



## ORBIT - Online Repository of Birkbeck Institutional Theses

---

Enabling Open Access to Birkbeck's Research Degree output

### Economic measurement and modelling with large datasets: theory, application and policy implications

<https://eprints.bbk.ac.uk/id/eprint/40452/>

Version: Full Version

**Citation: Simonsen, Lasse de la Porte (2019) Economic measurement and modelling with large datasets: theory, application and policy implications. [Thesis] (Unpublished)**

© 2020 The Author(s)

---

All material available through ORBIT is protected by intellectual property law, including copyright law.

Any use made of the contents should comply with the relevant law.

---

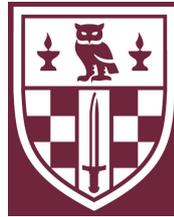
[Deposit Guide](#)  
Contact: [email](#)

Economic Measurement and Modelling with Large  
Datasets: Theory, Application and Policy  
Implications

*Lasse de la Porte Simonsen*

THESIS SUBMITTED FOR THE DEGREE OF  
DOCTOR OF PHILOSOPHY

DEPARTMENT OF ECONOMICS, MATHEMATICS AND STATISTICS  
BIRKBECK COLLEGE, UNIVERSITY OF LONDON



Supervised by PROFESSOR STEPHEN H. WRIGHT

November, 2019

# Declaration

I hereby declare that, except where specific reference is made to the work of others, the contents of this dissertation is original and has not been submitted in whole or in part for consideration of any other degree or qualification at Birkbeck College, University of London, or any other university. Chapter 2, titled “Residential Land Supply in 27 EU Countries: Pigovian Controls or Nimbyism?”, is joint work with Stephen H. Wright. I certify that I am one of the lead authors of the research, having designed and carried out the empirical analysis and data work. Further, Chapter 3, titled “Inflation Dynamics and Price Flexibility in the UK”, is joint work with Ivan Petrella and Emiliano Santoro. I certify that I am one of the lead authors of the research, having dealt with the non-trivial task of managing and understanding the 27.5 million observations in our dataset, and the technical details of the methodology used for the construction of the Consumer Price Index.

Lasse de la Porte Simonsen  
London, November, 2019

# Abstract

In this thesis I debate the merits of using large disaggregated datasets to drive economic theory, and its implications for economic policy making. In chapter 2 we use a previously unexplored dataset of 1/4 million photo-interpreted points to measure the supply of residential land across Europe. Using this dataset, we estimate the reduced-form parameters of a Pigovian externality model for land controls. We interpret the results of these estimates as evidence that across Europe the supply of residential land use is detrimentally restricted to below the social optimum. In chapter 3 the underlying dynamics of the UK consumer price inflation are explored from 27.5 million underlying price quotes in the Consumer Price Index. We find evidence of secular trends in the pricing mechanism of firms in the UK, as well as support for the theory of state dependent pricing. Further, our results indicate that neither the Bank of England nor professional forecasters are taking into account the information embedded in a flexibility index which could improve their inflation forecasts. Lastly, in chapter 4 I explore a framework for decomposing the inflation rate into missing observations, product replacements and regular matched inflation rates. Using this framework I explore a potential source of bias in the inflation measurement of a particular narrow category of good, “Women’s top, long sleeved, not blouse”, due to uncaptured quality change. The results are preliminary, and difficult to interpret as the bias found could be explained by fashion cycles rather than being a measurement error. Finally, in chapter 5 I provide my concluding remarks, finding that there is indeed a benefit to exploring large disaggregated datasets, as they can uncover features of economic fundamentals that are not readily observable in aggregated data.

# Acknowledgements

I would like to thank Birkbeck College for having taught me (almost) everything I know about economics. My thanks also goes to my fellow PhD students at Birkbeck (past and present) especially to Marco Pellicia, Andrea Nocera, Angelos Diamantopoulos, and Rory MacQueen. I acknowledge and thank the Economics and Social Research Council for financial support in order to undertake my PhD. My thanks also goes to the Structural Economic Analysis Division at the Bank of England for hosting me for a summer internship in 2015, during which the seeds for my last two chapters on price inflation were first sown. I thank the ONS, and especially Rowan Kelsoe, for all their patience and help in answering all my questions. Likewise, my thanks also goes to my current employers at Macrosynergy Partners and all to my amazing colleagues there for your help, patience and support. Thanks to Emiliano Santoro for being a brilliant co-author (together with Ivan Petrella) with a relentless energy for redrafting our paper and making us strive for better, and to Søren Hove Ravn for making me feel welcome whilst teaching at the summerschools at Copenhagen University, and lending me a desk to work from during the summer of 2014. My deepest of gratitude and thanks goes to my supervisors Stephen Wright, Ivan Petrella, and Ron Smith. Birkbeck is the place where I was taught everything I know about economics and econometrics, and you are the people who taught me, and still continue to do so. For all your help and guidance, I thank you. A huge thanks goes to my family and friends for all the support you have given me and the patience you have shown throughout the last 5 years. Lastly, I would like to thank my wife Christine, without whom I would never have been able to get through this journey. This thesis is dedicated to you my love.

# Contents

<b>List of Figures</b>	<b>6</b>
<b>List of Tables</b>	<b>8</b>
<b>1 Introduction</b>	<b>9</b>
<b>2 Residential Land Supply in 27 EU Countries: Pigovian Controls or Nimbyism?</b>	<b>12</b>
2.1 Introduction . . . . .	12
2.2 The dataset . . . . .	16
2.3 A simple model of Pigovian land supply . . . . .	28
2.4 Estimation . . . . .	34
2.5 Pigovian Land Supply? Attempting to make sense of the econometric reduced form. . . . .	43
2.6 Conclusions . . . . .	46
2.7 Appendices . . . . .	47
<b>3 Inflation Dynamics and Price Flexibility in the UK</b>	<b>53</b>
3.1 Introduction . . . . .	53
3.2 Microdata on consumer prices . . . . .	57
3.3 Analytical framework . . . . .	67
3.4 A generalized Ss model . . . . .	71
3.5 Implications for aggregate price adjustment . . . . .	76
3.6 Inflation dynamics and state dependence in price flexibility . . . . .	83
3.7 Concluding remarks . . . . .	88
3.8 Appendices . . . . .	89
<b>4 Adding up the UK's consumer price inflation data: Quality change, missing observations and the importance of sales</b>	<b>106</b>
4.1 Introduction . . . . .	106
4.2 The Basic Building Blocks: Index Calculations . . . . .	109
4.3 Decomposing the inflation rate . . . . .	118
4.4 An alternative pure price index: No quality change . . . . .	130
4.5 Conclusion . . . . .	130
4.6 Appendix . . . . .	133
4.7 Higher Aggregates: COICOP Classifications . . . . .	135
<b>5 Concluding Remarks</b>	<b>138</b>



# List of Figures

2.1	The LUCAS dataset . . . . .	16
2.2	Shares of residential land in total land in 27 EU countries . . . . .	18
2.3	Residential Shares by EU Region: Empirical CDFs . . . . .	18
2.4	Estimates of residential land per capita in 27 EU countries, with 95% confidence intervals . . . . .	20
2.5	Residential Land Per Capita by Region: Empirical CDFs . . . . .	21
2.6	Composition of building stock . . . . .	22
2.7	Main components of residential land per capita in 27 EU countries .	23
2.8	Residential vs. Total land: National Data . . . . .	24
2.9	Residential vs. Total land: Regional Data . . . . .	24
2.10	Residential land per capita vs. consumption per capita . . . . .	25
2.11	Housing Expenditure per $m^2$ vs Consumption per Capita in 27 EU Countries . . . . .	26
2.12	Composition of Total Land . . . . .	27
2.13	Housing Expenditure per $m^2$ and the Opportunity Cost of Land . . .	27
2.14	Comparative Statics . . . . .	33
2.15	Share Aggregation of homogeneous equations . . . . .	42
2.16	Calibrations: Realistic Calibration . . . . .	45
2.17	Calibrations: Unrealistic Calibration . . . . .	45
2.18	The LUCAS dataset . . . . .	49
2.19	Observation 44503638 . . . . .	50
3.1	Frequency of Adjustment and Average Price Changes . . . . .	62
3.2	Moments of the Distribution of Price Changes . . . . .	63
3.3	Analytical Framework . . . . .	69
3.4	A combined increase in $\phi$ and $\lambda$ . . . . .	70
3.5	Dispersion of Price Gaps and the Inaction Region . . . . .	75
3.6	Price Flexibility and Different Margins of Price Adjustment . . . . .	77
3.7	Event Study: VAT Changes . . . . .	81
3.8	Price Flexibility and Inflation Persistence . . . . .	84
3.9	Representativeness . . . . .	90
3.10	Aggregate vs Disaggregated Moments . . . . .	91
3.11	Monthly GDP and Detrended Coincident Indicator . . . . .	94
3.12	Estimated Price Gap Distributions and Hazard Functions . . . . .	97
3.13	Parameters of the Price Gap Distribution . . . . .	98
3.14	Parameters of the Hazard Function . . . . .	99
3.15	Probability of a High-flexibility Regime . . . . .	104
3.16	Price Flexibility and Inflation Persistence . . . . .	105

3.17	Price Flexibility and Inflation Volatility . . . . .	105
4.1	A Price Trajectory . . . . .	112
4.2	Elementary Aggregate Index . . . . .	115
4.3	Item Index (unchained) . . . . .	116
4.4	Chain-Linked Item Index . . . . .	118
4.5	Decomposition of the inflation rate . . . . .	129
4.6	Weighted index of prices . . . . .	131
4.7	Distribution of sampled prices . . . . .	132
4.8	Fit of the inflation rate . . . . .	134

# List of Tables

2.1	Summary table of the dataset . . . . .	19
2.2	The probability that any LUCAS survey point is classified as residential: Spatial Correlation and national vs regional dummy variables	36
2.3	The probability that any LUCAS survey point is classified as residential: Spatial Correlation and implicit model of regional land supply	41
2.4	Direct estimation of reduced form model of land supply on regional and national data . . . . .	43
2.5	Datasources . . . . .	48
3.1	Summary Statistics . . . . .	57
3.2	Correlations of Pricing Moments with Macroeconomic Variables . . .	64
3.3	Correlations of Pricing Moments with Macroeconomic Variables: the Role of Asymmetry . . . . .	65
3.4	Flexibility in Price Adjustment: Correlation with Real Activity and Inflation . . . . .	78
3.5	VAT Changes: Actual and Counterfactual Statistics . . . . .	82
3.6	Forecast Errors and Price Flexibility . . . . .	86
3.7	Coincident Indicator - Model Fit . . . . .	93
3.8	Correlations of Pricing Moments with Macroeconomic Variables (Quadratic Trends) . . . . .	100
3.9	Correlations of Pricing Moments with Macroeconomic Variables: the Role of Asymmetry (Quadratic Trends) . . . . .	101
3.10	Flexibility in Price Adjustment: Correlation with Real Activity and Inflation (Quadratic Trends) . . . . .	102
3.11	Forecast Errors and Price Flexibility: Robustness (MSE) . . . . .	103
4.1	Price trajectories . . . . .	111
4.2	A numerical example of missing observations . . . . .	122
4.3	Example of Quality Change . . . . .	126
4.4	Regression fit of own calculations to the ONS implied inflation rate .	135
4.5	Moments of the series . . . . .	135

# Chapter 1

## Introduction

### Motivation and Contributions

In recent years an increasing attention has been given to what is termed “big data” and the statistical methods associated with such datasets (see for instance Varian, 2014 as well as Brynjolfsson and McAfee, 2014). However, the difficulty faced in economics, and especially macroeconomics, for using some of these big data statistical tools is that often the questions we want answered simply does not have abundance of data required for some of these methods – at least at first glance. Instead, macroeconomists mostly rely on small sized datasets where the history generally extends back less than 30 years, with a sampling frequency that is monthly at best. Further, the data regularly displays strong serial correlations, frequent structural breaks, and occurrences of key co-movement events, such as business cycle turning points, are rare observations – all of which makes inference difficult. As an example, if we look at the Consumer Price Index (CPI) inflation rate in the UK, it is a monthly series dating back to February 1996 which, at the time of this writing, is a “meagre” 269 months of observations,<sup>1</sup> and with only one recession in that period (2008 to 2009). When compared to some of the vast datasets, such as the ones found within today’s online social network platforms where we are talking million of observations, this is a small dataset.

To try and overcome these issues and incorporate “big data“ into economics, recent advances in macroeconomics have focused on incorporating larger datasets into the analysis by exploiting co-movements across the cross-section. An example of this is the use of factor models (see for instance Giannone et al., 2008 or Stock and Watson, 2002) which use this cross-sectional co-movement to infer information about the current state of the economy and common trends. As a complement to this, in this thesis I will set out an alternative use of large datasets<sup>2</sup> in macroeconomics by

---

<sup>1</sup>With the current last published inflation data for the UK being for June 2018.

<sup>2</sup>I see our data as being large datasets, but not truly big in the sense of the machine learning social media millions of observations, and hence the “large” data in the title of this thesis.

focusing on the micro data underlying the aggregated datasets we most often use. I will explore how non-standard large datasets can give us insight into questions where there is a sparsity of data, how the underlying distributions as well as trends of the micro data might differ from the aggregates, and lastly, if from these micro datasets we can infer anything about the quality of the measured aggregate data. Further, I will suggest that using these large disaggregated dataset can help give us insights into some of today's key challenges for economic policy.

## Outline of the thesis

The thesis proceeds in chapter 2, by exploring a new approach to estimating differences in residential land supply across Europe, using a dataset of photo-interpreted points. The datasets are purely cross-sectional consisting of single 2012 time-dimension ( $T = 1$ ), and a large cross-section ( $N \approx 250,000$ ) of points from all across Europe.<sup>3</sup> From these points of data, we are able to derive estimates of the share of residential land (out of total land) across Europe, and from these estimates a measure of per capita residential land. Using this data we estimate a model of Pigovian externality, and find evidence towards too restrictive residential land use policies across Europe. Chapter 2 shows that using unconventional datasets, that have previously been unexplored for economic measurement, can help us quantify economics questions that otherwise lack data, and help us put economic policy making on a more sound empirical footing.

In chapter 3, the underlying inflation dynamics of the UK are explored using the price quotes underpinning the Consumer Price Index. These price quote data consists of a medium time-dimension from February 1996 till today ( $T \approx 259$ ), and a large cross-section ( $N \approx 1.2 \text{ million}$ ). The data matrix has a very sparse structure due to a lot of missing observations making it an unbalanced “fat” panel. We find a significant time-variation in the distribution of the inflation rate for the underlying price quotes, with secular trends in the form of increased dispersions and changes to the frequency of price adjustments. A feature of these two trends is that in the aggregate data they cancel each other out, and hence become unobservable. Significant changes to the distribution of underlying price quotes inflation rates are also found following VAT change episodes which indicates the importance of coordination between fiscal and monetary policy. The VAT changes that took place during the financial crisis allowed more firms to reset their price to the optimal level. As a result the easing of monetary policy passed through quicker to price inflation, and hence rendering the impact of the policy changes less effective. Lastly, we find that

---

<sup>3</sup>In some of our cross checks for the quality of the data points, we do explore the time-dimension in the form of previous surveys from 2006 and 2009, thereby extending the sample to  $T = 3$ .

the flexibility index in terms of prices resetting to their optimal level, has a significant effect on the persistence of the inflation rate process. We find evidence that both the Bank of England and professional forecasters can improve their forecasting accuracy for their inflation forecasts by taking into account a flexibility index derived from the underlying micro price quote data.

In chapter 4 I explore the question of whether using the dataset utilised in chapter 3 of disaggregated micro price quotes from the Consumer Price Index can help us quality check the measurement of our inflation rates. To do this, I set out a linear decomposition of the inflation rate from the aggregate index into the individual parts of price quotes. I find a significant effect of missing observations and uncaptured quality adjustment for the inflation rate of a particular example of a highly disaggregated category goods, “Women’s top, long sleeved not blouse”. A feature of the inflation rate for this good is strong seasonal patterns, primarily reflecting the sales seasons as given by fashion cycles. These seasonal patterns are time-varying both in terms of timings and magnitude throughout the sample, which makes it difficult to capture this fashion cycle in the index, and complicates the interpretations of the missing observations and quality change. It also raises some important conceptual issues as to whether it is possible to sample these fashion goods consistently at all across time.

Finally, I will provide my concluding remarks for the thesis in chapter 5. I argue that the use of large and non-standard datasets does indeed enable us to answer macroeconomic questions that are not easily investigated from the aggregated series. These datasets also help us uncover interesting underlying dynamics and trends that are otherwise hidden in aggregate data.

