\Theta \approx 0.38$ (compared to $\Theta \approx 0.18$, in the last column of Table 1.6).

Table 1.6: Calibration of Parameters Related to Real Rigidities

Parameter	Interpretation	Baseline	↑ Real Rigidity	↓ Real Rigidity
ϵ	elasticity of subst. between varieties	9	11	7
σ	inverse of the EIS	0.5	0.5	0.5
φ	(Frisch) elasticity of labour supply	0.5	0.3	0.7

Notes: Different calibrations of the model. Each implies a different degree of strategic complementarities in price setting across sectors. The first (baseline) implies $\Theta \approx 0.14$, the second produces $\Theta \approx 0.10$ and the third, $\Theta \approx 0.18$.

Table 1.7 summarises our results for the aggregate slope in (19) based on estimations for the forward-looking and the hybrid models. Given the potential sensitivity to the way orthogonality conditions are constructed in the GMM, we once more present results under two normalisations⁷⁹. To evaluate the model, correlations between our estimates for λ_k and benchmark infrequencies from the micro evidence – Bils and P. Klenow (2004) – are shown in brackets. As shown in the table, all of our estimates are positive and statistically significant at 1%. The slope is quite low, though, with a clear majority of the point estimates agreeing on a value of 0.01. In addition, correlations between the

⁷⁵In addition, if that cross-sectional variation is constant over time, one could add fixed effects in the model – as in McLeay and Tenreyro (2019), for instance. As later shown, our estimations perform well given the complexity of the system, so that we do not include such additional controls in the model – what would increase even more its dimensionality.

⁷⁶A novel approach that exploits cross-sectional variation to identify the aggregate Phillips curve can be seen in Jorda and Nechio (2018). Their methodology benefits from the fact that countries that rely on a fixed exchange rate regime can not conduct an independent monetary policy.

⁷⁷The elasticity of labour supply in the case with higher real rigidities (second column) addresses a possibility indicated by the micro evidence. For instance, Pencavel (1986) surveys that literature on φ , showing that estimates are generally lower than 1/3. Our baseline, at 0.5, it is also not far from that value. Main results of this paper do not change when we set an EIS of 1.

⁷⁸Conservative in the sense that it should imply a higher degree of stickiness (diverging from the micro evidence)

⁷⁹For example, for sectoral NKPCs in (19), these normalisations consider error vectors

estimated reset probabilities ($\hat{\lambda}_k$) and the micro benchmarks are generally high, frequently above 60% and none of which below 50%. In the appendix, we show that all the results in Table 1.7 are closely maintained when the instrument set is constructed based on approaches II to IV.

Our findings are even more pertinent given the fact that the system in (19) should bias *against* finding a statistically significant slope. Note that sectoral inflation is affected by aggregate cost-push shocks due to the presence of the output gap y_t in the sectoral NKPCs⁸⁰. Thus, to some extent, it is also influenced by the endogenous response of monetary policy to offset those shocks. Therefore, although such framework helps to identify the aggregate equation by alleviating the simultaneity problem caused by monetary policy, it still biases the aggregate slope in (19) towards zero. *Contrasting with this, we do not estimate any aggregate slope in (19) that was not positive and significant at 1% throughout this paper, including a series of robustness checks we later present.* The bias may explain differences in point estimates for the slope when comparing Table 1.5 and Table 1.7. Nonetheless, both still suggest a very low coefficient and the divergence is minimal when it comes to the hybrid model.

Table 1.8 presents results for the implied stickiness of the heterogeneous economy, (20). 95% confidence intervals for that parameter and estimates of β and γ are also shown. The column “ $\text{corr}(\theta_k, \text{micro})$ ” exhibits correlations between sectoral Calvo-pricing probabilities estimated from the model ($1 - \hat{\lambda}_k$) and the benchmarks. The latter would imply $\theta^{\text{micro}} \approx 0.48$. The first thing to note is that this value lies inside *four* out of the *six* confidence sets for θ . These that take the form:

$$v_{k,t}^{[1]} = \pi_t - \frac{\beta}{1 + \beta\gamma} E_t \pi_{t+1} - \frac{\gamma}{1 + \beta\gamma} \pi_{t-1} - \int_0^1 f(k) \left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta\lambda_k}{(1 + \beta\gamma)} \right] dk \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon\varphi^{-1}} - \frac{1}{\epsilon} \right) y_t - \frac{1}{\epsilon} \int_0^1 f(k) \left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta\lambda_k}{(1 + \beta\gamma)} \right] y_{k,t} dk, \quad (21)$$

$$v_{k,t}^{[2]} = \frac{1 + \beta\gamma}{\beta} \pi_t - E_t \pi_{t+1} - \frac{\gamma}{\beta} \pi_{t-1} - \int_0^1 f(k) \left[\frac{\lambda_k}{(1 - \lambda_k)\beta} - \lambda_k \right] dk \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon\varphi^{-1}} - \frac{1}{\epsilon} \right) y_t - \frac{1}{\epsilon} \int_0^1 f(k) \left[\frac{\lambda_k}{(1 - \lambda_k)\beta} - \lambda_k \right] y_{k,t} dk, \quad (22)$$

where the latter multiplies the former by $\frac{1 + \beta\gamma}{\beta}$. Orthogonality conditions are $E_t[v_{k,t}^n z_t]$, for normalisation $n \in \{1, 2\}$ and each instrument, z_t . The normalisations for the aggregate NKPC follow the same pattern.

⁸⁰In addition, sectoral output gaps also correlate with the output gap, since $y_t = \int_0^1 f(k) y_k dk$.

Table 1.7: Implied Slope from Estimations of (19)

Model	Normalisation	Calibration		
		Baseline	\uparrow Real Rigidity	\downarrow Real Rigidity
$\gamma = 0$	(1)	0.008***	0.006***	0.007***
		(0.000)	(0.000)	(0.000)
		[0.63]	[0.66]	[0.52]
	(2)	0.009***	0.006***	0.007***
		(0.000)	(0.000)	(0.000)
		[0.63]	[0.69]	[0.55]
$\gamma \neq 0$	(1)	0.014***	0.009***	0.022***
		(0.000)	(0.000)	(0.001)
		[0.73]	[0.64]	[0.69]
	(2)	0.011***	0.005***	0.012***
		(0.000)	(0.000)	(0.000)
		[0.64]	[0.61]	[0.68]

Notes: Implied slope from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β , γ and each λ_k are allowed to vary in the range (0,1). The instrument set is constructed based on approach I. Standard errors from Delta method are presented in parentheses. Correlations between estimated and benchmark infrequencies ($1 - \lambda_k$) that come from the micro data in Bils and P. Klenow (2004) are shown in brackets. We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1. In addition, * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

intervals are also quite narrow. In the appendix, we show that approaches II to IV produce similar findings.

There seems to be more uncertainty for the true value of γ , with estimates being more sensitive to the way we normalise the orthogonality conditions in the GMM⁸¹. These are centred around 0.6 for the first normalisation, but more clearly agree on a value of 0.3 under the second. Generally, implications of an ad-hoc indexation scheme as in (12) for the heterogeneous economy are not trivial. Since under such rule firms that can not readjust index their prices based on sectoral rather than aggregate inflation, we can not compare our estimates of γ with the empirical literature (which assumes an identical-firms model). In addition, it could be the case that lower estimates of θ shown in Table 1.8 are a direct consequence of an upward bias in γ .

To evaluate the accuracy of our method regarding γ , – and also to reconfirm our estimates for θ in Table 1.8 –, we re-estimate (19) while setting $\gamma \in \{0.3, 0.5, 0.7\}$. These estimations are exhibited in the appendix and produce very similar findings. Furthermore, we shall evaluate the sensitivity of the approach to general parameters in the indexation scheme. To this

⁸¹That is something we also observe for the other approaches exhibited in the appendix.

regard, we derive a version of the model featuring indexation à la Gali and Gertler (1999). These authors assumed that a fraction ω of the firms in the economy use a backward-looking rule to set their prices. We model ω as the fraction of the firms *in the sector* that follow the same rule, but since it is a fraction and it is synchronised across sectors, our estimates for the parameter should coincide with theirs⁸². The resulting NKPCs in this variation of our model are considerably more complex, what should further complicate our estimation strategy⁸³. Nonetheless, the model performs noticeably well. We find an estimate of ω that is very similar to that in Gali and Gertler (1999) – ω is estimated as 0.30 in our model, while their baseline estimate gives 0.26⁸⁴ –, a slope of 0.004 (also significant at 1%) and an implied Calvo-pricing probability for the economy of $\hat{\theta} = 0.65$ (with a very low standard error of 5×10^{-4}). For conciseness reasons, the model and these results are exhibited in the appendix.

In practice, the motivation to introduce indexation schemes in the model is largely empirical. However, these ad-hoc rules are not innocuous, being often at odds with the complete model and are also commonly rejected by studies based on disaggregated data – e.g., Bils and P. Klenow (2004), Nakamura and Steinsson (2008) and Nakamura and Steinsson (2013). These studies generally provide little evidence that prices change at a constant rate between more significant resets, as implied by those rules.

In that sense, a model that succeeds in capturing the dynamics of inflation in the data would not need those ad-hoc assumptions to work. Table 1.9 shows that it seems to be the case for the heterogeneous economy depicted here. The implied infrequency from the model is again low and more aligned with the micro evidence, around 0.60 for two of the calibrations. In addition, correlations with the benchmarks from disaggregated data are still high, often above 60%. Without the indexation scheme affecting the channel related to strategic complementarities, it is more clear how the model produces lower

⁸²Presumably, assuming that the estimation method is appropriate and that their model is still approximately valid (even under the identical-firms hypothesis).

⁸³For example, the aggregate NKPC in that economy follows:

$$\begin{aligned} \pi_t = E_t \int_0^1 f(k) \frac{\beta(1-\lambda_k)}{\phi_k} \pi_{k,t+1} dk + \int_0^1 f(k) \frac{\omega}{\phi_k} \pi_{k,t-1} dk \\ + \int_0^1 f(k) \left[\frac{(1-\omega)(\lambda_k)(1-\beta(1-\lambda_k))}{\phi_k} \right] dk \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon\varphi^{-1}} - \frac{1}{\epsilon} \right) y_t \\ + \frac{1}{\epsilon} \int_0^1 f(k) \left[\frac{(1-\omega)(\lambda_k)(1-\beta(1-\lambda_k))}{\phi_k} \right] y_{k,t} dk, \end{aligned} \quad (23)$$

where $\phi_k = 1 - \lambda_k + \omega(1 - (1 - \lambda_k)(1 - \beta))$. See the appendix for more details.

⁸⁴See estimates using the GDP deflator (and the ordinary way to construct moment conditions, (1)) in table 2 of Gali and Gertler (1999).

Table 1.8: Implied Stickiness from the Hybrid Model ($\gamma \neq 0$)

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters		
			θ^{agg}	β	γ
\uparrow Real Rigidity	(1)	0.64	0.47 (0.44 - 0.50)	0.99 (0.017)	0.59 (0.011)
	(2)	0.61	0.58 (0.54 - 0.62)	0.98 (0.000)	0.33 (0.000)
Baseline	(1)	0.73	0.50 (0.48 - 0.52)	0.99 (0.019)	0.52 (0.012)
	(2)	0.64	0.54 (0.48 - 0.61)	0.99 (0.000)	0.37 (0.000)
\downarrow Real Rigidity	(1)	0.69	0.48 (0.44 - 0.52)	0.99 (0.073)	0.70 (0.055)
	(2)	0.68	0.59 (0.55 - 0.64)	0.99 (0.000)	0.34 (0.002)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β , γ and each λ_k are allowed to vary in the range (0, 1). The instrument set is constructed based on approach I. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies ($1 - \lambda_k$) that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1.

estimates of θ as we introduce more strategic complementarities across sectors, broadly in line with the theoretical predictions. Note that we already estimated a lower slope when imposing more real rigidities. This was the case regardless of the hypothesis for γ – see Table 1.7. In the appendix, we show that those results are again valid for the other approaches to instruments. They also do not change when we calibrate β .

1.4.3.3

Behind the Scenes

The model delivers sensible estimates of the slope and the aggregate Calvo-pricing parameter with different calibrations for strategic interactions in the economy, econometric settings and macroeconomic specifications. It also generates satisfactory correlations with micro-based benchmarks, from 50% to 70%. However, correlations are not necessarily a good measure of how accurately the model captures the stickiness of the economy. For example, it is perfectly possible to obtain sensible estimates of θ , with values of $\hat{\theta}_k = (1 - \hat{\lambda}_k)$ performing well compared to the corresponding benchmarks in the cross section, but at the same time presenting a poor correlation with the latter.

To confirm that the model captures essential information found in micro

Table 1.9: Implied Stickiness from the Forward-Looking Model ($\gamma = 0$)

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters	
			θ^{agg}	β
\uparrow Real Rigidity	(1)	0.66	0.58	0.99
			(0.49 - 0.68)	(0.000)
	(2)	0.69	0.60	0.99
			(0.50 - 0.70)	(0.000)
Baseline	(1)	0.63	0.63	0.99
			(0.57 - 0.69)	(0.000)
	(2)	0.63	0.62	0.99
			(0.55 - 0.68)	(0.000)
\downarrow Real Rigidity	(1)	0.52	0.70	0.98
			(0.68 - 0.73)	(0.000)
	(2)	0.55	0.70	0.99
			(0.67 - 0.73)	(0.000)

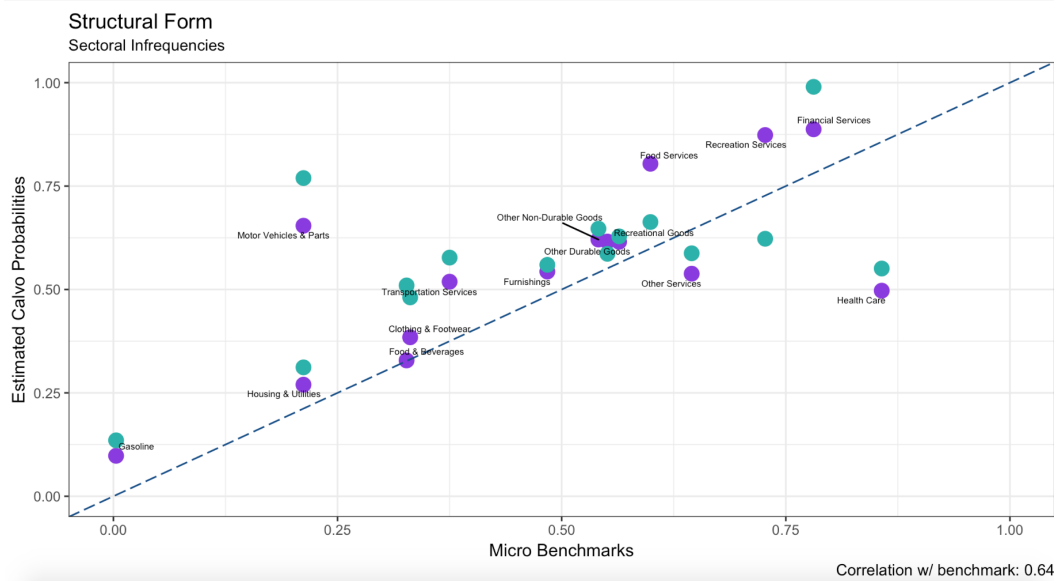
Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and each λ_k are allowed to vary in the range $(0, 1)$. We fix $\gamma = 0$ (purely forward-looking model). The instrument set is constructed based on approach I. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1.

data, Figure 1.1 compares $\hat{\theta}_k$ with micro benchmarks for the same parameters, based on disaggregated data in Bils and P. Klenow (2004). The calibration is the baseline in Table 1.6, but very similar findings are generated with other parameterisations of the model – see the appendix. Except for probably two sectors (“motor vehicles and parts” and “health care”), most part of the parameters are aligned with the benchmarks. Figure 1.2 shows how estimates under the two normalisation methods are similar. Individual standard errors are quite low, so that the confidence intervals are substantially narrow for each sector⁸⁵. Analogous charts are exhibited in the appendix, where we re-estimate the model using the alternative approaches to instruments.

1.4.4

Estimator Uncertainty

⁸⁵Normalisation (2) produces a $\hat{\theta}_k$ at the bound for “financial services”, what seems to be generating a higher standard error for that sector. Nonetheless, note that the point estimate does not lie far from the corresponding benchmark.

Figure 1.1: $\hat{\theta}_k$ vs. Micro Benchmarks

Notes: Estimated Calvo probabilities using the same econometric setting of Table 1.8. Results for normalisation 1 (2) are exhibited in purple (green). Benchmarks are implied probabilities from evidence in Bils and P. Klenow (2004) – see Table 1.1. Instrument selection follows approach I, the calibration used is the baseline and starting values are benchmark reset probabilities.

1.4.4.1

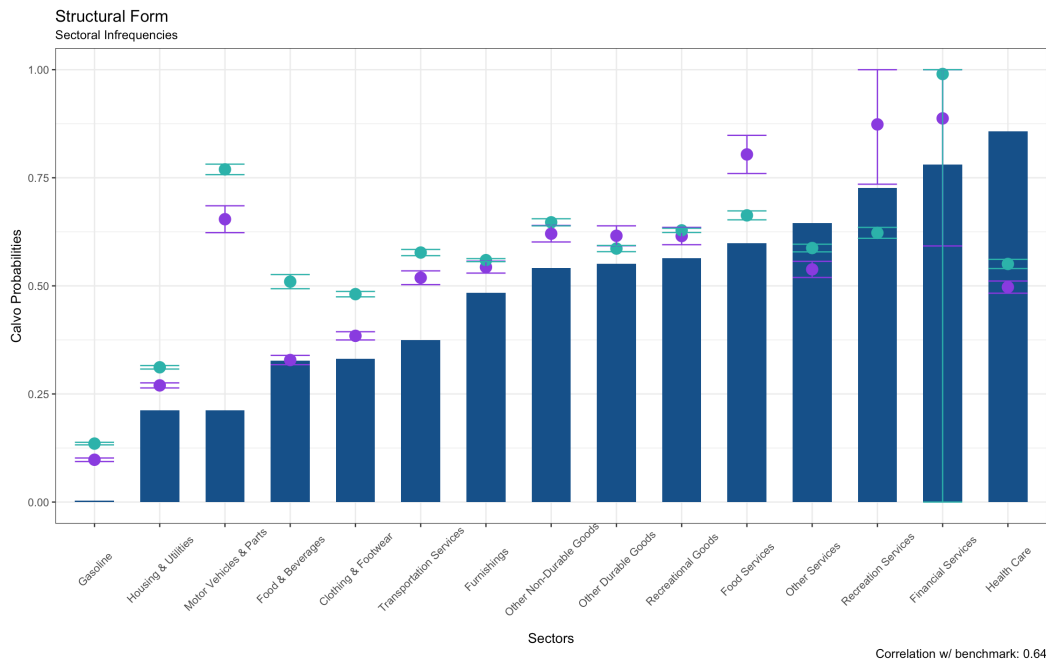
Parametric Stability

Previous estimations aimed to evaluate the sensitivity of our method to different specifications. However, the use of conventional first-order asymptotics in an environment where instruments are potentially weak can be misleading. This pitfall is even more important concerning structural models, especially with non-linear moment conditions as those derived from the system in (19). To address these uncertainties, weak-instrument-robust methods of inference are commonly recommended. Two techniques typically used construct robust confidence intervals for deep parameters by inverting the test statistics S (Stock and Wright (2000)) and K (Kleibergen (2005))⁸⁶. These tests are based on the empirical assessment that the estimations of parametric vectors are generally more meaningful than those of individual parameters. Hence, they seek to obtain a robust interval for one parameter while restricting the rest of the parametric space, usually assuming identification for them.

The problem with robust inference techniques in our setting is that, although they can be generalised to the presence of multiple endogenous variables, little is known about their implications⁸⁷. In such case, however, it

⁸⁶See Ma (2002), Andrews and Stock (2005), Nason and G. W. Smith (2008), Kleibergen and Mavroeidis (2009) and Mavroeidis, Plagborg-Møller, and H. Stock (2014), to cite just a few of the examples in the literature.

⁸⁷The use of robust inference in high-dimensional settings is not common. For example, Andrews and Stock (2005) analyse weak-instrument-robust methods covering a sample of

Figure 1.2: $\hat{\theta}_k$ vs. Micro Benchmarks – Confidence Intervals

Notes: Estimated Calvo probabilities using the same econometric setting of Table 1.8. Results for normalisation 1 (2) are exhibited in purple (green). Blue bars are micro-based benchmark probabilities implied from evidence in Bils and P. Klenow (2004) and presented in Table 1.1. For expository purposes, these are sorted by the degree of flexibility. 95% confidence intervals are also shown. Instrument selection follows approach I, the calibration used is the baseline and starting values are benchmark reset probabilities.

is well known that they suffer from poor power⁸⁸. Additionally, with seventeen parameters we would have to assume that a subvector of sixteen of them is identified to construct robust sets (unlikely). Lastly, in order to invert test statistics, the number of grid points at which we need to evaluate the null hypothesis grows exponentially with the dimension of the parametric vector.

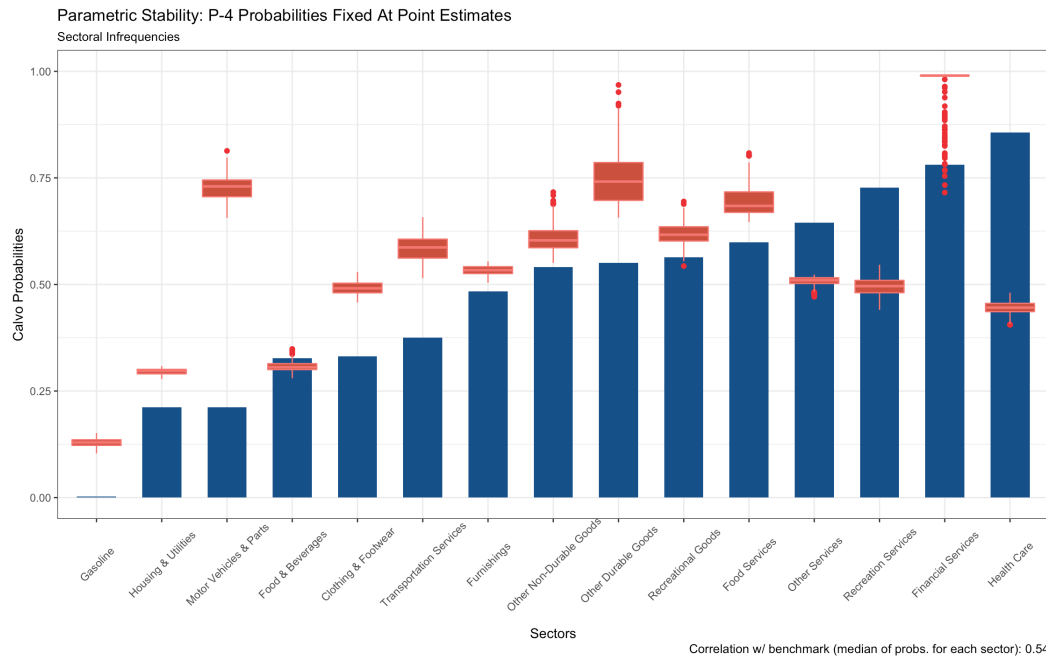
To circumvent those problems, we propose an alternative method to verify the model based on uncertainties involving the estimation. For each of the $P = 15$ sectoral probabilities, λ_k in (19), we fix $P - q$ of them at the original point estimates, re-estimating q parameters. For $q = 4$, this generates $C_4^{15} = 1365$ re-estimations of the model, with 365 estimates of λ_k for each sector. Figure 1.3 presents confidence intervals for sectoral Calvo-pricing probabilities ($1 - \lambda_k$), constructed with these values. Boxes for each sector represent the interquartile range. Vertical lines provide the 5%ile – 95%ile interval of the distribution. Note that results are quite in line with those exhibited in Figure 1.2. The correlation between the median estimate –

studies published in the American Economic Association journals. None of the 230 articles in their sample apply weak-IV methods using more than four endogenous variables in the estimations. Note that there are *seventeen* endogenous variables in our model.

⁸⁸In our case, both the S and the K tests hardly reject the null. This is true for most part of the points in the parametric space – likely because our GMM features too many moment conditions for such tests to be reliable.

horizontal line inside the boxes – of each sector and the corresponding micro-based benchmark is also high, at 54%. Repeating the exercise with $q > 4$ produces approximately the same chart. Cases when $q < 4$ are presented in the appendix, showing very similar findings.

Figure 1.3: Confidence Sets Constructed from Restricted Estimations



Notes: Parametric stability when 11 sectoral probabilities are fixed. Boxes represent the interval from the 25%ile to the 75%ile of distributions (for each sectoral probability). Horizontal lines are median estimates. 1365 restricted versions of the model are estimated, 365 estimates for each sector in total (vertically positioned). Instrument selection follows approach I. Baseline calibration. Starting values are benchmark reset probabilities.

1.4.4.2

Subsample Stability

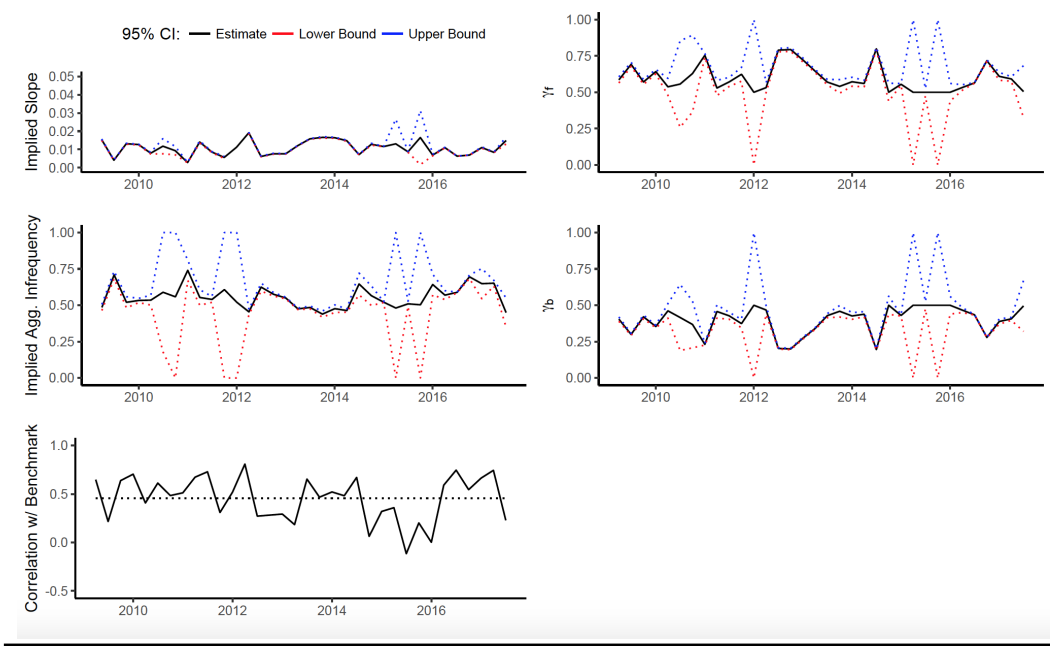
Next, we turn to uncertainties involving the sample. A typical practice has been to divide the sample into two or more periods, commonly splitting it around 1979 – e.g., Gali and Gertler (1999) and Clarida, Gali, and Gertler (2000). However, since our instrument sets are constructed based on AdaLASSO and the GMM using (19) has too many moment conditions, that separation is not feasible in our environment⁸⁹.

To test the subsample stability of our model, we conduct rolling-GMM estimations with $T = 180$ observations. In each step, instruments are re-selected based on (3), so that results also verify the degree of sensitivity of the instrument selection routine to the sample. Figure 1.4 presents the results for implied coefficients, such as the aggregate slope in (19) and the aggregate

⁸⁹The resulting number of observations when splitting the sample at 1979 restricts the reliability of the AdaLASSO routine for instruments.

infrequency in (20). Note that point estimates of both are substantially stable throughout the sample. This is also the case for forward and backward-looking coefficients in the aggregate NKPC. Although wider confidence intervals are produced in a few of our re-estimations, these are generally very low during the sample. Figure 1.4 also shows the correlations between estimated sectoral infrequencies $(1 - \hat{\lambda}_k)$ and those of the benchmark. They systematically lie around 0.5, approximately the average, in line with results of Table 1.8.

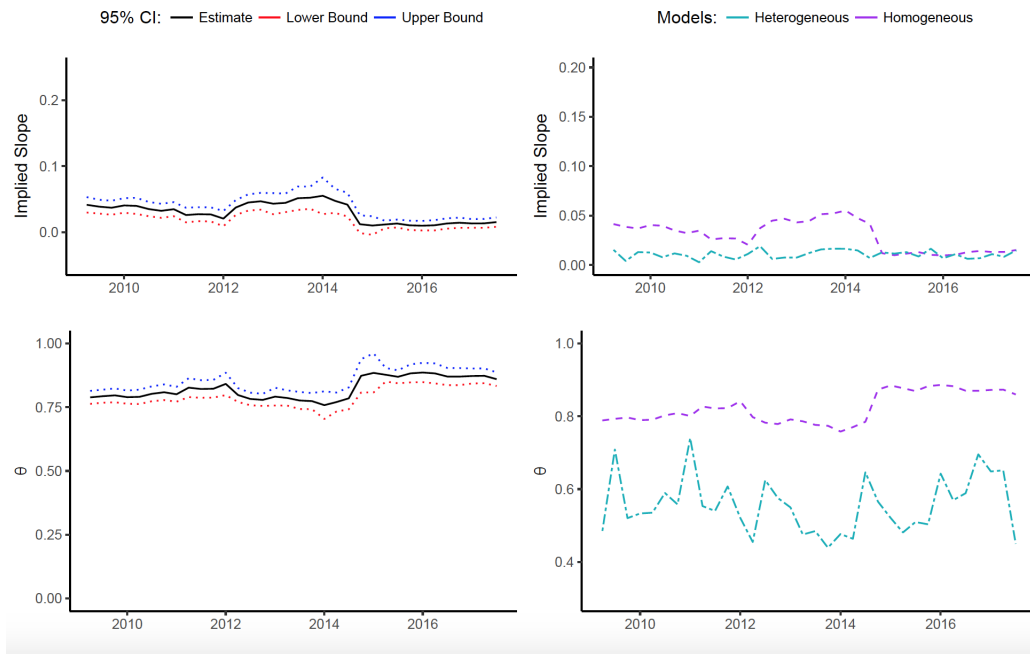
Figure 1.4: Subsample Stability: Implied Coefficients



Notes: Rolling-GMM estimates using the system in (19) and a HAC estimator for the covariance matrix. Confidence intervals are constructed by Delta method. Instruments are re-selected (based on the approach I) at each step. The horizontal axis measures the date of the last observation in the rolling subsample. “Implied aggregate infrequency” refers to (20). Results are for the normalisation (1) and starting values used are the benchmark reset probabilities based on micro data – see Bils and P. Klenow (2004) and Table 1.1 (appendix).

In Figure 1.5, we also compare estimations for the heterogeneous and the homogeneous economies. The same estimation setting is applied to the latter, exhibited in the first column (left). Consistent with previous results, note that the heterogeneous economy systematically implies a lower degree of stickiness (bottom right corner). For the homogeneous economy, the relationship between the slope and the degree of stickiness in the economy is trivial, so when estimates of the former are higher, the NKPC necessarily flattens (upper left). Such relationship becomes non-trivial when heterogeneity is introduced, due to the effects of strategic interactions across sectors. As a consequence, we estimate a slope that is low, positive – always significant at 1% – and stable when rolling the model (upper right), even though its implied stickiness moves in the interval from 0.50 to 0.70, exactly as in Table 1.8 and Table 1.9.

Figure 1.5: Subsample Stability: Homogeneous (Left) vs. Heterogeneous (Right)



Notes: Rolling-GMM estimates using the NKPC of the homogeneous economy (left) in (4) and the system for the heterogeneous economy, (19). For the latter, θ refers to (20). The horizontal axis measures the date of the last observation in the rolling subsample. Confidence intervals are constructed by Delta method. At each step, instruments are re-selected based on the approach I. Moment conditions are constructed using the baseline normalisation for both model (denoted as (1) in the previous tables). For the heterogeneous economy, starting values used are the benchmark reset probabilities based on micro data – see Bils and P. Klenow (2004) and Table 1.1.

A common discussion that has arisen in the empirical literature questions whether the Phillips curve for the US economy has flattened recently. Some studies do focus on how effective monetary policy may be attenuating the source of exogenous variation in the Phillips curve – e.g., Jorda and Nechio (2018) and McLeay and Tenreyro (2019), already mentioned. However, most of them emphasise the econometric setting instead – e.g., Stock and Watson (2007), Kuester, Müller, and Stölting (2009), Kleibergen and Mavroeidis (2009), Leduc, Wilson, et al. (2017) and Galí and Gambetti (2019). Results in Figure 1.5 can also shed light on this topic by suggesting that the Phillips curve has maintained its sensitivity to the output gap in recent years. One could adapt the model to evaluate the sensitivity to different measures of slack used in the literature, as the aggregate unemployment rate.

In the appendix, we reduce the number of moment conditions in the GMM by using approach II. We can then decrease the number of observations (to $T = 130$), increasing the number of re-estimations of the model. Analogous tables to Figure 1.4 and Figure 1.5 are presented, showing very similar results. There, we also show estimates of all of the deep parameters in (19). Deserving special attention, we find it interesting to compare the evolution of $\hat{\gamma}$ with that of $\hat{\theta}$. Estimates of the former lean downwards, whereas those of the

latter exhibit a clear upward trend. These findings are consistent with the already exposed assessment that the use of ad-hoc indexation schemes to improve the reliability of the model may mask essential dynamics in price setting. Heterogeneity likely mitigates the problem by the presence of strategic interactions between sectors, making the slope in its NKPC not particularly sensitive to γ . In contrast, in an identical-firms model, indexation rules can artificially inflate the slope. This may shed light on the reason why the real marginal cost commonly outperforms the output gap in the standard homogeneous models, since models with the former are often used while assuming some form of indexation.

1.4.5

Additional Robustness Checks

We conduct three additional robustness checks with the model. For the sake of conciseness, results are exhibited in the appendix.

First, we benefit from the fact that it is possible to identify the aggregate NKPC in (19) without directly estimating that equation. Since its parameters are also present in the remainder of system, the sectoral NKPCs in (16), it is possible to drop the aggregate equation. Results for the slope would still be biased towards zero – since y_t appears in the sectoral NKPCs. In addition, we would likely lose efficiency by removing important information otherwise exploited through cross-equation correlations in the errors. Nonetheless, all our findings are reconfirmed. Once more, results do not seem sensitive to the introduction of an indexation scheme in the economy.

Second, so far, *all* of our estimations point to a positive and statistically significant slope and to a degree of stickiness considerably more in line with the micro evidence. However, these findings depend on the calibration we assume for the parameters related to the degree of strategic complementarities across sectors – what defines Θ . We wonder how sensitive our results are to the latter, which, to some extent, is arbitrary.

Based on Carvalho (2006), we test a parameterisation that involves a lower elasticity of intertemporal substitution ($\sigma = 1$), a lower price elasticity of demand ($\epsilon = 5$) and a higher elasticity of the labour supply ($\varphi = 1.5$). This setting leads to $\Theta \approx 0.38$, further reducing the degree of real rigidities in the model, compared to the calibration in last column of Table 1.6, which implied $\Theta \approx 0.18$. Following the rationale, this should produce higher estimates of θ , as well as a higher slope. In the appendix, we show that $\hat{\theta}$ is still not as high as usually estimated in the literature – most values lie inside the interval (0.65, 0.70). The slope is also higher, but still very close to zero, around 0.04

for the hybrid model and 0.02 for the forward-looking one. Again, estimates are significant at 1%.

To test the sensitivity of a positive and statistically significant slope, we proceed with the following exercise. We increase the degree of real rigidities in the model until it generates a coefficient that is not significant at 5%. For the forward-looking model, this is first observed under an already quite unrealistic value of $\Theta \approx 0.05$ ⁹⁰. With the hybrid model, we could not generate a slope coefficient that was not significant⁹¹. The multi-sector heterogeneous economy works well in implying a positive dynamic relationship between inflation and the output gap regardless of the degree of real rigidity in the economy.

Finally, it is widely known that non-linear estimation methods can be quite sensitive to starting values. In the previous estimations, we relied on implied reset probabilities (λ_k) from the micro evidence as starting values. It could be that the seemingly reliable estimates we show are a direct consequence of that choice. For example, the vector of point estimates could be inherently copying the vector of starting values because of complexities in the moment conditions.

To verify the sensitivity of our model to initial values, we re-conduct structural estimates while relying on an “agnostic” routine to generate starting values. We individually estimate each sectoral equation, (16), and use the resulting $\hat{\lambda}_k$ as initial value for that sector in (19)⁹². Note that there are several complications involving the single estimation of sectoral NKPCs⁹³. Such procedure is *very* conservative, since nothing ensures that resulting estimates – and, then, starting values – are reliable. Table A.22 (appendix) shows that

⁹⁰In that case, we set $\sigma = 0.5$, $\epsilon = 20$ and $\varphi = 0.1$. Although not significant, the slope for the forward-looking model is still positive, at 0.001. The slope for the hybrid model was estimated at 0.003 (statistically significant).

⁹¹Even imposing $\Theta = 10^{-3}$ generates a slope of approximately 4×10^{-4} (still significant). We checked whether the model might be continuing to deliver a positive slope through distorted estimates of γ . It did not seem to be the case, as the latter was frequently around 0.50, with generally low standard errors ($< 10^{-2}$). We estimated an implied stickiness ($\hat{\theta}$) of 0.45 (or 0.30, depending on how we normalise the moment conditions). Standard errors on these values are also low (6×10^{-2}).

⁹²For each sectoral NKPC, instruments are the first three lags of the endogenous variables in the equation (as before).