## Chapter 2

# Residential Land Supply in 27 EU Countries: Pigovian Controls or Nimbyism?

### Abstract<sup>1</sup>

We exploit a previously unused dataset of around a quarter of a million survey points that allows us both to derive estimates of residential land on a per capita basis for 27 EU countries, and to model its supply. There is a fairly strong negative correlation between residential land per capita and population density, despite the fact that residential shares are typically very low. In the national data there is also a striking *lack* of correlation between residential land and per capita consumption, but with no indication that this reflects any true economic scarcity value. We model the spatial distribution of residential land assuming that planning policy restricts land supply to match its price to its perceived marginal social cost, allowing both for spatial correlation and the impact on land supply of a consumption externality from nearby housing. Our econometric results provide qualitative support for the model; but it is very hard to match our results to plausible structural parameters unless we assume a social planner who both gives a far greater weight to the impact of the externality than to the welfare gains from new housing, and perceives population density to be far larger than it actually is.

### 2.1 Introduction

Research into the housing market suffers from a paucity of data that allow direct inter-country comparisons of either quantity or absolute prices.<sup>2</sup> In this paper we

---

<sup>1</sup>This chapter is joint work with Professor Stephen Wright.

<sup>2</sup>Datasets such as those of Federal Reserve Bank of Dallas (2015) “*International House Price Database*”, the Bank for International Settlements (2015) “*Residential property statistics*”, the International Monetary Fund (2015) “*Global House Price Index*” or the OECD (2015) “*Focus on house prices*” provide valuable information on price changes, but do not enable direct comparisons of quantities or absolute prices

focus on a key input to housing, residential land, that can be measured directly, thus allowing inter-country comparisons. We use a dataset taken from the European Land Use and Cover Area-Frame Statistical Survey (LUCAS) (Eurostat, 2012) that (as far as we are aware) has not been previously exploited to analyse residential land. This allows us both to derive national and regional estimates of residential land on a per capita basis, and model its spatial distribution and economic determinants, on a consistent basis in 27 EU countries.

LUCAS provides us with a dataset of around one quarter of a million points from a stratified grid covering most of the inhabited geographical area of the 27 EU countries in the survey (which was carried out in 2012). The motivation for LUCAS was primarily to survey land use in agriculture and forestry; but as a by-product it tells us whether any given point in the survey was used for residential purposes, as well as providing some classification of its physical properties. This dataset allows us to carry out two complementary exercises: in measurement and modelling.

### 2.1.1 Measuring residential land in 27 EU countries

In Section 2.2 we use the percentage of points classified as residential to derive area estimates of residential land at both national and regional (Nomenclature of Units for Territorial Statistics level 2 or NUTS2) levels. At a national level these estimates have relatively tight confidence intervals (even after we account for non-trivial degrees of spatial correlation). However precision falls at a regional level; in smaller or more sparsely populated countries; and once we look at subcomponents of residential land.<sup>3</sup>

We focus on five key summary facts derived from the resulting estimates:

1. *Shares of residential land in total land are typically very low.* The median share at a national level is 2.4%, and at a (probably more representative) regional level it is 3.2%.<sup>4</sup> In only one country (Malta), and in fewer than 5% of the regions is the residential share above 20%. Shares of land that are actually built on are typically very much smaller.
2. *Residential land on a per capita basis has a very wide range of cross-sectional variation,* both at the national and regional level. At a national level its cross sectional log standard deviation is similar to that of consumption per

---

<sup>3</sup>In future drafts of this paper we plan to refine the dataset further using targeted photo-identification of characteristics of residential land in the neighbourhood of a subset of points identified in the survey. We are confident that this refinement (which is work in progress) will allow us to draw more confident conclusions on the subdivision of residential land by land cover, most notably identifying the footprint of residential buildings, as well deriving at least provisional estimates of floorspace.

<sup>4</sup>NUTS2 Regions: these are mostly more homogeneous in population terms than countries.

capita; while intra-country regional variations are typically even larger. Yet the *correlation of residential land per capita with national consumption per capita is close to zero*. Since land is a crucial input to housing consumption, and is fairly evidently an imperfect substitute for other inputs, this lack of correlation is quite striking.

3. Combining our land estimates with national accounts data, *estimates of housing expenditure per square metre of residential land display massive variation between EU countries*. For example, the average UK household has housing expenditure of around 22 Euros per square metre of residential land, roughly 10 times as much per square metre as an Hungarian or Polish households (for comparison, UK total consumption per capita is only around 3 times higher).<sup>5</sup>
4. We also combine LUCAS-based estimates of land used for non-residential purposes (predominantly agriculture and forestry) with estimates of value added for these sectors. This allows us to compare value added per square metre of land between housing and non-housing. While there is a modest positive cross-sectional correlation between the two, *the opportunity cost of land is extremely low in comparison to its value added in housing*. Thus it is very hard to explain the lack of correlation of residential land with consumption, noted in Fact 2, in terms of compelling economic demands in other uses. This provides a key motivation for our modelling work.
5. *Regional variations in residential land per capita within most countries are very large*, and in some countries distinctly larger than variation between countries. This variation is not simply a result of the large discrepancies between highly urban regions (such as congested cities) and sparsely populated rural regions.

### **2.1.2 Modelling the spatial distribution of residential land: an implicit model of land supply**

In the remainder of the paper we extend our analysis to the full dataset, by modelling the distribution of residential land at each of the roughly quarter million individual points in the survey. Our model incorporates both spatial correlation (captured by the share of neighbouring points that are classified as residential) and an implicit model of land supply, that is at least qualitatively consistent with a Pigovian equilibrium, that determines land supply (a framework similar in spirit to Cheshire and Sheppard (2002), but incorporating the impact of spatial correlation), which in turn

---

<sup>5</sup>We focus on comparisons where land estimates are reasonably precisely measured. The point estimate for Luxembourg is nearly 50 Euros per square metre, but the land estimates used in that calculation have a very wide margin of error.

determines the unconditional probability that any given point in a given country or region is residential.

In this framework, other people’s consumption of the housing services provided by residential land imposes a consumption externality on those living nearby. This is accentuated by spatial correlation, which, for a given share of residential land in total land, increases the conditional probability that housing will be consumed near other housing. We assume that planning policy deliberately restricts land supply to match its price to its perceived marginal social cost, in an attempt to mimic a Pigovian equilibrium.

In the absence of the externality, and with the opportunity cost of land in other uses assumed to be low and constant (consistent with Fact 4), the model would predict that the relationship between residential land per capita and consumption per capita would simply reflect the income elasticity of residential land. However, with Pigovian controls over land supply the impact of higher consumption will be much more limited. Residential land supply will be higher (and hence prices lower) where total land per capita is higher (since this lowers the unconditional probability of the consumption externality). Thus there will be a negative relationship between housing supply and population density (matching, at least qualitatively, the story told by Miles (2012)), but which would not arise in the absence of the externality. We estimate the model both in implicit form (as a determinant of the unconditional probability that a given point in any region will be residential, taking account of the impact of spatial correlation on the conditional probability) and directly using regional data. Results are very similar. Population density has a strong negative impact on land supply. We find that at a regional level consumption per capita does (as our theory would predict) also have a modest positive impact on supply of land, and hence on regional residential probabilities. However, this is largely obscured in data at a national level, due to large regional variations in population density.

While our econometric results provide qualitative support for our theoretical model, we show in Section 5 that it is very hard to match our reduced form results to plausible structural parameters unless we assume a social planner who is both “Nimbyist” (giving a far greater weight to the impact of the externality than to the welfare gains from new housing) *and* who ignores the impact of spatial correlation in determining the expected utility cost of the consumption externality.

### **2.1.3 Structure of the paper**

In Section 2.2 we describe our dataset and summarise its key features, which provides the basis for our Facts 1 to 5 as outlined above. Section 2.3 sets out our model of Pigovian land supply. Section 2.4 presents our estimation results. In section 2.5 we

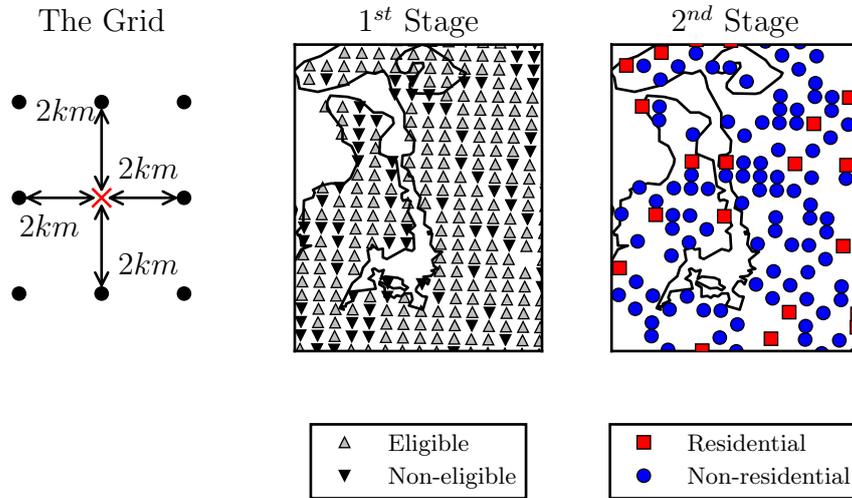
discuss the estimation of the reduced form parameters in relation to the structural parameters of the model. Finally Section 2.6 we draw conclusions from our analysis.

## 2.2 The dataset

### 2.2.1 The LUCAS Methodology

Eurostat’s “*Land Use and Cover Area frame Statistical Survey*” (LUCAS) is a two phase sample survey. The first phase is an equally spaced systematic grid of 1,078,764 observations (in the 2012 sample) in 27 EU countries, separated by 2 km in the four cardinal directions. Each of the points in the first-stage sample was photo-interpreted and classified in terms of land cover,<sup>6</sup> as well as eligibility (based on accessibility) for the second stage of the survey.<sup>7</sup> Together these two classifications give the stratifications of the first stage sample. For the second a subset of 270,277 eligible points from the first-stage sample were visited in person by surveyors. It is the dataset derived from this physical survey that we use in this paper.<sup>8</sup> Figure 1 illustrates for a particular region of the survey.

Figure 2.1: The LUCAS dataset



We consider two different definitions of residential land, using land use and cover definitions from the LUCAS survey (Eurostat, 2013). Our primary focus is on the first definition, of “*broad*” residential land, which uses all survey points classified as

<sup>6</sup>Using the “CORINE” classification. 1: Arable, 2: Permanent Crop, 3: Grassland, 4: Woodland and shrubland, 5: Bareland, 6: Artificial, 7: Water and Wetland

<sup>7</sup>Only points that were both below 1,500m in altitude and accessible by road were included in the second stage.

<sup>8</sup>As noted above, we plan in future drafts to complete a further stage of analysis of the dataset, using photo-identification.

residential by land use (LU “U370” in the dataset). We also consider an alternative “*narrow*” measure that uses only the subset of residential points that are also classified by land cover as artificial structures (land cover A11 “*Buildings with one to three floors*” and A12 “*Buildings with more than three floors*”).

Preliminary cross-checks of the second, narrow definition using photo-identification of a small number of points for a few countries lead us to be wary of data quality for this definition, due to a combination of the relatively small number of points sampled and some evidence of classification biases. For this reason we focus primarily on results using the more robust broad definition, based solely on land use.

We augment the LUCAS dataset by using Eurostat data on population and gross value added at a regional level (from NACE) as well as Consumption, Actual Rent, Implied Rent and Maintenance at a national level from national accounts (NAMA).

## 2.2.2 Estimates of Residential Land and its Components

The key features of the estimates we derive from the LUCAS dataset are summarised in Table 2.1.

### Estimated shares of residential land in total land

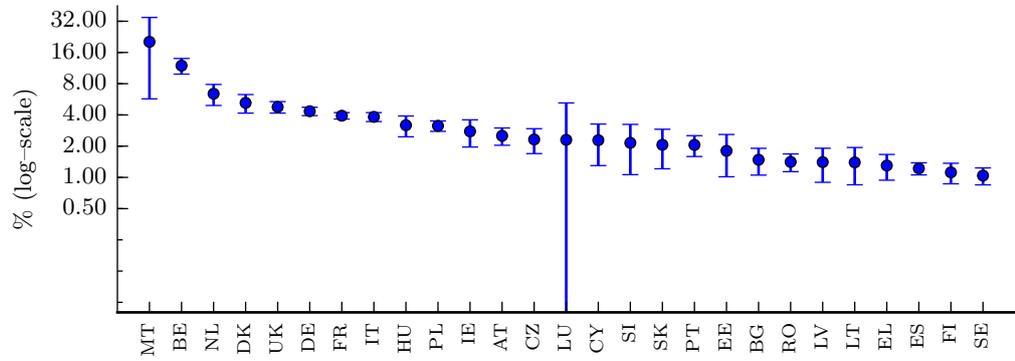
As documented in our Fact 1 and Table 2.1, Figure 2.2 show that for the great majority of EU countries residential land represents only a small proportion of the total land area. Only Malta and Belgium have residential shares in double figures.<sup>9</sup>

We can also do the same calculation for the 261 EU (NUTS2) regions covered by the survey. While individual regional point estimates are subject to nontrivial measurement error (discussed further in the next section) we can nonetheless derive some key features of the regional distribution. Figure 2.3 plots the empirical cumulative density function of residential shares at a regional level. This provides further substantiation of our Fact 1: in the great majority of regions of the EU residential shares are low, whether using our broad or narrow definition of residential land. The median regional residential shares are 3.2% on our broad definition, (less than 1% on the narrow definition). Only 10% of EU regions have broad residential shares in double figures; and only 4% of regions (all of which are major urban areas) have shares above 20%.

---

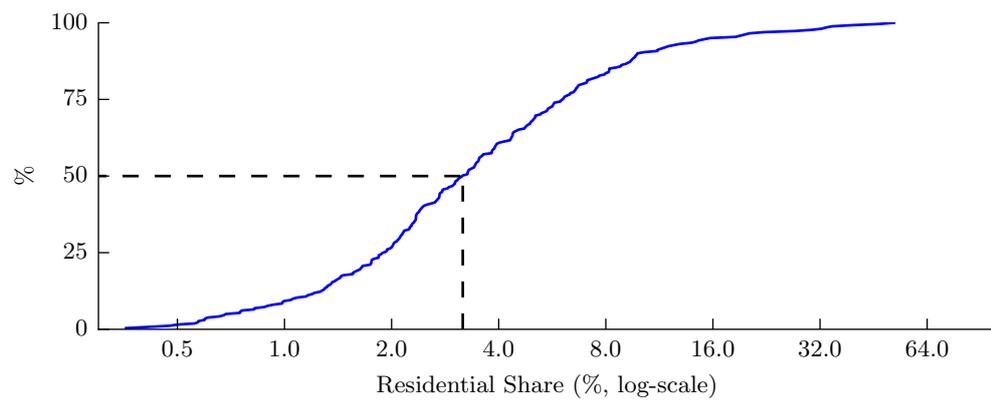
<sup>9</sup>Figure 2.2 shows point estimates together with estimated 95% confidence intervals. These are derived from the estimation procedure described below in Section 4.1, in which the point estimate of the unconditional probability that a given point in the sample is residential is estimated jointly with an estimate of the impact of spatial correlation on the conditional probability. Spatial correlation decreases the precision with which the unconditional probability can be estimated, so that the confidence intervals shown here are typically significantly wider than those that assume points are spatially independent (as in Gallego & Delincé, 2010, equation 12.2). This is a particular problem for a few very small countries (Luxembourg, Cyprus and Malta), for which there are relatively few observations.