⁹³First, inflation expectations may not vary much for some of the sectors. Second, we are not using instruments outside of the model to improve the estimations (what could be substantially important for some sectors – e.g., lags of commodity indices instrumenting the NKPC of “gasoline and other energy goods”). Third, issues of measurement errors might be quite relevant for sectors that are less representative of the entire economy (lower $f(k)$), what could produce distorted estimates of λ_k for this sector. Fourth, we find that estimates of λ_k are very sensitive to the econometric method when sectoral NKPCs are considered individually, what suggests to be sceptical on the reliability of these values. Lastly, note that a single distorted estimate of $\hat{\lambda}_k$ for some sector could further complicate the estimation of the system, if sensitivity to initial values is an issue for the latter.

our main findings are again maintained. Correlations with the benchmark are slightly lower, but still support the model⁹⁴.

1.5 Conclusion

This paper addresses a central complication in the empirical literature on the NKPC equation. It is well known that even when estimates for the coefficient of the slack variable are positive, the implied stickiness exhibited by the economy is too high to be consistent with the micro evidence. We propose a richer framework of a multi-sector economy with heterogeneity in pricing behaviour, which creates a high dimensional setting, exploited through a data-driven instrument selection routine.

The method delivers more reliable and precise estimations of the NKPC. These are less prone to common pitfalls evidenced in the literature and substantially more robust to underlying assumptions regarding the model, the estimation strategy, as well as arbitrary choices for calibrated parameters. We also show that ad-hoc workarounds typically applied in the model based on empirical motivations – as the use of the labour share as driving variable and the assumption of indexations schemes – no longer seem necessary in the heterogeneous framework. The model systematically delivers stable, positive and statistically significant estimates of the slope while implying a degree of stickiness that approaches that of the micro evidence. Estimated degrees of nominal rigidities in the cross section of sectors are also consistent with implications from disaggregated data, presenting narrow standard errors. In addition, the reliability of our approach does not seem to be affected by estimator uncertainties. We test it perturbing the econometric setting, as the number and which parameters are estimated, their values, the sample, the orthogonality conditions and the instrument sets. The model exhibits no significant change in performance in any of those exercises. Such level of robustness is certainly at odds with the rest of the literature, which typically struggles with minor changes in the econometric setting.

⁹⁴ Additionally, we conduct tests to verify the underlying rank conditions of the system. For example, seventeen parameters are estimated in the system (19), when considering $\gamma \neq 0$. We fix fifteen of them, and generate 2×2 combinations with the precision of 10^{-3} for the remaining two in the $[0, 1]^2$ space. The step is continued until all the possible combinations involving deep parameters are exploited – the Jacobian matrix is mapped $C_2^{17} = 136$ times. We do not reject that the Jacobian matrices for those combinations have full rank with a tolerance of 2.331468×10^{-14} – this corresponds to the number of rows in the matrix times the default *epsilon* in *Matlab*®.

Abstract

Most macro series used in academic research are usually smoothed, filtered and interpolated by official data providers. This paper shows that the use of filtered consumption series may considerably distort estimates of the elasticity of intertemporal substitution (EIS) in consumption-based asset pricing models. Once we use unfiltered consumption, we find that point estimates become more similar and confidence intervals can become tighter across different settings, data frequencies, as well as for different types of consumption data – macro and micro. Results also seem less sensitive to the presence of weak instruments, as, for instance, the completely uninformative weak-IV-robust confidence intervals usually found in the literature become rarer. Generally, we find that the EIS is quite low for macro data, albeit not as close to zero as commonly suggested in the literature. In this case, we often obtain values in the interval $[0, 0.5]$. For micro data, we estimate Euler equations conditional on the consumption of asset vs. non-asset holders. With unfiltered consumption, we do not find enough evidence of a different EIS across these groups. In addition, our estimates for stock holders are positive, but not above 0.3. Estimates for bond holders are higher, but more uncertain, usually from 0.4 to 1. In contrast, reported consumption seem unreliable, consistently returning negative estimates across groups.

2.1

Introduction

The elasticity of intertemporal substitution (EIS, henceforth) plays a central role in models of dynamic choice in macroeconomics and finance, capturing the sensitivity of consumption to the level of expected returns. For example, long-run risks models – Bansal and Yaron (2004), Bansal, Kiku, et al. (2014), Bollerslev, Tauchen, and Zhou (2009) and Bansal and Shaliastovich (2013), for instance – depend on the key assumption of an EIS above 1 to imply a sizeable equity premium, a low and stable risk-free rate and a correct cyclical behaviour of dividend-price ratios. However, the empirical literature

has struggled to reach an agreement on a reasonable range for that parameter. Estimates seem heavily influenced by the specification, econometric method, different measures of returns used and the characteristics of the household considered in the studies¹. It is still not clear what value should be considered as a reasonable guess to calibrate representative-agent models, for example. Early evidence from R. Hall (1988) had pointed to a “strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero”, but follow-up papers generally found mixed results and apparently none of them can provide a useful bridge to reconcile its empirical findings with values usually adopted in some macro models.

This paper attempts to improve estimates of the EIS by considering the fact that official consumption data has been smoothed, filtered and interpolated before release. Henceforth, we refer to these series as *reported* consumption. Kroencke (2017) showed that, in order to mitigate measurement errors, these statistical procedures undermine those series for research purposes by lowering their covariances with returns and introducing non-existent persistence. The former finding can be important to the extent that it implies that estimates based on reported data will be downward biased, what could explain the fact that a reasonable fraction of the papers in the literature state that the EIS is not statistically different from 0. Indeed, in the absence of the effects of such transformations, estimates might be higher and potentially more precise.

Our main findings suggest that *unfiltered* consumption – i.e., adjusted series that eliminate those noisy statistical procedures present in reported data – can significantly improve econometric results when estimating the EIS. We confirm this fact for several econometric methods, types of consumption data (macro and micro), as well as for data at different frequencies. Unfiltered consumption is also substantially important to obtain more precise estimates of the EIS when considering specific groups of asset holders in the Euler equation, relevant issue when testing the limited asset market participation theory (henceforth, LAMP), for instance.

First, using aggregate expenditures data from the National Income and Product Accounts (NIPA), we show that estimates of the EIS tend to increase relatively to cases that use reported consumption instead. Point estimates across different classes of estimators and frameworks also become more similar,

¹See Gomes and Paz (2013) for an example of an alternative return measure constructed to capture the representative agent’s asset portfolio. Moreover, characteristics of the household are relevant to the extent that they participate differently in local markets – see Vissing-Jorgensen (2002) and Guvenen (2006), for instance. Besides that, Havranek et al. (2013) find that distinct estimates of the EIS for different regions or cohorts seem more related to the specific assets held by different groups and their income than to local preferences.

being roughly from 0 to 0.5 in our baseline specifications, compared to -0.2 to 0.2 , when using reported consumption.

Second, those estimates of the EIS seem less affected by the presence of weak instruments. We obtain more stable estimates across econometric methods regardless of whether unfiltered consumption produces lower first-stage F-statistics, compared to its reported analogue.

Third, while completely uninformative weak-instrument-robust confidence intervals² are quite frequent in the empirical literature of the EIS for macro consumption data, – Yogo (2004), Ait-Sahalia, Parker, and Yogo (2004), Ascari, Magnusson, and Mavroeidis (2016) and Gomes and Paz (2013) –, this paper shows that unfiltered consumption can transform these impractical intervals into more plausible sets. For instance, Yogo (2004) found uninformative sets in 66 percent of specifications estimated with reported macro expenditures data for the US economy. In contrast, for the same framework, we only obtain uninformative sets in 28 percent of our econometric approaches relying on unfiltered consumption. Consistent with Yogo (2004), the use of reported data generates uninformative sets in nearly half of our specifications based on macro data.

Lastly, empirical benefits of unfiltered consumption series are also confirmed using micro data from the Consumer Expenditures Survey (CEX, henceforth). With this data, it is also possible to test the limited participation hypothesis by obtaining distinct estimates of the EIS based on different groups of asset holders. We split households between stock and non-stock holders and between bond and non-bond holders. Unfiltered consumption is once again important. It produces estimates for stock holders that lie inside the interval from 0 to 0.3. The EIS for bond holders is more uncertain, albeit higher, generally from 0.3 to 1. In contrast, estimations based on reported consumption exhibit a less clear pattern, consistently returning negative values across different groups of asset holders.

To construct unfiltered consumption data, we adapt the so-called Filter model in Kroencke (2017), so that we can use it to estimate the EIS using several types of data at different frequencies, rather than annual macro data only, as in the original model. According to the method, government statisticians who *only* collect an admittedly noisy observation for consumption opt to use a Kalman filter to estimate the unobserved level of consumed goods as precisely as possible – henceforth, we refer to the latter as *state* consumption³. Consequently, reported (or filtered) consumption is then defined

²We define uninformative intervals as either empty sets or ones that cover the whole real line.

³Kroencke (2017) used the term *true* consumption instead. This reflects the fact that

by fitted values of this Kalman filter, while unfiltered consumption is obtained by reverse engineering to guess what their first (supposedly noisy) observations were before being subject to the procedure. We propose a modification to the original model in order to introduce serially correlated measurement error. Specifically, our variation of the model is possible through the solution of a parallel quasi-differenced Kalman filter. We then map this solution back onto the original model, without changing main assumptions. This modification is essential to the extent that serially correlated measurement error terms are more relevant when either using disaggregated data or data at higher frequencies – see Wilcox (1992), Bell and Wilcox (1993) and the online appendix of Kroencke (2017)⁴.

With the modified model in hands, we review identification approaches for the EIS usually adopted in the established literature, importantly those of Yogo (2004) and Vissing-Jorgensen (2002). The former addressed the empirical puzzle that estimates of the EIS are often statistically less than 1, while their reciprocal is not different from 1. He considered L. G. Epstein and Zin (1989) preferences and eleven developed countries while applying weak-identification-robust techniques. Although his final conclusions agree with R. Hall (1988), he finds point estimates that are rather imprecise across countries, mostly reflecting the presence of weak instruments. This inaccuracy was particularly true for the US economy, addressed in our paper⁵. A sensible explanation for this fact is that limited participation in asset markets may be plaguing results once Euler equations may no longer hold for the representative agent, possibility addressed by Vissing-Jorgensen (2002) and Guvenen (2006)⁶. The former used CEX panel data to verify how estimates of the EIS may differ taking into account different types of households in asset markets, as bond vs. stock holders vs. non-asset holders. The latter shows how lower estimates of the EIS are obtained when considering aggregates that ignore the facts that the majority of households do not participate in stock markets and that

unobserved consumption in the model represents what government statisticians classify as *true* consumption according to their beliefs. We prefer to use the term *state* consumption in order to clarify the fact that this variable is modelled as *true* consumption *only* by those statisticians and should not be confused with the true level of consumed goods in the economy.

⁴Consistent with that, we found significantly weaker results at higher frequencies when serially correlated measurement errors were not allowed. This was the case for NIPA consumption at quarterly frequency, for instance. Simply applying the original model of Kroencke (2017) provided such imprecise estimates of the EIS that even official data performed better in comparison.

⁵Yogo (2004) finds many empty and infinite weak-IV-robust confidence intervals for the EIS using US data independent from the data frequency, indicating that his baseline model is entirely rejected for this country. We use his framework in section 3 of this paper.

⁶Yogo (2004) also mentioned this possibility.

most of the wealth is held by a small fraction of population with a high EIS. Similarly, Ait-Sahalia, Parker, and Yogo (2004) review the so-called Equity Premium Puzzle – Mehra and Prescott (1985) – and estimate the EIS using not only the consumption of essential goods, but also that of luxury goods. While Vissing-Jorgensen (2002) finds that the EIS is not the same for stock and bond holders (0.3-0.4 and 0.8-1, respectively), results in Ait-Sahalia, Parker, and Yogo (2004) are somewhat inconclusive, albeit they do mention that the parameter is possibly higher for the consumption of luxury goods.

While none of the papers in the EIS empirical literature have considered that inaccurate estimates might be related to the fact that official data are filtered in order to mitigate measurement errors, there are a few papers in the asset pricing literature accounting for this fact. In addition to Kroencke (2017), Savov (2011) addressed the Equity Premium Puzzle and showed that reported consumption performs so poorly in asset pricing models that even the use of garbage data instead provides much better results. In a more complex framework and relying on Bayesian methods, Schorfheide, Song, and Yaron (2018) present a mixed-frequency approach that controls for measurement errors and time-varying volatilities⁷. In general, it is consensus that is quite hard to track true consumption in the data.

Our paper brings up the question about how the use of unfiltered consumption data may generate more reliable and precise estimates of the EIS. Indeed, we present evidence on how filtering out noisy elements present in official releases of consumption data (paradoxically, due to filtering of the original data) can help us to discipline econometric results in the estimation of that parameter. Furthermore, we evaluate our findings relying on weak-identification routines. The use of these techniques in the EIS empirical literature does not seem sufficiently disseminated yet. In addition to Yogo (2004), only a few papers address the subject. Ascari, Magnusson, and Mavroeidis (2016), Ait-Sahalia, Parker, and Yogo (2004), Gomes and Paz (2013) and J. C. Fuhrer and Rudebusch (2002) are examples, albeit the latter in a more macro-based framework.

This paper is organised as follows. In section 2.2, we present our modifications of the model in Kroencke (2017). The details on the complete model are available in the appendix. Section 2.3 evaluates how unfiltered consumption affects the estimates of the EIS in a framework of L. G. Epstein and Zin

⁷They log-linearise and estimate a state-space representation that simultaneously accounts for consumption and its corresponding measurement errors at different frequencies. We come back to this later, but for now have in mind that the frequency of consumption matters much for researchers interested in asset pricing models that attempt to track implicit/noiseless consumption data.

(1989) preferences and log-linearised Euler equations, in spirit of Yogo (2004). In the following section, we feed the methodology into consumption measured by CEX data to verify the potential effects on estimates of the EIS across different types of asset holders, testing the LAMP. The last section concludes.

2.2 Model

This section describes how we adapt the filter model in Kroencke (2017), allowing for serially correlated measurement error. As mentioned earlier, this adaptation makes the model suitable for different types of data at several frequencies – rather than just for macro annual data, as in the original model. For the sake of conciseness, in this section we only cover parts of the model which are modified and that are relevant to a smooth reading of this paper. The complete model and its derivation are presented in the appendix.

As in Kroencke (2017), we assume that statisticians who prepare the data for release collect a first (primitive) measure of consumed goods y_t , believed to be noisy. They conjecture that y_t is formed by *their belief* of true consumption c_t and an additive measurement error component ξ_t ⁸:

$$y_t = c_t + \xi_t \quad (1)$$

Henceforth, we refer to c_t as *state* consumption. Our paper contrasts with Kroencke (2017) at this point, who used the term true consumption instead. We adopt a different name to emphasise that the model captures what statisticians *believe* to be true consumption (c_t), rather than the correct measure of true consumption in the economy⁹.

These statisticians model state consumption by a random walk representation:

$$c_t = c_{t-1} + \mu_{c,t} + \sigma_{\eta,t}\eta_t, \quad (2)$$

where $\eta_t \sim N(0, 1)$ and we assume $\mu_{c,t} = \mu_c = 0$. Equation (2) does *not* mean that true consumption follows a random walk process nor that it has a constant drift. Instead, it only implies that government statisticians filter the data considering that true consumption follows that stochastic process, while assuming a constant drift¹⁰. We later assume that $\sigma_{\eta,t}$ follows a GARCH

⁸You can see y_t as a first measure of consumption which has not been affected by filtering, smoothing and interpolation procedures. Alternatively, you can think of it as the garbage measure of Savov (2011).

⁹In addition, c_t will be the state variable in a Kalman filter, another reason for that name.

¹⁰Formally, we remove the mean of the series before calibrating the model, to then add it back to construct unfiltered consumption. These steps were also adopted in the original model and the use of data at different frequencies does not alter this part of the model.

process, as in the original model.

To generalise the model, we introduce persistent measurement error by relying on the following AR(1) representation:

$$\xi_t = \rho_\xi \xi_{t-1} + \sigma_\nu \nu_t, \quad (3)$$

where $\nu_t \sim N(0, 1)$. Generally speaking, it is trivial to expand a Kalman filter embedding (3)¹¹. Harvey, Ruiz, and Sentana (1992) discuss how to model and extend these filters while assuming ARCH or GARCH processes for variance terms. Nonetheless, the model with (3) can not be solved in terms of unfiltered consumption using typical procedures. In this regard, we follow E. Anderson et al. (1996), rewriting the state-space representation in terms of a “quasi-difference”:

$$\bar{y}_t = y_{t+1} - \rho_\xi y_t, \quad (4)$$

where \bar{y}_t represents the “quasi-differenced” counterpart of y_t ¹². Once the solution for \bar{y}_t is obtained, we then use (4) to map it back onto y_t . This final estimate of the latter (\hat{y}_t) is what we later call *unfiltered* consumption.

It is worth emphasising that we are not interested in fitted values of state variables after putting observable data into the filter. Instead, as in Kroencke (2017), we want to estimate what the original observed data (y_t) were once *all* we have are fitted values of a state variable (c_t) tracked by a Kalman filter. The quasi-differencing approach makes the reverse engineering we have to deal with when inverting the Kalman filter possible without imposing additional complications in the way we solve the model¹³. It also does *not* mean that statisticians consider (4) when filtering the data. Instead, they solely consider (1), (2) and (3).

Equations (1) to (4) form together a quasi-differenced Kalman filter whose solution can be written as¹⁴:

$$\hat{c}_t = \hat{c}_{t-1} + K_t^c (\bar{y}_{t-1} - (1 - \rho_\xi) \hat{c}_{t-1}), \quad (5)$$

$$K_t^c = \frac{P_t^c (1 - \rho_\xi) + \sigma_{\eta,t}^2}{P_t^c (1 - \rho_\xi)^2 + \sigma_{\eta,t}^2 + \sigma_\nu^2}, \quad (6)$$

$$P_t^c = P_{t-1}^c (1 - (1 - \rho_\xi) K_{t-1}^c) + (1 - K_{t-1}^c) \sigma_{\eta,t}^2, \quad (7)$$

where $\hat{c}_t = E_t[c_t]$ denotes *reported* consumption (conditional time- t estimate

¹¹Perhaps the simplest form is to expand the vector of latent variables, now including ξ_t .

¹²Typically, a “quasi-difference” involves lags of the variable. We are following the term used in E. Anderson et al. (1996) here. From (4), we have that $[y_{t+1}, y_t, \dots, y_0, \hat{c}_0]$ and $[\bar{y}_t, \bar{y}_{t-1}, \dots, \bar{y}_0, \hat{c}_0]$ span the same space. By construction, this implies that prediction errors in \bar{y}_t are actually innovations in y_{t+1} . See the appendix and E. Anderson et al. (1996) for more details.

¹³That is, we can solve the model following similar steps as in Kroencke (2017).

¹⁴Check the appendix for the derivation.

of true consumption), P_t^c is the conditional variance of c_t and $\sigma_{\eta,t}^2$ denotes the volatility parameter in (2). Importantly, K_t^c is the Kalman gain associated with true consumption, what directly governs the persistence of reported consumption. Let $(\bar{y}_{t-1} - (1 - \rho_\xi)\hat{c}_{t-1}) = u_t$ be the “re-scaled prediction error”, a surprise factor¹⁵. When K_t^c is relatively high, statisticians attribute more weight to the surprise factor than to their past estimate \hat{c}_{t-1} (reported consumption for the last period). Consequently, \hat{c}_t is less persistent. Kroencke (2017) had also shown that unfiltered consumption exhibits higher covariances with expected returns, what we later confirm in our results in terms of the EIS¹⁶.

One way to verify consistency of the filter is to check whether K_t^c increases in periods of economic turbulence (recessions, for example). Intuitively, when the variability of economic shocks is high relatively to the volatility of the measurement error, it is optimal for statisticians to adjust \hat{c}_t taking into account surprising data more than their past estimates, \hat{c}_{t-1} . Consequently, K_t^c is higher and reported consumption less persistent.

Algebraically, such mechanism comes from the analogue of (6) in the original model of Kroencke (2017)¹⁷. However, since equation (6) is not as simple as his, we must derive parametric conditions under which the derivatives of K_t^c for P_t^c and $\sigma_{\eta,t}^2$ are positive as well (so that more economic turbulence implies a higher Kalman gain). Fortunately, we find that our model behaves properly in this regard under reasonable parametric conditions. We present these conditions in proposition 1 below. For expository reasons, they are written in terms of a *homoscedastic* version of the model (when $\sigma_{\eta,t}^2 = \bar{\sigma}_\eta^2$, while \bar{K}^c and \bar{P}^c are also fixed to steady-state values)¹⁸. The formal proof as well as more details on the homoscedastic model are exhibited in the appendix. Henceforth, we refer to the baseline model when K_t^c , P_t^c and $\sigma_{\eta,t}^2$ are time-varying as *heteroscedastic*, but we will later present results for its homoscedastic analogue as well. Importantly, bear in mind that heteroscedasticity in the model does not imply the assumption of heteroscedasticity in our estimations, since the former encompasses *state* consumption, whereas the latter relates to error terms of models where *unfiltered* consumption is the regressor.

¹⁵This corresponds to the prediction error of the original model, but the term $(1 - \rho_\xi)$ adjusts it for the presence of the quasi-difference \bar{y}_t instead of y_t .

¹⁶When it is the other way around, reported consumption becomes a very persistent and predictable series and its correlation with asset returns normally lowers in comparison.

¹⁷See Kroencke (2017), p. 54, equation (5).

¹⁸The same conditions are valid point-to-point in time, but derivatives must hold at any single period. If the filter converges to steady-state values, that should not be a problem.

Proposition 1: if state consumption is homoscedastic and $4\frac{\sigma_\nu^2}{\sigma_\eta^2} > (1 + \rho_\xi)^2\bar{\sigma}_\eta^2 - (1 - \rho_\xi)^2$, then its unconditional variance and Kalman gain follow:

$$\begin{aligned}\bar{P}^c &= \frac{\bar{\sigma}_\eta^2}{2(1 - \rho_\xi)} \left([(1 - \rho_\xi)^2\bar{\sigma}_\eta^2 + 4\sigma_\nu^2]^{\frac{1}{2}} - (1 + \rho_\xi)\bar{\sigma}_\eta^2 \right), \\ \bar{K}^c &= \frac{\bar{P}^c(1 - \rho_\xi) + \bar{\sigma}_\eta^2}{\bar{P}^c(1 - \rho_\xi)^2 + \bar{\sigma}_\eta^2 + \sigma_\nu^2},\end{aligned}\tag{8}$$

while:

$$\frac{\partial \bar{K}^c}{\partial \bar{P}^c} > 0; \quad \frac{\partial \bar{K}^c}{\partial \bar{\sigma}_\eta^2} > 0 \iff \sigma_\nu^2 - (1 - \rho_\xi)\bar{P}^c\rho_\xi > 0.\tag{9}$$

Proof: See appendix.

In practice, a sufficiently small value of ρ_ξ ensures the second part in (9). In addition, we find that (9) is easily satisfied for different calibrations of the model.

Next, we need to derive a measure of unfiltered consumption that is compatible with the quasi-differenced filter in (5-7). Adapting developments in Kroencke (2017) for our model, one can isolate \bar{y}_t in (5) and conduct simple adjustments that account for time-aggregation bias to find¹⁹:

$$\hat{\bar{y}}_{t-1} = \frac{\hat{c}_t - (1 - (1 - \rho_\xi)\Omega_t)\hat{c}_{t-1}}{\Omega_t},\tag{10}$$

where $\Omega_t = \alpha K_t$ and we set $\alpha = 0.8$, as in the original model. Equation (10) above represents “quasi-differenced unfiltered consumption”. Once $\hat{\bar{y}}_t$ has been found, we need to transform it back into its primitive, unfiltered consumption, \hat{y}_t . Based on (4), we do this following:

$$\hat{y}_t = \hat{\bar{y}}_{t-1} + \rho_\xi \hat{y}_{t-1}\tag{11}$$

More details on how we use (11), as well as on how we initialise our model are presented in the appendix.

2.2.1 Consumption Volatility

As in the original Filter model, we consider a time-varying consumption volatility in (2), which follows a GARCH(1,1) stochastic process:

$$\sigma_{\eta,t}^2 = a_0 + a_1\eta_{t-1}^{*2} + a_2\sigma_{\eta,t-1}^2,\tag{12}$$

¹⁹Check the appendix for more details.

where $\eta_t^* = \sigma_{\eta,t}\eta_t$ ²⁰. With (12), we let the error term capture the inherent dynamics of the data, so that the random walk hypothesis is not an obstacle²¹. Furthermore, since we are ultimately interested in unfiltered consumption, it should not matter much if we have a random walk process with a data-driven model for its variance or a covariance-stationary model with constant volatility, as far as both generate unfiltered series whose moments are sufficiently similar. As in Kroencke (2017), we find that different calibrations for the GARCH specification generate very similar results if moments are relatively matched, so that (12) seems to perform well when applied to the data.

2.2.2

Adjusting Asset Returns

We also need time-aggregation-bias adjustments for returns, such that their timing is compatible with that of (10). These steps are identical to those in Kroencke (2017) when we use annual data and have nothing to do with a serially correlated measurement error, (4)²². However, we use a (necessary) slight adaptation when working with other data frequencies. Corrections are only performed on the return series when the econometric method applied uses unfiltered consumption. For reported consumption, asset returns need not be corrected for a different timing since we do not adjust that of reported data²³. It is worth emphasising that these adjustments are not essential to validate the main findings of this paper, since similar results are found using raw returns data²⁴.

2.2.3

Calibration

This paper follows parameterisation techniques in Kroencke (2017). However, the quasi-differencing approach demands an additional step. Specifically, with the presence of (3) we also need to calibrate ρ_ξ . To the best of our knowledge, Schorfheide, Song, and Yaron (2018) is the only paper that estimates something similar in the literature. Relying on Bayesian methods, the analogue of ρ_ξ with monthly consumption data is estimated as 0.06 in their paper,

²⁰Equation (12) is for the (baseline) *heteroscedastic* model. Presumably, $\sigma_{\eta,t}^2 = \bar{\sigma}_\eta^2$ for the *homoscedastic* version.

²¹By modelling state consumption growth as i.i.d., we let the GARCH component (12) absorb the dynamics of the data, so that the choice for the process itself becomes less fundamental.

²²Recall that the original model only handles annual data.

²³We follow Kroencke (2017) once more here. Results barely change when we repeat our estimations with reported consumption while correcting the timing of returns.

²⁴We repeat our main tables for NIPA consumption (macro data) using solely raw returns in the appendix.

albeit in a much more complex model. Discrepancies apart, when testing our filter on macro data (NIPA consumption) for different parametric combinations, we found out that it behaves properly for different values of ρ_ξ in a neighbourhood around 0.06 – such that K_t^c increases in recessions, in line with the intuition. Figure 2.1 below presents how K_t^c varies for ρ_ξ fixed around that neighbourhood, specifically at 0.03, 0.06 and 0.09. Besides, if remaining parameters are calibrated such that benchmarked moments of unfiltered consumption are sufficiently aligned, we find that different values of ρ_ξ in that neighbourhood simply do not matter much²⁵. Therefore, even when not using data at monthly frequency, we fix $\rho_\xi = 0.06$ throughout the paper, while adjusting remaining parameters following steps in Kroencke (2017)²⁶. The only exception is when we estimate the EIS using annual macro data. As already mentioned, serially correlated measurement error terms are less relevant for aggregate data at lower frequencies. Hence, we apply the original model ($\rho_\xi = 0$) in that case.

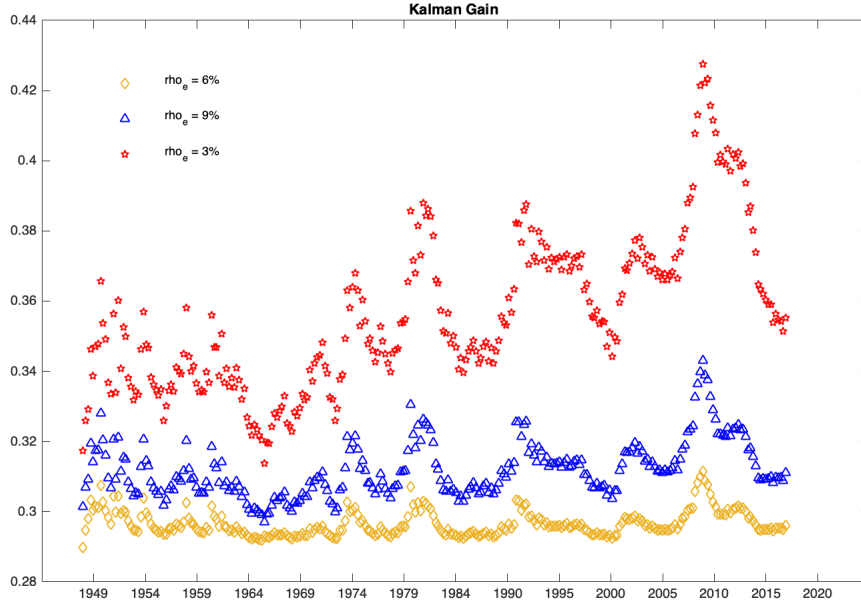
Since measurement error terms cancel out over longer horizons, – see Daniel and Marshall (1996) –, Kroencke (2017) used a value for $\bar{\sigma}_\eta$ that matched 6-year standard deviations of simulated and empirical data (garbage, reported and unfiltered as the latter)²⁷, calibrating the model based on post-war data. We calibrate $\bar{\sigma}_\eta$ under the exact same method when using annual NIPA data (as in his paper). When using quarterly NIPA data, we accumulate quarterly expenditures to get an implied annual consumption growth measure whose moments satisfy the same procedure. Since quarterly consumption data are benchmarked to annual counterparts before officially released, our calibration method is consistent with actual procedures conducted on the data²⁸. Based on Bansal and Yaron (2004), we find that values of $\bar{\sigma}_\eta$ in

²⁵In fact, correlations between different generated series of unfiltered NIPA consumption are similar once we change ρ_ξ while adjusting other parameters taking into account benchmark moments. Hence, the GARCH component of the model is perfectly adjustable to capture the consumption dynamics even when ρ_ξ is modified. Subsequent unfiltered consumption series are not econometrically distinguishable in terms of estimates of the EIS.

²⁶Later in section 4, data is at monthly frequency, even though consumption growth is semiannual. We transform the latter into monthly consumption growth in order to calibrate the model. Thus, scale and frequency of consumption follow Schorfheide, Song, and Yaron (2018) closely. Consequently, $\rho_\xi = 0.06$ is appropriate.

²⁷He chose that horizon as his benchmark based on considerations involving simulated data.

²⁸Monthly and quarterly official consumption data are based on the monthly retail trade survey (MRTS), while annual data comes from the annual retail trade survey (ARTS). Since issues of sampling error are more significant in the MRTS, data from the latter are used to mitigate these problems. Hence, calibrating our model for quarterly macro data based on corresponding (implied) moments for annual macro data is consistent with actual steps conducted on the data before release. For more details on how NIPA consumption is generated, see the online appendix of Kroencke (2017) or the official NIPA handbook: BEA (2017).

Figure 2.1: Kalman Gain Over Time for Different Values of ρ_ξ 

Note: Quarterly Kalman gain K_t^c in our sample for three different values of ρ_ξ (0.03, 0.06 and 0.09). The chart shows that consumption is less filtered (higher K_t^c) during the 2009 crisis, the early 2000's recession and early 90's and 80's recessions, for instance.

a neighbourhood of $\sqrt{3} \times 0.0078 \approx 1.4\%$ for quarterly NIPA consumption and $\sqrt{12} \times 0.0078 \approx 2.7\%$ for annual NIPA consumption match our moment requirements quite well²⁹. The latter is very similar to Kroencke's ($\bar{\sigma}_\eta = 2.5\%$, based on annual data as well), and this slight difference does not change moments substantially³⁰. Therefore, we set $\bar{\sigma}_\eta = 1.4\%$ and $\bar{\sigma}_\eta = 2.5\%$ for quarterly and annual data, respectively.

Next, σ_ν is the only parameter remaining to be calibrated. There is nothing particularly different here, and we choose its value such that the 6-year standard deviation of unfiltered consumption is approximately 1.2 times the value of its reported counterpart, as in the original Filter model. The motivation for this is that official statisticians should not make mistakes systematically when filtering the data, so that moments of reported consumption should not considerably exceed those of unfiltered consumption when measured over longer periods. For quarterly NIPA consumption, following this rule returns dif-

²⁹Those values represent the counterparts of $\bar{\sigma}_\eta$ in the model of Bansal and Yaron (2004), but adjusted for quarterly and annual data instead (they considered the value of 0.0078 at monthly frequency). The connection between their paper and the Filter model is not surprising. In fact, Kroencke (2017) used a modified version of their model to simulate state consumption (referred to as “true” consumption in that paper).

³⁰As in Kroencke (2017), we find that unfiltered consumption is not sensitive to different combinations of parameters in the GARCH(1,1) process for the conditional variance. We better specify how we treat these in the appendix.

Table 2.1: Calibrated Moments for NIPA Consumption of Nondurables and Services

[(Implied) Consumption Growth	$E(\Delta C_{year})$	$\sigma(\Delta C_{year})$	$\sigma(\sum_{year=1}^6 \Delta C_{year})/\sqrt{6}$	$Corr(\Delta C_{year}, \Delta C_{year-1})$
Reported (NIPA)	1.94%	1.41%	1.91%	32.01%
Simulated*	1.90%	2.48%	2.39%	1.54%
Garbage*	1.42%	2.86%	2.44%	-14.26%
Unfiltered - APWG* (1960-14)	1.85%	2.57%	2.44%	0.56%
Unfiltered - APWG* (1928-14)	1.79%	4.07%	3.08%	-10.89%
Unfiltered - Our Model (Quarterly Data)				
Homoscedastic (1960-14)	1.99%	2.88%	2.43%	-3.97%
Heteroscedastic (1960-14)	1.97%	2.30%	2.24%	2.69%
Homoscedastic (1947-17)	1.99%	3.26%	2.36%	-20.31%
Heteroscedastic (1947-17)	1.99%	2.49%	2.15%	-15.59%
Unfiltered - Our Model (Annual Data)				
Homoscedastic (1960-14)	1.30%	3.25%	2.45%	-8.22%
Heteroscedastic (1960-14)	1.91%	2.39%	2.26%	-0.71%
Homoscedastic (1930-17)	1.59%	5.11%	3.34%	-14.41%
Heteroscedastic (1930-17)	2.02%	3.71%	2.50%	-5.76%

Note: Moments of reported and unfiltered consumption (our model). We compare these moments with those of Kroencke (2017) as well: simulated consumption, garbage and unfiltered consumption (APWG stands for "Asset Pricing Without Garbage"). We have simply copied his results here, writing "*" next to variables presented in that paper. Reported and unfiltered consumption are for nondurables and services, from NIPA tables. We consider the quasi-differenced model with serially correlated measurement errors for quarterly data, setting $\rho_\varepsilon = 0.06$. For annual data, the model is the same as in Kroencke (2017). See section 2.3 for more details on calibration.

ferent calibrations across models: $\sigma_\nu = 3.8\%$ (heteroscedastic) and $\sigma_\nu = 2.5\%$ (homoscedastic). For annual NIPA data, statistical moments do not differ as much regarding the model and $\sigma_\nu = 2.8\%$ is set for both versions.

Table 2.1 compares moments of unfiltered and reported NIPA consumption based on nondurables and services. In the appendix, Table B.1 displays the same information for the consumption of nondurables only³¹. We present other relevant consumption measures shown in Kroencke (2017): *simulated* (he simulates state consumption using a long-run risk model built on Bansal and Yaron (2004)³²); *garbage* (as in Savov (2011)), and; *unfiltered* (for which we simply show results in Kroencke (2017)). Moments are displayed both for the complete sample (1930-2017 for annual and 1947:3-2017:4 for quarterly macro data) and for the post-war subsample used by Kroencke (2017) to calibrate moments, covering the period 1960-2014. For unfiltered NIPA consumption, we present results for both variants of our model: one where state consumption has constant volatility (homoscedastic version) and another with time-varying volatility (heteroscedastic version, the baseline model).

Regardless of the data frequency, our measures of unfiltered NIPA consumption can reproduce the mean-reversion behaviour exhibited by garbage. In addition, unfiltered NIPA consumption is more autocorrelated in the complete sample than in the period comprehending 1960 to 2014, consistent with results

³¹The calibration for this model is discussed in the appendix.

³²As mentioned above, he refers to this measure as "true" rather than state consumption.

found in Kroencke (2017) – see the first panel in Table 2.1. In contrast to its unfiltered analogue, reported NIPA consumption is quite persistent, consistent with the idea that the data are heavily filtered before release (lower K^c)³³.

Turning to micro (CEX) data, calibrated moments are exhibited in Table 2.2. The same procedures to calibrate the model are applied. The CEX data is subject to a number of statistical procedures before release, many of which relatively similar to those applied on NIPA data – see section 4 and the appendix for a discussion. It is also known that a significant fraction of the CEX consumption categories exhibit a similar behaviour compared to the NIPA analogues. Other categories do measure different things or have similar definitions but exhibit a CEX/NIPA ratio that is too low (high) over time. In terms of the estimation of the EIS, it is fundamental for the Filter model to be able to revert second moments and autocorrelations, as well as to exhibit higher covariation with returns. Inferring how much one source may be overestimating consumption growth relatively to the other is considerably less important. If overall there is no substantial change in how much CEX categories overestimate (underestimate) its NIPA analogues, then one can apply the same method to both sources when calibrating the model. In addition, it is a common procedure to aggregate consumption goods from the CEX taking the NIPA categories as reference³⁴.

As will become clear in section 4, the CEX data allows us to split households between different types of asset holders – stock holders vs. non-stock holders and bond holders vs. non-bond holders, for instance. Since official statistical procedures do not distinguish between different asset holders, we calibrate the model based on the consumption growth series of *all* households, imposing the resulting parameterisation to the consumption of the corresponding groups.