Figure 2.2: Shares of residential land in total land in 27 EU countries



Datasources: Eurostat *LUCAS* and *demo\_r\_d3area*

Figure 2.3: Residential Shares by EU Region: Empirical CDFs



Datasources: Eurostat *LUCAS* and *demo\_r\_d3area*

Table 2.1: Summary table of the dataset

	Total Land	Accessible Land	Residential Land	Expenditure on Housing	
				Share of Total Consumption	Per square metre of residential land
				(%)	( $\frac{EUR}{m^2}$ )
	( $m^2$ per capita)	( $m^2$ per capita)	( $m^2$ per capita)		
AT	9,976	5,623	256	14.9	11.2
BE	2,752	2,295	323	18.1	9.2
BG	15,135	10,399	223	16.0	2.6
CY	10,732	6,691	241	15.0	9.7
CZ	7,507	6,152	178	18.2	7.6
DE	4,364	3,187	187	17.6	16.2
DK	7,687	5,569	390	20.6	10.7
EE	34,128	24,166	600	12.7	1.3
EL	11,863	7,959	155	18.8	14.7
ES	10,808	7,291	130	18.8	17.8
FI	62,658	48,492	713	22.5	5.7
FR	9,692	6,797	380	20.3	8.9
HU	9,366	6,519	295	12.9	2.3
IE	15,231	11,810	426	18.8	7.0
IT	5,073	3,807	192	17.4	14.5
LT	21,740	15,103	304	8.9	1.9
LU	4,927	2,959	119	18.5	47.3
LV	31,574	20,105	453	14.3	1.8
MT	757	757	153	8.1	5.8
NL	2,483	1,714	162	17.5	17.5
PL	8,113	6,628	252	10.9	2.5
PT	8,747	5,898	177	13.7	8.1
RO	11,863	8,067	168	16.8	4.0
SE	46,249	29,189	473	20.5	7.7
SI	9,863	8,068	225	12.0	5.4
SK	9,073	6,905	184	12.4	4.7
UK	3,914	3,002	182	21.1	20.4
<i>Average</i>	13,936	9,820	279	16	9.9
<i>Median</i>	9,692	6,691	225	17	8
<i>CoV</i>	1.02	1.04	0.53	0.23	0.94
<i>Max</i>	62,658	48,492	713	22.5	47.3
<i>Min</i>	757	757	119	8.1	1.3

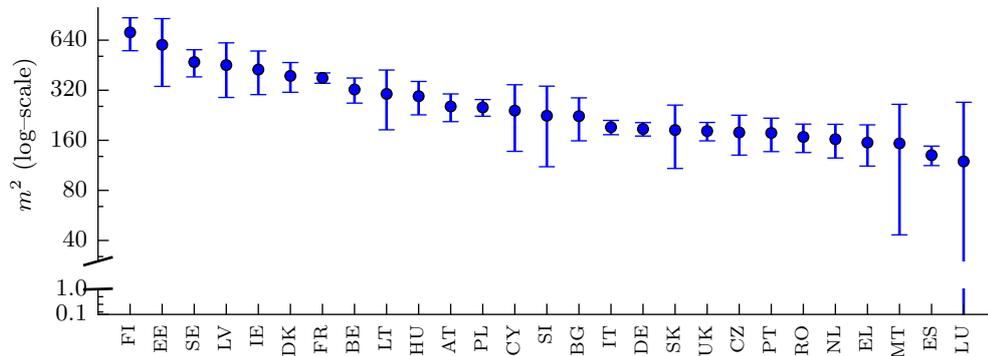
Sources: LUCAS; Eurostat, Accessible Land = area below 1,500 and accessible by road; Residential land = "broad" definition (land use = residential land, all forms of land cover); Housing Expenditure=Rent + Imputed Rent + Maintenance, EUR (Eurostat Table "nama\_10\_co3"); Total Consumption (Eurostat Table "nama\_10\_co3"); CoV: Coefficient of Variation.

## Estimates of residential land per capita

To document our Fact 2 we can also express our land estimates in per capita terms, which again show very large differences both at a national and regional level. Figure 2.4 shows estimates of per capital residential land (broad definition) at a national level, together with estimated 95% confidence intervals.<sup>10</sup>

For most countries the range of sampling uncertainty in the broad measure of residential land is small in comparison with the large differences between countries.<sup>11</sup>

Figure 2.4: Estimates of residential land per capita in 27 EU countries, with 95% confidence intervals



Datasources: Eurostat *LUCAS*, *demo\_r\_d3area* and *demo\_r\_d2jan*

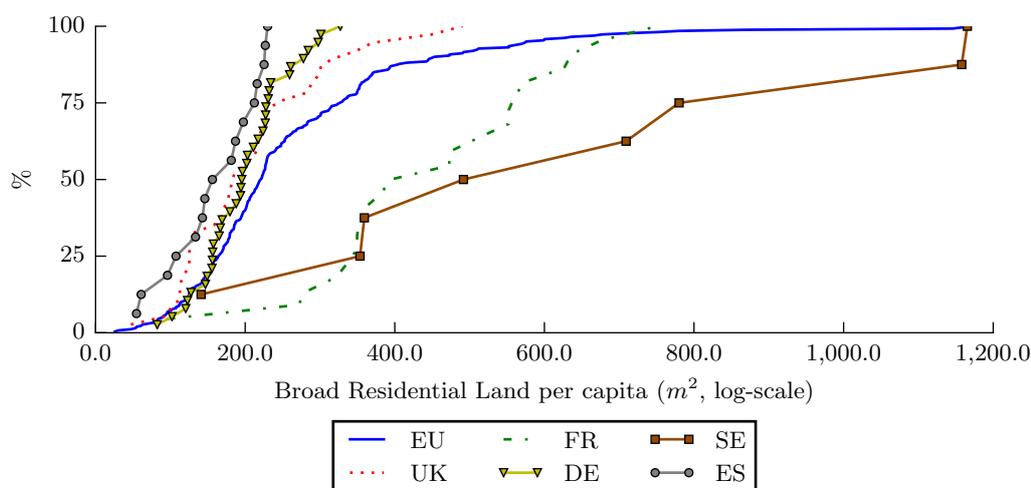
Since measurement error is decreasing in the number of points in any given area, which in turn is roughly proportional to its geographical size, precision of estimates of residential land falls off both for smaller countries and, a fortiori, for any given EU region. Nonetheless we can again derive some key features of the regional distribution, since these do not depend on the precision with which any given point is estimated. Figure 2.5 shows that the degree of variation in (broad) residential land per capita across EU regions, both within the EU as a whole, and within some individual countries, is very much greater than the variation between countries; but the chart also demonstrates the large differences between some countries even for the entire distribution across regions within that country.

Some features of these distributions are unsurprising. We would expect to see relatively low levels of residential land per capita in large cities, where many households live in apartment blocks that by their nature require little land per capita. And indeed the lowest levels of residential land per capita are typically found in cities in most EU countries. However, the large degree of dispersion is not simply

<sup>10</sup>Using the same methodology as outlined in footnote 9 for residential shares.

<sup>11</sup>For the narrow measure of residential land, the small number of observations results in a much wider range of uncertainty. Given the additional problems of the data quality concerns discussed above, in our econometric analysis we focus on results derived from the broad measure.

Figure 2.5: Residential Land Per Capita by Region: Empirical CDFs



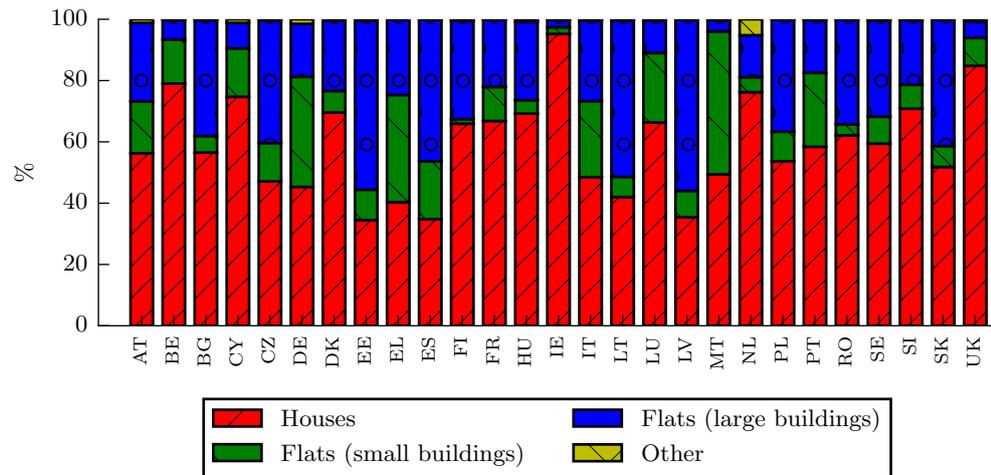
Datasources: Eurostat *LUCAS*, *demo\_r-d3area* and *demo\_r-d2jan*

driven by cities. Nor, in most cases, does it reflect any actual physical constraint, since, as shown in Figure 2.3, in only a handful of EU regions are residential shares of total land sufficiently high that there is simply “no space” for more residential land. In the great majority of EU regions the amount of residential land per capita thus reflects a policy choice, not a physical constraint.

Figure 2.5 also shows some very striking differences between countries. Thus two relatively sparsely populated countries, France and Spain, which Table 2.1 shows have very similar amounts of total land per capita, (the reciprocal of population density) have extremely different regional distributions of land per capita. France has a regional distribution that almost spans the entire EU regional distribution; whereas the range of values across the regions of Spain is very much smaller, with clear-cut dominance by the French distribution. In contrast, the distributions for Germany and the UK (both with similar and distinctly higher rates of population density at a national level) cross, with a considerably larger range of variation in the UK, but around a very similar average value.

To investigate this difference further, figure 2.6 plots the composition of the residential building stocks across the European Union using a Eurostat dataset (see Appendix 2.7.2. As can be seen there is a massive variation in the type of dwelling people live in, with, at one extreme, Ireland with 95.2% of the inhabitants living in single residence houses, whereas, for example, Spain has 65% of its residents living in apartment blocks. While the nature of residential buildings provides some insights into the differences in the distribution of residential land shown in Figure 6, it is, however, not of itself an explanation, since clearly the way in which residential land is utilised is also both endogenous to the price of land and to the restrictions on

Figure 2.6: Composition of building stock



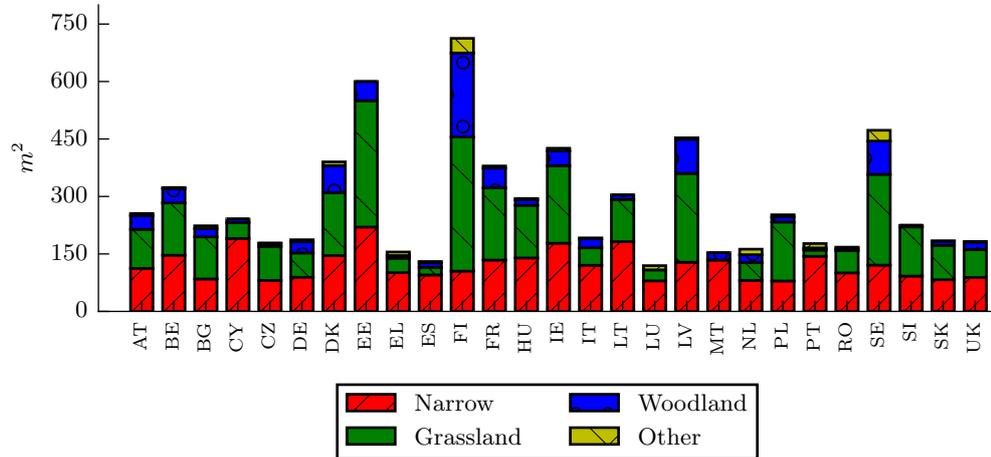
Datasources: Eurostat *ilc\_lwho02*

land supply and land use that planning policy imposes.

## Composition of residential land

Figure 2.7 uses the land cover classification provided by LUCAS to show the breakdown of residential land into its main components.

Figure 2.7: Main components of residential land per capita in 27 EU countries



Datasources: Eurostat *LUCAS* and *demo\_r\_d3area*

There are large differences in composition. At one extreme, in Malta almost all residential land consists of buildings; in contrast, in a number of Scandinavian and Baltic countries, a large proportion of residential land is made up of grass and woodland in gardens.

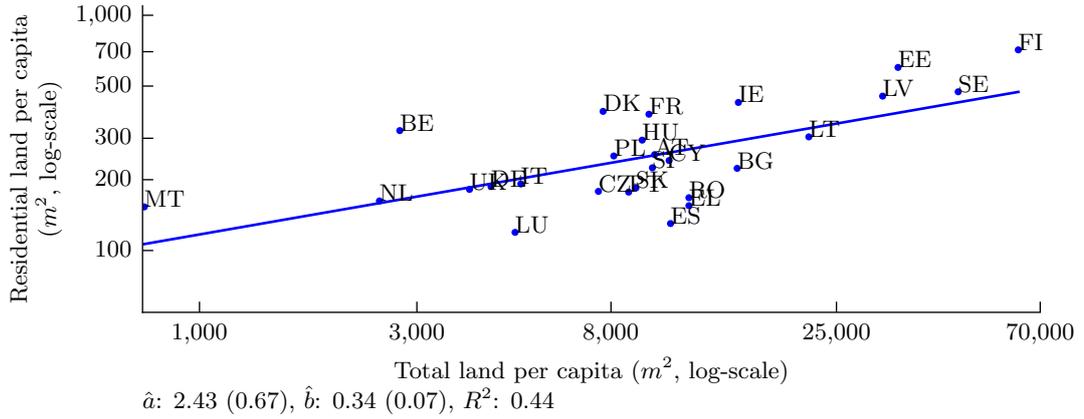
There is no clear-cut a priori case for choosing between the broad measure of residential land, which includes green space, and the narrow definition, which only focuses on buildings. In practice both are clearly subject to regulation (most gardens would, potentially, have space for at least one, often two or more additional houses, but in most countries regulation would not actually permit this additional building). We do, however, as discussed above, have grounds for scepticism about the quality of the data for the narrow measure. Thus, given the limitations of the dataset in its current form, we have considerably more confidence in inferences that can be drawn from the broad measure of residential land.

## Correlations

Figures 2.8 and 2.9 show the nature of the bivariate relationships between (broad) residential land per capita and total land per capita (the reciprocal of population density).

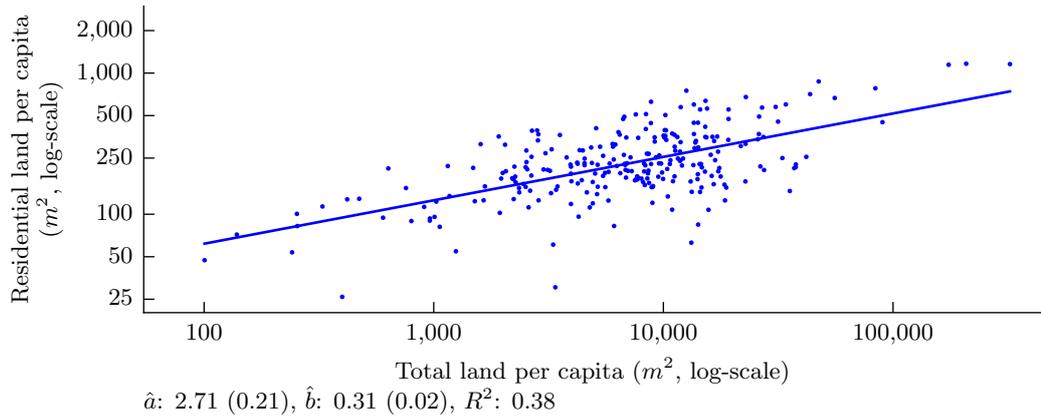
On the face of it, a positive association might seem unsurprising: more sparsely populated countries might be viewed as having “more space” for houses and gardens. However, a glance at Figure 2.2 should give pause for thought. For the overwhelming

Figure 2.8: Residential vs. Total land: National Data



Eurostat: Eurostat *LUCAS* and *demo\_r\_d3area*

Figure 2.9: Residential vs. Total land: Regional Data



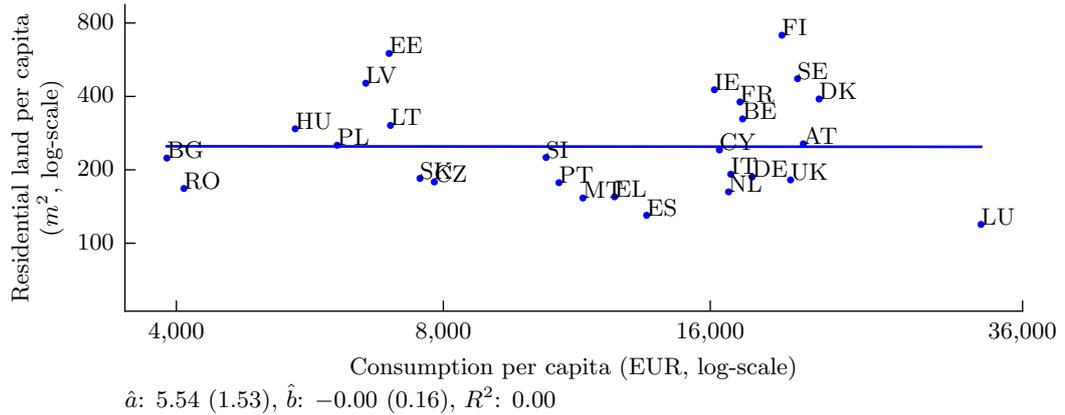
Eurostat: Eurostat *LUCAS* and *demo\_r\_d3area*

majority of countries and regions residential shares are so small that this argument is distinctly less plausible. We therefore conclude that we need to look for other explanations in our modelling.

We now turn to the (lack of) relationship with consumption per capita, as in our Fact 2.

Figure 2.10 show a distinct lack of any clear-cut bivariate relationship between residential land and consumption per capita across the cross section of 27 countries (our fact 2). This lack of relationship is itself very striking. As a key input to housing services, and indeed as a consumption good (“*space*”) in its own right, it appears unlikely on *a priori* grounds that consumption of the services of residential land is a borderline inferior good. Our model of Pigovian land supply set out below provides a rationale for such a weak correlation; our empirical results also show that the lack

Figure 2.10: Residential land per capita vs. consumption per capita



Eurostat: Eurostat *LUCAS*, *demo\_r\_d3area*, *demo\_r\_d2jan* and *nama\_a64*

of a simple correlation across the cross-section of countries in our sample conceals some (albeit fairly weak) impact of national consumption at a regional level, once we factor in the impact of other determinants.

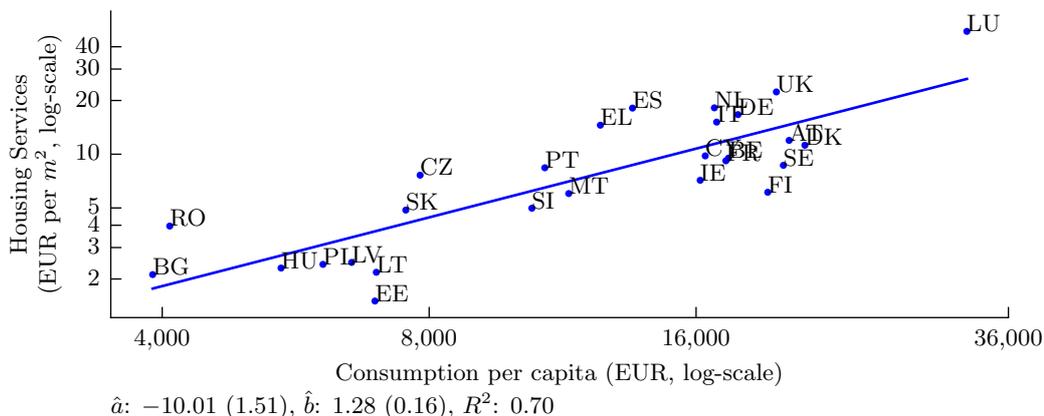
### 2.2.3 Estimates of housing expenditure per square metre and the opportunity cost of housing

We can also combine our land estimates derived from LUCAS with data from the national accounts (see Appendix 2.7.2) to derive estimates of housing expenditure, in Euros, per square metre of (broad) residential land, for each of the 27 countries in our sample. Table 2.1 shows the resulting figures, which, as noted in our Fact 3, show an extremely wide range of variation. As Figure 2.11 illustrates, in contrast to the lack of correlation of residential land with consumption, the single strongest explanatory factor for variation in housing expenditure per  $m^2$  is cross-sectional variation in total consumption per capita.

One possible supply-side-based explanation for the variation in housing expenditure (or, equivalently, value-added from the housing sector) per square metre of residential land might in principle be if there are also significant cross-sectional differences in the opportunity of land in other potential uses. Figure 2.12 shows that, in the majority of EU countries, agriculture (and to a lesser extent forestry) is the dominant alternative use.

We therefore construct two proxy measures of opportunity cost by dividing value added in forestry or agriculture (both measures from the national accounts, Appendix 2.7.2), by the respective areas used for these purposes. In figure 2.13 we plot the scatter plots of these measures of the opportunity cost of land against our measure of value added from housing, both measured on a comparable basis, per

Figure 2.11: Housing Expenditure per  $m^2$  vs Consumption per Capita in 27 EU Countries



Datasources: Eurostat *nama\_co* and *demo\_r\_d2jan*

square metre of land.

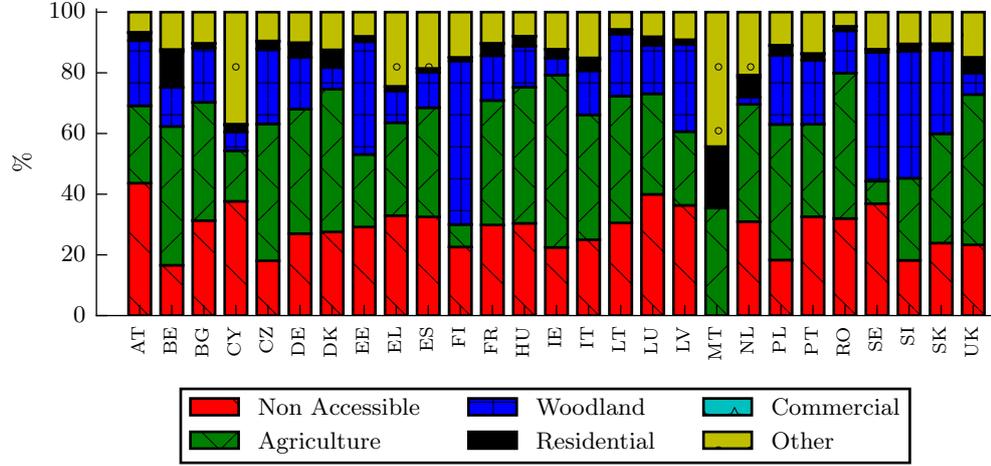
Figure 2.13 shows that there is indeed a positive correlation between our measure of value added from broad housing services and the opportunity cost of land used in agriculture (with a coefficient of 0.71 on the log values). However a closer look at the chart shows that this correlation cannot be viewed as representing any true economic relationship, since, as noted in Fact 4, the value added per  $m^2$  in the two sectors differ dramatically in magnitude. Where housing services has a value added per square metre of land between 1 and 50 euros, all our measures of value added from agricultural are below 1 euro.<sup>12</sup>

The correlation with the opportunity cost measure for forestry is not statistically significant and the differences in scales are even bigger.

Thus we conclude that there is no plausible explanation for the observed variation in housing expenditure per square metre in terms of the opportunity cost in competing uses of land.

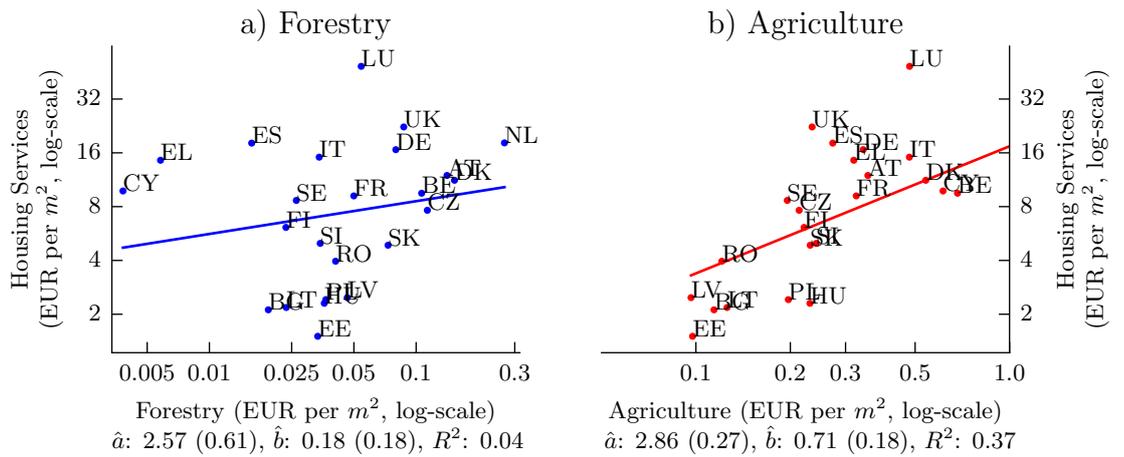
<sup>12</sup>We speculate that the correlation might be driven by the historical importance of agricultural productivity in determining land use in predominantly agrarian economies.

Figure 2.12: Composition of Total Land



Datasource: Eurostat *LUCAS*

Figure 2.13: Housing Expenditure per  $m^2$  and the Opportunity Cost of Land



Datasources: Eurostat *LUCAS*, Consumption and Regional Accounts

## 2.3 A simple model of Pigovian land supply

### 2.3.1 Distribution of Households

A set  $\mathcal{H}$  of households,  $h = 1, \dots, H$  live on a circle of points,  $i = 1, \dots, L \in \mathcal{L}$ . Nature allocates households to an “address” (a point on the circle) via a random mapping from  $h$  to  $i(h) \in \mathcal{I} \subset \mathcal{L}$ .

We assume that the households are distributed clockwise around the circle by a Markov chain, with the transition matrix given by:

$$\mathbf{M} = \begin{bmatrix} \Phi & 1 - \Phi \\ 1 - \gamma & \gamma \end{bmatrix} \quad (2.3.1)$$

with the state vector

$$\mathbf{x}_i \equiv \begin{bmatrix} 1_{(i \in \mathcal{I})} \\ 1_{(i \notin \mathcal{I})} \end{bmatrix} \quad (2.3.2)$$

where the first element of  $x_i$  is equal to 1 if the point  $i$  is a household address and zero otherwise. This implies that the (notional) law of motion for the states around the circle (which in turn determines the distribution of addresses) is given by:

$$\mathbf{x}_{i+1} = \mathbf{M}\mathbf{x}_i + \mathbf{u}_{i+1}. \quad (2.3.3)$$

Let  $l = L/H$ , be per capita land, hence population density,  $D = l^{-1}$ . To ensure the correct steady state distribution we must have:

$$\Phi = D + (1 - D)s \quad (2.3.4)$$

$$\gamma = 1 - D(1 - s) \quad (2.3.5)$$

where  $s \in (0, 1)$  is a parameter determining spatial correlation, with  $s = 0$  implying no spatial correlation. To simplify the algebra we assume that Nature repeats the allocation, until it reaches a finite sample with allocations that exactly match the steady-state distribution.

Conceptually the positive spatial correlation can be interpreted as a manifestation of either current (fixed) or historical positive externalities. Appendix 2.7.1 outlines a simple model of a historical building phase with positive externalities results in positive spatial correlations (clustering). However, in the overall perspective we argue that at the margin the negative externality most dominate, as the counterfactual of people being willing to pay more for living in ever more crowded living conditions seems counter-intuitive to us (space is still a premium).

### 2.3.2 Housing Technology

We assume a very simple housing technology that is intended to capture intensity of land use, which in turn determines the nature of the consumption externality.

We assume that households can only build houses to the left of their address, but also that the distance between household  $h$  and household  $h + 1$  (with addresses  $i(h)$  and  $i(h + 1)$ ) determines the nature of housing at address  $i(h + 1)$ . To simplify the algebra we assume that the housing technology available to household  $h$  is restricted to “bungalows” which are one storey high, and  $R_h$  points wide, or “high rises” which occupy a single point, but are  $R_h$  storeys high. This in turn determines the intensity of housing immediately to the right of household  $h$  (at the point  $i(h) + 1$ ) and hence the consumption externality imposed on household  $h$  by household  $h + 1$ .

The intensity of housing at the point to the right of household  $h$ 's address is given by

$$R_{i(h)+1} = \begin{cases} R_{h+1} & i(h + 1) - i(h) = 1 \\ 1 & i(h + 1) - i(h) = R_{h+1} \\ 0 & \text{Otherwise} \end{cases} \quad (2.3.6)$$

If the distance between the addresses is less than  $R_{h+1}$ , then the technology imposes that household  $h + 1$  must live in a  $R_{h+1}$  storey high rise. In contrast if the two households have addresses  $R_{h+1}$  or more points apart, household  $h + 1$  lives in a  $R_{h+1}$  point wide bungalow.

The impact on household  $h$  is that the neighbouring point,  $i(h) + 1$  will either have a high rise (if two addresses are on adjacent points) or a bungalow (if the two addresses are exactly  $R_{h+1}$  points apart), or no housing at all.<sup>13</sup>

### 2.3.3 Private Utility and Equilibrium

Each household is identical except in respect of the externality imposed by its neighbour. Household utility of household  $h$  (living at address  $i(h)$ ) is given by

$$\max_{G_h, R_h} U_h = (1 - \alpha) \ln(G_h - G^*) + \alpha \ln R_h + \ln(E - R_{i(h)+1}), \quad (2.3.7)$$

where  $G_h$  is nonresidential consumption, for which there may be a minimum subsistence level,  $G^* > 0$ . If so,  $R_h$  will be a superior good, with loglinearised income elasticity  $\frac{C_h}{C_h - G^*} > 1$ , where  $C_h$  is total consumption. We assume for simplicity that an  $R_h$  wide bungalow and an  $R_h$  storey high rise generate the same flow of housing consumption (for simplicity, measured in the same units). Thus congestion of hous-

---

<sup>13</sup>It is possible to allow for intermediate cases where houses may have varying heights, depending on the distance between addresses, but given the nature of the distribution of addresses this actually makes very little difference, while complicating the algebra nontrivially.

ing only matters to the extent that it may imply a greater intensity of housing at the adjacent point.<sup>14</sup>

The specification implies convex utility costs to household  $h$ , living at address  $i(h)$ , of housing in the neighbouring point  $i(h) + 1$ , which can be rationalised in terms of claims on finite resources in the local “*environment*” ( $E$ ). The stock of finite resources,  $E$  is treated as an endowment. We show below that the magnitude of  $E$  alone allows a sufficient parameterisation of the problem in terms of utility costs.<sup>15</sup>

Given the additive separable nature of the problem the externality has no direct impact on any household’s choices (although it will, as we shall show, have an indirect effect via the price of housing). The budget constraint faced by household  $h$  is given by:

$$G_h + QR_h = C_h, \quad (2.3.8)$$

where  $C_h$  is total consumption of other goods and housing services and  $Q$  is the price of housing services.

Optimising Equation 2.3.7 with respect to the constraint given in Equation 2.3.8 implies the following private demand function:

$$R_h = \alpha \frac{\widehat{C}_h}{Q}, \quad (2.3.9)$$

where  $\widehat{C}_h = C_h - G^*$  is surplus consumption.

Substituting back into the utility function and combining with the optimality condition, gives the indirect utility for household  $h$ :

$$V_h = V \left( \widehat{C}_h, Q, R_{i(h)+1} \right) = \ln \widehat{C}_h - \alpha \ln Q + \ln (E - R_{i(h)+1}) + \mathcal{C} \quad (2.3.10)$$

where  $\mathcal{C} = \ln((1 - \alpha)^{1-\alpha} \alpha^\alpha)$ . Hence we can straightforwardly calculate the consumption equivalent loss of utility (as a share of total consumption) for those with housing in neighbouring points using

$$V \left( (1 - \kappa) \widehat{C}_h, Q, 0 \right) = V \left( \widehat{C}_h, Q, R_{i(h)+1} \right), \quad (2.3.11)$$

which if we combine the indirect utility expressions yields:

$$\kappa(R_{i(h)+1}) = 1 - \frac{E - R_{i(h)+1}}{E} \quad (2.3.12)$$

---

<sup>14</sup>Allowing for this additional effect would complicate the algebra without changing the nature of the social planner’s problem, as set out below, since it would simply accentuate the expected utility impact of additional housing.

<sup>15</sup>We have also experimented with a specification in which household  $h$  also cares about intensity of housing at other neighbouring points but again this merely complicates the algebra without changing the nature of the problem.

If we investigate the limits of this expression, we have:

$$\lim_{E \rightarrow R_{i(h)+1}} \kappa(R_{i(h)+1}) = 1 \quad \text{and} \quad \lim_{E \rightarrow \infty} \kappa(R_{i(h)+1}) = 0,$$

the consumption equivalent impact of the externality is monotonically increasing in  $R_{i(h)+1}$  (the intensity of residential consumption at the next point) and decreasing in  $E$ , the scale of the exogenous “environment”, which thus provides a sufficient parameterisation for the magnitude of the externality.

### 2.3.4 Social Welfare Maximisation

The social planner acts behind a veil of ignorance and chooses the aggregate supply of residential land ( $R = \mathbb{E}_{\mathcal{H}} R_h$ ) to maximise a social welfare function given by the expected utility of a household chosen randomly from the set of  $\mathcal{H}$ . We thus deliberately restrict our social planner to make a policy choice only at the aggregate level, with the impact of the resulting externality on individual households being random.<sup>16</sup>

To simplify the analysis we assume that all households are identical, except for the externality.<sup>17</sup> The expected utility of the randomly chosen household that the social planner maximises is then given by:

$$\max_R W = \mathbb{E}_{\mathcal{H}} U_h = (1 - \alpha) \ln(G - G^*) + \alpha \ln R + \mathbb{E}_{\mathcal{H}} \ln(E - R_{i(h)+1}) \quad (2.3.13)$$

Substituting in the probability of the externality binding (and using the property of symmetric households), we can write this as

$$\begin{aligned} \max_R W &= (1 - \alpha) \ln(G - G^*) + (1 - \Phi) \ln E \\ &\quad + \alpha \ln R + \Phi \ln(E - R) + (1 - \Phi)(1 - \gamma) \gamma^{R-2} \ln\left(\frac{E - 1}{E}\right) \end{aligned}$$

where, clearly, we must assume  $E - R > 0$  to ensure a solution.