We have to make one small change to the calibration method when working with CEX data. The time series we construct in section 4 measures semi-annual consumption growth, but at monthly frequency. To equalise scale and frequency, we transform the data into monthly consumption growth before calibrating the model, to then revert the scale back into semiannual consumption growth. Since we calibrate the model based on monthly consumption growth at the same frequency, we fix $\bar{\sigma}_\eta = 0.0078$, following Bansal and Yaron (2004), who used monthly data. For the same reason, $\rho_\xi = 0.06$ is set, motivated by

³³In addition, observable means of all variables are similar. This is intuitive since both measurement errors and treatment procedures made on the data shall cancel over time. We also present results for the consumption of nondurables only. See Table 2.1 (nondurables and services) and Table B.1 (nondurables) in the appendix for more details.

³⁴See Attanasio and Weber (1995) and Vissing-Jorgensen (2002), for instance.

Table 2.2: Calibrated Moments for CEX Consumption: 1982 to 2013

Consumption Growth (Annual, Implied)	$E(\Delta C_{year})$	$\sigma(\Delta C_{year})$	$\sigma(\sum_{year=1}^6 \Delta C_{year})/\sqrt{6}$	$Corr(\Delta C_{year}, \Delta C_{year-1})$	Observations per Month (Mean)
NIPA Consumption					
Reported	1.65%	1.23%	1.78%	60.57%	-
Unfiltered	1.65%	2.13%	2.22%	11.51%	-
CEX: All Households					
Reported	1.98%	5.65%	2.64%	-36.40%	246
Unfiltered	2.10%	6.63%	3.17%	-55.45%	-
CEX: Stock Holders					
Reported	3.07%	7.36%	4.37%	-22.30%	49
Unfiltered	3.10%	10.30%	5.22%	-56.66%	-
CEX: Non-Stock Holders					
Reported	1.57%	5.22%	2.50%	-35.42%	197
Unfiltered	1.74%	6.34%	2.97%	-52.36%	-
CEX: Bond Holders					
Reported	3.14%	6.87%	4.01%	-22.69%	70
Unfiltered	3.25%	10.01%	4.96%	-58.49%	-
CEX: Non-Bond Holders					
Reported	1.27%	5.03%	2.18%	-39.34%	176
Unfiltered	1.45%	6.11%	2.65%	-54.42%	-

Note: Moments of reported and unfiltered based on CEX data (1982-2013). We also exhibit moments of reported and unfiltered consumption based on macro data (NIPA consumption), calculated for the CEX period. The original CEX data measures semi-annual consumption growth at monthly frequency. We convert these series into monthly consumption growth to calibrate the model but aggregate the data to obtain moments for (implied) annual consumption growth – first column – so that these are comparable with moments in Kroenke (2017). In order to account for the fact that most likely statisticians do not adjust the data splitting by households, we calibrate moments based on all households. The last column provides the mean number of (cross-sectional) observations for each month, measured over the sample. Our CEX sample consists of 90,080 households.

results in Schorfheide, Song, and Yaron (2018). Finally, the same rule for $\bar{\sigma}_\nu$ applies, which establishes that the long-run standard deviation of unfiltered consumption is not higher than 1.2 times that of its reported analogue. It returns $\bar{\sigma}_\nu = 2.7\%$.

From Table 2.2, it can be seen that unfiltered CEX consumption (for all households) repeats the same patterns in Table 2.1, for NIPA data. In fact, unfiltered data are again more volatile, exhibiting more mean reversion than reported consumption (it is also the case regardless of the group of asset holders considered). Check the appendix for more details on how we calibrate the model for CEX data.

2.3

EIS Estimates with Unfiltered Consumption Data

In this section, we repeat the estimation approach of Yogo (2004), using unfiltered and reported consumption based on nondurables and services. Our results are also evaluated based on weak-identification methods. Alternative estimations are presented in the appendix, broadly reconfirming our main

findings³⁵.

Under Epstein-Zin preferences, it is possible to derive typical Euler Equations usually used in the literature to estimate the EIS. These connect consumption growth with returns of an asset class i ³⁶:

$$\Delta c_{k,t+1} = \tau_{i,t} + \psi r_{i,t+1} + \epsilon_{i,t+1}, \quad (13)$$

$$r_{i,t+1} = \zeta_{i,t} + \Theta \Delta c_{k,t+1} + \varrho_{i,t+1}, \quad (14)$$

where $\tau_{i,t}$ and $\zeta_{i,t}$ encompass mainly terms of second order (conditional variances and covariances), related to consumption growth and returns and $\epsilon_{i,t+1}$ and $\varrho_{i,t+1}$ also include expectational error terms. These are correlated with regressors in (13) and (14), so that an IV model must be adopted to properly identify their slopes. A standard log-linearisation shows that theoretically we must have $\Theta = 1/\psi$ ³⁷. Nonetheless, it is often hard to show this result when relying on IV methods, regardless of the specification. Yogo (2004) addresses this puzzle, testing whether $\hat{\psi} = 1$ and $\hat{\Theta} = 1$, when individually estimating (13) and (14). He shows a rejection of the null in the first but not in the second estimation, pointing out that the presence of weak instruments may be substantially affecting these results.

We index consumption growth with k in (13-4) to emphasise the consumption series used. In the tables that follow we consider $k \in \{\text{Reported, Unf-Hom, Unf-Het}\}$, where the last two refer to unfiltered consumption, constructed from the homoscedastic and heteroscedastic (baseline) models, respectively. The identification approach of this section does not depend on the hypothesis for heteroscedasticity nor on the asset type i ³⁸. Specifically, we conduct estimations with both stocks and risk-free returns. Lags of the nominal interest rate, inflation, consumption growth (the measure relevant in the estimation, either reported or unfiltered) and log dividend-price ratio are used instruments. One could use the real interest rate rather than the nominal and inflation as instruments, but we prefer to follow Yogo (2004) strictly to eluci-

³⁵For instance, those estimations include the consumption of nondurables only, raw data for returns while using unfiltered consumption or applying the quasi-differenced model with serially correlated measurement error for annual data as well. In this section, we remove the first three observations (aiming to exclude the filter's training period) for estimations that use quarterly data. With annual data, we opt to use the entire sample, given the limited number of observations available. In the appendix, we also present additional results when the sample is restricted to 1960:1–2017:4 for quarterly and to 1940–2017 for annual data.

³⁶We present the recursive form of L. G. Epstein and Zin (1989) preferences and their associated non-linear Euler Equations in the appendix. Structural forms of (13) and (14) can be seen in Yogo (2004).

³⁷See Yogo (2004).

³⁸If i is the risk-free sovereign bond, for example, only $\tau_{i,t}$ and $\zeta_{i,t}$ change and some of their second-order terms become null. See Yogo (2004) for more details.

date the comparison. As in that paper, we lag all instruments twice to mitigate concerns of invalid moment conditions under conditional heteroscedasticity in (13) and (14). In addition, instruments that are lagged at least twice ensure that problems involving time-aggregation in consumption do not affect estimates, as advised by R. Hall (1988).

2.3.1

Homoscedastic Framework

Here we present estimates for the EIS (ψ) and its reciprocal ($1/\psi$) using equations (13) and (14), respectively, assuming conditional homoscedasticity. We apply three K-class estimators: TSLS, Fuller-K and LIML³⁹. First-stage F-statistics to infer about the relevance of instruments are also reported⁴⁰. Critical values for them under the null hypotheses in Stock and Yogo (2002) are presented in the appendix. In general terms, F-statistics above 10 ensure that the TSLS bias is low enough to be reliable, while the Fuller-K bias is not high enough when that number is above 6. Under conventional first-order asymptotics, all those three estimators should converge to the same limit distribution, with the TSLS being the efficient one. In contrast, under weak instruments, the Fuller-K and LIML are more robust estimators⁴¹. Here we present estimations for reported and unfiltered consumption series which are constructed from the consumption of nondurables and services component, found in the NIPA tables. Additional estimations for the consumption of nondurables only are provided in the appendix. They produce similar results.

It is worth reemphasising the difference between a homoscedastic filter and conditionally homoscedastic error in the Euler Equation. The former only implies that statisticians filter the data assuming a constant volatility parameter for *state* consumption. The latter refers to conditionally homoscedastic errors in the Euler equation. The filtering process and the presence of measurement error prevent us from concluding that the homoscedastic model is the preferred choice when errors in the Euler equation are homoscedastic. The more generalised heteroscedastic Filter model is our baseline, even though we exhibit results under both settings, for completeness.

³⁹Check the appendix for a better description of those estimators.

⁴⁰If error terms are not serially correlated and homoscedastic, the first-stage F-statistic is a sample analogue of the so-called concentration parameter, that captures how relevant instruments are. When the F-statistic (and the concentration parameter) is sufficiently high, the TSLS is reliable, approximately unbiased and its t-statistic exhibits a proper convergence towards a standard normal.

⁴¹Under weak instruments, the TSLS can be severely biased, compared with the Fuller-K. Additionally, the Wald test that corresponds to the LIML estimator is less size-distorted than that of the TSLS – see Stock and Yogo (2002), Stock, Wright, and Yogo (2002) and Murray (2006) for more details.

Table 2.3 below displays results for quarterly data, where unfiltered consumption considers the quasi-differenced Filter model ($\rho_\xi = 0.06$). The first thing to note is that there is much more agreement on estimates of the EIS (ψ) across different estimators when unfiltered consumption is used. Moreover, negative point estimates are completely absent, broadly in line with economic intuition. Our point estimates for the EIS under the heteroscedastic model lie in the range 0.15-0.38, roughly in line with R. Hall (1988), Yogo (2004) and L. Epstein and Zin (1991)⁴². In addition, higher point estimates can be obtained using the homoscedastic filter, from 0.20 to 0.57. In contrast, reported consumption shows quite a different picture, with point estimates from -0.20 to 0.07. On the one hand, negative estimates are frequent with reported consumption when we use stocks, and results indicate that a lower first-stage predictability may be explaining this fact. On the other hand, unfiltered consumption produces a very narrow interval of positive point estimates across estimators, roughly from 0.16 to 0.19 under the heteroscedastic model, but with a similar first-stage F-statistic. Standard errors are also more equalised using unfiltered consumption and stocks, suggesting that the former alleviates problems related to weak instruments. However, it is likely that such issue is still partially plaguing estimates, given the low first-stage F-statistic, slightly above 4. When the risk-free is used, these statistics generally exceed critical values of Stock and Yogo (2002), indicating that weak identification is not an issue⁴³.

The lower part of Table 2.3 confirms that the same improvements observed for the EIS under unfiltered consumption are valid for its reciprocal. In fact, our estimates of $1/\psi$ return a notably wide range using reported consumption (from -5.01 to 18.98), regardless of the asset class used. The story is once again quite different for unfiltered consumption, with estimates in the narrower range 0.35-5.25 under the heteroscedastic model, implying an EIS from 0.19 to 2.86⁴⁴. Considering the more robust Fuller-K and LIML estimators, implied EIS estimates using that model and quarterly data are in the range 0.19-0.70, basically in line with direct estimates. Although exhibiting lower first-stage F-statistics, the use of the homoscedastic model returns $1/\psi$ from 0.33 to 4.25. Under Fuller-K and LIML, this implies an EIS in the range from 0.17 to 0.97. Differences between the TSLS and the other two

⁴²Our results lie in the lower end of estimates in the latter (0.17-0.87), albeit their results restricted to consumption of nondurables and services are very similar to ours.

⁴³Gomes and Paz (2013) find the same result, conducting similar estimations for the risk-free rate, but using an alternative measure of returns, there argued to better capture the portfolio of the representative agent.

⁴⁴Table B.6 (appendix) exhibits implied estimates of the EIS (ψ) from estimates of its reciprocal ($1/\psi$) using (14).

Table 2.3: Estimates of the EIS Using K-Class Estimators and Quarterly Data

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	0.067*** (0.078)	0.053*** (0.093)	0.053*** (0.093)	22.172
	ψ	Unf-Hom	0.527 (0.467)	0.566 (0.481)	0.573 (0.484)	22.335
	ψ	Unf-Het	0.346* (0.336)	0.379* (0.350)	0.385* (0.352)	22.474
Stocks	ψ	Reported	0.006*** (0.017)	-0.101*** (0.096)	-0.199*** (0.213)	4.575
	ψ	Unf-Hom	0.204*** (0.108)	0.221** (0.114)	0.235** (0.120)	4.462
	ψ	Unf-Het	0.156*** (0.080)	0.178*** (0.088)	0.191*** (0.093)	4.310
Risk Free	$\frac{1}{\psi}$	Reported	0.438* (0.311)	4.953 (4.475)	18.979 (33.660)	6.630
	$\frac{1}{\psi}$	Unf-Hom	0.331*** (0.154)	1.031 (0.683)	1.745 (1.473)	1.890
	$\frac{1}{\psi}$	Unf-Het	0.349*** (0.168)	1.434 (0.997)	2.599 (2.380)	2.268
Stock	$\frac{1}{\psi}$	Reported	0.795 (2.724)	-4.150 (4.936)	-5.014 (5.346)	6.630
	$\frac{1}{\psi}$	Unf-Hom	2.884 (1.372)	3.559 (1.739)	4.251 (2.159)	1.890
	$\frac{1}{\psi}$	Unf-Het	3.427 (1.601)	4.491 (2.137)	5.247* (2.565)	2.268

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and quarterly data. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

estimators are broadly expected since, as mentioned above, the Fuller-K and LIML estimators are more robust to weak instruments. Since first-stage F-statistics are decreased by a factor of three when unfiltered consumption is used, this justifies the gap⁴⁵. Generally, our estimates for $1/\psi$ under (14) agree with what we obtain for the EIS (ψ) using (13), considering the more robust estimators⁴⁶.

Table 2.4 below presents similar results but for annual data, for which we simply re-calibrate the original Filter model in Kroencke (2017)⁴⁷. The big picture is very similar to that of quarterly data. By using unfiltered consumption, we once again get rid of negative estimates of the EIS, obtaining point values in the limited range 0.09-0.23. Such stability is even more surprising considering the fact that first-stage F-statistics are considerably lower with unfiltered consumption and stock returns (slightly above 1), compared to reported consumption in the same situation (around 5). The former generates $\hat{\psi}$ in the narrow interval 0.18-0.22. A higher first-stage predictability does not ensure positive estimates of the EIS using reported consumption: from -0.07 to 0.11. Counter-intuitively, note that estimations with stocks imply values around -0.05, statistically significant at 1%. Our findings for the reciprocal $1/\psi$ approximately repeat those for quarterly data. Unfiltered consumption gives much more precise results, even with first-stage F-statistics that are noticeably lower than those of Table 2.3. Using stocks, these estimates imply an EIS in the very tight range 0.22-0.25 – see Table B.6 (appendix). Even with substantially low first-stage F-statistics (lower than 2), note that standard errors are quite aligned across estimators. This suggests that the three methods converge to the same limit distribution, as in the case of conventional first-order asymptotics. For the risk-free, implied EIS estimates from $1/\psi$ are from 0.09 to 0.52, when excluding the less robust TSLS estimator, not that far from results for stocks.

Generally, it seems that the connection between first-stage predictability and more precise estimates is not that relevant with unfiltered consumption. Hence, it could be that a sizeable proportion of the econometric difficulties

⁴⁵The lower first-stage F-statistic for unfiltered consumption makes sense, once Table 2.1 (appendix) shows that unfiltered consumption is not as serially correlated as its reported analogue and consumption growth is the endogenous regressor in (14).

⁴⁶That being said, we could not revert the puzzle that ψ is generally statistically different from 1 but not its reciprocal. In this regard, Table 2.3 provides unclear results, what can indicate that weak-instruments are still affecting estimates when unfiltered consumption is used, even though not as heavily as with reported consumption.

⁴⁷Recall that serially correlated measurement error is not necessary at annual frequency, so we do not use (3) and (4). Consequently, our model is no longer quasi-differenced, being exactly that of Kroencke (2017). We still present results for annual data imposing the quasi-differenced model ($\rho_\xi = 0.06 \neq 0$) in the appendix. We show that we once more can improve estimates of the EIS, albeit with somewhat weaker results.

usually attributed to weak instruments corresponds instead to weaknesses involving the consumption time series⁴⁸. In addition, point estimates of the EIS are generally more close to 1, although still not statistically higher than it. Overall, under unfiltered consumption the improvement is expressive enough both quantitatively (higher and more equalised estimates across estimators, none with the wrong sign) and qualitatively (closer to usual choices of values, applied to macro models).

The next step is to evaluate how unfiltered compares with reported consumption using robust inference. For this, we invert Moreira (2003) and T. W. Anderson, Rubin, et al. (1949) test statistics, creating 95 percent weak-identification-robust confidence intervals. Table 2.5 below summarises results. When it comes to quarterly data, the first thing to note is that unfiltered consumption effectively reverts an empty set under the Anderson-Rubin statistic. Using the risk-free, the robust confidence interval with the heteroscedastic model in this case is in line with estimates of L. Epstein and Zin (1991): from -0.07 to 0.87. Based on the S-test of Stock and Wright (2000), Ascari, Magnusson, and Mavroeidis (2016) also found an empty interval, using a baseline Euler Equation as (13) and reported consumption. The S-test is a generalisation of the Anderson-Rubin test to a GMM setting, being not only robust to weak instruments but also to heteroscedasticity of arbitrary form. Ascari, Magnusson, and Mavroeidis (2016) test several Euler Equations, derived from many different assumptions, and confidence intervals similar to ours are only obtained relying on internal habit formation. Since habit formation tends to create inertia⁴⁹, implicitly flattening the relationship between consumption and returns, our finding is pertinent to the extent that it brings a similar confidence interval to a much simpler Euler Equation, without creating doubts about how habit formation might be implicitly lowering estimates of the EIS. In addition, as emphasised in Yogo (2004), uninformative robust sets are a natural consequence of a very weak IV setting, so that once more our results suggest that unfiltered consumption significantly improves the identification of the EIS. Even though a little wider, our confidence intervals generated by the conditional likelihood ratio test tell a similar story. Using stocks broadly confirms our results with the risk-free, with the additional benefit that it produces narrower intervals and that the homoscedastic and

⁴⁸We still can not rule out that weak instruments are affecting our estimation since estimates and standard errors – even though more equalised – are still different across estimators (recall that in the absence of weak instruments, limit distributions under the three estimators should be approximately the same). This is especially the case for annual data in Table 2.4, as well as for our estimates of $1/\psi$ in both Table 2.3 and Table 2.4, for which first-stage F-statistics are essentially lower.

⁴⁹J. C. Fuhrer (2000).

Table 2.4: Estimates of the EIS Using K-Class Estimators and Annual Data

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	0.112*** (0.105)	0.109*** (0.111)	0.108*** (0.113)	11.836
	ψ	Unf-Hom	0.099*** (0.190)	0.091*** (0.205)	0.089*** (0.208)	10.727
	ψ	Unf-Het	0.097*** (0.188)	0.090*** (0.203)	0.089*** (0.205)	10.727
Stocks	ψ	Reported	-0.049*** (0.034)	-0.061*** (0.038)	-0.065*** (0.040)	5.057
	ψ	Unf-Hom	0.186*** (0.081)	0.197*** (0.088)	0.226*** (0.108)	1.307
	ψ	Unf-Het	0.184*** (0.080)	0.194*** (0.086)	0.222*** (0.105)	1.312
Risk Free	$\frac{1}{\psi}$	Reported	1.364 (0.724)	4.592 (3.364)	9.246 (9.634)	1.893
	$\frac{1}{\psi}$	Unf-Hom	0.388 (0.382)	1.883 (1.890)	11.147 (25.868)	1.826
	$\frac{1}{\psi}$	Unf-Het	0.394 (0.387)	1.916 (0.475)	11.293 (26.190)	1.822
Stock	$\frac{1}{\psi}$	Reported	-6.808** (3.885)	-11.773* (6.818)	-15.285* (9.365)	1.893
	$\frac{1}{\psi}$	Unf-Hom	4.056* (1.835)	4.161 (1.915)	4.415 (2.106)	1.826
	$\frac{1}{\psi}$	Unf-Het	4.143* (1.863)	4.241* (1.937)	4.502 (2.131)	1.822

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

heteroscedastic models return more similar results.

Table 2.5 also shows that our results for annual data are not as impressive. Unfiltered consumption does increase the upper end of intervals, but confidence sets are generally wider. This result is more evident when stock returns are used. In this case, EIS values along the whole real line are possible. It is difficult to infer the reason for this, even though the small sample size for annual data could be a possible explanation⁵⁰. Moreover, recall that completely uninformative robust intervals for the EIS are frequent in the literature – Yogo (2004), Ait-Sahalia, Parker, and Yogo (2004), Ascari, Magnusson, and Mavroeidis (2016) and Gomes and Paz (2013). Indeed, our results are not particularly surprising in this respect.

2.3.2 Heteroscedastic Framework

Recall that the econometric approach mentioned above is still valid in a heteroscedastic setting. In this framework, GMM is the efficient method. Thus, we now turn to this estimator, using (13). In addition to the conventional two-step GMM (2S-GMM), we also present estimates using the continuously updated GMM estimator (CUE-GMM) – L. P. Hansen, Heaton, and Yaron (1996). The latter is less biased, provides confidence intervals with better coverage rates and performs better under weak instruments – L. P. Hansen, Heaton, and Yaron (1996), Stock, Wright, and Yogo (2002) and W. K. Newey and R. J. Smith (2004).

With (13), we conduct estimates of the EIS using 2S-GMM and CUE-GMM, with the risk-free as our measure of returns. In this more general GMM setting, ψ can also be identified in joint estimation using both the risk-free and stocks:

$$\Delta c_{k,t+1} = \tau_{f,t} + \psi r_{f,t+1} + \epsilon_{f,t+1}, \quad \Delta c_{k,t+1} = \tau_{m,t} + \psi r_{m,t+1} + \epsilon_{m,t+1}, \quad (15)$$

where indices f and m denote risk-free and market returns, respectively⁵¹. The system estimation can improve efficiency from exploiting cross-equation

⁵⁰In the appendix, we present results when unfiltered consumption at annual frequency is generated by the Filter model with serially correlated measurement errors instead. General findings are broadly in line with those of Table 2.5, suggesting that our hypothesis for measurement error is not causing that problem. Additionally, there we also re-estimate the Euler equations while further restricting the sample (so that substantially more observations are removed for early years, related to the Filter's training period). Results also do not seem sensitive to such choice.

⁵¹Drift terms must be allowed to differ across equations (given different second-order terms in $\tau_{i,t}$ depending on the asset class i) while slopes are restricted to the same value (EIS). Check the appendix for complete specifications in (15).

correlations in expectational errors included in both innovations. Additionally, weak-instrument-robust confidence intervals constructed by inverting the K-test statistic – Kleibergen (2005) – are presented. This test is robust to weak identification, as well as to autocorrelation and heteroscedastic error terms. It is similar to the S-test – Stock and Wright (2000) – mentioned above, albeit more computationally involved. We choose the K-test against the S-test based on several factors. First, the former applies in the context of non-linear moment conditions. Second, W. K. Newey and Windmeijer (2009) show that the K-test is valid under many weak moment conditions. Third, Andrews and Stock (2005) and Kleibergen and Mavroeidis (2009) specifically recommend it against available alternatives when dealing with heteroscedasticity of arbitrary form.

Table 2.6 below summarises our estimates for the EIS. First, it generally confirms our previous findings by showing higher point values for unfiltered consumption. Second, results with reported consumption are now more in line with those with unfiltered consumption, relatively to the previous tables. This is especially true for quarterly data, where the former no longer generates negative estimates: from 0.0 to 0.2. In this case, estimates with unfiltered consumption under the heteroscedastic model, for instance, do not differ much: from 0.0 to 0.5. In addition, results with unfiltered consumption and the risk-free rate (first two columns) are broadly in line with those of Table 2.3, suggesting that the homoscedasticity assumption may be less restrictive when stocks are not considered. This is not the case for reported data, with which estimates of the EIS are higher when allowing for a heteroscedastic environment. Unfiltered consumption constructed from the homoscedastic model once more presents higher estimates of the EIS, even though at the cost of higher standard errors for quarterly data. Results for annual data are a little weaker. Comparatively, point estimates lie closer to zero, both for reported and unfiltered consumption. In spite of that, the joint estimation using the former once again reaches a negative value, significant at 10%. Finally, 95% robust intervals still return uninformative sets, suggesting that identification issues related to weak instruments are possibly more relevant at annual frequency, regardless of the homoscedasticity hypothesis for the errors in the Euler equations⁵². In contrast, we again revert completely uninformative robust sets for quarterly data: $[0.18, 9.64]$ based on the homoscedastic model and $[0.05, 8.91]$ for the heteroscedastic analogue⁵³.

⁵²Recall the previous tables and the lower first-stage F-statistics obtained for annual data.

⁵³In the appendix, we present similar results for the consumption of nondurables only. We again can revert totally uninformative sets into more plausible ones using unfiltered consumption on quarterly data, albeit those intervals still exhibit negative values: $[-0.35, 0.84]$ relying on the homoscedastic model and $[-0.62, 0.26]$ for the heteroscedastic version.

Table 2.5: Weak-IV-Robust CIs for the EIS

Quarterly Data			
Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	\emptyset	$[-0.136, 0.235]$
	Unf-Hom	$[-0.319, 1.525]$	$[-0.377, 1.591]$
	Unf-Het	$[-0.077, 0.867]$	$[-0.307, 1.122]$
Stocks	Reported	\emptyset	$(-\infty, +\infty)$
	Unf-Hom	$[-0.045, 0.921]$	$[0.014, 0.650]$
	Unf-Het	$[-0.009, 0.710]$	$[0.023, 0.536]$
Annual Data			
Risk Free	Reported	$[-0.104, 0.316]$	$[-0.131, 0.341]$
	Unf-Hom	$[-0.272, 0.442]$	$[-0.357, 0.523]$
	Unf-Het	$[-0.270, 0.438]$	$[-0.352, 0.516]$
Stock	Reported	$[-0.245, 0.019]$	$[-0.199, 0.007]$
	Unf-Hom	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Note: Weak-instrument-robust 95% confidence intervals. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Data used both for reported and unfiltered consumption refer to the consumption of nondurables and services. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$) while for annual data we use the canonical version – with no persistence for measurement errors. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

Table 2.6: Heteroscedasticity-Robust Estimates of the EIS

Quarterly Data				
Δc_k	Two-Step	CUE	SYS	95% CI
Reported	0.133	0.189**	0.001	$(-\infty, +\infty)$
	(0.082)	(0.085)	(0.000)	
Unf-Hom	0.601	0.678	0.007	$[0.182, 9.639]$
	(0.523)	(0.525)	(0.006)	
Unf-Het	0.448	0.512	0.006	$[0.053, 8.907]$
	(0.374)	(0.377)	(0.006)	
Annual Data				
Reported	0.056	0.022	-0.015*	$(-\infty, +\infty)$
	(0.088)	(0.087)	(0.008)	
Unf-Hom	0.122	0.136	0.067	$(-\infty, +\infty)$
	(0.142)	(0.142)	(0.045)	
Unf-Het	0.119	0.133	0.066	$(-\infty, +\infty)$
	(0.141)	(0.141)	(0.045)	

Note: 2S-GMM and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free rate. The third column presents estimates of the same coefficient under the joint estimation (15), where market returns are also used, while allowing for different drifts across equations. We present 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting. These are constructed by inverting the K-statistic of Kleibergen (2005). Consumption series are relative to nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

2.4

EIS, Limited Participation and Unfiltered Consumption

The last section demonstrated how important is to account for the fact that macro consumption series are heavily filtered before release, when estimating the EIS. In this section, we aim to verify whether that is again the case when dealing with other types of consumption data series. We use micro data at the level of the household to construct the consumption growth series used in our estimations. These data are extracted from the CEX, a large-scale survey, designed to represent characteristics of the entire US population. We construct measures of consumption growth following the same procedures in Vissing-Jorgensen (2002), albeit with considerably more observations – 1982-2013 here, compared to 1982-1996 in that paper.

The CEX data is subject to a number of statistical transformations, many of which similar to those applied to NIPA consumption series. To cite a few: topcoding; suppression; reallocation, and; imputation procedures. These are all present in the survey. Further, measurement error is so evident in the data that back in 2009 the Bureau of Labor Statistics (BLS) approved the so-called Gemini project to research and develop a complete redesign of the CEX (to be implemented in 2022), addressing measurement error and respondent burden issues. Mechanically, statistical procedures in the CEX produce the same effect on final reported data, lowering its variance and diminishing correlations with different measures of return. We detail the procedures used in the CEX in the appendix, also comparing with those performed on NIPA consumption.

An advantage of using CEX data is that we can separate households – and their corresponding measures of consumption growth – based on their asset-holding status. This is relevant to test the LAMP, since it is possible to obtain different estimates of the EIS conditional on different groups of households and their participation in specific asset markets. As noted by Yogo (2004), it is possible that weak instruments are not explaining the whole story of troublesome estimates of the EIS, usually obtained in the literature. Perhaps limited participation in asset markets is plaguing results due to Euler equations that do not hold for the representative agent, as addressed in Vissing-Jorgensen (2002) and Guvenen (2006). Indeed, estimations that rely on an Euler Equation for some asset, but that use households that do not hold a position in that asset, will likely bias estimates of the EIS downwards⁵⁴.

First, let us turn to how we use data from the survey. The CEX interviews more than 7500 households per quarter⁵⁵. Each household is interviewed five

⁵⁴See Vissing-Jorgensen (1998).

⁵⁵Before 2000, that number was slightly lower, around 5000 per quarter. The programme

times, but only the last four interviews are publicly available. Interviews with the same household occur every three months, when they report consumption for the previous three months. In the last (fifth) interview, households report their financial information. We use this information to separate households according to their asset-holding status. Formally, households report holdings for “stocks, bonds, mutual funds and other securities”, “US savings bonds”, “savings accounts” and “checking accounts, brokerage accounts, and other similar accounts”. We use responses for the first two categories to label households as stock (vs. non-stock) or bonds (vs. non-bond) holders.

We classify households by their asset holding status adopting the same criteria in Vissing-Jorgensen (2002). As noted in that paper, it is not possible to perfectly separate households solely by those two categories. There is some overlap between bond and stock holders, but not between asset and non-asset holders for each type of asset (bonds or stocks). Note that imperfect separation should bias *against* finding different estimates of the EIS across these groups. Formally, we refer to households with positive responses to “stocks, bonds, mutual funds and other securities” as stock holders, and to those with positive responses to that same category *or* to “US savings bonds” as bond holders.

Additionally, note that an Euler equation that measures consumption from t to $t + 1$ should hold at the beginning of first period. Therefore, we shall split households across groups based on their holdings at the beginning of t . To do so, we use two more items in the survey: one question that asks whether households have more, less or the same amount of the asset, relatively to a year ago; and another one which asks the estimated dollar difference in market value of that asset last month, compared to a year ago last month. As in Vissing-Jorgensen (2002), a household is classified as holding asset class i if it: (i) reports the same amount compared to a year ago, holding a positive position in i when interviewed for the last time; (ii) reports lower holdings of the asset, relative to a year ago, or; (iii) reports an increase in its holdings of the asset, but the dollar difference is less than the current value of holdings. Because some of those questions we use to separate households are no longer available after March 2013, this is the last month of consumption observations in our data set.

Our final sample consists of 90,080 households, spread over the period from 1982 to 2013⁵⁶. Amongst these households, 19.6% are classified as

contains two components, the Interview Survey and the Diary Survey. Each has its own sample. We compile our data set using the former.

⁵⁶The CEX data is available beginning in 1980. However, we follow Vissing-Jorgensen (2002) in dropping observations for 1980 and 1981. She argues that the quality of the CEX consumption data is considerably lower for that period.

stock holders and 29.1% as bond holders⁵⁷. On average, our sample has 246 households each month, of which 70 are bond holders and 49 are stock holders⁵⁸. See Table 2.2.

Our final data set encompasses semiannual consumption growth rates at monthly frequency. To construct consumption growth observations, CEX expenditure categories are carefully aggregated as to mimic definitions of the NIPA consumption of nondurables and services. We exclude three major categories: health care, education costs and housing expenses (except for housing operations). Cash contributions, personal insurance and pensions are also dropped. Categories in the first group are excluded because they exhibit a substantial durable component. In the second, for the same reason or because their definitions are considered out of scope relatively to NIPA consumption – see Garner et al. (2003, p. 12). Major categories in our consumption measure are food (at and away from home), beverages, apparel, tobacco, public and private transportation (including gasoline), personal care services, housing operations, miscellaneous and utilities. Our definition is broadly in line with that given in Attanasio and Weber (1995), being also similar to the one in Vissing-Jorgensen (2002).

For each household h , its consumption growth rate is:

$$\frac{C_{h,m+6} + C_{h,m+7} + C_{h,m+8} + C_{h,m+9} + C_{h,m+10} + C_{h,m+11}}{C_{h,m} + C_{h,m+1} + C_{h,m+2} + C_{h,m+3} + C_{h,m+4} + C_{h,m+5}}.$$

As in Vissing-Jorgensen (2002), the aggregate consumption growth observation is the average of this ratio in the cross section of households of the same group. Since consumption growth is semiannual, groups are classified based on their holdings at the beginning of the relevant period in the Euler equation – i.e., m . We refer to Vissing-Jorgensen (2002) for a formal discussion on how averaging households in the cross section of groups can generate consistent estimates of the EIS in this framework. In addition, closely following that paper, we drop: (i) extreme outliers (observations for which the consumption growth ratio is higher than 5 or less than 0.2); (ii) households that report a change in the age of the household head between two subsequent interviews different from zero or one; (iii) households living in student housing, and; (iv) non-urban households. To construct the semiannual consumption growth ratio above, we need consumption data for all interviews, 2 to 5. Therefore, we also

⁵⁷In addition, 2154 households report an increase in holdings of some asset, but not the current value. We classify them as asset holders for the corresponding category – 1593 as stock holders and 561 as bond holders. In addition, a few households report an increase in holdings that exceeds their response for current values. We consider them as non-asset holders in the corresponding category.

⁵⁸Hence, 197 are non-stock holders and 176 are non-bond holders.

drop households for which any of these interviews are missing. As our last step, we deflate nominal consumption growth observations by the urban CPI for nondurable goods.

We later use the Filter model on the consumption growth series that corresponds to that final sample of 90,080 households. There may be reasons to be sceptical about this, arguing that statistical procedures applied on CEX data probably take those households dropped from our sample into account. Nonetheless, very similar results are found when estimations in this section are repeated applying and calibrating the model for the complete sample (without dropping households), but imposing the calibration to our final data set (which excludes them). Results are also maintained when we calibrate the model *and* estimate the EIS based on the complete sample. Since consumption growth is semiannual but at the monthly frequency, there is an overlap of five months between observations. To calibrate the model, we equalise the scale of consumption growth to its frequency (as in the previous estimations), to then transform it back into semiannual. For a more complete discussion on how we apply the Filter model on the CEX data and different groups of asset holders, check the appendix.

Two types of returns are used when estimating our Euler equations in this section. When differentiating between stock holders vs. non-stock holders, we use the value-weighted return from NYSE, NASDAQ and AMEX. When applying to bond vs. non-bond holders, we use T-bill returns. An important issue relates to how we compute the relevant asset return used in the estimations when the consumption growth data is semiannual. We follow Vissing-Jorgensen (2002), using the middle six months from $(1 + R_m)$ to $(1 + R_{m+10})$: $(1 + R_{m+2})(1 + R_{m+3})\dots(1 + R_{m+6})(1 + R_{m+7})$. In addition, since $C_{h,m}$ is relevant in the consumption growth measure, it follows naturally that lagged instruments are constructed based on $(1 + R_{m-1})(1 + R_{m-2})\dots(1 + R_{m-5})(1 + R_{m-6})$, also using six months. When we estimate the Euler equations with unfiltered CEX consumption, we conduct similar adjustments as those of the last section on these return series. We better detail them in the appendix.

We use three instrument sets for the log stock return or the log T-bill return in the Euler equations: (i) dividend-price ratio; (ii) dividend-price ratio, lagged stock returns and lagged T-bill returns, and; (iii) dividend-price ratio, lagged corporate bond default premium and lagged government bond horizon premium. The last two take the form $\frac{1+R_t^{\text{long-term corporate bonds}}}{1+R_t^{\text{long-term government bonds}}}$ and $\frac{1+R_t^{\text{long-term government bonds}}}{1+R_t^{\text{short-term government bonds}}}$, respectively. We once more refer to the appendix for more details on these variables. Lastly, returns are deflated by using the urban CPI for total consumption.

As in the last section, we estimate log-linearised Euler equations. For stock holders, for example, the econometric approach we follow is:

$$\begin{aligned} \frac{1}{H_t^s} \sum_{h=1}^{H_t^s} \Delta \ln C_{t+1}^{h,s} = \psi^s \ln(1 + R_{s,t}) + \frac{1}{H_t^s} \sum_{h=1}^{H_t^s} \Delta \ln(\text{family size})_{t+1}^{h,s} \\ + \alpha_1^s D_2 + \dots + \alpha_{12}^s D_{12} + u_{t+1}^s, \end{aligned} \quad (16)$$

where, as before, ψ^s is the EIS for stock holders, D_m are seasonal dummies and H_t^s denotes the number of consumption growth observations for stock holders at time t . Euler equation (16) assumes that seasonality and the family size are multiplicative factors in the utility function⁵⁹. These two variables are included in all the three aforementioned instrument sets⁶⁰. Equation (16) holds under the Epstein-Zin framework of last section, as well as under CRRA preferences⁶¹.

2.4.1 Results

Results for two samples are presented. The first encompasses data from 1982 to 1996, the same used in Vissing-Jorgensen (2002). The second uses all the available data, from 1982 to 2013. We do so due to a change in methodology and in the sampling frame around 1996, when the CEX was redesigned⁶².

We begin with results for the period 1982-1996, exhibited in Table 2.7⁶³. First, note that *all* estimates with reported consumption are negative, many

⁵⁹Formally, the family size variable is defined as the change in the log average family size for the last two interviews (4 and 5), compared to the first two interviews (2 and 3).

⁶⁰Therefore, it is assumed that family size controls and seasonality factors are exogenous in our estimations.

⁶¹Regardless of the assumption for the utility function, generally α_m^s involves conditional variances (covariances) of (between) log consumption growth and log returns. In the case of Epstein-Zin preferences, there are also conditional second-order terms relative to wealth returns, based on the total portfolio of households. If some of those conditional second-order terms are not constant, stochastic terms enter u_{t+1}^s – which already included expectational and measurement errors (present in the consumption data). These stochastic terms do not imply inconsistent estimates, as long as they are uncorrelated with instruments used. See Vissing-Jorgensen (1998) and Vissing-Jorgensen (2002) for a formal treatment. Lastly, note that (16) can be estimated by instrumental variables methods even when using unfiltered consumption. Since returns are assumed uncorrelated with the measurement error, autocorrelation in the latter does not invalidate lags of the former as instruments.

⁶²However, we calibrate the model based on the entire sample (1982-2013). We need enough observations to calculate the long-run standard deviations of the series. Since these use a horizon of 6 years, calibrating based on the period 1982-1996 gives weaker results.

⁶³The use of semiannual consumption growth data at monthly frequency generates overlapping observations for two subsequent months of data. It follows that an MA(5) process enters the error term in (16). Therefore, we use Two-Step GMM with a heteroscedasticity and autocorrelation-consistent (HAC) estimator for the covariance matrix. Main results of this section do not change when using CUE-GMM, which is more robust to the presence of weak-instruments.

of them statistically significant. In contrast, unfiltered consumption reverts these into more sensible estimates, when considering the Euler Equations for stock and bond holders. It suggests an EIS from 0 to 0.3 for the former, and from 0.4 to 1 for the latter group. Differences between unfiltered and reported consumption are less substantial when considering non-asset holders, but the former still provides estimates of the EIS that are slightly less negative. This is also generally the case when the Euler equation for all households is estimated. Importantly, robust intervals lean towards positive values with unfiltered consumption. These are also substantially narrower when estimating the Euler equation for stock holders, suggesting that weak instruments affect these estimations to a lesser extent. For bond holders, results are more uncertain. Although point estimates seem more precisely estimated with unfiltered consumption, robust intervals are considerably wider.

Unfiltered consumption seems to magnify differences in terms of the EIS between asset and non-asset holders. Such result is consistent with Table 2.2, which shows, for instance, that unfiltering the CEX data introduces more mean reversion and considerably more volatility in the consumption series of bond and stock holders, compared to non-bond and non-stock holders. Robust sets with unfiltered consumption suggest that the EIS is not above 0.6 for stock holders, nor above 0.2 for non-stock holders. Nonetheless, there is considerable overlap between the intervals. Therefore, there seems to be little evidence favouring the limited asset market participation theory in our results. This finding contrasts with Vissing-Jorgensen (2002), who finds substantial differences across those groups for the same period. She does not apply weak-IV-robust methods, though⁶⁴.

Table 2.8 below presents results for the entire sample, from 1982 to 2013. Again, reported consumption produces negatives estimates of the EIS for *all* groups and instrument sets applied. In contrast, estimations with unfiltered consumption return positive estimates in *all but one* of the cases tested. For all households, the EIS is estimated from 0.05 to 0.1 using stock returns. These values shift to 0.4 to 1.2 with risk-free returns. Note that reported consumption provides counter-intuitive estimates in those cases, from -0.1 to -1.3 , depending on the type of return and the instrument set used.

In contrast to Table 2.7, using unfiltered consumption for the entire sample generates estimates of the EIS that are quite alike, comparing asset and non-asset holders. For stock holders, for example, we estimate a coefficient in the narrow interval from 0 to 0.1, not distant from estimates for non-stock

⁶⁴Although she does not apply robust methods, our point estimates are considerably distinct from those reported in Vissing-Jorgensen (2002) (for the same period). The data have been revised several times since then, so that these revisions may explain the differences.

Table 2.7: Estimates of the EIS – CEX Data: 1982 to 1996

		A. Estimation with Stocks		
Δc_k	Households	Instrument Set		
		I	II	III
Reported	All	−0.245* (0.145) [−0.894, −0.101]	−0.243** (0.122) [−1.048, −0.087]	−0.425*** (0.181) [−1.045, −0.248]
	Stock Holders	−0.007 (0.197) [−0.703, 0.221]	−0.224 (0.172) [−0.943, 0.033]	−0.341 (0.228) [−1.496, −0.085]
	Non-Stock Holders	−0.295* (0.155) [−0.760, −0.130]	−0.287** (0.131) [−0.807, −0.112]	−0.513** (0.203) [−0.816, −0.349]
	All	−0.066 (0.418) [−0.286, 0.050]	−0.189 (0.226) [−0.454, −0.102]	−0.135 (0.189) [−0.345, −0.051]
	Stock Holders	0.151 (0.609) [−0.038, 0.311]	0.068 (0.349) [−0.338, 0.290]	0.323 (0.363) [0.033, 0.663]
	Non-Stock Holders	−0.091 (0.501) [−0.368, 0.195]	−0.228 (0.274) [−0.426, 0.235]	−0.222 (0.222) [−0.391, 0.100]
		B. Estimation with Treasury Bills		
Reported	All	−1.065** (0.505) [−3.044, 0.821]	−0.933* (0.503) [−2.547, 1.573]	−1.680*** (0.491) [−4.141, 0.667]
	Bond Holders	−0.208 (0.690) [−2.969, 3.481]	−0.278 (0.693) [−2.867, 5.012]	−1.064 (0.676) [−6.079, 3.753]
	Non-Bond Holders	−1.293** (0.510) [−3.332, 0.604]	−1.219 (0.496) [−3.504, 1.523]	−1.966*** (0.471) [−4.947, 0.322]
	All	−0.328 (2.076) [−4.930, 4.314]	−1.017 (1.714) [−14.154, 12.089]	−0.855 (1.091) [−16.033, 8.073]
	Bond Holders	1.070 (2.444) [−3.719, 8.085]	0.417 (2.228) [−19.829, 36.569]	0.765 (1.745) [−11.043, 16.187]
	Non-Bond Holders	−0.590 (2.600) [−5.682, 5.266]	−1.244 (2.199) [−40.931, 32.163]	−1.404 (1.377) [−28.711, 8.518]

Notes: Estimates of the EIS using Euler equation (16). The sample encompasses semi-annual consumption growth observations at monthly frequency, from 1982 to 1996. Unfiltered consumption is extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated. Here we assume that government statisticians filter the data based on our final sample – in which some households are dropped based on conditions described in the main text. Reported uses official CEX data. Unfiltered consumption growth is constructed from the heteroscedastic model. For this case, asset returns are adjusted for time-aggregation issues for any group of asset holders – see appendix. Instrument set I includes the log dividend-price ratio. Set II adds the lagged log real value-weighted return (from NYSE, NASDAQ and AMEX) and the lagged log real T-bill return. Set III replaces the last two by the lagged bond default premium and the lagged bond horizon premium. All these sets include the family size and seasonal controls as instruments (so that these are assumed exogenous). Standard errors are presented in parentheses. 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting are shown in brackets. We construct these intervals by inverting the K-test statistic in Kleibergen (2005). The null that the estimated coefficient equals 0 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

holders, around 0.15. Additionally, robust intervals are also similar. As in the previous table, these sets generally shift from showing negative to showing positive numbers, as we replace reported with unfiltered consumption. Once more, there seems to be more uncertainty involving estimations for bond and non-bond holders, as, for instance, robust intervals are again considerably wider. In spite of that, higher point estimates are still obtained with unfiltered consumption. For instrument sets II and III, it produces estimates around 0.4, the lower bound of results with analogous estimations in the previous table.

Generally, estimations in Table 2.8 once more produce limited evidence that the EIS differs substantially between asset and non-asset holders. This is the case even with unfiltered consumption. Recall that, in the previous section, we concluded that unfiltered consumption offered more reliable estimations, which also seemed less plagued by the presence of weak instruments. We then mentioned the possibility that commonly distorted estimates of the EIS may be a consequence of an Euler Equation that does not hold for the representative agent because of limited participation in asset markets. First, our estimations with the CEX data indicate that it is not the case. Second, even if it were, accounting for the fact that reported data are statistically treated before release can once again produce more precise estimates of the EIS.

2.5 Conclusion

The empirical evidence on the EIS is vast and several papers use different techniques attempting to provide better and more precise estimates of that parameter. Results seem heavily influenced by a number of factors as specification, econometric method, region and participation in asset markets. Nonetheless, there seems to be no academic research that takes into account the problematic consumption series used in the estimation yet.

The motivation of this paper is to fill this gap by conducting estimates of the EIS considering the fact that reported consumption series are filtered before release. These statistical procedures can be optimal to generate official data, but can also lead to undesirable consequences in terms of research. To unfilter reported consumption, we propose a modification to the so-called Filter model in Kroenke (2017), introducing serially correlated measurement error. With this adaptation, we can estimate the EIS with a number of types of data at different frequencies. First, we construct unfiltered consumption series for macro data (NIPA consumption), conducting estimations in the econometric framework of Yogo (2004). Second, we apply the model to disaggregated survey data (CEX), estimating the EIS based on the consumption of different

Table 2.8: Estimates of the EIS – CEX Data: 1982 to 2013

A. Estimation with Stocks				
Δc_k	Households	Instrument Set		
		I	II	III
Reported	All	−0.163** (0.065) [−0.246, −0.098]	−0.108* (0.060) [−0.158, −0.003]	−0.161*** (0.056) [−0.249, −0.094]
	Stock Holders	−0.107 (0.103) [−0.350, 0.070]	−0.117 (0.103) [−0.393, −0.039]	−0.175** (0.078) [−0.390, −0.045]
	Non-Stock Holders	−0.161*** (0.063) [−0.305, −0.084]	−0.101** (0.056) [−0.201, 0.048]	−0.149*** (0.056) [−0.292, −0.069]
	All	0.116 (0.241) [−0.022, 0.284]	0.044 (0.205) [−0.112, 0.186]	0.065 (0.210) [−0.094, 0.218]
	Stock Holders	0.000 (0.330) [−0.145, 0.145]	0.014 (0.294) [−0.128, 0.188]	0.035 (0.325) [−0.107, 0.196]
	Non-Stock Holders	0.173 (0.263) [0.029, 0.369]	0.107 (0.223) [−0.025, 0.269]	0.143 (0.239) [0.004, 0.331]
B. Estimation with Treasury Bills				
Reported	All	−1.340** (0.519) [−7.884, 2.655]	−0.389 (0.352) [−2.804, 7.158]	−0.578* (0.324) [−2.594, 7.302]
	Bond Holders	−0.916 (0.646) [−8.532, 9.719]	−0.349 (0.475) [−8.391, 7.065]	−0.532 (0.408) [−7.732, 7.839]
	Non-Bond Holders	−1.362*** (0.496) [−9.163, 2.868]	−0.450 (0.335) [−6.837, 4.205]	−0.687** (0.319) [−3.924, 5.542]
	All	1.181 (2.438) [−0.514, 9.875]	0.348 (1.287) [−2.057, 13.297]	0.418 (0.709) [−0.628, 9.138]
	Bond Holders	−0.337 (2.759) [−6.874, 45.645]	0.318 (1.570) [−39.949, 41.042]	0.423 (1.083) [−28.484, 47.735]
	Non-Bond Holders	2.154 (2.774) [−0.072, 31.714]	0.967 (1.479) [−1.769, 31.042]	0.566 (0.832) [−0.658, 26.812]

Notes: Estimates of the EIS using Euler equation (16). Our sample encompasses semi-annual consumption growth observations at monthly frequency, from 1982 to 2013. Unfiltered consumption is extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated. Here we assume that government statisticians filter the data based on our final sample – in which some households are dropped based on conditions described in the main text. Reported uses official CEX data. Unfiltered consumption growth is constructed from the heteroscedastic model. For this case, asset returns are adjusted for time-aggregation issues for any group of asset holders – see appendix. Instrument set I includes the log dividend-price ratio. Set II adds the lagged log real value-weighted return (from NYSE, NASDAQ and AMEX) and the lagged log real T-bill return. Set III replaces the last two by the lagged bond default premium and the lagged bond horizon premium. All these sets include the family size and seasonal controls as instruments (so that these are assumed exogenous). Standard errors are presented in parentheses. 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting are shown in brackets. We construct these intervals by inverting the K-test statistic in Kleibergen (2005). The null that the estimated coefficient equals 0 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

groups, according to their asset-holding status. With this separation, we revisit the LAMP using unfiltered consumption data. We also evaluate our econometric results for both the macro and the micro frameworks based on weak-identification-robust routines.

The main lesson from our paper is that unfiltered consumption can be substantially important when estimating the EIS. It provides considerably more stable estimates of that parameter across different econometric methods and settings. Furthermore, results based on macro consumption data seem less sensitive to the presence of weak instruments when we unfilter those series. In that case, we also show that unfiltered consumption can revert uninformative weak-IV-robust confidence intervals, which are quite frequent in the EIS empirical literature. Using micro data, improvements with unfiltered consumption are also substantial. Point estimates for the EIS are higher, positive, and corresponding robust intervals are commonly tilted to positive numbers. In contrast, reported micro data often produces negative values. Finally, our main findings for macro data suggest that the EIS is unlikely to be above 0.5. This is in line with a reasonable part of the literature, even though our point estimates with unfiltered consumption are higher in comparison. For micro data, our results imply an EIS from 0 to 0.3, based on the consumption of stock holders, and from 0.4 to 1, for bond holders.

References

- Ait-Sahalia, Yacine, Jonathan A Parker, and Motohiro Yogo (2004). “Luxury goods and the equity premium.” In: *The Journal of Finance* 59.6, pp. 2959–3004.
- Anderson, Evan et al. (1996). “Mechanics of forming and estimating dynamic linear economies.” In: *Handbook of Computational Economics*. Ed. by H. M. Amman, D. A. Kendrick, and J. Rust. 1st ed. Vol. 1. Elsevier. Chap. 04, pp. 171–252. URL: <https://EconPapers.repec.org/RePEc:eee:hecchp:1-04>.
- Anderson, Theodore W, Herman Rubin, et al. (1949). “Estimation of the parameters of a single equation in a complete system of stochastic equations.” In: *The Annals of Mathematical Statistics* 20.1, pp. 46–63.
- Andrews, Donald and James H Stock (2005). *Inference with weak instruments*.
- Aoki, Kosuke (2001). “Optimal monetary policy responses to relative-price changes.” In: *Journal of monetary economics* 48.1, pp. 55–80.
- Ascari, Guido, Leandro Magnusson, and Sophocles Mavroeidis (2016). “Empirical Evidence on the Euler equation for Consumption and Output in the US.” In: URL: https://editorialexpress.com/cgi-bin/conference/download.cgi?db_name=IAAE2016&paper_id=651.
- Attanasio, Orazio P and Guglielmo Weber (1995). “Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey.” In: *Journal of political Economy* 103.6, pp. 1121–1157.
- Ball, Laurence and David Romer (1990). “Real rigidities and the non-neutrality of money.” In: *The Review of Economic Studies* 57.2, pp. 183–203.
- Bansal, Ravi, Dana Kiku, et al. (2014). “Volatility, the macroeconomy, and asset prices.” In: *The Journal of Finance* 69.6, pp. 2471–2511.
- Bansal, Ravi and Ivan Shaliastovich (2013). “A long-run risks explanation of predictability puzzles in bond and currency markets.” In: *The Review of Financial Studies* 26.1, pp. 1–33.
- Bansal, Ravi and Amir Yaron (2004). “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles.” In: *Journal of Finance* 59.4, pp. 1481–1509. URL: <https://ideas.repec.org/a/bla/jfinan/v59y2004i4p1481-1509.html>.

- Bårdsen, Gunnar, Eilev S Jansen, and Ragnar Nymoen (2004). “Econometric evaluation of the New Keynesian Phillips curve.” In: *Oxford Bulletin of Economics and Statistics* 66, pp. 671–686.
- BEA (2017). *NIPA Handbook: Concepts and Methods of the U.S. National Income and Product Accounts*. U.S. Bureau of Economic Analysis.
- Bell, William R and David W Wilcox (1993). “The effect of sampling error on the time series behavior of consumption data.” In: *Journal of Econometrics* 55.1-2, pp. 235–265.
- Berriel, Tiago, Marcelo Medeiros, and Marcelo J. Sena (2016). “Instrument selection for estimation of a forward-looking Phillips Curve.” In: *Economics Letters* 145.C, pp. 123–125. URL: <https://EconPapers.repec.org/RePEc:eee:ecolet:v:145:y:2016:i:c:p:123-125>.
- Bils, Mark and Pete Klenow (2004). “Some Evidence on the Importance of Sticky Prices.” In: *Journal of Political Economy* 112.5, pp. 947–985. URL: <https://EconPapers.repec.org/RePEc:ucp:jpollec:v:112:y:2004:i:5:p:947-985>.
- Blinder, Alan et al. (1998). *Asking about prices: a new approach to understanding price stickiness*. Russell Sage Foundation.
- Bollerslev, Tim, George Tauchen, and Hao Zhou (2009). “Expected stock returns and variance risk premia.” In: *The Review of Financial Studies* 22.11, pp. 4463–4492.
- Breeden, Douglas T, Michael R Gibbons, and Robert H Litzenberger (1989). “Empirical tests of the consumption-oriented CAPM.” In: *The Journal of Finance* 44.2, pp. 231–262.
- Cagliarini, Adam, Tim Robinson, and Allen Tran (2011). “Reconciling microeconomic and macroeconomic estimates of price stickiness.” In: *Journal of Macroeconomics* 33.1, pp. 102–120. URL: <https://EconPapers.repec.org/RePEc:eee:jmacro:v:33:y:2011:i:1:p:102-120>.
- Calvo, Guillermo A. (1983). “Staggered prices in a utility-maximizing framework.” In: *Journal of Monetary Economics* 12.3, pp. 383–398. URL: <https://ideas.repec.org/a/eee/moneco/v12y1983i3p383-398.html>.
- Campbell, John Y (2003). “Consumption-based asset pricing.” In: *Handbook of the Economics of Finance* 1, pp. 803–887.
- Campbell, John Y, Luis M Viceira, Luis M Viceira, et al. (2002). *Strategic asset allocation: portfolio choice for long-term investors*. Clarendon Lectures in Economic.
- Carvalho, Carlos (2006). “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks.” In: *The B.E. Journal of Macroeconomics*

- 6.3, pp. 1–58. URL: <https://ideas.repec.org/a/bpj/bejmac/vfrontiers.2y2006i1n1.html>.
- Carvalho, Carlos and Jae Won Lee (2011). *Sectoral Price Facts in a Sticky-Price Model*. Departmental Working Papers 201133. Rutgers University, Department of Economics. URL: <https://ideas.repec.org/p/rut/rutres/201133.html>.
- Carvalho, Carlos and Fernanda Nechio (2016). “Factor specificity and real rigidities.” In: *Review of Economic Dynamics* 22, pp. 208–222.
- Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans (2005). “Nominal rigidities and the dynamic effects of a shock to monetary policy.” In: *Journal of political Economy* 113.1, pp. 1–45.
- Clarida, Richard, Jordi Gali, and Mark Gertler (2000). “Monetary policy rules and macroeconomic stability: evidence and some theory.” In: *The Quarterly journal of economics* 115.1, pp. 147–180.
- Cochrane, John H (1996). “A cross-sectional test of an investment-based asset pricing model.” In: *Journal of Political Economy* 104.3, pp. 572–621.
- Daniel, Kent and David Marshall (1996). *The equity premium puzzle and the risk-free rate puzzle at long horizons*. Working Paper Series, Issues in Financial Regulation WP-96-4. Federal Reserve Bank of Chicago. URL: <https://ideas.repec.org/p/fip/fedhfi/wp-96-4.html>.
- Del Negro, Marco and Frank Schorfheide (2008). “Forming priors for DSGE models (and how it affects the assessment of nominal rigidities).” In: *Journal of Monetary Economics* 55.7, pp. 1191–1208.
- Epstein, Larry G and Stanley E Zin (1989). “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework.” In: *Econometrica* 57.4, pp. 937–969. URL: <https://ideas.repec.org/a/ecm/emetrp/v57y1989i4p937-69.html>.
- Epstein, Lawrence and Stanley E Zin (1991). “Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis.” In: *Journal of political Economy* 99.2, pp. 263–286.
- Fuhrer, Jeff and George Moore (1995). “Inflation Persistence.” In: *The Quarterly Journal of Economics* 110.1, pp. 127–159. URL: <https://ideas.repec.org/a/oup/qjecon/v110y1995i1p127-159..html>.
- Fuhrer, Jeffrey C. (2000). “Habit formation in consumption and its implications for monetary-policy models.” In: *American Economic Review* 90.3, pp. 367–390.
- Fuhrer, Jeffrey C. and Glenn D. Rudebusch (2002). *Estimating the Euler equation for output*. Working Papers 02-3. Federal Reserve Bank of Boston. URL: <https://ideas.repec.org/p/fip/fedbwp/02-3.html>.

- Gali, Jordi and Luca Gambetti (2019). *Has the US wage phillips curve flattened? A semi-structural exploration*. Tech. rep. National Bureau of Economic Research.
- Gali, Jordi and Mark Gertler (1999). “Inflation dynamics: A structural econometric analysis.” In: *Journal of Monetary Economics* 44.2, pp. 195–222. URL: <https://ideas.repec.org/a/eee/moneco/v44y1999i2p195-222.html>.
- Gali, Jordi, Mark Gertler, and J. David López-Salido (2003). *Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve*. Working Papers 44. Barcelona Graduate School of Economics. URL: <https://ideas.repec.org/p/bge/wpaper/44.html>.
- Garner, Thesia I et al. (2003). “The Consumer Expenditure Survey in comparison: focus on personal consumption expenditures.” In: *Federal Economic Statistics Advisory Committee Meeting, March*. Vol. 17.
- Gomes, Fábio Augusto Reis and Lourenço S Paz (2013). “Estimating the elasticity of intertemporal substitution: Is the aggregate financial return free from the weak instrument problem?” In: *Journal of Macroeconomics* 36, pp. 63–75.
- Güvenen, Fatih (2006). “Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective.” In: *Journal of Monetary Economics* 53.7, pp. 1451–1472.
- Gwin, Carl R and David D VanHoose (2008). “Alternative measures of marginal cost and inflation in estimations of new Keynesian inflation dynamics.” In: *Journal of Macroeconomics* 30.3, pp. 928–940.
- Hale Shapiro, Adam (2008). “Estimating the new Keynesian Phillips curve: a vertical production chain approach.” In: *Journal of Money, Credit and Banking* 40.4, pp. 627–666.
- Hall, Robert (1988). “Intertemporal Substitution in Consumption.” In: *Journal of Political Economy* 96, pp. 339–57. DOI: 10.1086/261539.
- Hall, Robert E (1988). “Intertemporal substitution in consumption.” In: *Journal of political economy* 96.2, pp. 339–357.
- Hamilton, James D (2018). “Why you should never use the Hodrick-Prescott filter.” In: *Review of Economics and Statistics* 100.5, pp. 831–843.
- Hamilton, James Douglas (1994). *Time series analysis*. Vol. 2. Princeton university press Princeton, NJ.
- Hansen, Christian, Jerry Hausman, and Whitney Newey (2008). “Estimation with many instrumental variables.” In: *Journal of Business & Economic Statistics* 26.4, pp. 398–422.

- Hansen, Lars Peter, John Heaton, and Amir Yaron (1996). “Finite-sample properties of some alternative GMM estimators.” In: *Journal of Business & Economic Statistics* 14.3, pp. 262–280.
- Harvey, Andrew, Esther Ruiz, and Enrique Sentana (1992). “Unobserved Component Time Series Models With ARCH Disturbances.” In: *Journal of Econometrics* 52, pp. 129–157. DOI: 10.1016/0304-4076(92)90068-3.
- Havranek, Tomas et al. (2013). *Cross-Country Heterogeneity in Intertemporal Substitution*. Working Papers IES 2013/11. Charles University Prague, Faculty of Social Sciences, Institute of Economic Studies. URL: https://EconPapers.repec.org/RePEc:fau:wpaper:wp2013_11.
- Hooper, Peter, Frederic S Mishkin, and Amir Sufi (2019). *Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is It Just Hibernating?* Tech. rep. National Bureau of Economic Research.
- Imbs, Jean M, Eric Jondeau, and Florian Pelgrin (2007). “Aggregating Phillips Curves.” In: *Swiss Finance Institute Research Paper* 07-06.
- Imbs, Jean, Eric Jondeau, and Florian Pelgrin (2011). “Sectoral Phillips curves and the aggregate Phillips curve.” In: *Journal of Monetary Economics* 58.4, pp. 328–344. DOI: 10.1016/j.jmoneco.2011.07. URL: <https://ideas.repec.org/a/eee/moneco/v58y2011i4p328-344.html>.
- Jagannathan, Ravi and Yong Wang (2007). “Lazy investors, discretionary consumption, and the cross-section of stock returns.” In: *The Journal of Finance* 62.4, pp. 1623–1661.
- Jorda, Oscar and Fernanda Nechio (2018). “Inflation Globally.” In: Federal Reserve Bank of San Francisco.
- Kiley, Michael T (2015). “An evaluation of the inflationary pressure associated with short-and long-term unemployment.” In: *Economics Letters* 137, pp. 5–9.
- Kimball, Miles S (1995). *The quantitative analytics of the basic neomonetarist model*. Tech. rep. National Bureau of Economic Research.
- Kleibergen, Frank (2005). “Testing parameters in GMM without assuming that they are identified.” In: *Econometrica* 73.4, pp. 1103–1123.
- Kleibergen, Frank and Sophocles Mavroeidis (2009). “Weak instrument robust tests in GMM and the new Keynesian Phillips curve.” In: *Journal of Business & Economic Statistics* 27.3, pp. 293–311.
- Klenow, Peter J and Oleksiy Kryvtsov (2008). “State-dependent or time-dependent pricing: Does it matter for recent US inflation?” In: *The Quarterly Journal of Economics* 123.3, pp. 863–904.

- Klenow, Peter J and Jonathan L Willis (2006). “Real rigidities and nominal price changes.” In: *Available at SSRN 896325*.
- Kroencke, Tim A. (2017). “Asset Pricing without Garbage.” In: *Journal of Finance* 72.1, pp. 47–98. URL: <https://ideas.repec.org/a/bla/jfinan/v72y2017i1p47-98.html>.
- Kuester, Keith, Gernot J Müller, and Sarah Stölting (2009). “Is the New Keynesian Phillips Curve Flat?” In: *Economics Letters* 103.1, pp. 39–41.
- Leduc, Sylvain, Daniel J Wilson, et al. (2017). “Has the Wage Phillips Curve Gone Dormant?” In: *FRBSF Economic Letter* 30.
- Ma, Adrian (2002). “GMM estimation of the new Phillips curve.” In: *Economics Letters* 76.3, pp. 411–417.
- Mavroeidis, Sophocles (2004). “Weak identification of forward-looking models in monetary economics.” In: *Oxford Bulletin of Economics and statistics* 66, pp. 609–635.
- (2005). “Identification Issues in Forward-Looking Models Estimated by GMM, with an Application to the Phillips Curve.” In: *Journal of Money, Credit and Banking* 37.3, pp. 421–448. URL: <https://ideas.repec.org/a/mcb/jmoncb/v37y2005i3p421-48.html>.
- Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock (2014). “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve.” In: 52.
- McLeay, Michael and Silvana Tenreyro (2019). *Optimal inflation and the identification of the Phillips curve*. Tech. rep. National Bureau of Economic Research.
- Medeiros, Marcelo C and Eduardo F Mendes (2016). “1-regularization of high-dimensional time-series models with non-Gaussian and heteroskedastic errors.” In: *Journal of econometrics* 191.1, pp. 255–271.
- Mehra, Rajnish and Edward C Prescott (1985). “The equity premium: A puzzle.” In: *Journal of monetary Economics* 15.2, pp. 145–161.
- Moreira, Marcelo J (2003). “A conditional likelihood ratio test for structural models.” In: *Econometrica* 71.4, pp. 1027–1048.
- Murray, Michael P (2006). “Avoiding invalid instruments and coping with weak instruments.” In: *Journal of economic Perspectives* 20.4, pp. 111–132.
- Nakamura, Emi and Jón Steinsson (2008). “Five facts about prices: A reevaluation of menu cost models.” In: *The Quarterly Journal of Economics* 123.4, pp. 1415–1464.
- Nakamura, Emi and Jón Steinsson (2013). *Price Rigidity: Microeconomic Evidence and Macroeconomic Implications*. Working Paper 18705. National

- Bureau of Economic Research. DOI: 10.3386/w18705. URL: <http://www.nber.org/papers/w18705>.
- Nason, James M and Gregor W Smith (2008). “Identifying the new Keynesian Phillips curve.” In: *Journal of Applied Econometrics* 23.5, pp. 525–551.
- Newey, Whitney K and Richard J Smith (2004). “Higher order properties of GMM and generalized empirical likelihood estimators.” In: *Econometrica* 72.1, pp. 219–255.
- Newey, Whitney K and Frank Windmeijer (2009). “Generalized method of moments with many weak moment conditions.” In: *Econometrica* 77.3, pp. 687–719.
- Pencavel, John (1986). “Labor supply of men: a survey.” In: *Handbook of labor economics* 1, pp. 3–102.
- Piazza, Mr Roberto (2018). *Is There a Phillips Curve? A Full Information Partial Equilibrium Approach*. International Monetary Fund.
- Roberts, John M (1995). “New Keynesian economics and the Phillips curve.” In: *Journal of money, credit and banking* 27.4, pp. 975–984.
- Rudd, Jeremy and Karl Whelan (2005a). “Does labor’s share drive inflation?” In: *Journal of Money, Credit, and Banking* 37.2, pp. 297–312.
- (2005b). “New tests of the new-Keynesian Phillips curve.” In: *Journal of Monetary Economics* 52.6, pp. 1167–1181. URL: <https://ideas.repec.org/a/eee/moneco/v52y2005i6p1167-1181.html>.
- (2007). *Modeling inflation dynamics: a critical review of recent research*. Open Access publications. School of Economics, University College Dublin. URL: <https://EconPapers.repec.org/RePEc:ucn:oapubs:10197/201>.
- Savov, Alexi (2011). “Asset Pricing with Garbage.” In: *Journal of Finance* 66.1, pp. 177–201. URL: <https://EconPapers.repec.org/RePEc:bla:jfinan:v:66:y:2011:i:1:p:177-201>.
- Sbordone, Argia M. (2002). “Prices and unit labor costs: a new test of price stickiness.” In: *Journal of Monetary Economics* 49.2, pp. 265–292. URL: <https://ideas.repec.org/a/eee/moneco/v49y2002i2p265-292.html>.
- (2005). “Do expected future marginal costs drive inflation dynamics?” In: *Journal of Monetary Economics* 52.6, pp. 1183–1197.
- (2007). *Globalization and inflation dynamics: the impact of increased competition*. Tech. rep. National Bureau of Economic Research.
- Schorfheide, Frank (2008). “DSGE model-based estimation of the New Keynesian Phillips curve.” In: *FRB Richmond Economic Quarterly* 94.4, pp. 397–433.

- Schorfheide, Frank, Dongho Song, and Amir Yaron (2018). “Identifying Long-Run Risks: A Bayesian Mixed-Frequency Approach.” In: *Econometrica* 86.2, pp. 617–654. DOI: 10.3982/ECTA14308. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA14308>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA14308>.
- Stock, James H and Mark W Watson (2007). “Why has US inflation become harder to forecast?” In: *Journal of Money, Credit and banking* 39, pp. 3–33.
- Stock, James H and Jonathan H Wright (2000). “GMM with weak identification.” In: *Econometrica* 68.5, pp. 1055–1096.
- Stock, James H, Jonathan H Wright, and Motohiro Yogo (2002). “A survey of weak instruments and weak identification in generalized method of moments.” In: *Journal of Business & Economic Statistics* 20.4, pp. 518–529.
- Stock, James H and Motohiro Yogo (2002). *Testing for weak instruments in linear IV regression*.
- Tibshirani, Robert (1996). “Regression shrinkage and selection via the lasso.” In: *Journal of the Royal Statistical Society: Series B (Methodological)* 58.1, pp. 267–288.
- Vermeulen, Philip et al. (2007). “Price setting in the euro area: some stylised facts from individual producer price data.” In: *National Bank of Belgium Working Paper* 111.
- Vissing-Jorgensen, Annette (1998). “Limited stock market participation and the equity premium puzzle.” In:
- (2002). *Limited Asset Market Participation and the Elasticity of Intertemporal Substitution*. NBER Working Papers 8896. National Bureau of Economic Research, Inc. URL: <https://EconPapers.repec.org/RePEc:nbr:nberwo:8896>.
- Wilcox, David W (1992). “The construction of US consumption data: some facts and their implications for empirical work.” In: *The American economic review*, pp. 922–941.
- Woodford, M (2003). “Interest and Prices: Foundation of a Theory of Monetary Policy.” In:
- Working, Holbrook (1960). “Note on the correlation of first differences of averages in a random chain.” In: *Econometrica: Journal of the Econometric Society*, pp. 916–918.
- Yogo, Motohiro (2004). “Estimating the Elasticity of Intertemporal Substitution When Instruments Are Weak.” In: *The Review of Economics and*

Statistics 86.3, pp. 797–810. URL: <https://ideas.repec.org/a/tpr/restat/v86y2004i3p797-810.html>.

Zou, Hui (2006). “The adaptive LASSO ad its oracle properties.” In: *Journal of the American Statistical Association* 101, pp. 1418–1429. DOI: 10.1198/016214506000000735.

A

Appendix for Inflation Dynamics in a Multi-Sector Framework

This appendix refers to the first chapter. It includes several robustness checks, additional details, estimated specifications, the complete model and its underlying variations.

A.1

Additional Details on the Data Set

Table A.1 below summarises our data set, as well as applied transformations. Sectoral information is presented in the main paper – see Table 1.1.

Table A.1: Data Set for the Heterogeneous Economy

Variable	Source	Aggregate	Sectoral	Literature	HP Filter
Output	BEA	yes	no	yes	yes
Consumption	BEA	no	yes	no	yes
Non-Farm Labour Share	BLS	yes	no	yes	yes
PCE Inflation	BEA	yes	yes	yes	no
PCE Chain-Type Price Index	BEA	yes	yes	no	no
PPI Commodities Inflation	BLS	yes	no	yes	no
Effective Fed Funds	FED	yes	no	yes	no
5-Year Treasury Spread	FED	yes	no	yes	no
Avg. Hourly Earnings Inflation	BLS	yes	no	yes	no

Notes: Types of aggregate and sectoral data compiled in our data set. Output and consumption are measured in per capita terms to account for a model with no population growth. The PPI, the 5-year spread, the Fed Funds rate, the labour share as well as the wage inflation are taken from the Federal Reserve Bank of St. Louis' FRED economic database. The second column provides the sources of the variables: BEA stands for U.S. Bureau of Economic Analysis, BLS for U.S. Bureau of Labor Statistics and FED for the Federal Reserve System. The third and fourth columns give the level of aggregation of the data. The fifth column indicate whether that variable is usually used in the empirical literature or not. Finally, the last column shows for which of those variables we extract cyclical components based on the HP filter.

A.2

Heterogenous Economy: Completely Specified Model

Additional details on the heterogeneous economy of the first chapter.

The first order solution to the household problem gives:

$$\begin{aligned}\frac{W_{kj,t}}{P_t} &= \frac{L_{kj,t}^{\varphi^{-1}}}{C_t^{-\sigma}} \\ E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} I_t \frac{P_t}{P_{t+1}} \right] &= \frac{1}{\beta} \\ C_{k,t} &= f(k) C_t \left(\frac{P_{k,t}}{P_t} \right)^{-\epsilon} \\ C_{kj,t} &= f(k)^{-1} C_{k,t} \left(\frac{P_{kj,t}}{P_{k,t}} \right)^{-\epsilon}\end{aligned}$$

The pricing kernel between period t and $t + s$ of this economy is defined as:

$$Q_{t,t+s} = \beta \left(\frac{C_{t+s}}{C_t} \right)^{-\sigma} \frac{P_t}{P_{t+s}}$$

The result of the first order conditions as well as the law of movement for sectoral price indices are in the main text, equations (13) and (14) in the first chapter, respectively. Note that the specification for the interest rate rule is irrelevant for the purposes of this paper¹.