---

<sup>16</sup>A discussant of an earlier draft of this paper pointed out, correctly, that in our modelling framework a social planner with sufficient discretionary powers could in principle increase aggregate housing with *no* increase in the aggregate consumption externality, by careful choice of where new housing is supplied (i.e., by only allowing increased housing at addresses that are sufficiently far away from the neighbouring address). However, such a discretionary policy at a micro level would have clear distributional impacts; it is also far from clear what the implementable equivalent of policy of this kind would be.

<sup>17</sup>It is easy to extend the model to allow for differences in incomes between households, but this does not change any of the results, since our social planner has no tools to deal with inequality. Note also that we do not allow for *new* houses; housing supply in our framework is increased simply by all households having larger houses.

We can then define a function for the expected marginal disutility of the externality as

$$\mathcal{F}(R; \Phi, \gamma, E) = \overbrace{\Phi \left( \frac{1}{E-R} \right)}^{>0} - \overbrace{(R-2)(1-\Phi)(1-\gamma)\gamma^{R-3} \ln \left( \frac{E-1}{E} \right)}^{>0} \quad (2.3.14)$$

where the first term captures the marginal disutility of a larger high rise at the adjacent point (with constant probability  $\Phi$ ), while the second term captures the resulting higher probability of having the edge of a bungalow at the next point (with associated constant disutility).<sup>18</sup>

Assuming a unit marginal rate of technical substitution between housing and non-housing we have (using Equation 2.3.8 and 2.3.9)

$$\frac{W'_R}{W'_G} = \frac{\alpha}{1-\alpha} \frac{\widehat{G}}{R} - \frac{\widehat{G}}{1-\alpha} \mathcal{F}(R; \Phi, \gamma, E) = 1 = \text{MRT} \quad (2.3.15)$$

implying

$$\frac{R}{\widehat{C}} = \frac{\alpha}{1 + \mathcal{F}(R; \Phi, \gamma, E) \times \widehat{C}}. \quad (2.3.16)$$

Since we have from the private optimisation that  $\frac{QR}{C} = \alpha$ , this implies that the price is given by:

$$Q = 1 + \mathcal{F}(R; \Phi, \gamma, E) \times \widehat{C}, \quad (2.3.17)$$

i.e., the expenditure share of housing in surplus consumption ( $\widehat{C} = C - G^*$ ) is constant and equal to  $\alpha$  (implying a rising share in total consumption) but Pigovian land supply implies that the share of real housing in surplus consumption (determined implicitly by (2.3.16) is decreasing in the externality term  $\mathcal{F}(R; \Phi, \gamma, E)$ , requiring an increase in the price of land,  $Q$ .

In contrast, in the absence of an externality the marginal social cost function would simply be an invariant horizontal line, so that a 1% shift in aggregate consumption would simply cause a  $\left(\frac{C}{C-G^*}\right)$  % rise in residential land supply.  $Q$  is increasing in  $C$ , and negative externality implies a higher  $Q$ , but if the absent of externalities  $F = 0$  and  $Q = 1$  for a given  $C$ . Gives indirect evidence for negative externality to dominate which is also found in the value added per square meters. If positive externality dominated there would be no reason to add value. This is in contrast to Lucas (2001) who in his paper has an implicit assumption of a fixed supply of land and focuses on the demand side only, whereas our focus is on both

---

<sup>18</sup>Given that we measure housing in integer values  $\mathcal{F}(R; \Phi, \gamma, E)$  is strictly a first-order Taylor series approximation for the disutility impact of increasing  $R$  by one point. For sufficiently large  $R$  this approximation becomes arbitrarily good.

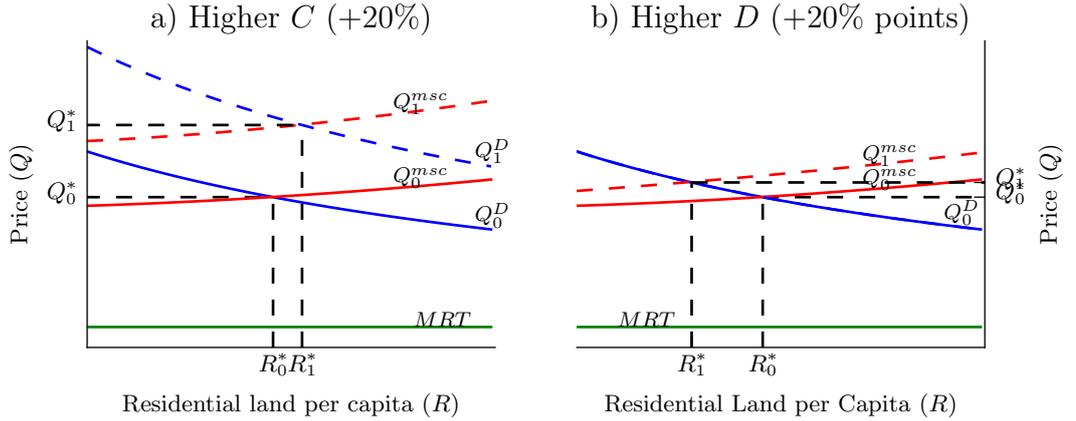
the demand and supply side. Further the only driver of land demand in Lucas (2001) is transportation costs (for households) and agglomeration effects for producers, without any social planner to intervene and restrict supply, which at least for a European context is unrealistic as we know that a large part of what determines land use in is planning controls.

### 2.3.5 Geometry

The model is given by three simple elements above (1) the private demand function from Equation 2.3.9 (2) Marginal Rate of Transformation (set equal to one) and (3) The Pigovian marginal social cost, found from the planner problem in Equation 2.3.17.

Figure 2.14 plots comparative statics for higher consumption ( $\Delta C > 0$ ) and higher population density ( $\Delta D > 0$ ):

Figure 2.14: Comparative Statics



In response to a shift in consumption both demand curve and marginal social cost functions shift, but the latter shifts by strictly less, so the new equilibrium implies increases in both price and quantity; higher population density shifts the MSC curve only, thus implies a higher land price but lower land supply per capita.

### 2.3.6 Log-Linearisation of the Model

In the Appendix we show that we can log-linearise the model as

$$\tilde{r} = \lambda_c \tilde{c} - \lambda_d \tilde{d} \quad (2.3.18)$$

where  $\tilde{r}$ ,  $\tilde{c}$  and  $\tilde{d}$  are log deviations around an equilibrium where  $\hat{C} = \bar{C}$  and  $D = \bar{D}$  are some mean values (eg across our cross-section) but land supply ignores

the externality. We then show that we can write

$$\lambda_c = \lambda_c(\bar{D}, s, \alpha, \kappa_\alpha) \quad \text{and} \quad \lambda_d = \lambda_d(\bar{D}, s, \alpha, \kappa_\alpha),$$

with  $\lambda_c, \lambda_d \in (0, 1)$ , where  $\kappa_\alpha$  is the consumption equivalent cost of the externality (as defined in (2.3.12) to household  $h$  if  $R_{i(h)+1} = \bar{R}$ , in a non-Pigovian equilibrium such that  $Q = 1$  (which straightforwardly implies  $\bar{R} = \alpha \bar{C}$ ). It is straightforward to show that in this restricted case we would then have

$$\lambda_c(\bar{D}, s, \alpha, 0) = \eta_c = \frac{\bar{C}}{\bar{C} - G^*} \quad \text{and} \quad \lambda_d(\bar{D}, s, \alpha, 0) = 0,$$

that is, in the absence of the externality population density would have no impact on land supply, and land supply would simply be determined by the income elasticity of residential land services.

In the next section we show that we can estimate the reduced form parameters  $\lambda_c$  and  $\lambda_d$  econometrically, as determinants of the probability that any given point in our dataset will be residential, and that this simple log-linear reduced form has very good explanatory power for national and regional differences in this probability. In Section 2.5, we ask whether we can make sense of these reduced form estimates in terms of plausible structural parameters.

## 2.4 Estimation

We estimate reduced form parameters on point level data: we show below that it is straightforward to translate the reduced form of our model to the unconditional probability that any given point in our sample will be residential. This allows us to use point-level regressors (in particular to capture spatial correlation) alongside regional and national regressors where available. As a cross-check we also estimate the reduced form using regional and national estimates of residential land as described in Section 2.2.<sup>19</sup>

We first consider the impact of spatial correlation, which is generic to all our point-level estimation methods.

### 2.4.1 Capturing spatial correlation

We have a set of points,  $\{q_{ijk}\}$ ,  $i = 1, \dots, 269, 328$ ,<sup>20</sup> which are zero (non-residential) or 1 (residential), in  $j = 1, \dots, 261$  regions, and  $k = 1, \dots, 27$  countries.

<sup>19</sup>Note that regional and national specifications are not exact aggregations of the point level specification, but they do provide a cross-check.

<sup>20</sup>We lose a few numbers observations, from the points without any neighbouring points close it to

We model the conditional probability that a given point is residential as

$$\mathbb{P}(q_{ijk} = 1 \mid s_{ijk}) = \phi_{jk}s_{ijk} + (1 - \phi_{jk})p_{jk}, \quad (2.4.1)$$

where  $s_{ijk}$  is the spatial autoregressive term (defined precisely below) which can be interpreted straightforwardly as the local residential share in the neighbourhood of a given point,  $q_{ijk}$ . The parameter  $\phi_{jk}$  captures the strength of spatial correlation: the higher is the local residential share,  $s_{ijk}$  the higher is the probability that point  $q_{ijk}$  will be residential. We allow for the possibility of heterogeneity in spatial correlation at a regional level, at a national level ( $\phi_{jk} = \phi_k$ ); or homogeneity across the dataset ( $\phi_{jk} = \phi$ ).<sup>21</sup>

$p_{jk} = \mathbb{P}(q_{ijk} = 1)$  is the unconditional probability that a given point will be residential. Initially we model  $p_{jk}$  using dummy variables for region  $j$ , country  $k$ ; we then proceed to derive it from the reduced form of our theoretical model, using regional and country level regressors.

In both cases our model is a Spatial Autoregressive model (SAR) as in Vega and Elhorst (2013) and (Pesaran, 2015, pp.797–816) and the spatial autoregressive term can be written as:

$$\mathbf{s} = \mathbf{W}\mathbf{q},$$

where  $\mathbf{q}$  is the vector of individual points stacked ( $q_{ijk}$ ) and  $\mathbf{W}$  is the spatial lag matrix given by:

$$\mathbf{W} = \begin{bmatrix} w_{11} & w_{12} & \cdots & w_{1n} \\ w_{21} & w_{22} & & \\ \vdots & & \ddots & \\ w_{n1} & & & w_{nn} \end{bmatrix}$$

with the restrictions that the individual elements in the matrix is given by:

$$w_{ij} = \begin{cases} 1/N & \forall i \neq j \text{ and } j \text{ is part of nearest } N \text{ points to } i \\ 0 & \text{Otherwise} \end{cases}$$

The unconditional probability,  $p_{jk}$ , in region  $j$  of country  $k$  is given by:

$$p_{jk} = \mathbb{E}_i(q_{ijk}) = \frac{R_{jk}}{L_{jk}} \quad (2.4.2)$$

---

<sup>21</sup>There is clearly a conceptual link between  $\Phi$  in our 1-dimensional theoretical model (the probability that there will be houses at two adjacent addresses) and its empirical counterpart  $\phi$  in two dimensions; but we cannot draw a precise analytical link, other than to note that we would expect a higher  $\phi$  to correspond to a higher  $\Phi$ .

where  $R_{jk}$  is residential land (not directly observable), and  $L_{jk}$  is total land in region  $j$  of country  $k$  (which we can take to be perfectly observable). Since  $R_{jk}$  is not directly observable, nor can  $p_{jk}$  be. We assume the expected value of the spatial correlation term to be equal to the unconditional probability.<sup>22</sup>

In Table 2.2 we first show the results of estimating Equation 2.4.1, modeling the unconditional probabilities as either constants or by national or regional dummies, together with the effect of spatial correlation, either homogeneous or with country regional heterogeneity:

$$q_{ijk} = \hat{\phi}_{jk}s_{ijk} + (1 - \hat{\phi}_{jk})\hat{p}_{jk} + \hat{u}_{ijk} \quad (2.4.3)$$

Table 2.2: The probability that any LUCAS survey point is classified as residential: Spatial Correlation and national vs regional dummy variables

	Broad Residential					
	B1	B2	B3	B4	B5	B6
<i>Dummy variables for unconditional probabilities</i>						
Constant: $\hat{p}_{jk} = \hat{p}$	✓					
National: $\hat{p}_{jk} = \hat{p}_k$		✓	✓			
Regional: $\hat{p}_{jk}$				✓	✓	✓
<i>Group Spatial Correlation Parameters</i>						
Homogenous: $\hat{\phi}$	0.497 (0.01)	0.425 (0.014)		0.309 (0.087)		
National: $\bar{\hat{\phi}}_k$			0.345 (0.179)		0.264 (0.173)	
Regional: $\bar{\hat{\phi}}_{jk}$						0.155 (0.377)
$R^2$	0.022	0.025	0.026	0.031	0.032	0.035
$SSR$	8773	8744	8735	8691	8683	8655
$AIC$	-3.42	-3.43	-3.43	-3.43	-3.43	-3.43
$BIC$	-3.42	-3.42	-3.42	-3.41	-3.41	-3.39
<i>Parameters</i>	2	28	54	262	288	522

Observations: 269, 238. Standard errors in brackets. Estimation of Equation 2.4.3:  $q_{ijk} = \hat{\phi}_{jk}s_{ijk} + (1 - \hat{\phi}_{jk})\hat{p}_{jk} + \hat{u}_{ijk}$ , with the unconditional probabilities ( $\hat{p}_{jk}$ ) specified by either a constant term, national or regional dummy.  $q_{ijk}$  is equal to 1 if a given point is classified as residential, and 0 otherwise. Coefficients with a bar above are the mean group estimates (Pesaran et al., 1996) given by:  $\bar{\hat{\phi}}_k = \frac{1}{K} \sum_k \hat{\phi}_k$  and  $\bar{\hat{p}}_k = \frac{1}{K} \sum_k \hat{p}_k$ . The standard errors are calculated using (Pesaran et al., 1996) formula, and reflects, especially for the model B4-B6 the large skewness and kurtosis in the data as demonstrated in Figure 2.3.

As can be seen from the table, this specification requires a large number of para-

<sup>22</sup>This very close to holding in the dataset. Small differences only arise when the adjacent points are in other regions but in practice this only marginally changes the results.

meters, to capture heterogeneity both in unconditional probabilities and in spatial correlation (522 parameters are required at a regional level). Point estimates of regional and local probabilities (not shown in table) are always extremely close to simple averages of the share of residential points.<sup>23</sup> The Mean Group estimates of the spatial autoregressive coefficients are calculated as outlined by (Pesaran et al., 1996, pp.155–156).

As is to be expected, given the relatively low frequency of residential points in most regions, the notional fit of the specification is very poor. As would also be expected, the more dummy variables in the model, and the greater the degree of heterogeneity allowed in the spatial correlation parameter, the lower are the resulting estimates of the average degree of spatial correlation. There is a clear analogy here with estimation of autoregressive parameters in time series data, where it is a standard result that a sufficiently large number of dummy variables for different time periods will, in the limit, capture a large proportion of serial correlation, but typically only with the benefit of hindsight. In our case, while we treat regional dummies as exogenous, they are clearly not: the definition of regions post dates the emergence of population clusters.

Strikingly, however, the key determinant of improved fit appears to be heterogeneity in the unconditional probability: heterogeneity in spatial correlation has much more marginal impact, and, particularly in regions and smaller countries, heterogeneous spatial coefficients are not well-determined. For this reason in deriving the confidence intervals shown in Figures 2 and 3 we using standard errors for unconditional probabilities from the model with homogeneity of spatial correlation imposed (model B2).

While the spatial autoregressive parameters in Table are not large (certainly in comparison with their temporal correlation equivalents) they are strongly significant. It should also be borne in mind that the average distance to the set of 20 nearest points used in calculating the spatial term  $s_{ijk}$  is typically around 7km (given only partial sampling of points on the grid in the first stage sample that are each 2km apart) - well beyond the distances at which significant externalities are likely to occur. By implication the degree of spatial correlation at the shorter distances at which externalities are likely to occur must be very much greater.<sup>24</sup>

---

<sup>23</sup>These in turn are very close to, but not identical to, the residential shares shown in Figures 2 and 3, which are calculated on the basis of total, rather than accessible land, using a weighting scheme on individual points reflecting the frequency of sampling in that region of the grid. In practice however the differences are typically very small.