One can log-linearise the equations of the model, in terms of deviations from steady-state values. The law of movement becomes:

$$p_{k,t} = \lambda_k (x_{k,t} - p_{k,t-1}) + (1 - \lambda_k) p_{k,t-1} + \gamma (1 - \lambda_k) \pi_{k,t-1}, \quad (1)$$

while the price-setting problem in (9) and (10) of the first chapter results:

$$x_{k,t} = (1 - \beta(1 - \lambda_k)) E_t \sum_{s=0}^{\infty} \beta^s (1 - \lambda_k)^s [p_{t+s} + \Theta y_{t+s} - \gamma \pi_{k,t+s-1,t-1}], \quad (2)$$

where Θ and $\pi_{\tau,t}$, $\tau > t$, are defined in the main text. Other parts of the model

¹It is important that it does not consider sectoral shocks, though. See section 4 of the first chapter.

are the same as in Carvalho (2006), and can be linearised as follows:

$$\begin{aligned}
 y_t &= E_t y_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1}) \\
 y_t &= \int_0^1 f(k) y_{k,t} dk \\
 y_{k,t} &= y_t - \epsilon(p_{k,t} - p_t) \\
 y_{kj,t} &= y_{k,t} - \epsilon(p_{kj,t} - p_{k,t}) \\
 y_{k,t} &= \int_0^1 y_{kj,t} dj \\
 p_t &= \int_0^1 f(k) p_{k,t} dk \\
 p_{k,t} &= \int_0^1 p_{kj,t} dj \\
 w_{kj,t} - p_t &= \varphi^{-1} l_{kj,t} + \sigma c_t \\
 b_t &= 0 \\
 y_{kj,t} &= c_{kj,t} = n_{kj,t} = l_{kj,t}
 \end{aligned}$$

Each firm kj that resets its price, fixing $x_{kj,t}$ at period t , will have a future demand for its variety that takes the form $y_{kj,t+s} = y_{t+s} - \epsilon(x_{kj,t} - p_{t+s})$. Using this expression in the log-linearised first-order condition for the labour market above yields:

$$w_{kj,t+s} = (1 + \varphi^{-1}\epsilon)p_{t+s} + (\varphi^{-1} + \sigma)y_{t+s} - \epsilon\varphi^{-1}x_{kj,t} \quad (3)$$

A.2.1 Deriving the NKPCs

We now derive the generalised NKPC, equation (15) of the main text. As usual, it is possible to write the reset prices as a discounted stream of future marginal costs (with an additional term, due to indexation):

$$x_{kj,t} = (1 - \beta(1 - \lambda_k))E_t \sum_{s=0}^{\infty} \beta^s (1 - \lambda_k)^s [w_{kj,t+s} - \gamma \pi_{k,t+s-1,t-1}] \quad (4)$$

Write $mc_{kj,t+s}^r = w_{kj,t+s} - p_{k,t+s}$, the real marginal cost of a firm kj in $t + s$, which reset its price at t . Using this, subtracting $p_{k,t-1}$ from (4), writing in terms of all the firms in the same sector, and rearranging:

$$x_{k,t} - p_{k,t-1} = (1 - \beta(1 - \lambda_k))E_t \sum_{s=0}^{\infty} \beta^s (1 - \lambda_k)^s [mc_{k,t+s}^r + p_{k,t+s} - p_{k,t-1} - \gamma \pi_{k,t+s-1,t-1}] \quad (5)$$

It is possible to expand this equation to:

$$\begin{aligned}
 x_{k,t} - p_{k,t-1} = & (1 - \beta(1 - \lambda_k))E_t \sum_{s=0}^{\infty} \beta^s (1 - \lambda_k)^s [mc_{k,t+s}^r] + \\
 & (1 - \beta(1 - \lambda_k))E_t \{ \beta^0 (1 - \lambda_k)^0 (p_{k,t} - p_{k,t-1}) + \\
 & \beta(1 - \lambda_k)(p_{k,t+1} - p_{k,t-1} + \gamma \pi_{k,t,t-1}) + \\
 & \beta^2 (1 - \lambda_k)^2 (p_{k,t+2} - p_{k,t-1} + \gamma \pi_{k,t+1,t-1}) + [\dots] \}
 \end{aligned} \tag{6}$$

Cancelling terms and writing recursively:

$$x_{k,t} - p_{k,t-1} = [1 - \beta(1 - \lambda_k)]mc_{k,t} + [1 - \beta\gamma(1 - \lambda_k)]\pi_{k,t} + [1 - \beta(1 - \lambda_k)]E_t \{ x_{k,t+1} - p_{k,t} \} \tag{7}$$

Using the law of movement, (1) in this appendix:

$$x_{k,t} - p_{k,t-1} = \frac{1}{\lambda_k} [\pi_{k,t} - \gamma(1 - \lambda_k)\pi_{k,t-1}] \tag{8}$$

One period forward:

$$x_{k,t+1} - p_{k,t} = \frac{1}{\lambda_k} [\pi_{k,t+1} - \gamma(1 - \lambda_k)\pi_{k,t}] \tag{9}$$

Combining (7) and (9):

$$\begin{aligned}
 \pi_{k,t} [1 - \lambda_k(1 - \gamma(1 - \lambda_k)\beta) + (1 - \lambda_k)^2 \beta \gamma] = & \lambda_k [1 - (1 - \lambda_k)\beta] mc_{k,t}^r + \\
 (1 - \lambda_k)\beta E_t \pi_{k,t+1} + & \gamma(1 - \lambda_k)\pi_{k,t-1}
 \end{aligned} \tag{10}$$

Rearranging this equation, it is possible to derive the sectoral NKPC for with the real marginal cost as forcing variable:

$$\pi_{k,t} = \frac{\beta}{1 + \beta\gamma} E_t \pi_{k,t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{k,t-1} + \frac{\lambda_k(1 - (1 - \lambda_k)\beta)}{(1 - \lambda_k)(1 + \beta\gamma)} mc_{k,t}^r \tag{11}$$

To derive the sectoral NKPC with the output gap, combine (3) and (11)². Use the property that $\pi_t = \int_0^1 f(k)\pi_{k,t} dk$ to get to the generalised NKPC, (15) in the main paper.

We now derive the “expanded” NKPC, that features the marginal cost as the forcing variables, as well as endogenous sectoral terms as controls. First, write the sectoral slope coefficients as a function of their weighted average in the economy:

$$\psi_k(\lambda_k, \beta, \gamma) \equiv \left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta\lambda_k}{(1 + \beta\gamma)} \right] = \psi + \zeta_k \tag{12}$$

This is equation (17) in the first chapter. Since the model features constant returns to scale and and future marginal costs are synchronised across firms within a sector, it follows that:

²These are similar developments as in Woodford (2003, Appendix B.7, p. 666-669).

$$mc_t^r = \int_0^1 (mc_{k,t} - p_{k,t}) dk, \quad (13)$$

where $mc_{k,t}$ is the *nominal* marginal cost of firms in sector k . Next, integrate equation (11) above over the sectors, using the density function:

$$\begin{aligned} \pi_t = & \frac{\beta}{1 + \beta\gamma} E_t \pi_{t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{t-1} + \\ & \int_0^1 f(k) \left[\frac{\lambda_k}{(1 - \lambda_k)(1 + \beta\gamma)} - \frac{\beta\lambda_k}{(1 + \beta\gamma)} \right] (mc_{k,t} - p_{k,t}) dk, \end{aligned} \quad (14)$$

to then use (12) and (13) above, getting to:

$$\begin{aligned} \pi_t = & \frac{\beta}{1 + \beta\gamma} E_t \pi_{t+1} + \frac{\gamma}{1 + \beta\gamma} \pi_{t-1} + \psi mc_t^r \\ & + \int_0^1 f(k) \zeta_k (mc_{k,t} - p_{k,t}) dk \end{aligned} \quad (15)$$

Combine (3) and (15) to find the expanded NKPC, (18) in the main text.

A.3

Instrument Sets

Table A.2 below provides additional details on how we construct the instrument sets used in this paper. Recall that the first approach (I) is the baseline. Results for this are exhibited in the main text. We repeat the estimations using II, III and IV in this appendix.

Table A.2: Instrument Sets

Approach	Type	Candidates			Choice	
		Literature	Sectoral Data	Selection	Reduced Form	Structural Form
I	adaLASSO	yes	output gaps	lag 1	lags 1 to 3	lags 1 to 2
II	adaLASSO	yes	output gaps	lag 1	lags 1 to 2	lag 1
III	adaLASSO	yes	output gaps relative prices	lag 2	lag 2	lag2
IV	ad-hoc	yes	output gaps	-	lags 1 to 3	lags 1 to 2

Notes: “Literature” indicates whether variables outside of the model usually used in the literature are present as candidates. These are the Fed Funds rate, the Treasury spread and inflation in commodities and in wages. As before, the output gap and the non-farm labour share are always candidates. “Sectoral data” shows sectoral variables whose lags are considered as instruments (“relative prices” are given by $p_{k,t} - p_t$). The following columns gives the lag considered when selecting based on (3). The last two columns present the choice of instruments based on candidates selected.

A.4

Additional Results for the Heterogeneous Economy

A.4.1**Reduced-Form Estimations**

The following provides complementary results for Table 1.5 in section 4. Here we perform robustness checks considering approaches II to IV to construct instruments.

Table A.3: Reduced-Form NKPCs for the Heterogeneous Economy – Approach II

Coefficient	Generalised		Expanded	Homogeneous
	$\gamma = 0$	$\gamma \neq 0$	$\gamma \neq 0$	$\gamma \neq 0$
γ_f	1.016 (0.049)	0.610 (0.006)	0.650 (0.012)	0.746 (0.018)
γ_b	-	0.405 (0.002)	0.408 (0.015)	0.563 (0.013)
Output Gap	0.147*** (0.055)	-0.002 (0.005)	0.001 (0.008)	-
Marginal Cost	-	-	0.016*** (0.005)	0.146*** (0.009)
$\beta = 0.99$ and $\gamma = 0.5$				
Output Gap	0.067*** (0.020)	0.002 (0.003)	-0.011** (0.005)	-
Marginal Cost	-	-	-0.005*** (0.002)	-0.006*** (0.001)

Notes: Estimates of the NKPCs in (15) and (18) using CUE-GMM. For the former we present results for both the forward-looking ($\gamma = 0$) and the hybrid ($\gamma \neq 0$) versions of the model. The instrument set is constructed based on approach II (see the main text). The lower end of the table presents estimates while fixing $\beta = 0.99$ and $\gamma = 0.5$. HAC standard errors are presented in parentheses. The hypotheses of statistically insignificant coefficients for the output gap and the labour share are tested: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A.4: Reduced-Form NKPCs for the Heterogeneous Economy – Approach III

Coefficient	Generalised		Expanded	Homogeneous
	$\gamma = 0$	$\gamma \neq 0$	$\gamma \neq 0$	$\gamma \neq 0$
γ_f	0.998 (0.013)	0.488 (0.074)	0.526 (0.151)	0.532 (0.019)
γ_b	-	0.491*** (0.002)	0.454*** (0.155)	0.439*** (0.022)
Output Gap	0.027 (0.069)	0.035 (0.088)	-0.012 (0.132)	-
Marginal Cost	-	-	0.011 (0.056)	0.042*** (0.004)
$\beta = 0.99$ and $\gamma = 0.5$				
Output Gap	0.032 (0.080)	-0.010 (0.052)	-0.017 (0.067)	-
Marginal Cost	-	-	0.045 (0.033)	0.013*** (0.002)

Notes: Estimates of the NKPCs in (15) and (18) using CUE-GMM. For the former we present results for both the forward-looking ($\gamma = 0$) and the hybrid ($\gamma \neq 0$) versions of the model. The instrument set is constructed based on approach III (see the main text). The lower end of the table presents estimates while fixing $\beta = 0.99$ and $\gamma = 0.5$. HAC standard errors are presented in parentheses. The hypotheses of statistically insignificant coefficients for the output gap and the labour share are tested: *p<0.1; **p<0.05; ***p<0.01.

Table A.5: Reduced-Form NKPCs for the Heterogeneous Economy – Approach IV

Coefficient	Generalised		Expanded	Homogeneous
	$\gamma = 0$	$\gamma \neq 0$	$\gamma \neq 0$	$\gamma \neq 0$
γ_f	1.027 (0.002)	0.719 (0.002)	0.660 (0.003)	1.125 (0.006)
γ_b	-	0.320 (0.002)	0.356 (0.003)	-0.195 (0.006)
Output Gap	0.066*** (0.005)	-0.055 (0.088)	0.007*** (0.002)	-
Marginal Cost	-	-	0.021*** (0.002)	0.012*** (0.004)
$\beta = 0.99$ and $\gamma = 0.5$				
Output Gap	0.030*** (0.006)	0.074*** (0.002)	0.015*** (0.003)	-
Marginal Cost	-	-	0.013*** (0.002)	-0.018*** (0.002)

Notes: Estimates of the NKPCs in (15) and (18) using CUE-GMM. For the former we present results for both the forward-looking ($\gamma = 0$) and the hybrid ($\gamma \neq 0$) versions of the model. The instrument set is constructed based on approach IV (see the main text). The lower end of the table presents estimates while fixing $\beta = 0.99$ and $\gamma = 0.5$. HAC standard errors are presented in parentheses. The hypotheses of statistically insignificant coefficients for the output gap and the labour share are tested: *p<0.1; **p<0.05; ***p<0.01.

A.4.2

Structural Estimations: Implied Slope

We also provide complementary material for Table 1.7, exhibited in the main text (section 4). The following tables consider approaches II to IV.

Table A.6: Implied Slope from Estimations of (19) – Approach II

Model	Normalisation	Calibration		
		Baseline	↑ Real Rigidity	↓ Real Rigidity
$\gamma = 0$	(1)	0.008***	0.007***	0.009***
		(0.000)	(0.000)	(0.000)
		[0.67]	[0.47]	[0.78]
	(2)	0.010***	0.007***	0.010***
		(0.000)	(0.000)	(0.000)
		[0.59]	[0.48]	[0.55]
$\gamma \neq 0$	(1)	0.010***	0.006***	0.012***
		(0.000)	(0.000)	(0.000)
		[0.64]	[0.62]	[0.61]
	(2)	0.010***	0.007***	0.010***
		(0.000)	(0.000)	(0.000)
		[0.82]	[0.84]	[0.69]

Notes: Implied slope from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and γ are allowed to vary. The instrument set is constructed based on approach II. Standard errors from Delta method are presented in parentheses. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in brackets. We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix). In addition, * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A.7: Implied Slope from Estimations of (19) – Approach III

Model	Normalisation	Calibration		
		Baseline	↑ Real Rigidity	↓ Real Rigidity
$\gamma = 0$	(1)	0.015***	0.006***	0.020***
		(0.002)	(0.000)	(0.005)
		[0.58]	[0.42]	[0.64]
	(2)	0.015***	0.010***	0.020***
		(0.003)	(0.001)	(0.004)
		[0.62]	[0.47]	[0.63]
$\gamma \neq 0$	(1)	0.012***	0.007***	0.015***
		(0.001)	(0.001)	(0.001)
		[0.59]	[0.58]	[0.73]
	(2)	0.012***	0.010***	0.017***
		(0.000)	(0.000)	(0.000)
		[0.52]	[0.41]	[0.54]

Notes: Implied slope from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and γ are allowed to vary. The instrument set is constructed based on approach III. Standard errors from Delta method are presented in parentheses. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in brackets. We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix). In addition, *p<0.1; **p<0.05; ***p<0.01.

Table A.8: Implied Slope from Estimations of (19) – Approach IV

Model	Normalisation	Calibration		
		Baseline	↑ Real Rigidity	↓ Real Rigidity
$\gamma = 0$	(1)	0.007***	0.007***	0.007***
		(0.000)	(0.000)	(0.000)
		[0.61]	[0.66]	[0.44]
	(2)	0.007***	0.007***	0.006***
		(0.000)	(0.001)	(0.000)
		[0.64]	[0.72]	[0.22]
$\gamma \neq 0$	(1)	0.009***	0.009***	0.017***
		(0.000)	(0.001)	(0.000)
		[0.66]	[0.62]	[0.73]
	(2)	0.005***	0.005***	0.013***
		(0.000)	(0.000)	(0.000)
		[0.72]	[0.54]	[0.68]

Notes: Implied slope from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and γ are allowed to vary. The instrument set is constructed based on approach IV. Standard errors from Delta method are presented in parentheses. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in brackets. We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix). In addition, *p<0.1; **p<0.05; ***p<0.01.

A.4.3

Structural Estimations: Implied Stickiness

We now present robustness tables for Table 1.8 and Table 1.9, exhibited in the fourth section of the main text. The following tables consider approaches II to IV.

Table A.9: Implied Stickiness from the Hybrid Model ($\gamma \neq 0$) – Approach II

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters		
			θ^{agg}	β	γ
\uparrow Real Rigidity	(1)	0.62	0.58	0.99	0.37
			(0.49 - 0.66)	(0.013)	(0.007)
	(2)	0.84	0.63	0.99	0.22
			(0.32 - 0.94)	(0.001)	(0.006)
Baseline	(1)	0.64	0.58	0.99	0.37
			(0.57 - 0.60)	(0.012)	(0.007)
	(2)	0.82	0.66	0.99	0.21
			(0.46 - 0.87)	(0.000)	(0.005)
\downarrow Real Rigidity	(1)	0.61	0.62	0.99	0.36
			(0.59 - 0.64)	(0.011)	(0.006)
	(2)	0.69	0.65	0.99	0.23
			(0.54 - 0.67)	(0.000)	(0.002)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β , γ and each λ_k are allowed to vary in the range $(0, 1)$. The instrument set is constructed based on approach II. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

Table A.10: Implied Stickiness from the Hybrid Model ($\gamma \neq 0$) – Approach III

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters		
			θ^{agg}	β	γ
\uparrow Real Rigidity	(1)	0.58	0.57 (0.30 - 0.85)	0.99 (0.013)	0.73 (0.253)
	(2)	0.41	0.54 (0.33 - 0.75)	0.99 (0.001)	0.10 (0.012)
Baseline	(1)	0.59	0.54 (0.39 - 0.69)	0.99 (0.211)	0.72 (0.162)
	(2)	0.52	0.65 (0.31 - 0.98)	0.99 (0.001)	0.14 (0.010)
\downarrow Real Rigidity	(1)	0.73	0.59 (0.52 - 0.66)	0.99 (0.080)	0.48 (0.050)
	(2)	0.54	0.65 (0.38 - 0.91)	0.99 (0.001)	0.11 (0.010)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β , γ and each λ_k are allowed to vary in the range (0, 1). The instrument set is constructed based on approach III. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies ($1 - \lambda_k$) that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

Table A.11: Implied Stickiness from the Hybrid Model ($\gamma \neq 0$) – Approach IV

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters		
			θ^{agg}	β	γ
\uparrow Real Rigidity	(1)	0.63	0.50 (0.44 - 0.55)	0.96 (0.011)	0.57 (0.007)
	(2)	0.54	0.58 (0.56 - 0.61)	0.98 (0.000)	0.33 (0.001)
Baseline	(1)	0.73	0.52 (0.49 - 0.56)	0.98 (0.017)	0.49 (0.011)
	(2)	0.64	0.56 (0.50 - 0.62)	0.98 (0.001)	0.31 (0.003)
\downarrow Real Rigidity	(1)	0.73	0.56 (0.46 - 0.66)	0.99 (0.012)	0.46 (0.007)
	(2)	0.68	0.61 (0.56 - 0.65)	0.99 (0.001)	0.27 (0.002)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β , γ and each λ_k are allowed to vary in the range (0, 1). The instrument set is constructed based on approach IV. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies ($1 - \lambda_k$) that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

Table A.12: Stickiness from the Forward-Looking Model ($\gamma = 0$) – Approach II

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters	
			θ^{agg}	β
\uparrow Real Rigidity	(1)	0.47	0.58	0.99
			(0.51 - 0.66)	(0.001)
	(2)	0.48	0.57	0.99
			(0.49 - 0.65)	(0.001)
Baseline	(1)	0.67	0.63	0.99
			(0.57 - 0.70)	(0.000)
	(2)	0.59	0.62	0.99
			(0.55 - 0.69)	(0.000)
\downarrow Real Rigidity	(1)	0.78	0.70	0.99
			(0.61 - 0.78)	(0.000)
	(2)	0.55	0.67	0.99
			(0.61 - 0.72)	(0.000)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and each λ_k are allowed to vary in the range $(0, 1)$. We fix $\gamma = 0$ (purely forward-looking model). The instrument set is constructed based on approach II. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

Table A.13: Stickiness from the Forward-Looking Model ($\gamma = 0$) – Approach III

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters	
			θ^{agg}	β
\uparrow Real Rigidity	(1)	0.42	0.65	0.99
			(0.53 - 0.78)	(0.000)
	(2)	0.47	0.62	0.99
			(0.00 - 1.00)	(0.004)
Baseline	(1)	0.58	0.65	0.99
			(0.00 - 1.00)	(0.006)
	(2)	0.63	0.68	0.99
			(0.00 - 1.00)	(0.011)
\downarrow Real Rigidity	(1)	0.64	0.69	0.99
			(0.00 - 1.00)	(0.000)
	(2)	0.63	0.70	0.99
			(0.00 - 1.00)	(0.007)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and each λ_k are allowed to vary in the range $(0, 1)$. We fix $\gamma = 0$ (purely forward-looking model). The instrument set is constructed based on approach III. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

Table A.14: Stickiness from the Forward-Looking Model ($\gamma = 0$) – Approach IV

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters	
			θ^{agg}	β
\uparrow Real Rigidity	(1)	0.66	0.60	0.97
			(0.53 - 0.68)	(0.000)
	(2)	0.72	0.59	0.99
			(0.50 - 0.68)	(0.000)
Baseline	(1)	0.61	0.66	0.99
			(0.62 - 0.71)	(0.000)
	(2)	0.64	0.65	0.99
			(0.59 - 0.71)	(0.000)
\downarrow Real Rigidity	(1)	0.44	0.71	0.98
			(0.69 - 0.73)	(0.000)
	(2)	0.22	0.73	0.98
			(0.68 - 0.78)	(0.000)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and each λ_k are allowed to vary in the range $(0, 1)$. We fix $\gamma = 0$ (purely forward-looking model). The instrument set is constructed based on approach IV. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

A.4.4

Structural Estimations: β and γ Calibrated

Tables below give complementary information to the results of section 4. We re-estimate our models while imposing $\beta = 0.99$ and, if applicable (hybrid model), $\gamma \in \{0.3, 0.5, 0.7\}$. All estimations use the baseline routine (I) to construct the instrument set.

Table A.15: Implied Slope – $\beta = 0.99$ and $\gamma \in \{0, 0.3, 0.5, 0.7\}$

Model	Normalisation	Calibration		
		Baseline	\uparrow Real Rigidity	\downarrow Real Rigidity
$\gamma = 0$	(1)	0.007***	0.004***	0.005***
		(0.000)	(0.000)	(0.000)
		[0.29]	[0.34]	[-0.04]
	(2)	0.006***	0.004***	0.006***
		(0.000)	(0.000)	(0.000)
		[0.27]	[0.47]	[0.15]
$\gamma = 0.3$	(1)	0.010***	0.005***	0.018***
		(0.000)	(0.000)	(0.000)
		[0.72]	[0.49]	[0.76]
	(2)	0.008***	0.003***	0.010***
		(0.000)	(0.000)	(0.000)
		[0.60]	[-0.08]	[0.64]
$\gamma = 0.5$	(1)	0.012***	0.007***	0.021***
		(0.000)	(0.000)	(0.001)
		[0.67]	[0.44]	[0.69]
	(2)	0.010***	0.004***	0.011***
		(0.000)	(0.000)	(0.000)
		[0.57]	[0.18]	[0.58]
$\gamma = 0.7$	(1)	0.013***	0.008***	0.021***
		(0.000)	(0.000)	(0.000)
		[0.58]	[0.50]	[0.58]
	(2)	0.011***	0.005***	0.012***
		(0.000)	(0.000)	(0.000)
		[0.53]	[0.10]	[0.51]

Notes: Implied slope from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. Each λ_k is allowed to vary in the range (0, 1). We fix $\beta = 0.99$ and $\gamma \in \{0.3, 0.5, 0.7\}$ in the hybrid version of the model. The instrument set is constructed based on approach I. Standard errors provided by Delta method are presented in parentheses. Correlations between estimated and benchmark infrequencies ($1 - \lambda_k$) that come from the micro data in Bils and P. Klenow (2004) are shown in brackets. We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix). In addition, *p<0.1; **p<0.05; ***p<0.01.

Table A.16: Implied Stickiness – $\beta = 0.99$ and $\gamma \in \{0, 0.3, 0.5, 0.7\}$

Calibration	Normalisation	$\gamma = 0$		$\gamma = 0.3$		$\gamma = 0.5$		$\gamma = 0.7$	
		$Corr(\theta_k, Micro)$	θ^{agg}	$Corr(\theta_k, Micro)$	θ^{agg}	$Corr(\theta_k, Micro)$	θ^{agg}	$Corr(\theta_k, Micro)$	θ^{agg}
↑ Real Rigidity	(1)	0.33	0.77 (0.64 - 0.90)	0.49	0.59 (0.36 - 0.82)	0.44	0.52 (0.24 - 0.80)	0.50	0.46 (0.18 - 0.75)
		0.47	0.71 (0.62 - 0.81)	-0.08	0.66 (0.15 - 1.00)	0.18	0.58 (0.32 - 0.83)	0.18	0.52 (0.42 - 0.62)
	(2)	0.29	0.70 (0.56 - 0.84)	0.73	0.58 (0.45 - 0.71)	0.67	0.52 (0.37 - 0.67)	0.58	0.50 (0.36 - 0.68)
		0.27	0.70 (0.56 - 0.84)	0.60	0.59 (0.34 - 0.84)	0.57	0.52 (0.23 - 0.81)	0.53	0.49 (0.16 - 0.82)
↓ Real Rigidity	(1)	-0.04	0.75 (0.69 - 0.81)	0.76	0.56 (0.46 - 0.66)	0.69	0.51 (0.40 - 0.63)	0.59	0.50 (0.37 - 0.64)
		0.15	0.73 (0.71 - 0.76)	0.64	0.61 (0.42 - 0.81)	0.58	0.56 (0.34 - 0.78)	0.51	0.52 (0.50 - 0.54)
	(2)								

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of (19) using SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. Each λ_k is allowed to vary in the range $(0, 1)$. We set $\beta = 0.99$ and $\gamma = 0.3$. The instrument set is constructed based on approach I. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the columns “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

A.4.5

Alternative Form of Indexation

Here, we estimate a variation of the model that features an indexation scheme à la Gali and Gertler (1999). Specifically, a fraction ω of the firms in each sector are “backward-looking”. These readjust their prices based on their competitors in the same sector³:

$$p_{k,t}^b = x_{k,t-1} + \pi_{k,t-1} \quad (16)$$

Forward-looking firms continue to behave as in the model above, optimising their prices looking to a discounted stream of expected nominal marginal costs (denote $p_{k,t}^f$ as their reset prices). It follows that, under this setting, the sectoral index for newly set prices takes the form:

$$x_{k,t} = (1 - \omega)p_{k,t}^f + \omega x_{k,t}^b$$

Sectoral prices take the form:

$$p_{k,t} = \lambda_k x_{k,t} + (1 - \lambda_k)p_{k,t-1} \quad (17)$$

Write $p_{k,t}^f$ as a discounted sequence of future *nominal* marginal costs:

$$\begin{aligned} p_{k,t}^f &= (1 - \beta(1 - \lambda_k))E_t \sum_{s=0}^{\infty} \beta^s (1 - \lambda_k)^s [mc_{k,t+s}] = \\ &= (1 - \beta(1 - \lambda_k))m_{k,t} + \beta(1 - \lambda_k)E_t p_{k,t+1}^f \end{aligned} \quad (18)$$

Use (17) to write:

$$x_{k,t} = \frac{1}{\lambda_k}(p_{k,t} - (1 - \lambda_k)p_{k,t-1}), \quad (19)$$

and:

$$x_{k,t-1} = \frac{1}{\lambda_k}(p_{k,t-1} - (1 - \lambda_k)p_{k,t-2}) \quad (20)$$

Combine the latter with (16) to find:

$$\begin{aligned} p_{k,t}^b &= \frac{1}{\lambda_k}(p_{k,t-1} - (1 - \lambda_k)p_{k,t-2} + \pi_{k,t-1}) \\ &= \frac{1}{\lambda_k}(\pi_{k,t-1} + \lambda_k p_{k,t-1}) \\ &= \frac{1}{\lambda_k}\pi_{k,t-1} + p_{k,t-1} \end{aligned} \quad (21)$$

And it follows that:

$$p_{k,t}^b - p_{k,t} = -\pi_{k,t} + \frac{1}{\lambda_k}\pi_{k,t-1} \quad (22)$$

³It is assumed that the backward-looking firm can not infer which of its competitors in the same sector are also backward-looking.

Combining (17) and the sectoral index for newly set prices:

$$\frac{1}{\lambda_k} p_{k,t} - \frac{1 - \lambda_k}{\lambda_k} p_{k,t-1} = (1 - \omega) p_{k,t}^f + \omega p_{k,t}^b \quad (23)$$

Subtracting $p_{k,t}$ from this:

$$\frac{1 - \lambda_k}{\lambda_k} \pi_{k,t} = (1 - \omega)(p_{k,t}^f - p_{k,t}) + \omega(p_{k,t}^b - p_{k,t}) \quad (24)$$

Isolating $p_{k,t}^f$ in this equation, subtracting $p_{k,t}$ from (18), and combining both resulting expressions:

$$p_{k,t}^f = p_{k,t} + \frac{1 - \lambda_k}{\lambda_k(1 - \omega)} \pi_{k,t} - \frac{\omega}{1 - \omega} (p_{k,t}^b - p_{k,t}) \quad (25)$$

Next, from this and (22) forward:

$$E_t p_{k,t}^f = E_t p_{k,t} + \frac{1 - \lambda_k}{\lambda_k(1 - \omega)} E_t \pi_{k,t+1} + \frac{\omega}{1 - \omega} E_t \pi_{k,t+1} - \frac{\omega}{(1 - \omega)\lambda_k} \pi_{k,t} \quad (26)$$

Combine this expression with $p_{k,t}^f - p_{k,t}$, where $p_{k,t}^f$ uses (18):

$$\begin{aligned} p_{k,t}^f - p_{k,t} &= [1 - \beta(1 - \lambda_k)] m c_{k,t}^r - \beta(1 - \lambda_k) p_{k,t} + \beta(1 - \lambda_k) E_t p_{k,t+1} \\ &\quad + \frac{\beta(1 - \lambda_k)(1 - \lambda_k + \omega \lambda_k)}{\lambda_k(1 - \omega)} E_t \pi_{k,t+1} \\ &\quad - \frac{\beta(1 - \lambda_k)\omega}{\lambda_k(1 - \omega)} \pi_{k,t} \end{aligned} \quad (27)$$

Finally, by combining this expression with (22) and (24), it is possible to show that the aggregate NKPC in this economy is:

$$\begin{aligned} \pi_t &= E_t \int_0^1 f(k) \frac{\beta(1 - \lambda_k)}{\phi_k} \pi_{k,t+1} dk + \int_0^1 f(k) \frac{\omega}{\phi_k} \pi_{k,t-1} dk \\ &\quad + \int_0^1 f(k) \left[\frac{(1 - \omega)(\lambda_k)(1 - \beta(1 - \lambda_k))}{\phi_k} \right] dk \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon \varphi^{-1}} - \frac{1}{\epsilon} \right) y_t \\ &\quad + \frac{1}{\epsilon} \int_0^1 f(k) \left[\frac{(1 - \omega)(\lambda_k)(1 - \beta(1 - \lambda_k))}{\phi_k} \right] y_{k,t} dk, \end{aligned} \quad (28)$$

while sectoral NKPCs take the form:

$$\begin{aligned} \pi_t &= E_t \frac{\beta(1 - \lambda_k)}{\phi_k} \pi_{k,t+1} + \frac{\omega}{\phi_k} \pi_{k,t-1} \\ &\quad + \left[\frac{(1 - \omega)(\lambda_k)(1 - \beta(1 - \lambda_k))}{\phi_k} \right] \left(\frac{\varphi^{-1} + \sigma}{1 + \epsilon \varphi^{-1}} - \frac{1}{\epsilon} \right) y_t \\ &\quad + \frac{1}{\epsilon} \left[\frac{(1 - \omega)(\lambda_k)(1 - \beta(1 - \lambda_k))}{\phi_k} \right] y_{k,t}, \end{aligned} \quad (29)$$

where $\phi_k = 1 - \lambda_k + \omega(1 - (1 - \lambda_k)(1 - \beta))$.

Note that the system comprised of (28) and (29) is far more complex than that in (19) of the main text. First, non-linearities are more evident. Second, the aggregate NKPC has even more variables, since we not obtain the aggregate inflation rate by aggregating its sectoral counterparts. Even so, Table A.17 shows that the model performs remarkably well in reconfirming results of section 4 in the main paper. In addition, ω can be compared with that in Galí and Gertler (1999). While we generally estimate $\hat{\omega} = 0.30$, they obtained 0.27 in their baseline specification.

Table A.17: Estimations of the Model with Indexation à la Galí and Gertler (1999)

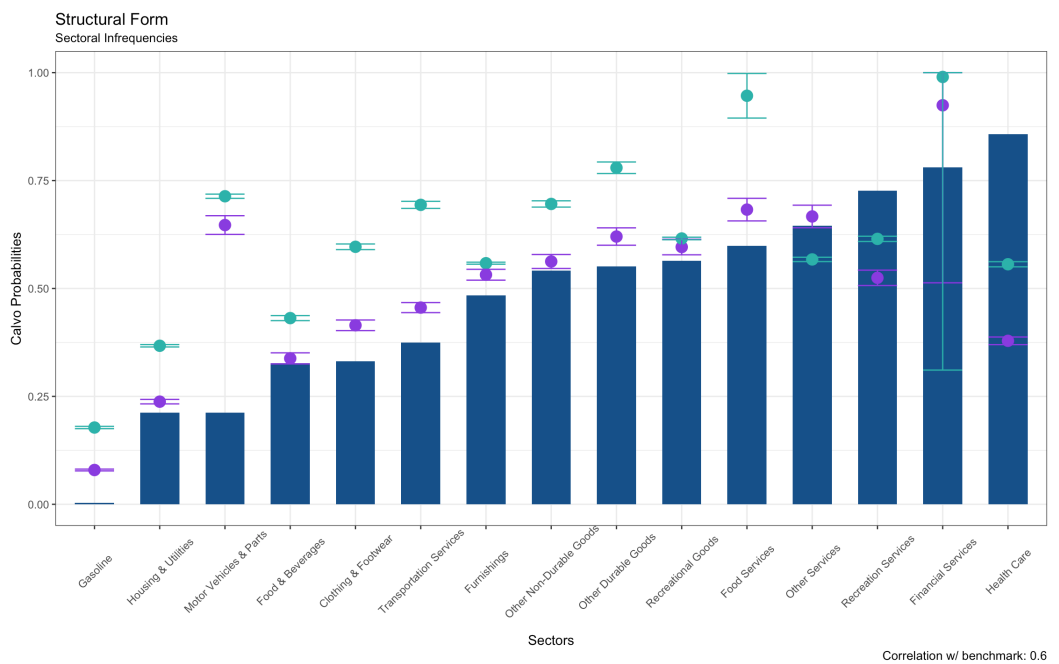
Calibration	$Corr(\theta_k, Micro)$	Parameters			
		θ^{agg}	β	ω	Slope ^{agg}
↑ Real Rigidities	0.25	0.74 (0.74 - 0.74)	0.98 (0.000)	0.30 (0.000)	0.001*** (0.000)
Baseline	0.03	0.65 (0.65 - 0.66)	0.99 (0.000)	0.30 (0.000)	0.004*** (0.000)
↓ Real Rigidities	0.00	0.69 (0.69 - 0.70)	0.99 (0.000)	0.32 (0.000)	0.005*** (0.000)

Notes: Implied slope in (25) and aggregate Calvo-pricing probability in (20) from estimations of the system comprised on (25) and (26). We rely on SYS-GMM with a HAC estimator for the covariance matrix. The calibrations used are shown in Table 1.6. β , ω and each λ_k are allowed to vary in the range (0,1). The instrument set is constructed based on approach I. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. **The micro benchmark implies $\theta^{\text{micro}} \approx 0.48$.** Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix). The null of a statistically insignificant slope is tested: *p<0.1; **p<0.05; ***p<0.01.

A.4.6
Sectoral Infrequencies: Estimated vs. Implied from Micro Evidence

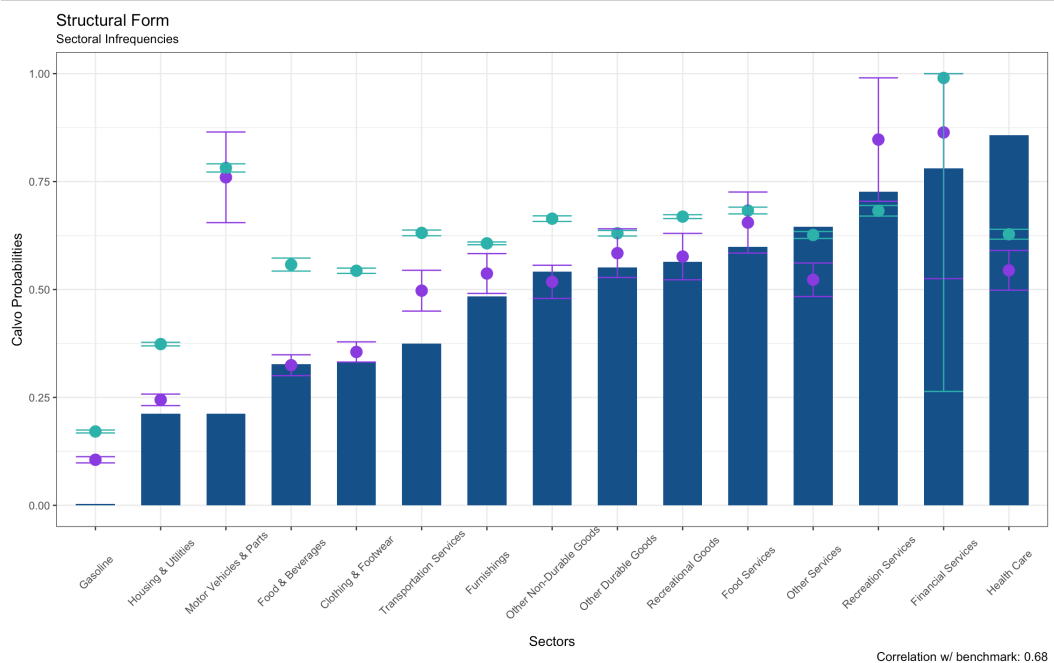
Here we present additional charts that compare estimated infrequencies $(1 - \hat{\lambda}_k)$ with micro-based analogues. First, we present tables that use the remaining two calibrations in Table 1.6 not shown in the main text. Second, we also show results under the other three approaches (II to IV) and relying on the baseline calibration.

Figure A.1: Implied θ_k vs. Micro Benchmarks – \uparrow Real Rigidities



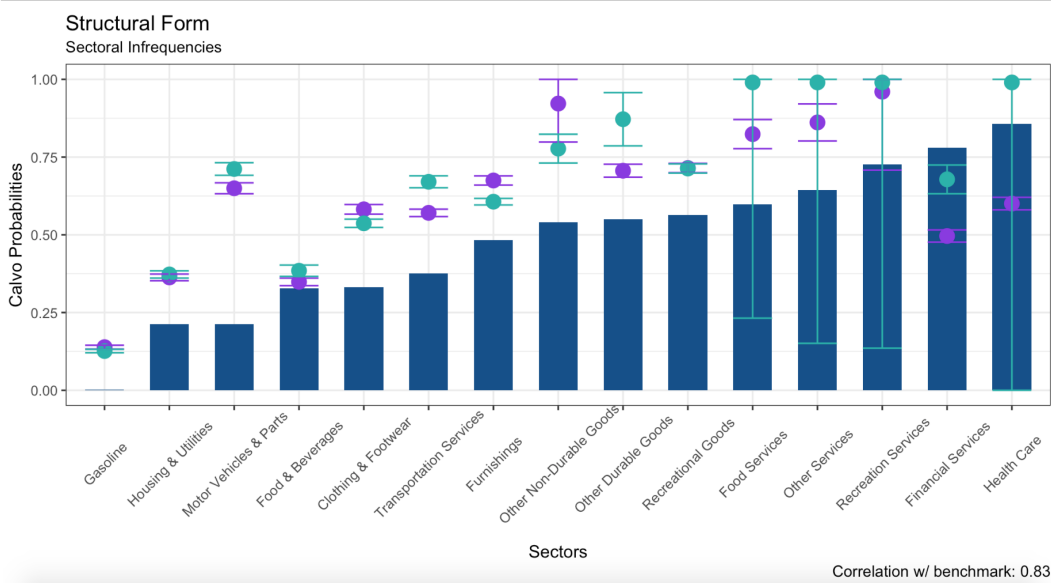
Notes: Estimated Calvo probabilities using SYS-GMM – based on (19) – and a HAC estimator for the covariance matrix. β and γ are also estimated. We set the calibration that implies more real rigidities – see Table 1.6. Results for normalisation 1 (2) are exhibited in purple (green). Blue bars are micro-based benchmark probabilities implied from evidence in Bils and P. Klenow (2004) and presented in Table 1.1 (appendix). For expository purposes, these are sorted by the degree of flexibility. 95% confidence intervals are also shown. Instrument selection follows approach I, the calibration used is the baseline and starting values are benchmark reset probabilities.

Figure A.2: Implied θ_k vs. Micro Benchmarks – \downarrow Real Rigidities



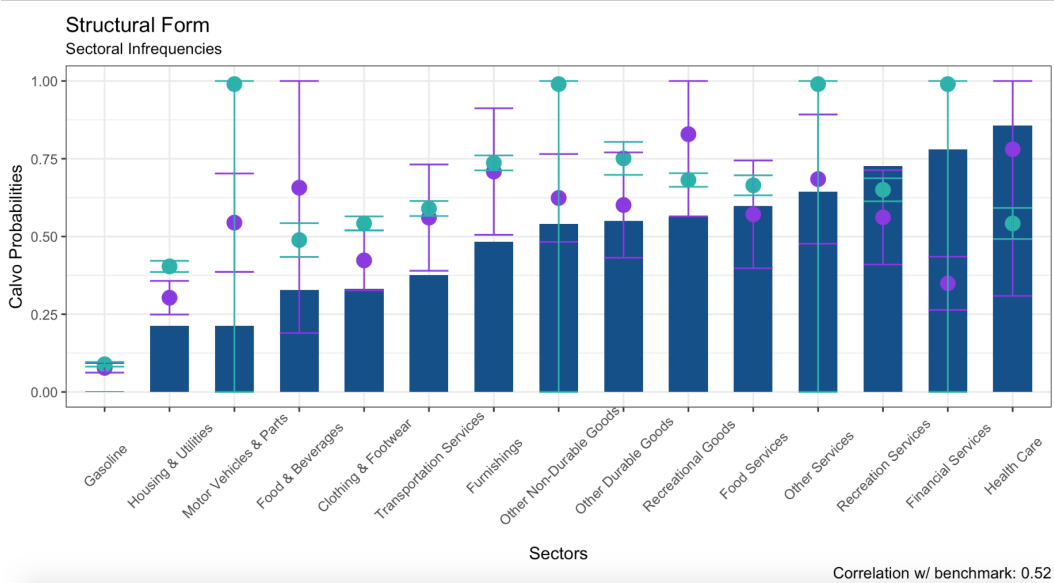
Notes: Estimated Calvo probabilities using SYS-GMM – based on (19) – and a HAC estimator for the covariance matrix. β and γ are also estimated. We set the calibration that implies less real rigidities – see Table 1.6. Results for normalisation 1 (2) are exhibited in purple (green). Blue bars are micro-based benchmark probabilities implied from evidence in Bils and P. Klenow (2004) and presented in Table 1.1 (appendix). For expository purposes, these are sorted by the degree of flexibility. 95% confidence intervals are also shown. Instrument selection follows approach I, the calibration used is the baseline and starting values are benchmark reset probabilities.

Figure A.3: Implied θ_k vs. Micro Benchmarks – Approach II



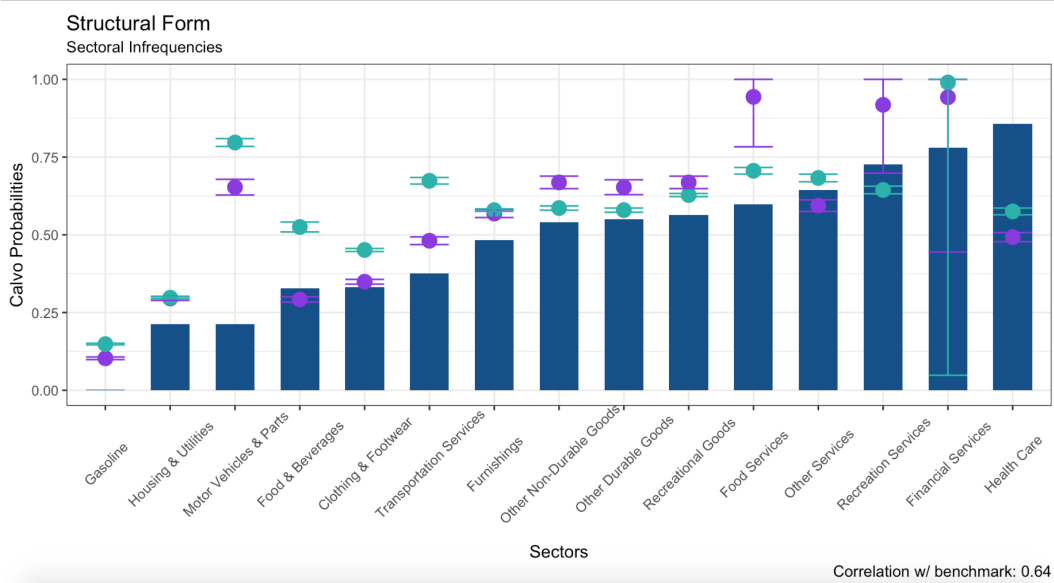
Notes: Estimated Calvo probabilities using SYS-GMM – based on (19) – and a HAC estimator for the covariance matrix. β and γ are also estimated. We rely on the baseline calibration in Table 1.6. Results for normalisation 1 (2) are exhibited in purple (green). Blue bars are micro-based benchmark probabilities implied from evidence in Bils and P. Klenow (2004) and presented in Table 1.1 (appendix). For expository purposes, these are sorted by the degree of flexibility. 95% confidence intervals are also shown. Instrument selection follows approach II, the calibration used is the baseline and starting values are benchmark reset probabilities.

Figure A.4: Implied θ_k vs. Micro Benchmarks – Approach III



Notes: Estimated Calvo probabilities using SYS-GMM – based on (19) – and a HAC estimator for the covariance matrix. β and γ are also estimated. We rely on the baseline calibration in Table 1.6. Results for normalisation 1 (2) are exhibited in purple (green). Blue bars are micro-based benchmark probabilities implied from evidence in Bils and P. Klenow (2004) and presented in Table 1.1 (appendix). For expository purposes, these are sorted by the degree of flexibility. 95% confidence intervals are also shown. Instrument selection follows approach III, the calibration used is the baseline and starting values are benchmark reset probabilities.

Figure A.5: Implied θ_k vs. Micro Benchmarks – Approach IV

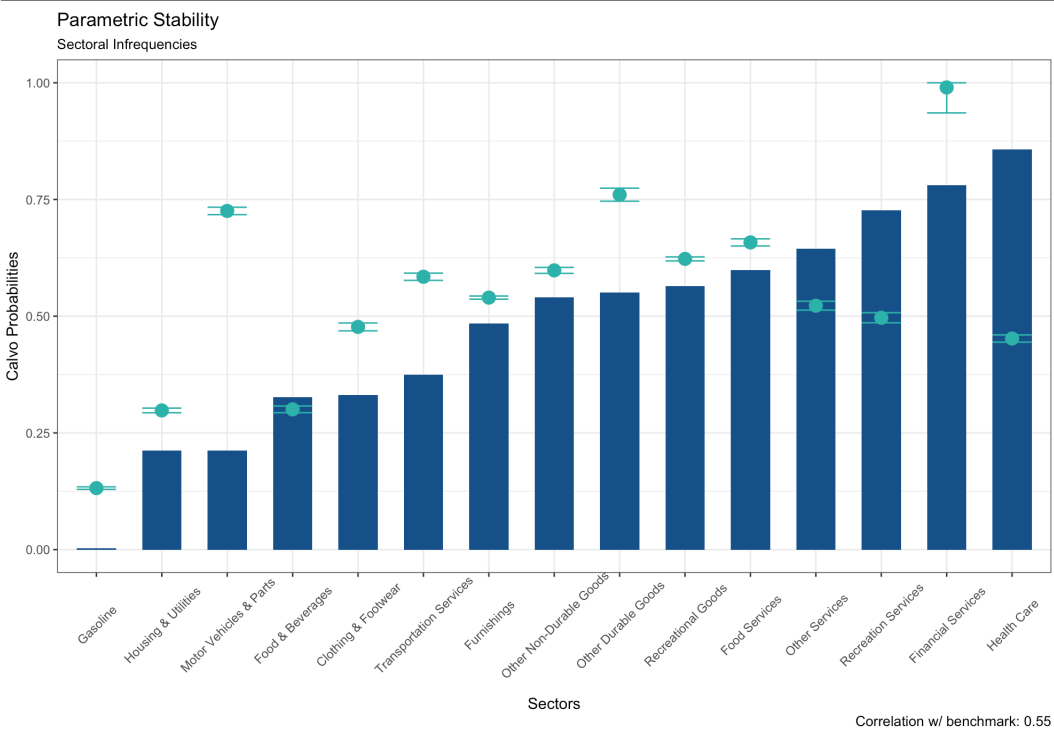


Notes: Estimated Calvo probabilities using SYS-GMM – based on (19) – and a HAC estimator for the covariance matrix. β and γ are also estimated. We rely on the baseline calibration in Table 1.6. Results for normalisation 1 (2) are exhibited in purple (green). Blue bars are micro-based benchmark probabilities implied from evidence in Bils and P. Klenow (2004) and presented in Table 1.1 (appendix). For expository purposes, these are sorted by the degree of flexibility. 95% confidence intervals are also shown. Instrument selection follows approach IV, the calibration used is the baseline and starting values are benchmark reset probabilities.

A.4.7
Parametric Stability

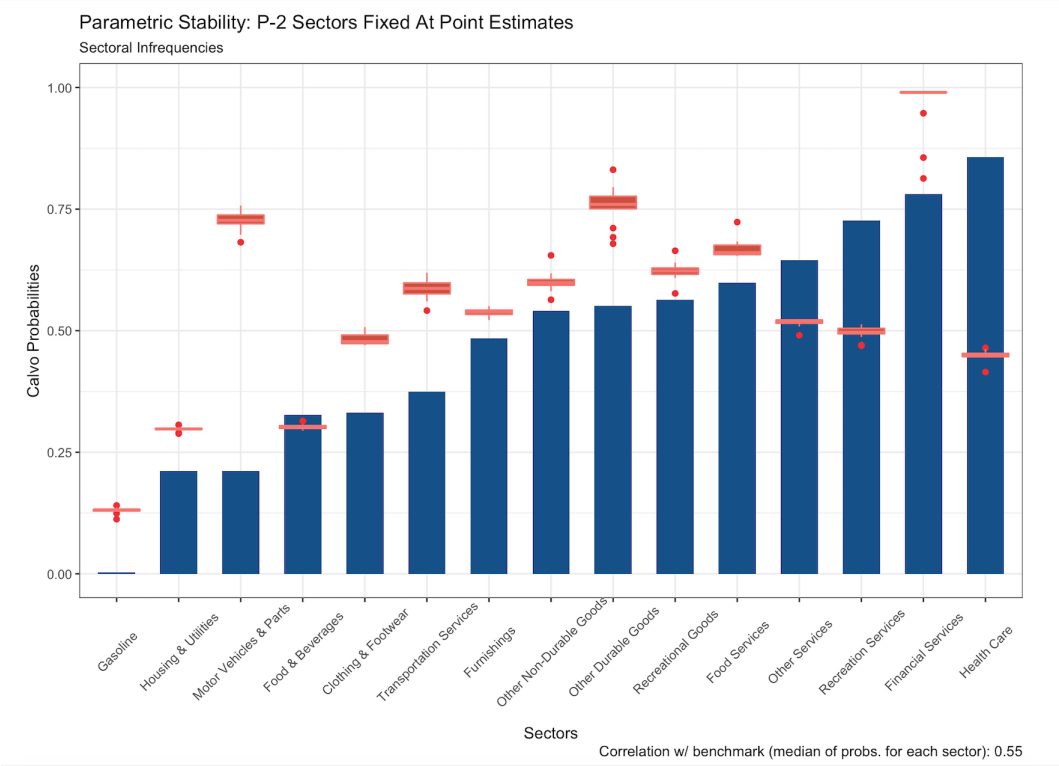
We show auxiliary charts that complement Figure 1.3. Specifically, here we exhibit results for cases when $q \in \{1, 2, 3\}$, the number of sectoral reset probabilities fixed in the estimation.

Figure A.6: Confidence Sets for Each λ_k (Estimated Individually)



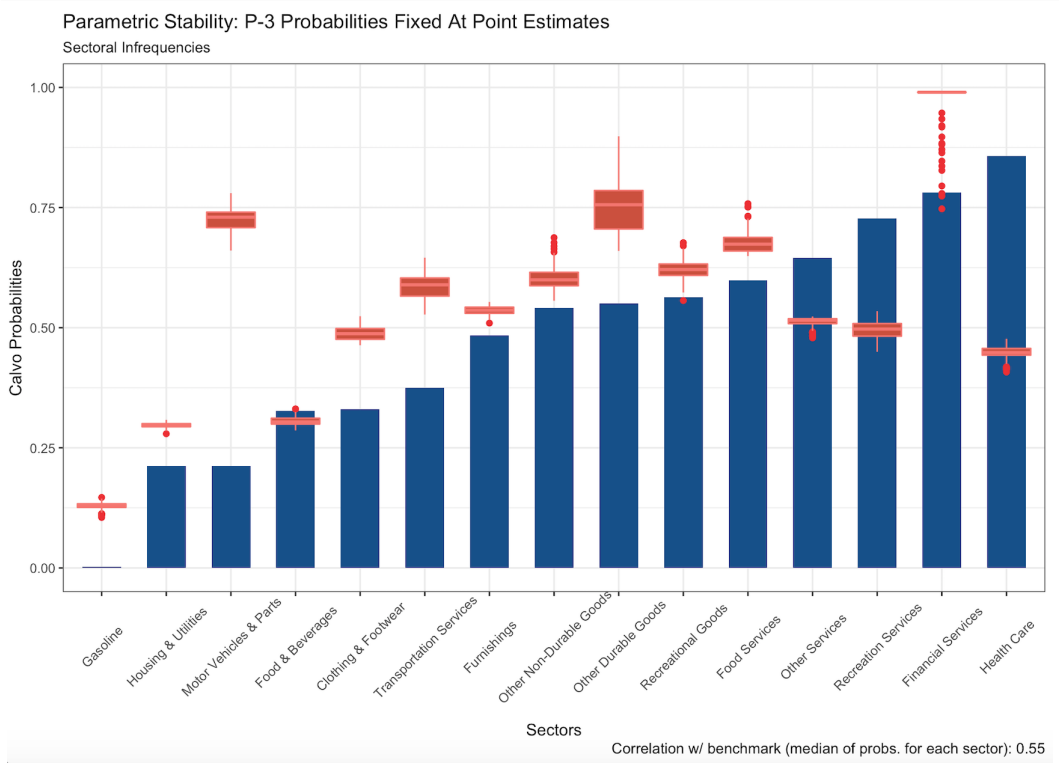
Notes: Parametric stability when 14 sectoral probabilities are fixed at point estimates (only one being estimated). Boxes represent the interval from the 25%ile to the 75%ile of distributions (for each sectoral probability). Horizontal lines are median estimates. Instrument selection follows approach I. Baseline calibration. Starting values: benchmark reset probabilities.

Figure A.7: CIs Constructed Based on Restricted Estimations



Notes: Parametric stability when 13 sectoral probabilities are fixed (two being estimated). Boxes represent the interval from the 25%ile to the 75%ile of distributions (for each sectoral probability). Horizontal lines are median estimates. 105 restricted versions of the model are estimated. Estimates for each of these versions are vertically positioned for each sector. Instrument selection follows approach I. Baseline calibration. Starting values: benchmark reset probabilities.

Figure A.8: CIs Constructed Based on Restricted Estimations

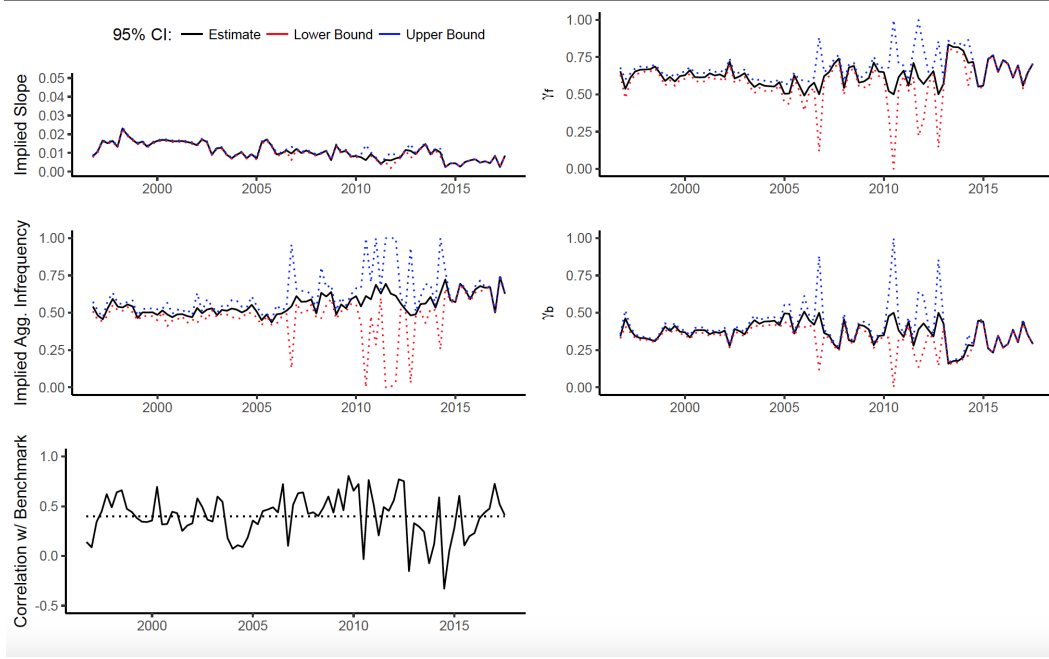


Notes: Parametric stability when 12 sectoral probabilities are fixed (three being estimated). Boxes represent the interval from the 25%ile to the 75%ile of distributions (for each sectoral probability). Horizontal lines are median estimates. 455 restricted versions of the model are estimated. Estimates for each of these versions are vertically positioned for each sector. Instrument selection follows approach I. Baseline calibration. Starting values: benchmark reset probabilities.

A.4.8

Rolling-GMM Estimations: Alternative Approach

Complementary results to Figure 1.4 and Figure 1.5, when we reduce the number of observations and moments conditions of the model. We use approach II and $T = 130$ observations.

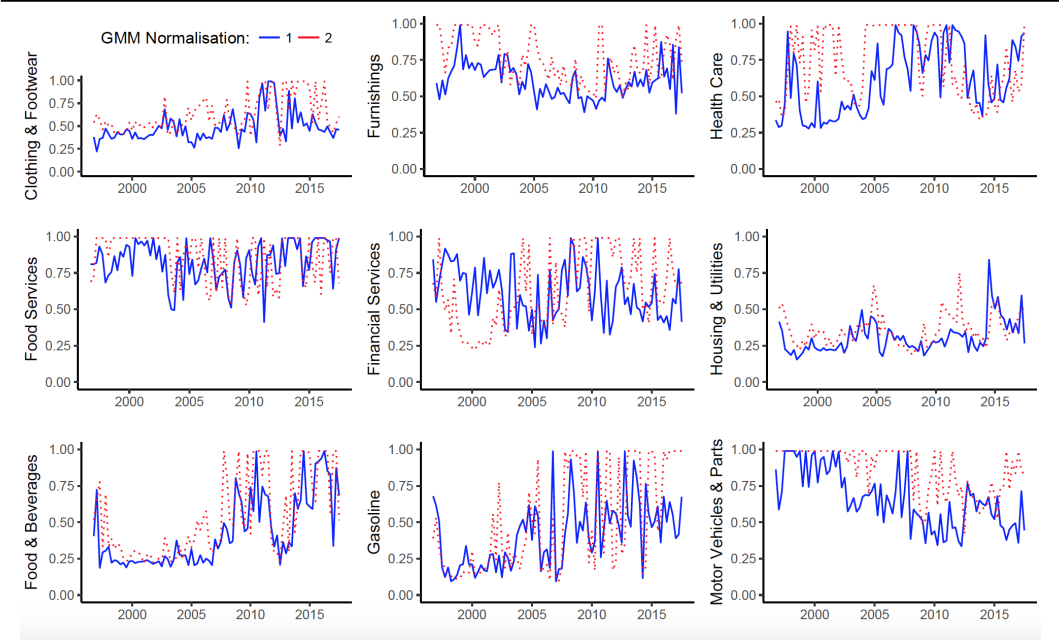
Figure A.9: Subsample Stability: Implied Coefficients ($T = 130$)

Notes: Rolling-GMM estimates using the system in (19) and a HAC estimator for the covariance matrix. Confidence intervals are constructed by Delta method. Instruments are re-selected (based on the approach II) at each step. The number of observations is $T = 130$. The horizontal axis measures the date of the last observation in the rolling subsample. “Implied aggregate infrequency” refers to (20). Results are for the normalisation (1) and starting values used are the benchmark reset probabilities based on micro data – see Bils and P. Klenow (2004) and Table 1.1 (appendix).

Figure A.10: Subsample Stability: Homogeneous vs. Heterogeneous ($T = 130$)

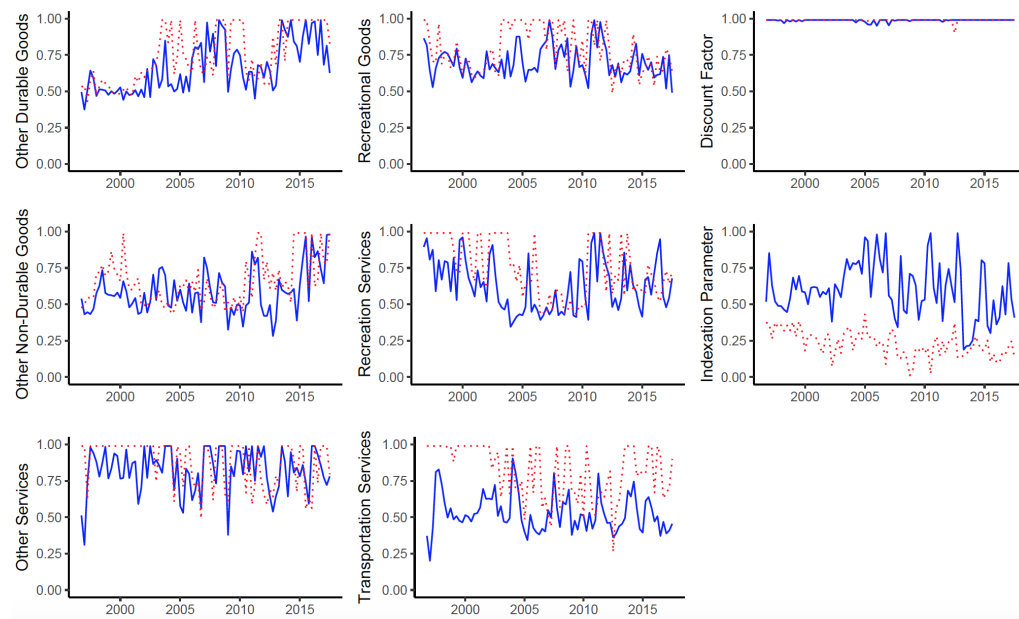
Notes: Rolling-GMM estimates using the NKPC of the homogeneous economy, (4), and the system for the heterogeneous economy, (19). The first column exhibits results for the homogeneous economy. “Implied Slope” refers either to the aggregate slope in (19) or to the one in (4). θ refers to (20). The horizontal axis measures the date of the last observation in the rolling subsample. Confidence intervals are constructed by Delta method. Instruments are re-selected (based on the approach II) at each step. Moment conditions are constructed using the baseline normalisations for both model (denoted as (1) in the previous tables). For the heterogeneous economy, starting values used are the benchmark reset probabilities based on micro data – see Bils and P. Klenow (2004) and Table 1.1 (appendix).

Figure A.11: Subsample Stability: Deep Parameters (Part I)



Notes: Rolling-GMM estimates using the NKPC of the homogeneous economy, (4), and the system for the heterogeneous economy, (19). Estimates of sectoral reset probabilities, $\hat{\lambda}_k$ are shown (we only exhibit the name of the sectors in the axes). The horizontal axis measures the date of the last observation in the rolling subsample. “Discount factor” and “indexation parameter” refer to β and γ , respectively. Confidence intervals are constructed by Delta method. Instruments are re-selected (based on the approach II) at each step. Moment conditions are constructed using the baseline normalisations for both model (denoted as (1) in the previous tables). For the heterogeneous economy, starting values used are the benchmark reset probabilities based on micro data – see Bils and P. Klenow (2004) and Table 1.1 (appendix).

Figure A.12: Subsample Stability: Deep Parameters (Part II)



Notes: Rolling-GMM estimates using the NKPC of the homogeneous economy, (4), and the system for the heterogeneous economy, (19). Estimates of sectoral reset probabilities, $\hat{\lambda}_k$ are shown (we only exhibit the name of the sectors in the axes). The horizontal axis measures the date of the last observation in the rolling subsample. “Discount factor” and “indexation parameter” refer to β and γ , respectively. Confidence intervals are constructed by Delta method. Instruments are re-selected (based on the approach II) at each step. Moment conditions are constructed using the baseline normalisations for both model (denoted as (1) in the previous tables). For the heterogeneous economy, starting values used are the benchmark reset probabilities based on micro data – see Bils and P. Klenow (2004) and Table 1.1.

A.4.9

Structural Estimations: Dropping the Aggregate NKPC in (19)

Additional robustness checks when the aggregate NKPC, (15), is dropped from the system in (19).

Table A.18: Implied Slope from (19) *Without* the Aggregate NKPC

Model	Normalisation	Calibration		
		Baseline	↑ Real Rigidity	↓ Real Rigidity
$\gamma = 0$	(1)	0.012***	0.017***	0.007***
		(0.000)	(0.000)	(0.000)
		[0.17]	[0.16]	[0.38]
	(2)	0.012***	0.007***	0.016***
		(0.000)	(0.000)	(0.000)
		[0.19]	[0.65]	[0.60]
$\gamma \neq 0$	(1)	0.011***	0.007***	0.015***
		(0.000)	(0.000)	(0.000)
		[0.78]	[0.84]	[0.81]
	(2)	0.005***	0.006***	0.007***
		(0.000)	(0.000)	(0.000)
		[0.45]	[0.81]	[0.50]

Notes: Implied slope from estimations of the system in (19) *without* the aggregate NKPC. We use SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and γ are allowed to vary. The instrument set is constructed based on approach I. Standard errors from Delta method are presented in parentheses. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in brackets. We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix). In addition, *p<0.1; **p<0.05; ***p<0.01.

Table A.19: Implied Stickiness from (19) *Without* the Aggregate NKPC

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters		
			θ^{agg}	β	γ
\uparrow Real Rigidity	(1)	0.84	0.47 (0.44 - 0.50)	0.99 (0.018)	0.40 (0.011)
	(2)	0.81	0.58 (0.54 - 0.62)	0.99 (0.000)	0.22 (0.004)
Baseline	(1)	0.78	0.63 (0.52 - 0.74)	0.99 (0.017)	0.37 (0.009)
	(2)	0.45	0.76 (0.66 - 0.87)	0.99 (0.000)	0.18 (0.001)
\downarrow Real Rigidity	(1)	0.81	0.59 (0.55 - 0.63)	0.99 (0.037)	0.50 (0.022)
	(2)	0.50	0.70 (0.68 - 0.72)	0.99 (0.000)	0.20 (0.000)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of the system in (19) *without* the aggregate NKPC. We use SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β , γ and each λ_k are allowed to vary in the range $(0, 1)$. The instrument set is constructed based on approach I. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

Table A.20: Implied Stickiness from (19) *Without* the Aggregate NKPC ($\gamma = 0$)

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters	
			θ^{agg}	β
\uparrow Real Rigidity	(1)	0.16	0.49	0.99
			(0.43 - 0.54)	(0.000)
	(2)	0.65	0.60	0.99
			(0.42 - 0.77)	(0.000)
Baseline	(1)	0.17	0.60	0.99
			(0.46 - 0.75)	(0.000)
	(2)	0.19	0.59	0.99
			(0.44 - 0.75)	(0.000)
\downarrow Real Rigidity	(1)	0.38	0.60	0.97
			(0.57 - 0.63)	(0.000)
	(2)	0.60	0.62	0.99
			(0.53 - 0.71)	(0.000)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of the system in (19) *without* the aggregate NKPC. We use SYS-GMM and a HAC estimator for the covariance matrix. Different calibrations of the model are used, described in Table 1.6. β and each λ_k are allowed to vary in the range (0, 1). We fix $\gamma = 0$ (forward-looking model). The instrument set is constructed based on approach I. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies ($1 - \lambda_k$) that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

A.4.10

Structural Estimations When $\Theta \approx 0.38$

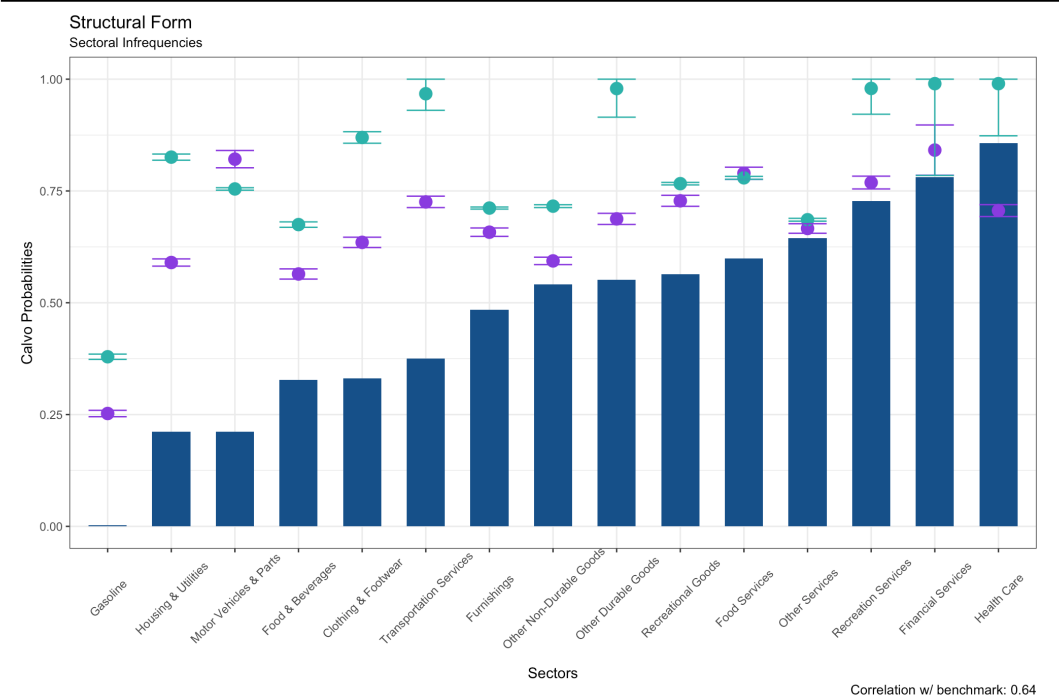
Here we present estimations when we impose $\epsilon = 5$, $\varphi = 1.5$ and $\sigma = 1$ ($\Theta \approx 0.38$). Such economy features much less real rigidities, and we should estimate a higher degree of stickiness from the model. Nonetheless, we find that $\hat{\theta}$ is still not as high as usually estimated in the literature. Most estimates lie inside the interval $(0.65, 0.70)$. Additionally, the estimated slope is not too different from that in our estimations in the main text (still very low and statistically significant at 1%).

Table A.21: Estimations of (19) With $\Theta \approx 0.38$

Model	Normalisation	$Corr(\theta_k, Micro)$	Parameters			Slope ^{agg}
			θ^{agg}	β	γ	
$\gamma = 0$	(1)	0.84	0.66	0.99	-	0.048***
			(0.62 - 0.70)	(0.001)		(0.000)
	(2)	0.81	0.69	0.99	-	0.047***
			(0.64 - 0.75)	(0.001)		(0.001)
$\gamma \neq 0$	(1)	0.64	0.66	0.96	0.49	0.034***
			(0.65 - 0.67)	(0.010)	(0.007)	(0.000)
	(2)	0.64	0.80	0.98	0.21	0.016***
			(0.78 - 0.82)	(0.003)	(0.001)	(0.000)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – from estimations of the system in (19) when we impose $\epsilon = 5$, $\varphi = 1.5$ and $\sigma = 1$ ($\Theta \approx 0.38$). We use SYS-GMM and a HAC estimator for the covariance matrix. β , γ and each λ_k are allowed to vary in the range $(0, 1)$ in the lower half of the table. The upper half gives estimates under the forward-looking model (fixing $\gamma = 0$). The instrument set is constructed based on approach I. Standard errors are presented in parentheses. We present 95% confidence intervals for θ obtained by Delta method. Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are benchmark reset probabilities implied from the micro evidence – Bils and P. Klenow (2004). See Table 1.1 (appendix).

Figure A.13: Implied θ_k vs. Micro Benchmarks ($\Theta \approx 0.38$)



Notes: Estimated Calvo-setting probabilities using SYS-GMM – based on (19) – and a HAC estimator for the covariance matrix. β and γ are also estimated. We fix $\epsilon = 5$, $\varphi = 1.5$ and $\sigma = 1$ ($\Theta \approx 0.38$). Results for normalisation 1 (2) are exhibited in purple (green). Blue bars are micro-based benchmark probabilities implied from evidence in Bils and P. Klenow (2004) and presented in Table 1.1 (appendix). For expository purposes, these are sorted by the degree of flexibility. 95% confidence intervals are also shown. Instrument selection follows approach I, the calibration used is the baseline and starting values are benchmark reset probabilities.

A.4.11
Structural Estimations: Sectoral NKPCs to Generate Initial Values

Robustness checks when starting values are obtained estimating sectoral NKPCs, individually, without instrument selection routines.

Table A.22: Estimations of (19): Initial Values Generated by Sectoral NKPCs

Calibration	Normalisation	$Corr(\theta_k, Micro)$	Parameters			
			θ^{agg}	β	γ	$Slope^{agg}$
↑ Real Rigidity	(1)	0.14	0.47 (0.44 - 0.50)	0.91 (0.001)	0.49 (0.001)	0.009*** (0.000)
	(2)	0.26	0.58 (0.54 - 0.62)	0.99 (0.000)	0.35 (0.001)	0.005*** (0.000)
Baseline	(1)	0.26	0.65 (0.65 - 0.66)	0.94 (0.001)	0.47 (0.001)	0.004*** (0.000)
	(2)	0.22	0.69 (0.65 - 0.74)	0.99 (0.000)	0.37 (0.001)	0.005*** (0.000)
↓ Real Rigidity	(1)	0.07	0.74 (0.74 - 0.74)	0.93 (0.001)	0.45 (0.001)	0.004*** (0.000)
	(2)	0.33	0.71 (0.70 - 0.71)	0.99 (0.000)	0.37 (0.000)	0.006*** (0.000)

Notes: Implied aggregate Calvo-pricing probability θ – see (20) – and aggregate slope from estimations of the system in (19). We use SYS-GMM and a HAC estimator for the covariance matrix. The baseline calibration in Table 1.6 is applied. β , γ and each λ_k are allowed to vary in the range (0, 1). The instrument set is constructed based on approach I. Standard errors are presented in parentheses, constructed by Delta method when applicable. We present 95% confidence intervals for θ . Correlations between estimated and benchmark infrequencies $(1 - \lambda_k)$ that come from the micro data in Bils and P. Klenow (2004) are shown in the column “ $Corr(\theta_k, Micro)$ ”. **The micro benchmark implies $\theta^{micro} \approx 0.48$.** We use two normalisations for moment conditions. These are (21) and (22) in the main text, shown respectively as (1) and (2) here for expository reasons. Starting values used in the algorithm are first-stage coefficients, $\hat{\lambda}_k$. These estimates are obtained by estimating each sectoral NKPC individually.

B

Appendix for EIS with Unfiltered Consumption

B.1

Data

B.1.1

Section 3

Returns data are compiled from the Kenneth French's online library, which uses CRSP data for stocks and the one-month Treasury bill rate (from Ibbotson Associates) for risk-free returns. We transform these series into real terms adjusting for CPI inflation (explained below). Adjustments described in the main paper are performed on each of these datasets – see more details below in this appendix. Ignoring observations lost by applying lagged instruments in the estimation, our initial dataset covers the period between 1931 to 2017 for annual and between 1947:3 to 2017:4 for quarterly data. Our original consumption data (or equivalently, reported consumption) come from NIPA tables, available on BEA's website. We use two time series: consumption of nondurables and services and consumption of nondurables (only). As explained in the main paper, unfiltered consumption has been constructed from these two series under distinct calibrations. For annual data, we construct unfiltered consumption from the original model in Kroencke (2017), which does not feature serially correlated measurement errors. For quarterly data, these are relevant, so that we introduce such form of persistence relying on the quasi-differenced filter described in section 2 of chapter 2. For more details on how we calibrate the model, see section "calibration" below.