<sup>24</sup>One objective of our proposed further stage of photo-identification of the dataset will be to investigate the nature of spatial correlation at much shorter distances.

## 2.4.2 Estimating reduced form parameters

### Deriving the econometric specification

In Section 4.1 we estimated regional and national residential probabilities directly, using dummy variables, to allow us to focus on the estimates of spatial correlation, but at the cost of a very large number of parameters.

By definition the (true) probability of a point being residential is given by

$$p_{jk} = \frac{R_{jk}}{L_{jk}} = \frac{R_{jk}}{H_{jk}} \frac{H_{jk}}{L_{jk}} = \frac{R_{jk}}{H_{jk}} \exp(d_{jk}), \quad (2.4.4)$$

where  $\exp(d_{jk}) = D_{jk} = \frac{H_{jk}}{L_{jk}}$  is the population density. We now attempt to estimate the determinants of these regional probabilities, using the reduced form of the model of Pigovian land supply set out in Section 2.3, which can be written as the log of the implied predictor

$$\ln\left(\widehat{\frac{R_{jk}}{H_{j,k}}}\right) = \beta + \lambda_c \ln\left(\frac{C_{jk}}{H_{jk}}\right) - \lambda_d \ln\left(\frac{H_{jk}}{L_{jk}}\right) \quad (2.4.5)$$

where  $H_{jk}$  is population,  $C_{jk}$  is aggregate consumption,  $R_{jk}$  is residential land (which cannot be measured directly from our dataset) and  $L_{jk}$  is total land, in region  $j$ , country  $k$ . Defining  $c_{jk} = \ln\left(\frac{C_{jk}}{H_{jk}}\right)$  as the log of consumption per capita, and combining the two equation with an approximation error they can be transformed into

$$\ln(\hat{p}_{jk}) = \ln\left(\widehat{\frac{R_{jk}}{H_{jk}} \frac{H_{jk}}{L_{jk}}}\right) = \beta + \lambda_c c_{jk} + (1 - \lambda_d) d_{jk}, \quad (2.4.6)$$

where, again  $p_{jk} = \mathbb{E}_i(q_{ijk})$  is the implied unconditional probability that any given point in region  $j$  of country  $k$  is residential. Further we have that the true probability for region  $j$  and country  $k$  is given by

$$p_{jk} = \hat{p}_{jk} + u_{jk}. \quad (2.4.7)$$

where  $\hat{p}_{jk}$  is the log of the predictor and  $u_{jk}$  is the approximation error.

Using the above the process for whether a point ( $q_{ijk}$ ) is residential we find that it can be modelled as

$$\begin{aligned} q_{ijk} &= \phi_{jk} s_{ijk} + (1 - \phi_{jk}) p_{jk} + \epsilon_{ijk} \\ &= \phi_{jk} s_{ijk} + (1 - \phi_{jk}) \hat{p}_{jk} + (1 - \phi_{jk}) u_{jk} + \epsilon_{ijk} \\ &= \hat{\phi}_{jk} s_{ijk} + (1 - \phi_{jk}) \exp(\beta + \lambda_c c + (1 - \lambda_d) d) + e_{ijk} \end{aligned} \quad (2.4.8)$$

where  $p_{jk}$  is the true probability and  $\hat{p}_{jk}$  is the log of the predictor, and  $s_{ijk}$  is the spatial correlation term. Further  $e_{ijk} = \epsilon_{ijk} + (1 - \phi_{jk}) u_{jk}$ . All variable on the right-

hand-side of the last line apart from the error term are measurable, thus, we can again estimate an equation of the same form as 2.4.3, but where, instead of using national and regional dummy variables, we now model the unconditional residential probability by

$$\hat{p}_{jk} = \exp \left( \beta + \lambda_c \ln \left( \frac{C_{jk}}{H_{jk}} \right) + (1 - \lambda_d) \ln \left( \frac{H_{jk}}{L_{jk}} \right) \right) \quad (2.4.9)$$

which gives an implementable regression with regional/country regressors, that we can estimate by non-linear least squares. Note that this specification automatically imposes a lower bound of zero for the implied probability (since both terms are bounded below by zero).<sup>25</sup> There is no necessary upper bound of unity, but in practice the implied estimate of the residential share is always so far below unity that this issue is immaterial. Thus while in principle we might need to follow Angrist and Pischke (2009) and estimate by restricted least squares, in practice there is no need to impose any restrictions in estimation.<sup>26</sup>

The specification in equation 2.4.9 in principle allows for all regressors to be measured at a regional level. In practice at present we measure land,  $L_{jk}$  and population  $H_{jk}$  at a regional level and aggregate consumption  $C_{jk} = C_k$  at a national level.

The error for the  $\{ijk\}^{th}$  point is given by:

$$e_{ijk} = q_{ijk} - \mathbb{P}(q_{ijk} = 1 \mid s_{ijk})$$

We need to allow for heteroscedasticity in the errors, since the estimated conditional probability, and hence the variance of the point-wise errors, varies considerably across the sample.

## Estimation Results

The results are summarised in table 2.3. We estimate a sequence of models. In the first specification we ignore spatial correlation; in the remaining models we augment the model with spatial regressors analogous to those shown in Table 2, with spatial autoregressive coefficients that are, respectively, homogeneous and then heterogeneous at a national and regional level.

As can be seen from the above table, all the regressors are highly significant. Just as in Table 2, the  $R^2$ s are quite low, but, strikingly, our simple economic model with a very small number of parameters comes quite close to replicating the results

---

<sup>25</sup>Recall that our measure of spatial correlation is bounded between zero and one ( $s_{ijk} \in [0, 1]$ )

<sup>26</sup>We also deliberately eschew alternative estimation procedures for binary dependent variables such as probit and logit, since we have a clear rationale for our particular specification in terms of the underlying model.

given by the equivalent models in Table 2 which all have a much larger number of parameters.

	<b>Broad Residential</b>			
	B1	B2	B3	B4
<i>Group Spatial Correlation Parameters</i>				
Homogenous: $\hat{\phi}$		0.380 (0.007)		
National: $\hat{\phi}_k$			0.331 (0.0)	
Regional: $\hat{\phi}_{jk}$				0.256 (0.0)
<i>Reduced Form Parameters: Unconditional Probabilities</i>				
Intercept: $\hat{\beta}$	-9.33 (0.238)	-9.33 (0.378)	-9.33 (0.395)	-9.34 (0.378)
Consumption: $\hat{\lambda}_c$	0.3600 (0.0255)	0.3600 (0.0404)	0.3599 (0.0423)	0.3599 (0.0412)
Pop. Density: $\hat{\lambda}_d$	0.4404 (0.0073)	0.4399 (0.0115)	0.4398 (0.0116)	0.4400 (0.0168)
$R^2$	0.017	0.028	0.029	0.032
$SSR$	8,817	8,723	8,713	8,685
$AIC$	-3.42	-3.43	-3.43	-3.43
$BIC$	-3.42	-3.43	-3.43	-3.41
<i>Parameters</i>	3	4	30	264

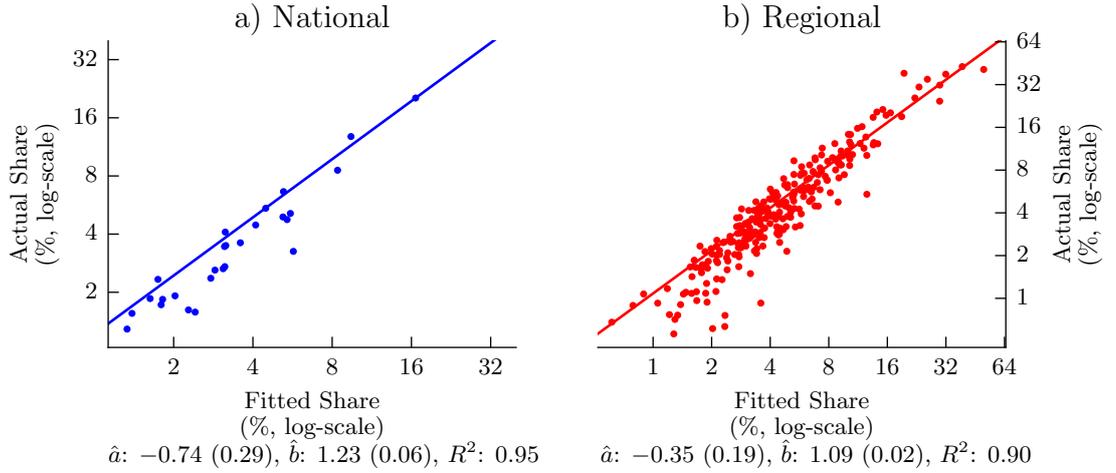
Observations: 269,238. Standard errors in brackets. Estimation of Equation 2.4.3:  $q_{ijk} = \hat{\phi}_{jk}s_{ijk} + (1 - \hat{\phi}_{jk})\hat{p}_{jk} + \hat{u}_{ijk}$ , with the unconditional probabilities ( $\hat{p}_{jk}$ ) given by:  $\exp\left(\hat{\beta} + \hat{\lambda}_c c_k + (1 - \hat{\lambda}_d)d_{jk}\right)$ , as in Equation 2.4.9.  $q_{ijk}$  is equal to 1 if a given point is classified as residential, and 0 otherwise.  $s_{ijk}$  is the spatial correlation term,  $c_k$  is the log national consumption per capita, and  $d_{jk}$  is the log regional population density. Models B1 only provides an estimate of the unconditional probabilities, ignoring spatial correlation. Models B2 imposes a homogenous spatial correlation terms, whereas models B3 and B4 allow spatial correlation terms to vary at a national and regional level. The table shows mean groups estimates (Pesaran et al., 1996) of heterogenous coefficient estimates.

Table 2.3: The probability that any LUCAS survey point is classified as residential: Spatial Correlation and implicit model of regional land supply

## Predictive Power for National and Regional Data

As noted above, neither of the specifications in Tables 2 and 3 has much explanatory power at a level of individual points, since what we are predicting is a sequence of 1s that typically occur with very low frequency, even once we allow for spatial correlation. However, as a cross-check we can aggregate the implied predicted values from Table 3 at national and regional levels and compare with observed residential shares.<sup>27</sup> Figure 2.15 shows the results.

Figure 2.15: Share Aggregation of homogeneous equations



While clearly there remains a nontrivial component of variation at both national and regional levels, it is evident at the same time a large proportion of the variation in residential shares is captured by our equation. Since the impact of spatial correlation essentially averages away, it is the reduced form parameters,  $\hat{\lambda}_c$  and  $\hat{\lambda}_d$  that are doing most of the work.

## Estimation of regional and national aggregates

A cross-check we can also estimate the following equation at a national or regional level:

$$\ln \left( \frac{R_{jk}}{H_{jk}} \right) = \beta + \lambda_c \ln \left( \frac{C_{jk}}{H_{jk}} \right) - \lambda_d \ln \left( \frac{H_{jk}}{L_{jk}} \right) + u_{jk} \quad (2.4.10)$$

Where  $R_{jk}$  are estimates of residential land constructed as outlined in Section 2.1. In the absence of spatial correlation (or if the aggregated spatial correlation term was orthogonal to regional regressors) this would be equivalent to a regional aggregation of the point-wise equation, thus reducing the dataset to 261 observations. We can also aggregate further to a national level, setting  $R_{jk} = R_k$ , which reduces to just 27 observations.

<sup>27</sup>Here, for consistency with estimation, we simply use the unweighted share of residential points in accessible land covered by LUCAS.

Table 2.4: Direct estimation of reduced form model of land supply on regional and national data

	<b>Broad Residential</b>				
	National data		Regional data		
	BN1	BN2	BR1	BR2	BR3
Intercept ( $\beta$ )	1.053 (1.464)	2.433 (0.672)	-0.043 (0.627)	0.411 (0.594)	2.708 (0.214)
Consumption ( $\lambda_c$ )					
National	0.130 (0.123)		0.270 (0.058)		
Regional				0.199 (0.048)	
Pop. Density ( $\lambda_d$ )					
National	0.356 (0.074)	0.337 (0.073)			
Regional			0.329 (0.024)	0.343 (0.025)	0.308 (0.024)
$R^2$	0.463	0.441	0.431	0.422	0.384
$SSR$	3.12	3.25	48.87	49.67	52.91
$AIC$	-1.94	-1.97	-1.65	-1.64	-1.58
$BIC$	-1.43	-1.63	-1.55	-1.53	-1.51
<i>Parameters</i>	3	2	3	3	2
<i>Observations</i>	27	27	261	261	261

Standard errors in brackets. Estimate of Equation 2.4.10:  $r_{jk} = \hat{\beta} + \hat{\lambda}_c c_{jk} - \hat{\lambda}_d d_{jk} + \hat{u}_{jk}$ . Where  $r_{jk}$  is our LUCAS estimate of the log regional (or national) residential land per capita,  $c_{jk}$  is the log consumption per capita (regional or national), and  $d_{jk}$  is the log population density. BN1 shows the estimation using national aggregates where BN2 is at the national level but excludes consumption.

Table 4 shows that, using only national data, the lack of a cross-sectional bivariate correlation between consumption per capita and residential land per capita noted in Section 2.2.4 is also evident in a multivariate framework: the implied estimate of  $\lambda_c$  is positive but insignificant. However, on regional data, once we condition on land per capita (the reciprocal of population density) at a regional level, the estimate of  $\lambda_c$  is larger, and significantly different from zero, whether we use national data on consumption or a regional proxy, regional GDP per capita.

## 2.5 Pigovian Land Supply? Attempting to make sense of the econometric reduced form.

In setting out our theoretical model in Section 3 we derived a log-linear reduced form from a model of Pigovian land supply, which we have shown is at least qualitatively

consistent with our econometric estimates of reduced form coefficients. However, on closer inspection it proves distinctly harder to get even an approximate *quantitative* match that allows us to rationalise what we observe with a truly Pigovian equilibrium.

We have seen that the link between the theoretical model and the observable reduced form can be reduced to the impact of four key magnitudes:  $\alpha$ , the weight of housing in total consumption;  $\kappa_\alpha$ , the consumption equivalent value of the externality (evaluated in the absence of any attempt to mitigate it by Pigovian policies);  $D$ , population density; and  $s$ , a parameter determining spatial correlation.

Clearly a truly Pigovian policy would be required to trade off the cost of the externality,  $\kappa_\alpha$  against the utility gain of higher housing, captured by  $\alpha$ . We can get at least a ballpark value for  $\alpha$  by looking at shares of housing expenditure in total consumption, as given in Table 1. On national data (which is all that we have) these have a cross-sectional average around 16%. Under the maintained assumption of log utility (hence unit price elasticity) this magnitude will be invariant to the price of land; but must clearly be a significant over-estimate of  $\alpha$ , since only a fraction of housing expenditure is on land *per se*. We start by setting  $\alpha = 0.05$ . Since  $\kappa_\alpha$  is inherently un-knowable, we allow it to vary over its full range of  $[0, 1]$ .

Spatial correlation, for which we have shown there is nontrivial econometric evidence, matters in our model because, for any given value of population density,  $D$ , it increases  $\Phi$ , the probability that the neighbouring point will be residential. In the absence of spatial correlation, this would reduce to the unconditional probability of a given point being residential: but we have seen that observed residential shares are so low that this would make it very unlikely that the externality will occur, thus reducing its impact on the social planner, who maximises the expected utility of a randomly chosen household. Thus higher spatial correlation accentuates the impact of the externality. Since, *ceteris paribus*, residential consumption would rise with aggregate consumption, this means that it will dampen the impact of higher consumption on land supply (ie,  $\partial\lambda_c/\partial s < 0$ ).

However at the same time, higher spatial correlation will, *ceteris paribus*, make population density *less* important, simply by inspection of equation 3.2, which determines  $\Phi$ , the Markov probability that point  $i(h) + 1$  will be residential, which can be re-written as  $\Phi = s + D(1 - s)$ .

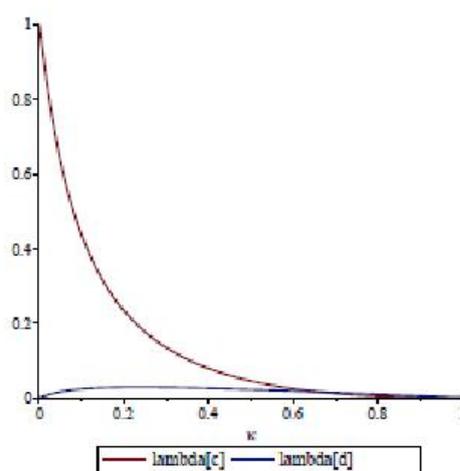
Figures 2.16 and 2.17 illustrate the difficulties of reconciling the reduced form with plausible structural parameters.