For other series, inflation rate uses quarter-over-quarter and year-over-year CPI data, the nominal interest rate is the same used above (from Kenneth French's website) and the dividend-price ratio has been taken from Robert Shiller's online data source. Recall that we take logs of the latter.

B.1.2

Section 4

For stocks, we use value-weighted returns that consider NYSE, NASDAQ and AMEX. For T-bill returns, we rely on the same dataset of section 3

of chapter 2. To calculate the bond default premium, we use the Moody's Seasoned BAA Corporate Bond Yield (which is based on bonds with maturities of 20 years and above) as the long term corporate yield. To compute the bond horizon premium, we rely on 20-year and 1-month T-bill rates (Federal Reserve). From January 1984 to September 1993, the 20-year data are not available, so we use the 10-year analogue instead.

As indicated in the main paper, consumption growth data is constructed from the CEX interviews. These are deflated using the CPI deflator for nondurables considering urban households. To deflate return series mentioned above we use the CPI for total consumption (also for urban households).

B.1.3

A Few Notes on Data Sources involving Consumption, Frequencies and Measurement Errors

The CEX encompasses two major data sources, the interview and the diary surveys. They do not share the same sample. We use the former to construct our consumption data series for the groups of asset holders. Data in the diary survey is much more detailed, and likely more prone to measurement error and misreporting issues. However, the interview survey is also prone to such problems. Indeed, data that are probably misreported is easily verified for several of the questions used to construct consumption in section 4. For instance: households that report two distinct quantities for the consumption of the same item, in the same month of the same year; households that report negative values for some item; or even households that do not report consumption for some quarter (or, some interview), but that report numbers to other questions, not related to consumption (these households are included in our sample).

The BLS has been systematically attempting to change the methodology, so that respondent burden and measurement errors are less significant. The Gemini project of 2009, mentioned in the main paper, is an example.

B.1.3.1

Data Sources Comparison: NIPA vs. CEX

There are many methodological differences between consumption measured by the PCE (or NIPA, from BEA) and the CEX (from BLS).

First, consumption measured by the BLS takes into account that the data source varies with the sampling frequency of consumption data. For example, monthly and quarterly data in the PCE are based on the monthly retail trade survey (MRTS), while annual data comes from the annual retail trade survey

(ARTS). It is known that sampling errors are more problematic in the former, and the BEA takes this into account. In so far as we can tell, adjustments applied by the BLS are not specific to any major data source in the CEX. Hence, even though the interview and the diary survey have distinct samples, the Filter model could be also applied to the latter.

Second, the BEA benchmarks quarterly and annual data to “the best available source data”, which happens to be the quinquennial – see BEA (2017). It follows that quarterly and annual frequency data are interpolated, so that they are compatible with benchmark years. In a second step, quarterly estimates are benchmarked to their annual counterparts. Bell and Wilcox (1993) argue that benchmark procedures reduce measurement errors, inherently affecting the autocorrelation in the consumption series. Contrasting with the PCE, the CEX does not benchmark the data. Instead, the survey is continuously redesigned to circumvent issues of measurement error.

Imputation is present in both the PCE and the CEX data. BEA and BLS rely on statistical models for non-response to predict missing values. These are considerably in the ARTS, about 8%. The CEX did not apply imputation to asset data, but began to do in 2004. Allocation routines are also common across both sources. The CEX applies tabulation corrections *before* and *after* other adjustment routines.

Residual methods are applied by the BEA to measure the consumption of some categories. They use “residuals” from government expenditures to do so. As far as we can tell, these routines are absent in the CEX.

Finally, the BLS applies smoothing techniques over the data. Direct forms of smoothing are absent in the CEX, albeit topcoding routines are present, generating similar effects. Topcoding techniques modify the consumption, positions in assets and the income data of outliers, so that these can not be identified from the public micro-files. Thresholds applied are also constantly under revision, aiming to correctly disentangle “true” outliers from misrespondents or coding errors.

B.2

Completely-Specified Quasi-Differenced Filter Model

As mentioned before, when we referred to annual consumption in section 3 of the main text we were using the canonical filter model in Kroenke (2017). The only difference in that case relates to calibration (since different time series are used). Specifically, measurement errors followed a simple white noise stochastic process with no form of persistence introduced. The quasi-differenced filter model (that accounts for such persistence) is only used when

handling quarterly data in section 3, and semi-annual data (but at monthly frequency), in section 4. This follows for the reasons described in the main text. In this section, we provide the complete specification of this quasi-differenced version and better detail how we modify the original Filter model in Kroencke (2017).

Recall that a Filter model without persistence in measurement error is not suitable for data at monthly or quarterly frequencies. The more frequent the data (or the smaller the level of aggregation) the more likely accounting for a serially correlated error becomes essential since sampling errors are probably autocorrelated – see the online appendix of Kroencke (2017). In our estimations, unfiltered consumption performed poorly in terms of estimates of the EIS when serially correlated measurement error terms are not considered in the model and the data frequency is either monthly or quarterly. Comparatively, even reported consumption provided more precise estimates in those cases.

Turning to the model, assume a simple state-space representation:

$$x_{t+1} = Fx_t + R\eta_{t+1}, \quad (1)$$

$$y_t = Hx_t + \xi_t, \quad (2)$$

those representing the state and measurement equations of a simple Kalman filter, respectively. x_t represents a vector of state variables and the last term is its corresponding disturbance. y_t is observed consumption (it can be the garbage measure in Savov (2011), for instance). Unfiltered consumption is our estimate of this time series, while ξ_t represents measurement errors.

By permitting serially correlated measurement error terms we introduce (3) of the main text in the system:

$$\xi_t = \rho_\xi \xi_{t-1} + \nu_t,$$

where ν_t is a simple white noise process. Generally, the usual assumption is $E[\xi_t \eta'_{t+h}] = E[\eta_t \eta'_s] = E[\xi_t \xi'_s] = 0$, for $t \neq s$ and $h > 0$ – see James Douglas Hamilton (1994), for example. We relax this hypothesis in order to introduce the possibility of (3). Particularly, let us assume that the innovation of state consumption and that of measurement errors are conditionally normally distributed:

$$\begin{bmatrix} \eta_{t+1} \\ \xi_t \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} R\sigma_{\eta,t+1}^2 R' & R\omega_{\eta,\nu} \\ \omega_{\nu,\eta} R' & \sigma_\nu^2 \end{bmatrix} \right), \quad (3)$$

where we are not assuming any zeros in the covariance matrix, but we allow for a time-varying element in its upper left corner.

Next, define:

$$P_{t+1|t} = E[(x_{t+1} - \hat{x}_{t+1|t})(x_{t+1} - \hat{x}_{t+1|t})' | \mathcal{F}_t], \quad (4)$$

where the information set \mathcal{F}_t tracks data realisations conditional on period t and $P_{t+1|t}$ is the covariance of prediction errors conditional on the same period (*a priori*). With a small abuse of notation, we denoted its first element (relative to consumption) by P_t^c in section 2 of chapter 2 – we turn back to this below. In addition, $\hat{x}_{t+1|t} = E[x_{t+1} | \mathcal{F}_t]$, as usually. Note that $x_{t+1|t-1} \sim N(\hat{x}_{t+1|t-1}, P_{t+1|t-1})$, and:

$$\hat{x}_{t+1|t-1} = F\hat{x}_{t|t-1}, \quad (5)$$

by (1). In a similar vein and using $E[x_t \eta'_{t+1}] = E[\hat{x}_{t|t-1} \eta'_{t+1}] = 0$:

$$P_{t+1|t-1} = E[(F(x_t - \hat{x}_{t|t-1}) + R\eta_{t+1})(F(x_t - \hat{x}_{t|t-1}) + R\eta_{t+1})' | \mathcal{F}_{t-1}] = FP_{t|t-1}F' + R\sigma_{\eta,t+1}^2 R', \quad (6)$$

Likewise:

$$\hat{y}_{t|t-1} = H\hat{x}_{t|t-1}, \quad S_{t|t-1} = HP_{t|t-1}H' + \sigma_\nu^2, \quad (7)$$

where $S_{t|t-1} = E[(y_t - \hat{y}_{t|t-1})(y_t - \hat{y}_{t|t-1})' | \mathcal{F}_{t-1}]$, the covariance of pre-fit prediction errors. Finally, if $\Sigma_{t|t-1} = E[(y_t - \hat{y}_{t|t-1})(x_{t+1} - \hat{x}_{t+1|t})' | \mathcal{F}_{t-1}]$ is the cross-correlation matrix between state and observable variables, then:

$$\Sigma_{t|t-1} = FP_{t|t-1}H' + \omega_{\nu,\eta}R', \quad (8)$$

so that:

$$\begin{bmatrix} x_{t+1} | \mathcal{F}_{t-1} \\ y_t | \mathcal{F}_{t-1} \end{bmatrix} \sim N \left(\begin{bmatrix} F\hat{x}_{t|t-1} \\ H\hat{x}_{t|t-1} \end{bmatrix}, \begin{bmatrix} P_{t+1|t-1} & \Sigma_{t|t-1} \\ \Sigma'_{t|t-1} & S_{t|t-1} \end{bmatrix} \right). \quad (9)$$

From (8), one can thus express the distribution of $x_{t+1} | \mathcal{F}_t$ by marginalising $x_{t+1} | \mathcal{F}_{t-1}$ in terms of $y_t | \mathcal{F}_{t-1}$. Hence, relying on the multivariate normal:

$$\begin{aligned} \hat{x}_{t+1|t} &= F\hat{x}_{t|t-1} + \Sigma_{t|t-1}S_{t|t-1}^{-1}(y_t - H\hat{x}_{t|t-1}) \\ &= F\hat{x}_{t|t-1} + (FP_{t|t-1}H' + R\omega_{\eta,\nu})(HP_{t|t-1}H' + \sigma_\nu^2)^{-1}(y_t - H\hat{x}_{t|t-1}), \end{aligned} \quad (10)$$

using (6) and (7). In a similar fashion:

$$\begin{aligned} P_{t+1|t} &= P_{t+1|t-1} - \Sigma_{t|t-1}S_{t|t-1}^{-1}\Sigma'_{t|t-1} \\ &= FP_{t|t-1}F' + R\sigma_{\eta,t+1}^2 R' - (FP_{t|t-1}H' + R\omega_{\eta,\nu})(HP_{t|t-1}H' + \sigma_\nu^2)^{-1}(FP_{t|t-1}H' + \omega_{\eta,\nu}R')' \end{aligned} \quad (11)$$

Finally, the Kalman gain is simply:

$$K_t = \Sigma_{t|t-1} S_{t|t-1}^{-1} = (F P_{t|t-1} H' + R \omega_{\eta, \nu}) (H P_{t|t-1} H' + \sigma_\nu^2)^{-1} \quad (12)$$

One can derive the original filter in Kroencke (2017) with equations (9-11). However, up to this point, we have not allowed for serially correlated measurement errors – equation (3) of the main text – yet. Perhaps the easiest way to introduce this possibility in a Kalman filter is to expand the model so that measurement errors are defined as a new state variable. Harvey, Ruiz, and Sentana (1992) present different forms of modelling that while still relying on ARCH or GARCH processes to express variances (as we do here). However, note that we actually aim to *invert* a Kalman filter. This is, we are not interested in tracking state variables given observables. Instead, we aim to infer what would those observables were once we (researchers) *only* know estimates of state variables that supposedly took those into account to be constructed. As in Kroencke (2017), we do *not* use simple recursions of a Kalman filter but instead their reverse counterparts. With that in mind, re-scaling the system is much simpler than developing alternative methods of “reverse engineering” that support an expanded system that includes (3). We then opt to re-scale the original Kalman filter above, expressing it in terms of a quasi-differenced system. This re-scaled filter does not modify the standard interpretation given in Kroencke (2017).

In the following, we blend findings of E. Anderson et al. (1996) with the original Filter model, establishing our quasi-differencing approach. Assume that $\nu_t \sim N(0, \sigma_\nu^2)$, the error term in equation (3). Next, we define observables in terms of a quasi-difference:

$$\bar{y}_t \equiv y_{t+1} - \rho_\xi y_t, \quad (13)$$

where \bar{y}_t is referred here as “quasi-differenced observable consumption”. Note that the state equation (1) does not change with this modification, but the measurement equation is transformed into¹:

$$\bar{y}_t = (HF - \rho_\xi H)x_t + HR\eta_{t+1} + \nu_{t+1} \equiv \bar{H}x_t + \bar{\xi}_t \quad (14)$$

It is worth mentioning that by rewriting the Filter model in terms of a quasi-differenced observable we are *not* assuming that the raw data first observed by official statisticians is \bar{y}_t . Instead, *we are only re-scaling the system in order to solve it, to then mapping back \bar{y}_t onto y_t* . This will probably become more clear below.

¹Write the measurement equation one period forward. Use the state equation in x_{t+1} . Then, subtract $\rho_\xi \times y_t$ from this, using the measurement equation in the current period as y_t .

By following similar developments, we can rewrite covariances in terms of the new composite error term in (13). Particularly, if $R\eta_{t+1} \equiv \bar{\eta}_{t+1}$, it follows that $\sigma_{\bar{\xi}}^2 = HR\sigma_{\eta,t+1}^2R'H' + \sigma_{\nu}^2$ and $\omega_{\bar{\eta},\bar{\xi}} = R\sigma_{\eta,t+1}^2R'H'$. Applying the same algebra as above:

$$\begin{aligned}\hat{x}_{t+1|t} = & F\hat{x}_{t|t-1} \\ & + (FP_{t|t-1}\bar{H}' + R\sigma_{\eta,t+1}^2R'H')(\bar{H}P_{t|t-1}\bar{H}' + HR\sigma_{\eta,t+1}^2R'H' + \sigma_{\nu}^2)^{-1}(\bar{y}_t - \bar{H}\hat{x}_{t|t-1})\end{aligned}\quad (15)$$

$$\begin{aligned}P_{t+1|t} = & FP_{t|t-1}F' + R\sigma_{\eta,t+1}^2R' \\ & - (FP_{t|t-1}\bar{H}' + R\sigma_{\eta,t+1}^2R'H')(\bar{H}P_{t|t-1}\bar{H}' + HR\sigma_{\eta,t+1}^2R'H' + \sigma_{\nu}^2)^{-1} \\ & (\bar{H}P_{t|t-1}F' + HR\sigma_{\eta,t+1}^2R'),\end{aligned}\quad (16)$$

with the Kalman gain vector following:

$$K_t = (FP_{t|t-1}\bar{H}' + R\sigma_{\eta,t}^2R'H')(\bar{H}P_{t|t-1}\bar{H}' + HR\sigma_{\eta,t}^2R'H' + \sigma_{\nu}^2)^{-1} \quad (17)$$

Note that $[\bar{y}_t, \bar{y}_{t-1}, \dots, \bar{y}_0, \hat{x}_0]$ and $[y_{t+1}, y_t, \dots, y_0, \hat{x}_0]$ span the same space since:

$$\begin{aligned}\bar{y}_t - E[\bar{y}_t|\bar{y}_{t-1}, \dots, \bar{y}_0, \hat{x}_0] &= (y_{t+1} - Dy_t) \\ & - E[y_{t+1} - Dy_t|y_t - Dy_{t-1}, y_{t-1} - Dy_{t-2}, \dots, y_0, \hat{x}_0] \\ &= y_{t+1} - Dy_t + Dy_t - E[y_{t+1}|y_t, y_{t-1}, y_{t-2}, \dots, y_0, \hat{x}_0] \\ &= y_{t+1} - E[y_{t+1}|y_t, y_{t-1}, y_{t-2}, \dots, y_0, \hat{x}_0]\end{aligned}\quad (18)$$

Equations (14-16) express the algorithm in the quasi-differenced Kalman filter. The next step is to invert those equations to isolate quasi-differenced unfiltered consumption. Before, let's identify our system in terms of our findings above. State variables are defined as:

$$x_t = \begin{bmatrix} c_t \\ c_{t-1} \\ \eta_t^* \end{bmatrix}, \quad F = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad R = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \quad \eta_t^* = \sigma_{\eta,t}^2\eta_t, \quad (19)$$

while we assume, as in Kroencke (2017), that its variance follows a GARCH(1,1) stochastic process given by:

$$\sigma_{\eta,t}^2 = a_0 + a_1\eta_{t-1}^{*2} + a_2\sigma_{\eta,t-1}^2 \quad (20)$$

We can express the variance of prediction errors as:

$$\begin{aligned} S_{t|t-1} &= \begin{bmatrix} 1 - \rho_\xi & 0 & 0 \end{bmatrix} P_{t|t-1} \begin{bmatrix} 1 - \rho_\xi \\ 0 \\ 0 \end{bmatrix} + \sigma_{\eta,t}^2 + \sigma_\nu^2 \\ &= (1 - \rho_\xi)^2 P_{t|t-1}^c + \sigma_{\eta,t-1}^2 + \sigma_\nu^2, \end{aligned} \quad (21)$$

where $P_{t|t-1}^c$ denotes the element (1,1) – related to the state variable c_t – of covariance matrix $P_{t|t-1}$, and we have used the fact that $H = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}$ and then $\bar{H} = \begin{bmatrix} 1 - \rho_\xi & 0 & 0 \end{bmatrix}$. One can then obtain the first component of the Kalman gain in (16):

$$\begin{aligned} K_t &= \left\{ \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \sigma_{\eta,t}^2 + \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} P_{t|t-1} \begin{bmatrix} 1 - \rho_\xi \\ 0 \\ 0 \end{bmatrix} \right\} \\ &\quad [(1 - \rho_\xi)^2 P_{t|t-1}^c + \sigma_{\eta,t}^2 + \sigma_\nu^2]^{-1}, \end{aligned} \quad (22)$$

where the matrix pre-multiplying $P_{t|t-1}$ is F . Solving (22) for its first element, we obtain the Kalman gain – relative to consumption – described in the main text, equation (6):

$$K_t^c = \frac{(1 - \rho_\xi) P_{t|t-1}^c + \sigma_{\eta,t}^2}{(1 - \rho_\xi)^2 P_{t|t-1}^c + \sigma_{\eta,t}^2 + \sigma_\nu^2},$$

and in section 2 we used the notation P_t^c instead than $P_{t|t-1}^c$.

Next, use (16), writing it in terms of our model:

$$\begin{aligned} P_{t|t-1} &= \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} P_{t-1|t-2} \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \sigma_{\eta,t}^2 \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 1 \end{bmatrix} \\ &\quad - K_{t-1} \left\{ \begin{bmatrix} 1 - \rho_\xi & 0 & 0 \end{bmatrix} P_{t-1|t-2} \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 1 \end{bmatrix} \sigma_{\eta,t}^2 \right\}, \end{aligned} \quad (23)$$

whose first element is:

$$P_{t|t-1}^c = P_{t-1|t-2}^c (1 - (1 - \rho_\xi) K_{t-1}^c) + (1 - K_{t-1}^c) \sigma_{\eta,t}^2,$$

Recall that $P_t^c \equiv P_{t|t-1}^c$, in our notation of section 2.

Our quasi-differenced Filter model is almost completely described now.

The next step is to obtain our best guess of \bar{y}_{t-1} . In order to do that, write (15) for the update phase ($\hat{x}_{t|t}$). Substituting relevant matrices (19), the first element of $\hat{x}_{t|t}$ is:

$$\hat{c}_t = \hat{c}_{t-1} + K_t^c(\bar{y}_{t-1} - (1 - \rho_\xi)\hat{c}_{t-1}), \quad (24)$$

where $\hat{c}_t \equiv E[c_t|\mathcal{F}_t]$. Isolating \bar{y}_{t-1} :

$$\hat{\bar{y}}_{t-1} = \underbrace{\frac{\hat{c}_t - (1 - (1 - \rho_\xi)K_t^c)\hat{c}_{t-1}}{K_t^c}}_{\text{Unfiltered Quasi-Differenced Consumption}} \quad (25)$$

Equation (25) is *not* ready to be mapped back onto original unfiltered consumption yet. We still need a few adjustments, described in the following.

B.2.1

Adjusting Unfiltered Consumption for Time-Aggregation Bias

First, we adapt (25), so that it accounts for time-aggregation bias. R. Hall (1988) addressed this using an AR(2) representation. In a similar fashion and adapting adjustments in Kroencke (2017) to your model, we have²:

$$\Delta c_t^{adTA} = \frac{[\Delta c_t^{TA} - (1 - (1 - \rho_\xi)\alpha)\Delta c_{t-1}^{TA}]}{\alpha}, \quad (26)$$

where Δc_t^{adTA} denotes the time-aggregation bias-adjusted estimate of consumption growth and Δc_t^{TA} represents its time-aggregated counterpart. The parameter α guarantees that the second moment of Δc_t^{adTA} is the same of point-to-point consumption. It is set to the same value of Kroencke (2017), 0.8³. Adapting (25) to our model gives:

$$\hat{\bar{y}}_{t-1} = \frac{\hat{c}_t - (1 - (1 - \rho_\xi)\Omega_t)\hat{c}_{t-1}}{\Omega_t}, \quad (27)$$

where $\Omega_t = \alpha K_t$.

²See Kroencke (2017) for more details.

³It solves $Var(\Delta c_t^{TA}) = \frac{\alpha}{2-\alpha}Var(\Delta c_t^{adTA}) = \frac{2}{3}Var(\Delta c_t^{adTA})$. The former equality is implied from (34) setting $\rho_\xi = 0$ while the latter uses results in Working (1960) and Breeden, Gibbons, and Litzenberger (1989). The approximation for ρ_ξ does not distort results sensibly since we have set that parameter to a value very close to zero in our estimations. The first equality does not change in comparison with Kroencke (2017) since our model still relies on a random-walk representation for state consumption whose conditional variances are modelled through a GARCH(1,1).

B.2.2

Adjusting the Timing of Asset Returns

By construction, the variable in (27) has its second moment perfectly compatible with consumption measured point-to-point in time. However, the timing of asset returns is misaligned. We correct for this in section 3 by adapting adjustments in Kroencke (2017) – which was based on Cochrane (1996) – for the use of quarterly series.

First, we sum end-of-month levels $\Pi_{i,m,t+1}$ to obtain a measure of quarterly time-aggregated stock returns:

$$\Delta R_{i,t+1}^{TA} = \frac{\sum_{m=1}^3 \Pi_{i,m,t+1}}{\sum_{m=1}^3 \Pi_{i,m,t}} - 1, \quad (28)$$

where i represents the asset class and m is the corresponding month of quarter t .

Second, we bring first and second moments of this series back to point-to-point counterparts to make it compatible with (27). We conduct this adjustment using returns for the last quarter of the year. The motivation is in Jagannathan and Wang (2007), who argued that investors are more prone to adjust their investment portfolios in the fourth quarter of the year.

$$\Delta R_{i,t+1}^{adTA} = \frac{\Delta R_{i,t+1}^{TA} - E(\Delta R_{i,t+1}^{TA})}{\sigma(\Delta R_{i,t+1}^{TA})} \sigma(\Delta R_{i,t+1}^{Q4-Q3}) + E(\Delta R_{i,t+1}^{Q4-Q3}), \quad (29)$$

where $\Delta R_{i,t+1}^{Q4-Q3} = \Pi_{i,12,t+1}/\Pi_{i,9,t+1}$, returns measured for the last quarter⁴.

Step (27) is not necessary in section 4 since data were at monthly frequency. This simply implies that $\Delta R_{i,t+1}^{TA}$ is equal to (raw) monthly returns used. When it comes to (27), we do the following, for section 4:

$$\Delta R_{i,t+1}^{adTA} = \frac{\Delta R_{i,m+1}^{TA} - E(\Delta R_{i,m+1}^{TA})}{\sigma(\Delta R_{i,m+1}^{TA})} \sigma(\Delta R_{i,December}) + E(\Delta R_{i,December}), \quad (30)$$

where sub-index “December” denotes monthly returns in December of the relevant year. Results of section 4 are exactly the same if we use October or November instead.

⁴Kroencke (2017) conducted a similar adjustment but for annual series, using the first two moments of December-to-December consumption growth as his correction in a similar fashion. Note that by using moments of the last quarter components $E(\Delta R_{i,t+1}^{Q4-Q3})$ and $\sigma(\Delta R_{i,t+1}^{Q4-Q3})$ change every 4 observations (or equivalently, every 4 quarters) - contrasting with corrections in Kroencke (2017), that change for each observation (since described in annual terms). This could imply an unnecessary persistence for the series $\Delta R_{i,t+1}^{adTA}$. However, by comparing return series generated by (28) and (29) with their raw analogues – and repeating the same experiment for annual series, using the method in Kroencke (2017) – we have found that the impact of those modifications for our estimates is minimal. Hence, we evaluate that our adjustments for quarterly data do not perform considerably different from those of the original model. For complete results involving raw returns data, see section 6.4.5 below.

Finally, $\Delta R_{i,t+1}^{adTA}$ in (29) is aggregated to represent semi-annual returns as described in the main text – this is, we used $R_{i,t+1}$ instead of $\Delta R_{i,t+1}^{adTA}$ for simplicity reasons in section 4.

B.2.3

Mapping Unfiltered Quasi-Differenced Consumption Back onto Unfiltered Consumption

The last step is the simplest one. In order to map \hat{y}_{t-1} back onto \hat{y}_t – unfiltered consumption – we perform:

$$\underbrace{\hat{y}_t}_{\text{Unfiltered Consumption}} = \hat{y}_{t-1} + \rho_\xi \hat{y}_{t-1} \quad (31)$$

The fact that \hat{y}_{t-1} appears in the right hand side of (30) implies we can not identify y_0 . We turn back to this and how we initialise the model below.

B.2.4

Consumption Volatility

Motivated by results in Harvey, Ruiz, and Sentana (1992), we use the approximation $\eta_{t-1}^{*2} \approx E_{t-1}(\eta_{t-1}^{*2})$. It follows that $E_t(\eta_t^{*2})$ – and hence $E_{t-1}(\eta_{t-1}^{*2})$ – is obtained by:

$$\begin{aligned} E_t(\eta_t^{*2}) &= P_{t|t}^\eta + \eta_{t|t}^{*2} \\ &= \left(1 - \frac{\sigma_{\eta,t-1}^2}{P_{t|t-1}^c (1 - \rho_\xi)^2 + \sigma_{\eta,t-1}^2 + \sigma_\nu^2} \right) \sigma_{\eta,t-1}^2 \\ &\quad + \left(\frac{\sigma_{\eta,t-1}^2}{P_{t|t-1}^c (1 - \rho_\xi)^2 + \sigma_{\eta,t-1}^2 + \sigma_\nu^2} \right)^2 u_t^2, \end{aligned} \quad (32)$$

where $u_t \equiv \bar{y}_{t-1} - (1 - \rho_\xi)\hat{c}_{t-1}$, the re-scaled prediction error.

B.2.5

Homoscedastic Counterpart and Proof of Proposition 1

It is simple to derive a version of the model that features homoscedasticity in state consumption. This not only implies $\sigma_{\eta,t}^2 = \bar{\sigma}_\eta^2$, but also $K_t^c = \bar{K}^c$ and $P_{t|t-1}^c = \bar{P}^c$.

From (7) in the main text, we have:

$$\bar{P}^c = \bar{P}^c (1 - (1 - \rho_\xi)\bar{K}^c) + (1 - \bar{K}^c)\bar{\sigma}_\eta^2 \quad (33)$$

And from (6), also from the main paper:

$$\bar{K}^c = \frac{(1 - \rho_\xi)\bar{P}^c + \bar{\sigma}_\eta^2}{(1 - \rho_\xi)^2\bar{P}^c + \bar{\sigma}_\eta^2 + \sigma_\nu^2}, \quad (34)$$

By plugging (34) in (33) and after some algebraic manipulations, one can finally find the second order equation:

$$\{\bar{P}^c\}^2 + \bar{\sigma}_\eta^2 \left(\frac{1 + \rho_\xi}{1 - \rho_\xi} \right) \bar{P}^c - \frac{\bar{\sigma}_\eta^2 \sigma_\nu^2}{(1 - \rho_\xi)^2} = 0 \quad (35)$$

After rewriting terms, it is possible to show that the roots of (35) are given by:

$$\bar{P}^c = \frac{\bar{\sigma}_\eta^2}{2(1 - \rho_\xi)} \left\{ \pm \left[(1 - \rho_\xi)^2 \bar{\sigma}_\eta^2 + 4\sigma_\nu^2 \right]^{\frac{1}{2}} - (1 + \rho_\xi) \bar{\sigma}_\eta^2 \right\}, \quad (36)$$

so that the only sensible solution is (recall that $1 - \rho_\xi > 0$):

$$\bar{P}^c = \frac{\bar{\sigma}_\eta^2}{2(1 - \rho_\xi)} \left\{ \left[(1 - \rho_\xi)^2 \bar{\sigma}_\eta^2 + 4\sigma_\nu^2 \right]^{\frac{1}{2}} - (1 + \rho_\xi) \bar{\sigma}_\eta^2 \right\} \quad (37)$$

Given (37), our model only makes sense if:

$$4 \frac{\sigma_\nu^2}{\bar{\sigma}_\eta^2} > (1 + \rho_\xi)^2 \bar{\sigma}_\eta^2 - (1 - \rho_\xi)^2 \quad (38)$$

Equations (33) and (34) characterise the homoscedastic model. Those are also useful when initialising the heteroscedastic Filter model – see more details below, in this appendix. For our parameterisation, it follows that the right hand side of (38) is negative and even if otherwise a sizeable σ_ν^2 compared with $\bar{\sigma}_\eta^2$ would do the trick. Generally, we have found that (37) is met under an ample set of realistic parameterisations.

As mentioned in the main paper, one can evaluate how well a Filter model behaves by checking whether its corresponding Kalman gain increases during a period of economic turbulence (recessions, for instance). This generates an unfiltered series less persistent and probably more connected with movements in the assets market.

The Kalman gain in our model is more complicated than that in Kroencke (2017), but we can still see that a more volatile measurement error lowers \bar{K}^c and expectations of the state of consumption do not change much, in line with the intuition. Now we must ensure that reported consumption is less filtered when \bar{P}^c and(or) $\bar{\sigma}_\eta^2$ are higher. By taking derivatives of the Kalman gain, it

is possible to show that:

$$\frac{\partial \bar{K}^c}{\partial \bar{P}^c} = \left\{ (1 - \rho_\xi) \Phi - (1 - \rho_\xi)^2 [(1 - \rho_\xi) \bar{P}^c + \bar{\sigma}_\eta^2] \right\} \Phi^{-2}, \quad (39)$$

$$\frac{\partial K_t}{\partial \bar{\sigma}_\eta^2} = \left\{ \Phi - (1 - \rho_\xi) \bar{P}^c + \bar{\sigma}_\eta^2 \right\} \Phi^{-2}, \quad (40)$$

where $\Phi = (1 - \rho_\xi)^2 \bar{P}^c + \bar{\sigma}_\eta^2 + \sigma_\nu^2 > 0$. It is straightforward to see that (39) is always positive while (40) is positive for ρ_ξ small enough, so that $\sigma_\nu^2 - (1 - \rho_\xi) \bar{P}^c \rho_\xi > 0$. Our calibration – when $\rho_\xi = 0.06$ – easily meets this condition. The fact that (39) and (40) hold in our model then ensures its validity. We have also found that both conditions are also met in the heteroscedastic model – when derivatives those are time-dependent.

B.2.6

Initialising the Model

We start the model using its homoscedastic analogue described above, so that $P_{t=1}^c = \bar{P}^c$ and $K_{t=1}^c = \bar{K}^c$. In addition, based on the long term representation of the GARCH process in (12): $\sigma_{\eta,t=1}^2 = \alpha_0 / (1 - \alpha_1 - \alpha_2)^5$. As mentioned earlier, by construction we are not able to identify \hat{y}_0 when using the quasi-differenced Filter model. Therefore, we initialise the filter assuming that $\Delta y_{t=1} = 0$, then burning the first observations for which consumption growth seems to exhibit an abrupt and unrealistic movement. It follows that we burned the first three observations when dealing with quarterly data in section 3 but none when using annual data – the latter justified by the low number of observations available. Below we repeat methods in section 3, while restricting the sample to 1960:1–2017:4 (1940–2017) for quarterly (annual) data. In section 4, we burned the first 7 months of data – recall that consumption growth is semi-annual but the frequency is monthly.

B.2.6.1

Notes on Calibration

Recall that for quarterly data in section 3, we use the quasi-differenced Filter model shown in section 2. In that case we fix $\bar{\sigma}_\eta = 0.0078 * \sqrt{3} \approx 1.4\%$, adapting similar results with monthly data presented in Bansal and Yaron (2004) for that frequency. Kroencke (2017) noted that his model seems little sensitive to choices of a_1 and a_2 (GARCH process), once remaining parameters are correctly calibrated to the same moments. We confirmed the same finding

⁵Being more specific, we do $\sigma_{\eta,t=1}^2 = \bar{\sigma}_\eta^2$ and calibrate a_1 and a_2 based on bechmark moments of section 2 such that a_0 is uniquely determined.

for our model. We choose $a_1 = 0.22$ and $a_2 = 0.5$ ⁶. Since the services component of consumption is quite more imprecise (and volatile) than its nondurables analogue, we fix different values for $\bar{\sigma}_\nu$ based on each type of consumption: nondurables and services or nondurables only. Specifically, we use $\bar{\sigma}_\nu = 3.8\%$ for the former and $\bar{\sigma}_\nu = 2.2\%$ for the latter when applying the heteroscedastic model. For its homoscedastic analogue, we use $\bar{\sigma}_\nu = 2.5\%$ and $\bar{\sigma}_\nu = 2.0\%$, respectively. These values not only match our benchmark moments quite well but also the difference itself makes sense, given the imprecision of the services component mentioned above (implying a lower value of $\bar{\sigma}_\nu$ when that group is removed from the measure). Finally, recall that we establish $\rho_\xi = 0.06$ regardless of the consumption type.

For results involving annual data in section 3 our model does not feature persistent measurement error. Therefore, we follow the exact same steps in Kroencke (2017), with mere alternations in calibration to account for an updated time series (until 2017 instead of 2014)⁷. There we set $\bar{\sigma}_\eta = 2.5\%$, $a_1 = 0.01$ and $a_2 = 0.85$. The former is the same value used in Kroencke (2017). We do not adapt it to our time series since by following the same logic we use for quarterly data would give a very similar value ($0.0078 * \sqrt{12} \approx 2.7\%$) and very similar results. In addition, we fix $\bar{\sigma}_\nu = 2.8\%$ (nondurables and services) and $\bar{\sigma}_\nu = 1.9\%$ (nondurables only – see section 6.4.3 below) when using annual data. These values do not change depending on the model used to unfilter consumption.

As in Kroencke (2017), we ensure that the long-term standard deviation (over 6 years) of annual unfiltered consumption is not much more than 1.2 times that of reported consumption (we impose this condition when calibrating $\bar{\sigma}_\nu$). It is intuitive that this gap should not be considerably high since: (i) measurement errors should cancel out when consumption is measured over longer horizons, – see Daniel and Marshall (1996) –, and; (ii) filtering procedures should be smart enough such that in reported data is not considerably more volatile than unfiltered data in the long run – presumably, the implicit algorithm would be otherwise corrected to take new evidence and perceived errors into account. That being said, it is clear that the intuition behind that rule does not change regardless of the nature of the stochastic

⁶That implies $a_0 = 0.00005$.

⁷Technically, we start our Filter model in 1930 and its first observation generated for unfiltered consumption relates to 1931. Kroencke (2017) expanded the original time series provided in NIPA tables to the period that 1927-30, so that it encompasses the Great Depression period. He used data available in Robert Shiller's website for that, with the implicit assumption that the representative statistician does not change the hypothetical Filter model across different datasets. We do not use data from that period, so that all our consumption observations come from BEA (NIPA).

process for the measurement error. For example, for nondurables and services we have that the ratio between long-run standard deviations of unfiltered and reported consumption are 1.12 and 1.23 for quarterly and annual data, respectively. For completeness, in Table B.1 below we present similar results as those shown in Table 2.1 but for consumption of nondurables only. The 1.2 times rule is still valid.

Table B.1: Calibrated Moments for NIPA Consumption of Nondurables Only

[(Implied) Consumption Growth	$E(\Delta C_{year})$	$\sigma(\Delta C_{year})$	$\sigma(\sum_{year=1}^6 \Delta C_{year})/\sqrt{6}$	$Corr(\Delta C_{year}, \Delta C_{year-1})$
Reported (NIPA)	1.38%	2.60%	2.34%	32.09%
Unfiltered - APWG* (1960-14)	-	2.68%	2.30%	0.77%
Unfiltered - APWG* (1928-14)	-	4.15%	3.12%	-11.31%
Unfiltered - Our Model (Quarterly Data)				
Homoscedastic (1960-14)	1.40%	3.38%	2.13%	-29.93%
Heteroscedastic (1960-14)	1.34%	2.16%	1.88%	-11.53%
Homoscedastic (1947-17)	1.34%	3.75%	2.10%	-32.34%
Heteroscedastic (1947-17)	1.29%	2.30%	1.78%	-21.80%
Unfiltered - Our Model (Annual Data)				
Homoscedastic (1960-14)	1.29%	2.62%	2.23%	-0.66%
Heteroscedastic (1960-14)	1.29%	2.63%	2.23%	-0.66%
Homoscedastic (1930-17)	1.51%	4.10%	2.89%	-7.22%
Heteroscedastic (1930-17)	1.51%	4.02%	2.87%	-6.48%

Note: Moments of reported and unfiltered consumption (our model). We compare these moments with those of unfiltered consumption in Kroencke (2017) as well (APWG stands for "Asset Pricing Without Garbage"). We have simply copied his results here, writing "*" next to variables presented in that paper. Reported and unfiltered consumption are for nondurables only, from NIPA tables. We consider the quasi-differenced model with serially correlated measurement errors of section 2 for quarterly data, while setting $\rho_\xi = 0.06$. For annual data, the model is the same as in Kroencke (2017). See section 2.3 for more details on calibration.