In figure 2.16 we pick what appear, a priori, relatively plausible values of  $\alpha = D = 0.05$ , and  $s = 0.5$ , and then plot both reduced form parameters as a function of  $\kappa_\alpha$ , the consumption equivalent value of the externality.<sup>28</sup> We work in deliberately

---

<sup>28</sup>We assume that  $D$  is best captured by the residential share since this captures the probability

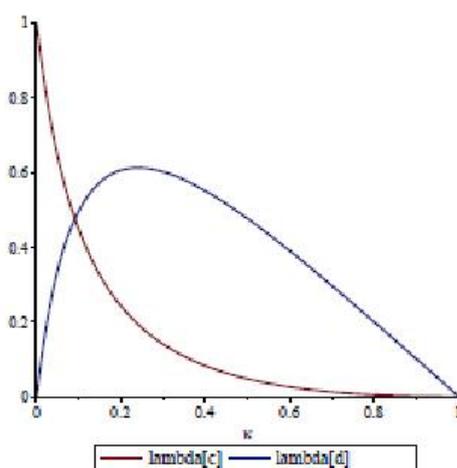
Figure 2.16: Calibrations: Realistic Calibration



round values since our purpose is purely to illustrate the puzzle, rather than seek a precise match.

The two key features illustrated in this first calibration are, first, that the externality needs to be large (with an impact of the order of 10% to 20% of the consumption of affected households) to bring down the reduced form consumption elasticity to anything close to the observed value of around 0.35; but, second, more crucially, with this calibration, population density has only a very modest impact on land supply, whatever the consumption equivalent cost of the externality.

Figure 2.17: Calibrations: Unrealistic Calibration



In Figure 2.17 we can get at least an approximate match for the two reduced form coefficients, at a relatively modest (but still high) value of  $\kappa_\alpha$  but only by making two very significant changes.

---

of a given point having an address on it.

The first is to pick an arbitrarily low value of  $\alpha$ , which we set to one tenth of its value in Figure 2.16. We can crudely characterise this as a “Nimbyist” outcome, in which the social planner sets a very low weight on the utility gains from new housing. But this alone will not provide a match: we also need to assume (against the strong evidence in the data) that there is no spatial correlation (and thus set  $s = 0$ ), which means that the conditional and unconditional probability of the externality are equalised. Without *both* of these features, we cannot even get close to matching both the reduced form coefficients.

We thus conclude that, at least on the basis of the simple model that we devise to analyse the spatial distribution of land, it is very hard to characterise land supply policies as truly Pigovian in nature. The very weak observed impact of higher consumption on land supply requires either that the externality be very costly, or that its costs are given excessive weight in the social planner’s problem (i.e., Nimbyism). But the strength of the observed negative impact of population density is also a puzzle - despite the apparent intuition that less populous regions and countries will have “more space” for residential land. It is actually very difficult to rationalise the strength of this relationship, given that, as shown in Section 2, so little land is actually used for residential purposes.

## 2.6 Conclusions

In this paper we have analysed a new dataset of around  $1/4$  million survey points, taken from the European Land Use and Cover Area-Frame Statistical Survey (LUCAS), covering 27 EU countries. This allows us both to derive national and regional estimates of residential land on a per capita basis, and model its spatial distribution and economic determinants, in light of a theoretical model in which restrictions on land supply attempt to mimic a Pigovian optimum.

Our econometric results show that supply of residential land per capita is affected rather weakly by higher consumption per capita, but somewhat more strongly (and negatively) by population density. While this is qualitatively in line with what would be predicted by a truly Pigovian land supply, we show that it is very hard to rationalise the magnitude of these effects with plausible structural parameters.

## 2.7 Appendices

### 2.7.1 A motivating model of spatial correlation

Assume that the distribution of addresses is generated by an initial historic phase of building. The cost function for building a house at address  $i + 1$  takes the form

$$c(i + 1) = c_0 \exp(-c_1 \times \mathbf{1}_{(i \in \mathcal{I})} + q) \quad (2.7.1)$$

where  $q \sim \mathcal{N}(0, \sigma_q^2)$  is a random normal (with variance  $\sigma_q^2$ ) component in building costs due to differences in terrain, etc. Costs of building a house at address  $i + 1$  are reduced if there is already a house at address  $i$  (due to positive externalities in building costs because of shared overheads such as roads, sewers, electricity etc.). Assume for simplicity that in this initial phase all houses are identical except to the extent that positive consumption externalities lead to a premium and thus sell at price given by

$$p(i + 1) = p_0 \exp(p_1 \times \mathbf{1}_{(i \in \mathcal{I})}) \quad (2.7.2)$$

where  $p_1$  implies a positive consumption externality or agglomeration effect, that gives benefits from living in close proximity to other houses. Assume that supply of houses in the building phase is competitive. This implies the decision rule that  $i + 1$  is an address ( $i + 1 \in \mathcal{I}$ ) if and only if

$$\begin{aligned} c(i + 1) &< p(i + 1) \\ \Leftrightarrow \exp(-c_1 \times \mathbf{1}_{(i \in \mathcal{I})} + q) &< \frac{p_0}{c_0} \exp(p_1 \times \mathbf{1}_{(i \in \mathcal{I})}) \\ \Leftrightarrow \frac{q}{\sigma_q} &< \frac{\ln\left(\frac{p_0}{c_0}\right) + (p_1 + c_1) \times \mathbf{1}_{(i \in \mathcal{I})}}{\sigma_q} \end{aligned} \quad (2.7.3)$$

which in turn implies

$$\begin{aligned} \Phi &= \mathbb{P}(c(i + 1) < p(i + 1) \mid i \in \mathcal{I}) = F\left(\frac{\ln\left(\frac{p_0}{c_0}\right) + c_1 + p_1}{\sigma_q}\right) \\ 1 - \gamma &= \mathbb{P}(c(i + 1) < p(i + 1) \mid i \notin \mathcal{I}) = F\left(\frac{\ln\left(\frac{p_0}{c_0}\right)}{\sigma_q}\right) \end{aligned} \quad (2.7.4)$$

where in the above  $F(\cdot)$  is the standard normal cumulative density function. By inspection, spatial correlation may arise from positive externalities in terms of either cost savings ( $c_1 > 0$ ) or positive consumption externalities ( $p_1 > 0$ ), or both. Notice

also that  $\gamma \leq \Phi$ , with the equality is only binding if  $c_1 = p_1 = 0$ .<sup>29</sup>

Combining the last property that  $1 - \gamma \leq \Phi$  with the definitions of the definition of the two Markov probability terms as given in equations 2.3.4 and 2.3.5. Recall that  $\Phi = D + (1 - D)s$  and  $1 - \gamma = D(1 - s)$ , where by construction the spatial correlation term is restricted to be between zero and one ( $s \in [0, 1]$ ) and the population density given the unit of lands in square meters is also at least empirically between zero and ones ( $D \in [0, 1]$ ). Combining this we find that

$$\begin{aligned} 1 - \gamma &\leq \Phi \\ \Leftrightarrow D(1 - s) &\leq D + (1 - D)s && (2.7.5) \\ \Leftrightarrow 0 &\leq s. \end{aligned}$$

The implication of our historical building phase is that we require a non-negative spatial correlation term as expected from the positive externalities and which is in line with our empirical observations.

## 2.7.2 Datasources

Table 2.5: Datasources

Data	Eurostat code
LUCAS	
Consumption	<i>nama_10_co3_p3_1_data</i>
Housing Expenditures	<i>nama_10_co3_p3_1_data</i>
Agricultural Output	<i>nama_10_a64</i>
Forestry Output	<i>nama_10_a64</i>
Regional Gross Value Added	
Population	<i>demo_r_d2jan</i>
Regional Area	<i>demo_r_d3area</i>
National Area	<i>demo_r_d3area</i>

## 2.7.3 The LUCAS Methodology

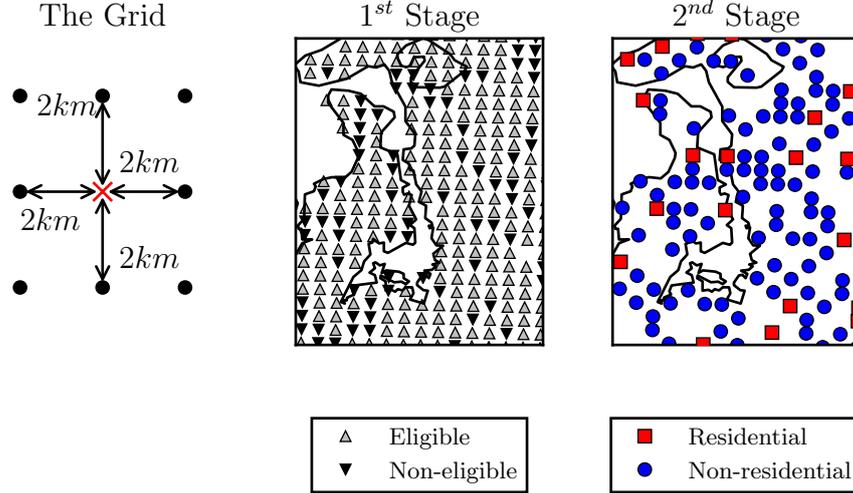
Eurostat’s “*Land Use and Cover Area frame Statistical Survey*” (LUCAS) is a two phase sample survey. The first phase is an equally spaced systematic grid of 1,078,764 observations (in the 2012 sample) in 27 EU countries, separated by 2 km in the four cardinal directions. Each of the points in the first-stage sample are photo-interpreted and classified in terms of land cover,<sup>30</sup> as well as eligibility (based on

<sup>29</sup>By assumption all parameters are positive e.g.  $c_0 > 0$ ,  $c_1 \geq 0$ ,  $p_0 > 0$ , and  $p_1 \geq 0$ .

<sup>30</sup>Using the “CORINE” classification. 1: Arable, 2: Permanent Crop, 3: Grassland, 4: Woodland and shrubland, 5: Bareland, 6: Artificial, 7: Water and Wetland

accessibility) for the second stage of the survey.<sup>31</sup> Together these two classifications give the stratifications of the first stage sample. For the second a subset of 270,277 eligible points from the first-stage sample were visited in person by a surveyor. It is the dataset derived from this physical survey that we use in this paper.

Figure 2.18: The LUCAS dataset



The individual points are then visited as in the figure below and interpreted by the photo taken:

The area estimates of land use and cover classification for the individual NUTS2 regions are then estimated following the methodology of Eurostat (following the methodology of Cochran (1977, chpt. 12)) as:

$$L_{jk,c} = A_{jk} \sum_{h \in \psi} \frac{T_{jk,h}}{T_{jk}} \frac{t_{jk,hc}}{t_{jk}}, \quad (2.7.6)$$

where  $A_{jk}$  is the total area of region  $j$  in country  $k$ ,  $T_{ijk,h}$  is the number of points with stratum classification  $h$  and  $\psi$  is the set of stratum classifications (1 to 7) which is part of the second phase survey and  $T_{jk}$  is the total number of points in the first phase for region  $j$  and country  $k$ .  $t_{jk}$  is the total number of points for the region in the second phase survey and  $t_{jk,hc}$  is the number of points with land use/cover classification  $c$  within stratum  $h$ .

For robustness we proceed with two different definitions of residential land, using land use and cover definitions from the LUCAS survey (Eurostat, 2013). The first, “*broad*” residential land, uses all survey points classified as residential by land use (LU “*U370*” in the dataset). The alternative “*narrow*” measure uses only the subset of residential points that are also classified by land cover as artificial structures (land

<sup>31</sup>Only points that were both below 1,500m in altitude and accessible by road were included in the second stage.

Figure 2.19: Observation 44503638



55°49'46.5" N  
12°02'27.0" E

cover A11 “Buildings with one to three floors” and A12 “Buildings with more than three floors”).

We augment the LUCAS dataset by using Eurostat data on population and gross value added at a regional level (from NACE) as well as Consumption, Actual Rent, Implied Rent and Maintenance at a national level from national accounts (NAMA).

#### 2.7.4 Log–Linearisation of the Model

In log terms, our equilibrium condition of the model in Equation 2.3.16 can be written as

$$\ln R - \ln \hat{C} = \ln(\alpha) - \underbrace{\ln \left( 1 + e^{\ln \mathcal{F}(R,D|s,E) + \ln \hat{C}} \right)}_{\Gamma(R,D,\hat{C})}.$$

Recall that the probabilities is a function of the population density, e.g.:

$$\Phi(D, s) = D + (1 - D)s$$

$$\gamma(D, s) = 1 - D(1 - s)$$

and that the marginal disutility of the externality is given by:

$$\begin{aligned}\mathcal{F}(R; \Phi, \gamma, E) &= \Phi \left( \frac{1}{E-R} \right) + (R-2)(1-\Phi)\gamma^{R-3}(1-\gamma) \ln \left( \frac{E}{E-1} \right) \\ &= (D + (1-D)s) \left( \frac{1}{E-R} \right) \\ &\quad + (R-2)(1-D - (1-D)s)(1-D(1-s))^{R-3}(D(1-s)) \ln \left( \frac{E}{E-1} \right)\end{aligned}$$

Which means that if we log-linearise the above equation, with respect to residential consumption ( $R$ ), total consumption ( $C$ ) and the population density ( $D$ ) and defining  $\tilde{r} = \frac{R-\bar{R}}{\bar{R}}$  for all the variables, we get the log-linearised model as:

$$\tilde{r} - \left( \frac{\bar{C}}{\bar{C} - G^*} \right) \tilde{c} = - \left( \frac{\Gamma_{\hat{C}} \times \bar{\hat{C}}}{\Gamma} \frac{\bar{C}}{\bar{\hat{C}}} \right) \tilde{c} - \left( \frac{\Gamma_R \times \bar{R}}{\Gamma} \right) \tilde{r} - \left( \frac{\Gamma_D \times \bar{D}}{\Gamma} \right) \tilde{d}$$

With the partial derivatives given by:

$$\begin{aligned}\Gamma_{\hat{C}} &= \left( \mathcal{F}(\cdot) \hat{C} \right) \frac{1}{\hat{C}} \\ \Gamma_{\hat{R}} &= \left( \mathcal{F}(\cdot) \hat{C} \right) \mathcal{F}_R \\ \Gamma_{\hat{D}} &= \left( \mathcal{F}(\cdot) \hat{C} \right) \mathcal{F}_D\end{aligned}$$

And the derivatives of the externalities marginal disutility term given by:

$$\begin{aligned}\mathcal{F}_D &= (1-s) \left( \frac{1}{E-R} \right) \\ &\quad + \ln \left( \frac{E}{E-1} \right) (R-2)(1-D(1-s))^{R-2} \\ &\quad \times \left( (1-s)^2(1-2D) - (R-3) \left( \frac{(1-D - (1-D)s)D(1-s)}{1-D(1-s)} \right) \right) \\ \mathcal{F}_R &= -\Phi \left( \frac{1}{E-R} \right)^2 + (1-\Phi)\gamma^{R-3}(1-\gamma) \ln \left( \frac{E}{E-1} \right) (1 + (R-3)\gamma^{-1})\end{aligned}$$

Which we can write as:

$$\tilde{r} = (1-\mu) \eta_c \tilde{c} - \mu \eta_r \tilde{r} - \mu \eta_d \tilde{d}$$

where

$$\begin{aligned}\mu &= \frac{\mathcal{F}(\cdot)\bar{\hat{C}}}{\Gamma(\cdot)} = \frac{\mathcal{F}(\cdot)\hat{C}}{1 + \mathcal{F}(\cdot)\hat{C}} \\ \eta_c &= \frac{\bar{C}}{\bar{C} - G^*} \\ \eta_r &= \frac{\mathcal{F}_R(\cdot)\bar{R}}{\mathcal{F}(\cdot)} \\ \eta_d &= \frac{\mathcal{F}_D(\cdot)\bar{D}}{\mathcal{F}(\cdot)}\end{aligned}$$

for some (reduced form equilibrium outcome)

$$\tilde{r} = \lambda_c \tilde{c} - \lambda_d \tilde{d} \quad (2.7.7)$$

where

$$\lambda_c = \left( \frac{1 - \mu}{1 + \mu\eta_r} \right) \eta_c \quad \text{and} \quad \lambda_d = \left( \frac{\mu}{1 + \mu\eta_r} \right) \eta_d \quad (2.7.8)$$

with  $\lambda_c, \lambda_d \in [0, 1]$ . Thus far, in line with our empirics.

However, while we get a qualitative match, it is by no means so easy to get a quantitative match, particularly for the magnitude of the coefficient on population density.