Table 2.2 summarises semi-annual consumption growth moments for CEX data – treated in section 4. Although we use data from that survey, we benchmark moments to results obtained for unfiltered NIPA consumption in section 3. We restrict moments of the latter to the available sample period of the former accordingly.

It is well known that a large fraction of CEX consumption categories reproduce a similar behaviour compared to NIPA counterparts. Other categories do measure different things or have similar definitions but exhibit a ratio CEX/NIPA that is too low (high) overtime. However, in terms of the estimation of the EIS, it is more central for the Filter model to be able to revert second moments and auto-correlations than to infer how much one source may be overestimating consumption compared to the other. Indeed, if overall there is no considerable change in how much CEX categories overestimate (underestimate) its NIPA analogues, then one can benchmark CEX aggregates to NIPA counterparts to calibrate the model.

Since official statistical procedures do not differentiate between different asset holders, we calibrate the model based on the consumption growth series for *all* households. Recall that, even though consumption growth is semi-annual, its frequency is monthly. Therefore, to suit its scale for its frequency we convert the series into monthly consumption growth, and then calibrate the model based on the latter. Once unfiltered consumption is obtained, we transform the series back into semi-annual growth terms. This series is then comparable with the original input in the model – reported (CEX) consumption. We repeat this procedure imposing the calibration of the model for *all* households to each type of asset holder. This generates unfiltered consumption growth series associated with each group.

Given that we calibrate the model based on monthly consumption growth at the same frequency, we fix $\bar{\sigma}_\eta = 0.0078$ in section 4 – the same value in Bansal and Yaron (2004), for the same scale and frequency. We once more parameterise $\bar{\sigma}_\nu$ such that the long-term (6-years) standard deviation is not greater than 1.2 times that of reported consumption. This rule gives us $\bar{\sigma}_\nu = 2.0\%$. We maintain $\rho_\xi = 0.06$ in section 4⁸. There we also establish $a_1 = 0.20$ and $a_2 = 0.30$ ⁹.

Note that patterns observed in Table 2.2 for CEX data are similar to those of Table 2.1 for NIPA consumption. Unfiltered consumption once more is more volatile and less auto-correlated than official data. In addition, note that this also holds true for most cases shown in Table 2.2 when we split the sample between different types of asset holders. Particularly, unfiltered consumption pictures a strong mean reversion pattern for stock and bond holders.

Our measures of consumption based on CEX data have a window of 10 periods of missing observations. This happens due to a methodological change, at the end of 1985. The BLS replaced the households IDs, so that we can not match households across that change. To construct unfiltered consumption, we need to impute values to those missing observations. We use the final consumption series for each group to estimate the observation for the next period (out of sample), based on a simple AR(1) model. We then apply and calibrate the model, but exclude those observations associated with the window. To isolate the effect of the training period around the imputation, we remove 6 observations before and after the period.

⁸We will test the stability of our model for different values of ρ_ξ in a future version of this paper.

⁹Such that $a_0 = 0.00003$.

B.3

Epstein-Zin Preferences Framework

This section gives auxiliary algebra and complementary results for section 3.

B.3.1

Euler Equations

Consider L. G. Epstein and Zin (1989) recursive preferences defined by:

$$U_t = \left[(1 - \delta) C_t^{\frac{1-\gamma}{\theta}} + \delta (E_t U_{t+1}^{1-\gamma})^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \quad (41)$$

where $\theta = (1 - \gamma)/(1 - \psi^{-1})$, δ is the discount factor, γ denotes the relative risk aversion coefficient and $C_{k,t}$ is real consumption of type k (unfiltered or reported) in period t . Denote w as the household's wealth and $1 + R_{w,t+1}$ as the gross real return on wealth. If the representative household combines it with the implicit inter-temporal budget constraint $W_{t+1} = (1 + R_{w,t+1})(W_t - C_{k,t})$, it is possible to show that the following Euler Equation holds¹⁰:

$$1 = E_t \left[\left(\delta \left(\frac{C_{k,t+1}}{C_{k,t}} \right)^{-\frac{1}{\psi}} \right)^{\theta} \left(\frac{1}{1 + R_{w,t+1}} \right)^{1-\theta} (1 + R_{f,t+1}) \right], \quad (42)$$

where $1 + R_{f,t+1}$ denotes the gross real returns on risk-free bonds.

Following Campbell (2003) but allowing for time-varying second-order variables as in Yogo (2004) and Campbell, Viceira, Viceira, et al. (2002), if we assume that returns and consumption are jointly log-normal, this implies that the riskless real interest rate is¹¹:

$$r_{f,t+1} = -\log(\delta) + \frac{1}{\psi} E_t[\Delta c_{k,t+1}] + \frac{\theta - 1}{2} \text{Var}_t[r_{w,t+1} - E_t r_{w,t+1}] - \frac{\theta}{2\psi^2} \text{Var}_t[\Delta c_{k,t+1} - E_t \Delta c_{k,t+1}], \quad (43)$$

Equation (42) above is used to estimate (14) of the main text with the risk-free rate. In addition, as in Yogo (2004), we obtain (14), of the main text, for market returns ($i = m$), from (56), by properly defining $r_{i,t+1} - E_t r_{i,t+1} - \frac{1}{\psi}(\Delta c_{k,t+1} - E_t \Delta c_{k,t+1}) \equiv \varrho_{i,t+1}$ under a similar log-linearisation:

$$\begin{aligned} r_{i,t+1} = & -\log(\delta) + \frac{1}{\psi} E_t[\Delta c_{k,t+1}] + \frac{\theta - 1}{2} \text{Var}_t[r_{w,t+1} - E_t r_{w,t+1}] - \frac{\theta}{2\psi^2} \text{Var}_t[\Delta c_{k,t+1} - E_t \Delta c_{k,t+1}] \\ & - \frac{1}{2} \text{Var}_t[r_{i,t+1} - E_t r_{i,t+1}] + \frac{\theta}{\psi} \text{Cov}_t[r_{i,t+1} - E_t r_{i,t+1}, \Delta c_{k,t+1} - E_t \Delta c_{k,t+1}] \\ & + (1 - \theta) \text{Cov}_t[r_{i,t+1} - E_t r_{i,t+1}, r_{w,t+1} - E_t r_{w,t+1}] \end{aligned} \quad (44)$$

¹⁰See L. Epstein and Zin (1991).

¹¹Where $r_{f,t} = \ln(1 + R_{f,t})$, for instance.

Finally, (13) of the main text is obtainable from (44) by rearranging terms while defining $\Delta c_{k,t+1} - E_t \Delta c_{k,t+1} - \psi(r_{i,t+1} - E_t r_{i,t+1}) \equiv \epsilon_{i,t}$.

B.4

K-Class Estimators and Critical Values

Consider the standard simultaneous equations system¹²:

$$y = Y\beta + X\gamma + u \quad (45)$$

$$Y = Z\Pi + X\Phi + V \quad (46)$$

As in Yogo (2004), the three K-class estimators used can be synthesised by¹³:

$$\hat{\beta} = [Y^\perp (I - kM_{z^\perp} Y^\perp) Y^\perp]^{-1} [Y^\perp (I - kM_{z^\perp} Y^\perp) y^\perp] \quad (47)$$

If $k = 1$, then we have TSLS. If k is the smallest root of $|\bar{Y}' M_x \bar{Y} - k \bar{Y}' M_z \bar{Y}|$, then (55) is the LIML estimator. Finally, the Fuller-K estimator is obtained when $k = k_{LIML} - 1/(T - K_1 - K_2)$.

For expository reasons, we repeat critical values of Stock and Yogo (2002) for first-stage F-statistics under the following null hypotheses:

- 1 TSLS bias is a fraction not greater than 10 percent that of the OLS: 10.27
- 2 Size of the TSLS t-test (5% significance) can not be greater than 10 percent: 24.58
- 3 Fuller-K bias as a fraction of the OLS bias is not greater than 10 percent: 6.37
- 4 Size of the LIML t-test (5% significance) can not be greater than 10 percent: 5.44

B.5

Consumption: Nondurables Only

In this subsection we repeat tables of section 3 but for consumption series constructed from nondurables only (excluding the services component). Our main results are maintained.

Two things stand out in results for the homoscedastic framework (Table B.2, Table B.3 and Table B.4). First, when we estimate $1/\psi$ using (14), first-stage F-statistics are actually twice as high for unfiltered consumption

¹²Where y denotes the dependent variable, Y is a matrix constructed from n endogenous variables, X is the matrix of K_1 exogenous regressors and Z has K_2 instruments. All variables have dimension T .

¹³Where $Y^\perp = M_x Y$, $\bar{Z} = [X, Z]$, $\bar{X} = [Y, X]$ and $\bar{Y} = [y, Y]$.

(heteroscedastic model) as for its reported analogue. In contrast, estimates with the former are again similar across estimators, once more suggesting that first-step predictability does not seem especially relevant in generating more sensible estimates¹⁴. Second, with unfiltered consumption our weak-instrument-robust confidence intervals are mostly in the positive region. Unfortunately, we still have uninformative robust intervals under the AR and LR tests and unfiltered consumption at annual frequency.

There is nothing particularly different in Table B.5. Although we obtain negative estimates for the EIS using quarterly data and unfiltered consumption, there is still improvement relative to reported consumption. Again, barely none of those estimates are statistically different from zero¹⁵. However, robust intervals from the J-K test are substantially narrower relative to Table 2.6.

¹⁴Other variables apart, Table B.1 would suggest that unfiltered consumption is a weaker instrument. In absolute terms, its auto-correlation diminishes roughly by a factor of five when we use unfiltered instead of reported consumption, jumping from 32.1% to mere -6.4% (for the complete sample). However, cross-correlations with asset returns are much more definite for unfiltered consumption. Most probably the second effect prevails on net, accounting for the more disciplined estimates obtained.

¹⁵The only exception is the estimate under SYS-GMM and reported consumption for annual data. However, the value of -0.003 is sufficiently small and not statistically significant at 5%.

Table B.2: Estimates of the EIS Using K-Class Estimators and Quarterly Data

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	−0.082***	−0.101***	−0.106***	22.120
			(0.127)	(0.131)	(0.132)	
	ψ	Unf-Hom	0.561	0.585	0.590	22.296
			(0.588)	(0.604)	(0.608)	
	ψ	Unf-Het	0.356**	0.378**	0.382**	22.308
			(0.281)	(0.294)	(0.296)	
Stocks	ψ	Reported	−0.002***	−0.016***	−0.020***	4.530
			(0.028)	(0.034)	(0.036)	
	ψ	Unf-Hom	0.246***	0.263***	0.279***	4.405
			(0.135)	(0.142)	(0.148)	
	ψ	Unf-Het	0.125***	0.158***	0.170***	4.306
			(0.066)	(0.079)	(0.084)	
Risk Free	$\frac{1}{\psi}$	Reported	−0.819***	−3.716	−9.473	1.636
			(0.539)	(3.100)	(11.839)	
	$\frac{1}{\psi}$	Unf-Hom	0.247***	0.831	1.694	1.669
			(0.122)	(0.609)	(1.745)	
	$\frac{1}{\psi}$	Unf-Het	0.435***	1.660	2.620	2.665
			(0.183)	(1.036)	(2.033)	
Stock	$\frac{1}{\psi}$	Reported	−0.311	−11.003	−49.778	1.636
			(3.446)	(13.680)	(90.231)	
	$\frac{1}{\psi}$	Unf-Hom	2.416	2.939	3.584	1.669
			(1.180)	(1.482)	(1.901)	
	$\frac{1}{\psi}$	Unf-Het	3.354	5.061*	5.889*	2.665
			(1.607)	(2.436)	(2.901)	

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and quarterly data. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). All consumption series refer to nondurables only. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.3: Estimates of the EIS Using K-Class Estimators and Annual Data

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	-0.088*** (0.084)	-0.123*** (0.094)	-0.128*** (0.096)	12.498
			0.094*** (0.208)	0.051*** (0.230)	0.043*** (0.234)	10.827
	ψ	Unf-Hom	0.091*** (0.204)	0.049*** (0.225)	0.042*** (0.229)	10.817
			0.030*** (0.024)	0.047*** (0.030)	0.051*** (0.032)	6.100
	ψ	Unf-Het	0.199*** (0.084)	0.228*** (0.102)	0.265*** (0.128)	1.469
			0.195*** (0.082)	0.222*** (0.099)	0.256*** (0.123)	1.482
Risk Free	$\frac{1}{\psi}$	Reported	-1.224*** (0.775)	-5.262* (3.399)	-7.792 (5.811)	2.947
			0.264** (0.307)	0.971 (1.195)	23.076 (124.367)	2.160
	$\frac{1}{\psi}$	Unf-Hom	0.266** (0.314)	0.973 (1.210)	24.042 (132.284)	2.147
			5.308 (3.448)	15.157 (8.954)	19.669 (12.325)	2.947
	$\frac{1}{\psi}$	Unf-Het	3.411 (1.531)	3.587 (1.675)	3.773 (1.823)	2.160
			3.521 (1.571)	3.703 (1.715)	3.901 (1.868)	2.147

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. All consumption series refer to nondurables only. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.4: Weak-IV-Robust CIs for the EIS

Quarterly Data			
Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	$[-0.383, 0.139]$	$[-0.392, 0.145]$
	Unf-Hom	$[-0.611, 1.838]$	$[-0.624, 1.851]$
	Unf-Het	$[0.143, 0.624]$	$[-0.204, 0.995]$
Stocks	Reported	$[-0.124, 0.040]$	$[-0.160, 0.051]$
	Unf-Hom	$[-0.087, 1.135]$	$[-0.004, 0.786]$
	Unf-Het	$[0.015, 0.539]$	$[0.022, 0.503]$
Annual Data			
Risk Free	Reported	$[-0.186, -0.075]$	$[-0.350, 0.045]$
	Unf-Hom	$[-0.236, 0.302]$	$[-0.482, 0.499]$
	Unf-Het	$[-0.238, 0.299]$	$[-0.473, 0.489]$
Stock	Reported	$[0.015, 0.102]$	$[-0.005, 0.152]$
	Unf-Hom	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Notes: Weak-instrument-robust 95% confidence intervals. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Data used for both reported and unfiltered consumption refer to the consumption of nondurables only. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$) while for annual data we use the canonical version - with no persistence for measurement errors. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. The calibration of other parameters described in this section apply in both cases.

Table B.5: Heteroscedasticity-Robust Estimates of the EIS

Quarterly Data				
Δc_k	Two-Step	CUE	SYS	95% CI
Reported	−0.136 (0.113)	−0.151 (0.114)	0.000 (0.000)	$(-\infty, \infty)$
Unf-Hom	0.125 (0.555)	0.152 (0.555)	0.015 (0.012)	$[-0.345, 0.841]$
Unf-Het	−0.034 (0.268)	−0.046 (0.268)	0.003 (0.003)	$[-0.616, 0.262]$
Annual Data				
Reported	−0.079 (0.091)	−0.220 (0.103)	−0.003* (0.001)	$(-\infty, \infty)$
Unf-Hom	0.089 (0.144)	0.094 (0.143)	0.039 (0.010)	$(-\infty, \infty)$
Unf-Het	0.085 (0.142)	0.091 (0.141)	0.039 (0.027)	$(-\infty, \infty)$

Note: 2S-GMM and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free rate. The third column presents estimates of the same coefficient under the joint estimation (15), where market returns are also used, while allowing for different drifts across equations. We present 95% confidence intervals that are robust to both heteroscedasticity and weak-IV settings. These are constructed by inverting the K-statistic of Kleibergen (2005). Consumption series are relative to nondurables only. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

B.6**Implied ψ (EIS) from Estimates of $1/\psi$**

This section presents implied EIS estimates from $1/\psi$ (lower part of Table 2.3 and Table 2.4) using equation (14). All consumption data refer to nondurables and services and standard errors have been constructed by Delta method.

Table B.6: Implied EIS from the Estimation of (14) - Nondurables and Services

Asset	Δc_k	K-Class Estimator		
		TSLS	Fuller-K	LIML
Risk Free	Reported	2.285	0.202***	0.053***
		(1.621)	(0.182)	(0.093)
	Unf-Hom	3.021	0.970	0.573
		(1.406)	(0.643)	(0.484)
	Unf-Het	2.866	0.697	0.385*
		(1.376)	(0.484)	(0.352)
Stock	Reported	1.258	−0.241***	−0.199***
		(4.308)	(0.287)	(0.213)
	Unf-Hom	0.550	0.242***	0.168***
		(0.690)	(0.218)	(0.145)
	Unf-Het	0.292***	0.223***	0.191***
		(0.136)	(0.106)	(0.093)
Annual Data				
Risk Free	Reported	0.733	0.218***	0.108***
		(0.389)	(0.160)	(0.113)
	Unf-Hom	2.577	0.531	0.090***
		(2.537)	(0.533)	(0.208)
	Unf-Het	2.539	0.522***	0.089***
		(2.496)	(0.129)	(0.205)
Stock	Reported	−0.147***	−0.085***	−0.065***
		(0.084)	(0.049)	(0.040)
	Unf-Hom	0.247***	0.240***	0.227***
		(0.112)	(0.111)	(0.108)
	Unf-Het	0.241***	0.236***	0.222***
		(0.109)	(0.108)	(0.105)

Note: Implied ψ (EIS) estimates from (14). Consumption series have been constructed taking into account the consumption of nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard values are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.7: Implied EIS from the Estimation of (14) - Nondurables Only

Quarterly Data				
Asset	Δc_k	K-Class Estimator		
		TSLS	Fuller-K	LIML
Risk Free	Reported	-1.221***	-0.269***	-0.106***
		(0.803)	(0.225)	(0.132)
	Unf-Hom	4.049	1.203	0.590
		(2.000)	(0.882)	(0.608)
	Unf-Het	2.298	0.602	0.382**
		(0.965)	(0.376)	(0.296)
Stock	Reported	-3.220	-0.091***	-0.020***
		(35.721)	(0.113)	(0.036)
	Unf-Hom	0.414***	0.340***	0.279***
		(0.202)	(0.172)	(0.148)
	Unf-Het	0.298***	0.198***	0.170***
		(0.143)	(0.095)	(0.084)
Annual Data				
Risk Free	Reported	-0.817***	-0.190***	-0.128***
		(0.518)	(0.123)	(0.096)
	Unf-Hom	3.788	1.030	0.043***
		(4.405)	(1.267)	(0.234)
	Unf-Het	3.766	1.028	0.042***
		(4.455)	(1.278)	(0.229)
Stock	Reported	0.188***	0.066***	0.051***
		(0.122)	(0.039)	(0.032)
	Unf-Hom	0.293***	0.279***	0.265***
		(0.132)	(0.130)	(0.128)
	Unf-Het	0.284***	0.270***	0.256***
		(0.127)	(0.125)	(0.123)

Note: Implied ψ (EIS) estimates from (14). Consumption series have been constructed taking into account the consumption of nondurables only. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard values are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

B.7**Results Using Raw Returns for Both Reported and Unfiltered Consumption**

In the main text, we adjusted return series for potential time-aggregation bias when conducting estimates with unfiltered consumption – see section 2.2. Recall that returns are never adjusted when reported consumption is used. Here, we present results for the case when we use raw returns with both reported and unfiltered consumption. In general terms, our findings are broadly similar to those of section 3. Note that robust intervals constructed from inverting AR and LR statistics are no longer uninformative with unfiltered consumption and annual data. However, our results for stock returns are somewhat weaker in comparison. In addition, we can not revert uninformative sets in the heteroscedasticity-robust framework (J-K test).

Table B.8: EIS Using K-Class Estimators and Quarterly Data – Raw Returns

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	0.067*** (0.078)	0.053*** (0.093)	0.053*** (0.093)	22.172
	ψ	Unf-Hom	0.521 (0.461)	0.559 (0.475)	0.566 (0.478)	22.335
	ψ	Unf-Het	0.342** (0.332)	0.375* (0.346)	0.380* (0.348)	22.474
Stocks	ψ	Reported	0.006*** (0.017)	−0.101*** (0.096)	−0.199*** (0.213)	4.575
	ψ	Unf-Hom	0.215*** (0.114)	0.234*** (0.121)	0.249*** (0.127)	4.462
	ψ	Unf-Het	0.165*** (0.084)	0.189*** (0.093)	0.202*** (0.099)	4.310
Risk Free	$\frac{1}{\psi}$	Reported	0.438* (0.311)	4.953 (4.475)	18.979 (33.660)	6.630
	$\frac{1}{\psi}$	Unf-Hom	0.334*** (0.155)	1.043 (0.691)	1.765 (1.491)	1.890
	$\frac{1}{\psi}$	Unf-Het	0.353*** (0.170)	1.452 (1.009)	2.630 (2.408)	2.268
Stock	$\frac{1}{\psi}$	Reported	0.795 (2.724)	−4.150 (4.936)	−5.014 (5.346)	6.630
	$\frac{1}{\psi}$	Unf-Hom	2.725 (1.296)	3.362 (1.643)	4.015 (2.040)	1.890
	$\frac{1}{\psi}$	Unf-Het	3.237 (1.512)	4.242 (2.019)	4.956 (2.423)	2.268

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and quarterly data. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). Return series are *not* adjusted for time-aggregation issues as in section 2.2. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.9: EIS Using K-Class Estimators and Annual Data – Raw Returns

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	0.112*** (0.105)	0.109*** (0.111)	0.108*** (0.113)	11.836
	ψ	Unf-Hom	0.271*** (0.269)	0.289** (0.284)	0.293** (0.288)	11.699
	ψ	Unf-Het	0.121*** (0.172)	0.137*** (0.183)	0.140*** (0.186)	11.707
Stocks	ψ	Reported	-0.049*** (0.034)	-0.061*** (0.038)	-0.065*** (0.040)	5.057
	ψ	Unf-Hom	-0.037*** (0.058)	-0.024*** (0.063)	-0.022*** (0.064)	11.432
	ψ	Unf-Het	0.019*** (0.042)	0.034*** (0.046)	0.038*** (0.047)	9.026
Risk Free	$\frac{1}{\psi}$	Reported	1.364 (0.724)	4.592 (3.364)	9.246 (9.634)	1.893
	$\frac{1}{\psi}$	Unf-Hom	0.574 (0.328)	1.788 (1.303)	3.410 (3.345)	1.689
	$\frac{1}{\psi}$	Unf-Het	0.613 (0.462)	2.688 (2.297)	7.127 (9.426)	1.764
Stock	$\frac{1}{\psi}$	Reported	-6.808** (3.885)	-11.773* (6.818)	-15.285* (9.365)	1.893
	$\frac{1}{\psi}$	Unf-Hom	-1.587** (1.187)	-5.532 (5.335)	-46.488 (137.715)	1.689
	$\frac{1}{\psi}$	Unf-Het	1.715 (2.085)	10.520 (9.642)	26.585 (33.559)	1.764

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. Return series are *not* adjusted for time-aggregation issues as in section 2.2. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.10: Weak-IV-Robust CIs for the EIS – Raw Returns

Quarterly Data			
Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	\emptyset	$[-0.136, 0.235]$
	Unf-Hom	$[-0.315, 1.507]$	$[-0.372, 1.573]$
	Unf-Het	$[-0.076, 0.857]$	$[-0.303, 1.109]$
Stocks	Reported	\emptyset	$(-\infty, +\infty)$
	Unf-Hom	$[-0.048, 0.975]$	$[0.014, 0.688]$
	Unf-Het	$[-0.009, 0.751]$	$[0.024, 0.568]$
Annual Data			
Risk Free	Reported	$[-0.104, 0.316]$	$[-0.131, 0.341]$
	Unf-Hom	$[-0.304, 0.934]$	$[-0.292, 0.920]$
	Unf-Het	$[-0.203, 0.510]$	$[-0.233, 0.545]$
Stock	Reported	$[-0.245, 0.019]$	$[-0.199, 0.007]$
	Unf-Hom	$[-0.131, 0.107]$	$[-0.143, 0.125]$
	Unf-Het	$[-0.044, 0.148]$	$[-0.049, 0.158]$

Note: Weak-instrument-robust 95% confidence intervals. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Return series are *not* adjusted for time-aggregation issues as in section 2.2. Data used both for reported and unfiltered consumption refer to the consumption of nondurables and services. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$) while for annual data we use the canonical version – with no persistence for measurement errors. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

Table B.11: Heteroscedasticity-Robust Estimates of the EIS – Raw Returns

Quarterly Data				
Δc_k	Two-Step	CUE	SYS	95% CI
Reported	0.133 (0.082)	0.189** (0.085)	0.001 (0.000)	$(-\infty, +\infty)$
Unf-Hom	0.594 (0.517)	0.670 (0.519)	0.008 (0.006)	$(-\infty, +\infty)$
Unf-Het	0.443 (0.370)	0.506 (0.372)	0.007 (0.001)	$(-\infty, +\infty)$
Annual Data				
Reported	0.056 (0.088)	0.022 (0.087)	-0.015* (0.008)	$(-\infty, +\infty)$
Unf-Hom	0.067 (0.261)	0.047 (0.261)	-0.001 (0.000)	$(-\infty, +\infty)$
Unf-Het	0.033 (0.172)	0.025 (0.172)	-0.002 (0.001)	$(-\infty, +\infty)$

Note: 2S- and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free. The third column presents results under the joint estimation (15), where market returns are also used (allowing for different drifts across equations). We present 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting. These are constructed by inverting the K-statistic of Kleibergen (2005). Return series are *not* adjusted for time-aggregation issues as in section 2.2. Consumption series are relative to nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

B.8

Results with Restricted Sample

In this section, we repeat the estimations of section 3 in the main paper while restricting our sample. The motivation here is to remove first observations for which the Kalman filter is still learning.

Our estimations here use the period encompassing 1960:1 to 2017:4 for quarterly and 1940 to 2017 for annual data. Tables below broadly confirm our findings in the main paper.

Table B.12: EIS Using K-Class Estimators and Quarterly Data – 1960:2017

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	0.120*** (0.075)	0.083*** (0.108)	0.081*** (0.110)	22.810
	ψ	Unf-Hom	0.300* (0.392)	0.321* (0.404)	0.326* (0.406)	23.558
	ψ	Unf-Het	0.263** (0.284)	0.281** (0.293)	0.285** (0.294)	23.652
Stocks	ψ	Reported	0.017*** (0.016)	−0.011*** (0.069)	−3.946*** (46.662)	3.821
	ψ	Unf-Hom	0.135*** (0.082)	0.152*** (0.090)	0.163*** (0.094)	4.133
	ψ	Unf-Het	0.117*** (0.063)	0.126*** (0.066)	0.134*** (0.069)	4.025
Risk Free	$\frac{1}{\psi}$	Reported	0.498** (0.246)	4.978 (4.168)	12.292 (16.545)	12.118
	$\frac{1}{\psi}$	Unf-Hom	0.299*** (0.169)	1.216 (0.981)	3.071 (3.830)	1.658
	$\frac{1}{\psi}$	Unf-Het	0.472** (0.240)	1.720 (1.271)	3.514 (3.632)	1.768
Stock	$\frac{1}{\psi}$	Reported	1.467 (2.411)	−0.093 (2.944)	−0.253 (2.997)	12.118
	$\frac{1}{\psi}$	Unf-Hom	3.475 (1.819)	4.857 (2.652)	6.144 (3.560)	1.658
	$\frac{1}{\psi}$	Unf-Het	5.272* (2.569)	6.210* (3.090)	7.440* (3.843)	1.769

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and quarterly data. We restrict our sample to the period from 1960:1 to 2017:4. Unfiltered consumption extracted relying on the quasi-differenced Filter model whose measurement errors are serially correlated ($\rho_\xi = 0.06$). All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.13: EIS Using K-Class Estimators and Annual Data – 1940:2017

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	0.012*** (0.094)	0.014*** (0.098)	0.015*** (0.099)	16.600
	ψ	Unf-Hom	0.151*** (0.146)	0.149*** (0.147)	0.146*** (0.149)	11.846
	ψ	Unf-Het	0.151*** (0.146)	0.150*** (0.147)	0.147*** (0.149)	11.828
Stocks	ψ	Reported	−0.060*** (0.044)	−0.077*** (0.051)	−0.088*** (0.057)	2.691
	ψ	Unf-Hom	0.051*** (0.056)	0.039*** (0.066)	0.025*** (0.077)	1.500
	ψ	Unf-Het	0.051*** (0.056)	0.040*** (0.065)	0.027*** (0.076)	1.506
Risk Free	$\frac{1}{\psi}$	Reported	0.208 (0.713)	1.502 (2.376)	68.677 (468.18)	1.694
	$\frac{1}{\psi}$	Unf-Hom	2.636 (1.632)	3.370 (2.353)	6.848 (6.962)	0.620
	$\frac{1}{\psi}$	Unf-Het	2.622 (1.620)	3.370 (2.352)	6.826 (6.921)	0.626
Stock	$\frac{1}{\psi}$	Reported	−5.996* (4.000)	−8.842* (5.681)	−11.380 (7.443)	1.694
	$\frac{1}{\psi}$	Unf-Hom	4.800 (3.469)	6.468 (5.811)	39.723 (121.877)	0.620
	$\frac{1}{\psi}$	Unf-Het	4.867 (3.469)	6.616 (5.855)	37.384 (107.139)	0.626

Notes: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. We restrict our sample to the period from 1940 to 2017. Unfiltered consumption extracted relying on the Filter model whose measurement errors are not persistent. All consumption series refer to nondurables and services. We apply the same setting of Yogo (2004), using 3 types of K-class estimators and assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, asset returns have not been adjusted for time-aggregation. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.14: Weak-IV-Robust CIs for the EIS – Restricted Sample

Quarterly Data: 1960:1 – 2017:4			
Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	\emptyset	$[-0.136, 0.283]$
	Unf-Hom	$[-0.446, 1.134]$	$[-0.478, 1.168]$
	Unf-Het	$[-0.262, 0.860]$	$[-0.295, 0.896]$
Stocks	Reported	\emptyset	$(-\infty, +\infty)$
	Unf-Hom	$[-0.062, 0.691]$	$[-0.021, 0.509]$
	Unf-Het	$[-0.033, 0.609]$	$[0.005, 0.393]$
Annual Data: 1940 – 2017			
Risk Free	Reported	$[-0.164, 0.196]$	$[-0.185, 0.218]$
	Unf-Hom	$[-0.367, 0.599]$	$[-0.186, 0.452]$
	Unf-Het	$[-0.365, 0.600]$	$[-0.185, 0.453]$
Stock	Reported	$[-2.074, 0.034]$	$[-0.583, 0.015]$
	Unf-Hom	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Note: Weak-instrument-robust 95% confidence intervals inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. We restrict our estimations to the sample 1960:1-2017:4 (1940-2017) for quarterly (annual) data. Data used both for reported and unfiltered consumption refer to the consumption of nondurables and services. For quarterly data, we use our quasi-differenced Filter model ($\rho_\xi = 0.06$) while for annual data we use the canonical version – with no persistence for measurement errors. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

Table B.15: Heteroscedasticity-Robust Estimates of the EIS – Restricted Sample

Quarterly Data: 1960:1 – 2017:4				
Δc_k	Two-Step	CUE	SYS	95% CI
Reported	0.123 (0.077)	0.347 (0.093)	0.008 (0.002)	$(-\infty, +\infty)$
Unf-Hom	0.349 (0.333)	0.363 (0.333)	0.002 (0.001)	[0.022, 2.782]
Unf-Het	0.305 (0.240)	0.312 (0.240)	0.004 (0.001)	$[-0.173, 2.066]$
Annual Data: 1940 – 2017				
Reported	0.024 (0.087)	-0.029 (0.087)	-0.008 (0.005)	$(-\infty, +\infty)$
Unf-Hom	0.152 (0.141)	0.175 (0.141)	0.015 (0.009)	$(-\infty, +\infty)$
Unf-Het	0.152 (0.140)	0.173 (0.140)	0.015 (0.009)	$(-\infty, +\infty)$

Note: 2S-GMM and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free rate. The third column presents estimates of the same coefficient under the joint estimation (15), where market returns are also used, while allowing for different drifts across equations. 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting are also shown in the last column. These are constructed by inverting the K-statistic of Kleibergen (2005). We restrict our estimations to the sample 1960:1-2017:4 (1940-2017) for quarterly (annual) data. Consumption series are relative to nondurables and services. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

B.9

Results for Annual Data when Measurement Errors are Persistent

In the main paper, we showed results for annual data considering the original model in Kroencke (2017), which does not feature serially correlated measurement errors ($\rho_\xi = 0$). For completeness, in this section we exhibit results of section 3 for annual data while $\rho_\xi = 0.06 \neq 0$. It is worth reemphasising that we have found little sensitiveness of unfiltered consumption to different values of ρ_ξ once other parameters have been properly calibrated according to benchmark moments. Hence, we repeat $\rho_\xi = 0.06$ (as we did with quarterly data), but other parameters have been changed: $\bar{\sigma}_\eta = 0.0078 \times \sqrt{12} \approx 2.7\%$, $a_1 = 0.05$, $a_2 = 0.85$ and $\sigma_\nu = 3.3\%$ (heteroscedastic model) or $\sigma_\nu = 2.1\%$ (homoscedastic model)¹⁶. These parametric conditions ensure that moments of annual unfiltered consumption are not much different from those presented in Table 2.1. All consumption data refer to nondurables and services.

¹⁶We lower σ_ν for the homoscedastic model simply to ensure that the long-term standard deviation of unfiltered consumption is not greater than 1.2 times that of reported consumption.

Table B.16: Estimates of the EIS – Persistent M.E. and Annual Data

Asset	Estimate	Δc_k	K-Class Estimator			1S-F
			TSLS	Fuller-K	LIML	
Risk Free	ψ	Reported	0.112***	0.109***	0.108***	11.836
			(0.105)	(0.111)	(0.113)	
	ψ	Unf-Hom	0.097***	0.089***	0.088**	10.727
			(0.188)	(0.203)	(0.205)	
	ψ	Unf-Het	-0.030***	0.012***	0.020***	11.096
			(0.197)	(0.213)	(0.216)	
Stocks	ψ	Reported	-0.049***	-0.061***	-0.065***	5.057
			(0.034)	(0.038)	(0.040)	
	ψ	Unf-Hom	0.184***	0.194***	0.222***	1.312
			(0.080)	(0.086)	(0.105)	
	ψ	Unf-Het	0.059***	0.124***	0.168***	1.979
			(0.078)	(0.115)	(0.145)	
Risk Free	$\frac{1}{\psi}$	Reported	1.364	4.592	9.246	1.893
			(0.724)	(3.364)	(9.634)	
	$\frac{1}{\psi}$	Unf-Hom	0.394	1.910	11.293	1.822
			(0.387)	(1.916)	(26.190)	
	$\frac{1}{\psi}$	Unf-Het	-0.115***	0.275	49.474	1.739
			(0.358)	(1.043)	(528.708)	
Stock	$\frac{1}{\psi}$	Reported	-6.808**	-11.773*	-15.285*	1.893
			(3.885)	(6.818)	(9.365)	
	$\frac{1}{\psi}$	Unf-Hom	4.143*	4.241*	4.502	1.822
			(1.863)	(1.937)	(2.131)	
	$\frac{1}{\psi}$	Unf-Het	1.818	4.137	5.956	1.739
			(2.282)	(3.734)	(5.130)	

Note: Estimates of the EIS and its reciprocal using (13) and (14) and annual data. We use 3 types of K-class estimators while assuming that errors conditionally follow a martingale difference sequence. When reported consumption is used, we have not adjusted the timing of returns. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. The quasi-differenced model with $\rho_\xi = 0.06$ has been used, while adjusting other parameters to the dynamics and benchmark moments of annual data. All consumption measures refer to nondurables and services. Standard errors are presented in parentheses. The null that the estimated coefficient equals 1 has been tested: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.

Table B.17: Weak-IV-Robust CIs for the EIS – Persistent M.E. and Annual Data

Asset	Δc_k	Anderson-Rubin	Likelihood Ratio
Risk Free	Reported	$[-0.104, 0.316]$	$[-0.131, 0.341]$
	Unf-Hom	$[-0.269, 0.438]$	$[-0.352, 0.516]$
	Unf-Het	$[-0.334, 0.430]$	$[-0.395, 0.515]$
Stock	Reported	$[-0.245, 0.018]$	$[-0.199, 0.007]$
	Unf-Hom	$(-\infty, +\infty)$	$(-\infty, +\infty)$
	Unf-Het	$(-\infty, +\infty)$	$(-\infty, +\infty)$

Note: Weak-instrument-robust 95% confidence intervals for annual data. Sets constructed by inverting statistics of the Anderson-Rubin and Likelihood Ratio tests. Data used both for reported and unfiltered consumption refer to consumption of nondurables and services. We set $(\rho_\xi = 0.06)$ while adjusting other parameters to align benchmark moments. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively.

Table B.18: Het-Robust Estimates of the EIS – Persistent M.E. and Annual Data

Δc_k	Two-Step	CUE	SYS	95% CI
Reported	0.056	0.022	-0.015*	$(-\infty, \infty)$
	(0.088)	(0.087)	(0.008)	
Unf-Hom	0.119	0.133	0.066	$(-\infty, \infty)$
	(0.041)	(0.141)	(0.045)	
Unf-Het	-0.099	-0.399	0.000	$(-\infty, \infty)$
	(0.201)	(0.204)	(0.001)	

Note: 2S-GMM and CUE-GMM estimates of ψ (EIS) in equation (13) using the risk-free rate. The third column presents estimates of the same coefficient under the joint estimation (15), where market returns are also used, while allowing for different drifts across equations. We present 95% confidence intervals that are robust to both heteroscedasticity and a weak-IV setting. These are constructed by inverting the K-statistic of Kleibergen (2005). Consumption series are relative to nondurables and services. We set $\rho_\xi = 0.06$ while adjusting other parameters to align benchmark moments. Reported denotes official consumption data from NIPA tables. Unf-Hom and Unf-Het refer to unfiltered consumption, constructed by the homoscedastic and heteroscedastic models, respectively. Standard errors are presented in parentheses. The null that the estimated coefficient equals 0 has been tested using conventional t-statistics: ***, ** and * denote rejection of the null hypothesis at 1, 5 and 10 percent significance levels.