To explore further, for simplicity linearise around an equilibrium where the key ratio  $\frac{\hat{C}}{E-R}$  (determining  $\mu$ , and hence the  $\lambda_i$ ) is evaluated in an equilibrium where  $\hat{C} = \bar{\hat{C}}$  is some mean value (eg across our cross-section) and land supply ignores the externality, such that  $Q = 1$ , hence  $\bar{R} = \alpha\bar{C}$ . Using (2.3.12) we have  $\frac{R}{E-R} = \frac{\kappa}{1-\kappa}$ , and hence

$$\mu = \frac{\Phi \frac{\bar{\hat{C}}}{R} \frac{R}{E-R}}{1 + \Phi \frac{\bar{\hat{C}}}{R} \frac{R}{E-R}} = \frac{\frac{\Phi}{\alpha} \frac{\kappa_\alpha}{1-\kappa_\alpha}}{1 + \frac{\Phi}{\alpha} \frac{\kappa_\alpha}{1-\kappa_\alpha}} = \mu(\Phi(D, s), \alpha, \kappa_\alpha)$$

where  $\kappa_\alpha = \frac{E-\alpha\bar{C}}{E}$ . Substituting into (2.7.8) we have, using  $\eta_r = \frac{\kappa_\alpha}{1-\kappa_\alpha}$ ,

$$\begin{aligned}\lambda_c(\Phi, \alpha, 0) &= \lambda_c(0, \alpha, \kappa_\alpha) = 1 \\ \lambda_d(\Phi, \alpha, 0) &= \lambda_d(0, \alpha, \kappa_\alpha) = 0\end{aligned}$$

as expected. There is however a problem in finding parameter combinations that map to values similar to what we find in the data, for plausible values of  $\kappa_\alpha$ . A high value of  $s$ , and hence  $\Phi$ , lowers  $\lambda_c$  by enough to match our estimates but implies extremely low values for  $\lambda_d$ .

# Chapter 3

## Inflation Dynamics and Price Flexibility in the UK

### Abstract<sup>1</sup>

Using microdata underlying the UK consumer price index we study how the capacity of nominal demand shocks to stimulate the rate of inflation has evolved over the last two decades. To this end, we estimate a generalized  $(S, s)$  model of lumpy price adjustment, and document sizeable time variation in the behaviour of price flexibility. Most notably, the latter shoots up in the aftermath of the Great Recession and rapidly falls thereafter, with these sharp movements reflecting increased inflation volatility. These features map into a marked non-linearity of inflation dynamics with respect to the degree of price flexibility, with mean reversion being significantly faster when prices are relatively more flexible. State dependence plays a major role for price setting at the microeconomic level, and more so when inflation is particularly high and volatile. Neglecting these facts may severely bias our understanding of inflation dynamics.

### 3.1 Introduction

Over the last decade, the increasing availability of disaggregated data has allowed economists to attain a deeper understanding of consumer micro price behaviour and its implications for price flexibility at the macroeconomic level. The degree of aggregate price flexibility lies at the core of the monetary policy transmission mechanism, ultimately embodying Central Banks' capacity to stimulate output and inflation. As a result, a wide number of empirical contributions have been concerned with measuring the response of prices to nominal demand shocks. However, much less emphasis has been placed on the extent and characteristics of time variation in

---

<sup>1</sup>Joint work with Ivan Petrella and Emiliano Santoro.

aggregate price flexibility,<sup>2</sup> and how this information can be usefully employed to study inflation dynamics.

Using microdata underlying the UK consumer price index (CPI), we document how the distribution of price changes has evolved over the last two decades, and how that reflects into the behaviour of price flexibility. While in the first half of the sample the frequency of adjustment has been roughly stable, during the last decade it has displayed substantial variation, dropping markedly since after the Great Recession. Over the same period, the dispersion of price changes denotes a sustained increase. These facts stand in contrast with the behaviour of US microdata, where the cross-sectional standard deviation of price changes typically comoves positively with the frequency of adjustment (Vavra, 2014 and Berger and Vavra, 2017).

To contextualize these findings, we employ the menu cost model popularized by Barro (1972). Within this setting, diverging trends in the dispersion of price changes and the frequency of adjustment may emerge as the result of a persistent increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price: as long as the resulting expansion in the *inaction region* (i.e., the area where it is not worth adjusting prices) overcomes the effects of low frequency movements in the dispersion of *price gaps* (i.e., the wedge between the actual and the optimal price), the distribution of price changes becomes more dispersed and firms hit the adjustment bands less frequently. To test this prediction, we estimate the generalized *Ss* model developed by Caballero and Engel (2007), fitting the distribution of price gaps and the *hazard function* (i.e., the probability of individual price adjustment) over the price quotes available in each month. By the end of the sample, about five times as many firms appear inactive, as compared with the pre-2010 time window. In line with the framework employed to build our comparative-statics analysis, this implies that the expansion in the inaction region dominates the increase in the dispersion of price changes.

Changes in the distribution of price gaps and the hazard function inevitably reflect in the way shocks are propagated to the economy. To dig deeper into the connection between individual price adjustment and the response of aggregate inflation to nominal stimulus, we compute a measure of aggregate price flexibility, and track its behaviour over the last two decades. The response of aggregate inflation to nominal demand shocks increases substantially during the Great Recession—eventually reaching its (sample) peak in 2011—thus reverting and attaining its minimum in the first quarter of 2017. This implies that, over the last decade, the capacity of nominal stimulus to generate inflation has decreased markedly. More generally, changes in price flexibility tend to occur in correspondence of sizeable departures of CPI

---

<sup>2</sup>In this respect, Caballero and Engel (1993b) and Berger and Vavra (2017) represent some notable exceptions.

inflation from the Bank of England’s institutional target. In this respect, two facts stand out when examining inflation dynamics in the post-Great Recession sample: i) inflation has been outside the 1%-3% interval for a total of 22 out of 40 quarters, while the same has happened only in 11 quarters during the previous decade; ii) over the same period, inflation has shot above and below the target, reaching both its maximum (+4.8%) and minimum value (-0.1%) in the overall sample. In light of this, accounting for time variation in price flexibility may help us understand why hitting the inflation target may have proven to be rather difficult in the last decade.

Changes in price flexibility exert a major impact on the dynamics of aggregate price inflation. The half-life of the inflation response is twice as big in periods of relatively low flexibility, along with appearing remarkably close to the one obtained in the linear setting. In light of this, we posit that neglecting that inflationary shocks are propagated at different speeds depending on the overall degree of price flexibility may lead to overstating inflation persistence. We test this implication, and show that the Bank of England and other market participants do not appear to be taking into account changes in price flexibility when computing their inflation expectations. In fact, price flexibility accounts for roughly 25% of the variability in the absolute forecast error at a four-quarter horizon.

Taking a dynamic perspective is also shown to be important when contrasting the role of time-dependent protocols of price setting, for which the timing of all price changes is predetermined, with that of state-dependent models, for which the timing of price changes can itself respond to shocks. To this end, we decompose the time series of price flexibility into predetermined price adjustments—the so-called *intensive margin*—and adjustments triggered or cancelled by the shock—the *extensive margin*.<sup>3</sup> The latter appear rather relevant, and more so in periods of particularly volatile inflation. In fact, during these episodes the difference between actual inflation and its ‘Calvo counterfactual’—i.e., the inflation rate obtained by setting the period hazard function to a constant equal to the intensive margin—is particularly large. Looking at the behaviour of prices in the correspondence of changes in the value-added tax (VAT) allows us to quantify the importance of adjustments along the extensive margin (see also Gagnon et al., 2013 and Karadi and Reiff, 2014). Massive repricing occurring during these episodes does not emerge as a mere translation of the distribution of price gaps. In fact, many firms seize the opportunity to adjust their prices by more than the VAT change, which implies that inflationary/deflationary pressures from other sources are released in the process. By estimating the generalized *Ss* model in correspondence of a given VAT change, we are then able to devise some alternative counterfactual scenarios that disentangle

---

<sup>3</sup>Adjustments occurring over the intensive margin characterize both time- and state-dependent models. The extensive margin, instead, is a defining feature of state-dependent models.

changes in the hazard function from changes in the distribution of price gaps. All in all, state-dependent pricing plays a major role in amplifying the effects of a VAT change on aggregate inflation.

Our work relates to a number of studies that have examined the connection between microprice changes and aggregate inflation.<sup>4</sup> Among these, Berger and Vavra (2017) represents the contribution that is more in line with the spirit of the present paper. Compared with this study, we highlight the emergence of persistent movements in the distribution of UK price changes and, in this respect, we point to some distinctive patterns in the decade following the Great Recession. Moreover, we elaborate on the role of state dependence in price flexibility and its implications for predicting inflation dynamics. Our work also relates to a number of papers that devise and estimate specific structural models that connect changes in the distribution of price changes to price flexibility (see, e.g., Midrigan, 2011, Alvarez et al., 2016 and Vavra, 2014, among others). As discussed by Berger and Vavra (2017), an empirical limitation of this approach is to rely on specific shocks to the price-setting units, while our approach is more agnostic, in this sense. This represents a strategic advantage in the analysis of UK microdata, where the implied pattern of time variation in the distribution of price changes has been somewhat discontinuous, emerging at different points in time as the result of a different mix of first- and second-moment shocks, as well as persistent changes in the determinants of the inaction region of price setting. Finally, our work relates to Gagnon et al. (2013) in that we focus on the distinction between price adjustments that are determined ahead of shocks, and those that are triggered or cancelled by the shocks. Compared with this study, our empirical model allows us to examine the behaviour of the distribution of price gaps and that of the hazard function in connection with different episodes of VAT changes, thus highlighting important asymmetries over different margins of price setting.

Our paper also features some broad connection with recent empirical contributions employing individual UK consumer prices. In this respect, Bunn and Ellis (2012) have been among the first to appreciate the key characteristics of the frequency of price setting and the hazard functions implied by the microdata from the Office for National Statistics (ONS), while Dixon et al. (2014) have focused on the impact of the Great Recession on price setting. As compared with these papers, we pose particular emphasis on state dependence in price flexibility, as well as on its role for the transmission of nominal demand shocks. Moreover, our application underlines the importance of the selection effect for aggregate inflation (see, on this,

---

<sup>4</sup>See, among others, Bils and Klenow (2004), Dotsey and King (2005), Alvarez et al. (2006), Gertler and Leahy (2008), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2008), Gagnon (2009), Costain and Nakov (2011), Midrigan (2011), Nakamura et al. (2011), Alvarez and Lippi (2014), Karadi and Reiff (2014), Berardi et al. (2015), Alvarez et al. (2016), Nakamura et al. (2018).

Table 3.1: Summary Statistics

	Categories			
	COICOP	Unique	History	Regular
Price Quotes				
Total	27,479,532	27,314,761	23,258,171	19,954,005
Avg. per Month	106,099	105,462	89,800	77,042
Price Trajectories	4,333,302	4,314,903	3,196,697	2,880,332
Avg. CPI Weight	60.73%	60.37%	52.22%	46.48%
Sales and Recoveries				
Avg. per Month (Unweighted)	9.07%	9.10%	8.84%	
Avg. per Month (Weighted)	7.46%	7.49%	7.15%	
Product Substitutions				
Avg. per Month (Unweighted)	6.67%	6.67%	5.30%	
Avg. per Month (Weighted)	5.04%	5.05%	3.91%	

Notes: *COICOP* stands for the *Classification Of Individual CO*nsumption by Purpose price quotes used to calculate the CPI index; *Unique* indicates the COICOP price quotes for which we can uniquely identify a price trajectory; *History* refers to the subset of price quotes in the Unique category for which we can identify at least two consecutive price quotes; *Regular* refers to the price quotes in the History category that do not correspond to sales, product substitutions, or recovery prices. For each of these, we compute the total number of price trajectories, the weighted contribution of each category’s price quotes to the CPI index, as well as the relative number of price quotes corresponding to sales, recovery prices, and product substitutions. Whenever weighted, these statistics have been obtained by accounting for CPI, item-specific, stratum and shop (i.e., elementary aggregate) weights. Sample period: 1996:M2-2017:M8.

Carvalho and Kryvtsov, 2017 and references therein). Specifically, we highlight the versatility of the empirical approach proposed by Caballero and Engel (2007), and show how this can be followed to appreciate the importance of the extensive margin of price adjustment for inflation dynamics in the UK.

The rest of the paper is organized as follows. In Section 3.2 we discuss the key characteristics of the ONS microdata on consumer prices. Section 3.3 discusses the menu cost model that frames our empirical analysis. Section 3.4 reviews the generalized *Ss* model developed by Caballero and Engel (2007), and takes it to the data. Section 3.5 assesses time variation in price flexibility and identifies the relative contribution of adjustments along the intensive and the extensive margin. Section 3.6 discusses the implications of state dependence in price flexibility for inflation dynamics. Section 3.7 concludes.

## 3.2 Microdata on consumer prices

We use ONS microdata that underpin the UK CPI. Prices are collected on a monthly basis, for more than 1,100 categories of goods and services, and published with a

month-lag. Our sample covers the 1996:M2-2017:M8 time window, thus resulting into about 27.5 million observations (see Table 3.1). Each month around 106,000 prices are collected by a market research firm on behalf of the ONS. There are also about 140 items for which the corresponding price quotes are centrally collected. These are excluded from the publicly available dataset, as the structure of their market segment theoretically allows identification of some price setters, or because of the need to frequently adjust for quality changes.<sup>5</sup> The price quotes are recorded on or around the second or third Tuesday of the month, with the exact date being kept secret so as to avoid abnormal prices that, among other things, may be due to the collection of prices during bank-holiday weeks or to price manipulations by service providers and retailers. Furthermore, to make sure the collected price quotes are valid prices, the ONS has set various checks in place, both at the collection point and at later stages in the process. As a preliminary step in handling the dataset, we only employ price quotes that have been marked as being validated by the system or accepted by the ONS. Thus, any price quote that has been marked as missing, non-comparable, or temporarily out of stock is excluded from our sample. We refer to the remaining subset of prices—which make for approximately 60% of those included in the CPI—as *Classification Of Individual CO*nsumption by Purpose (COICOP) approved price quotes.

Each price quote is classified by region, location, outlet and item. The region refers to the geographical entity within the UK from which a given price quote is recorded. The location is intended as a shopping district within a given region: on price-collection days, 146 different locations are visited.<sup>6</sup> For a given location, the shop code is a unique but anonymized *id* associated with the outlet from which the quote is recorded. In turn, each shop is further classified according to whether it is independent (i.e., part of a group comprising less than 10 outlets at the national level) or part of a chain (i.e., more than 10 outlets). Due to a confidentiality agreement between the ONS and the individual shops, for each price quote only the region, outlet and item classifications are published. In light of this, some of the price quotes may not be uniquely identified. This is typically the case when the ONS samples the same item, in the same outlet, but for multiple locations within the same region. As an example, in March 2013 we pick an item with the following characteristics: ‘Women’s Long Sleeves Top’ (*id*: 510223) sold in multiple outlets (*shop type*: 1) within the region of London (*region*: 2). With these coordinates at hand we retrieve two different price quotes: one location sells the item for £22, and

---

<sup>5</sup>This is typically the case for personal computers, whose frequent model upgrades impose the use of hedonic regressions to enhance comparisons across time.

<sup>6</sup>Until August 1996, 180 different locations were being sampled. New locations are chosen every year, with about 20% of them being replaced. As a result, a location is expected to survive an average of about four years in the sample.

one for £26. In February 2013 the price quotes for the same goods were recorded at £25 and £26, respectively. The price quotes are so close that telling the two price trajectories apart may be challenging. To make sure that price trajectories can be univocally identified, we look at ‘base prices’, which are intended as the January’s price for each of the goods under scrutiny.<sup>7</sup> Given this information, we are able to uniquely identify the price trajectories for the two types of good. Even after conditioning on base prices, though, a small portion of price trajectories are still not uniquely identified (about 0.1%, on average): we opt for discarding these. In Table 3.1 the column labelled ‘History’ refers to the price quotes with an identifiable history that spans at least two consecutive periods. Following the criteria outlined above, we drop about 12,000 quotes per month.<sup>8,9</sup>

To aggregate the individual price quotes into a single price we also make use of the following weights produced by the ONS:<sup>10</sup> the *shop* weights, which are employed to account for the fact that a single item’s price is the same in different shops of the same chain (e.g., a pint of milk at a Tesco branch);<sup>11</sup> the *stratification* weights, which reflect the fact that purchasing patterns may differ markedly by region or type of outlet;<sup>12</sup> finally, the *item* and *COICOP* weights are used to reflect consumers’ expenditure shares in the national accounts.

### 3.2.1 Variable definition

After deriving our price quotes in line with the criteria set out above, it is important to make a distinction between regular and temporary price changes. We start by dealing with sales, whose behaviour tends to be significantly different from that of regular prices (see Eichenbaum et al., 2011 and Kehoe and Midrigan, 2015). To this end, we first exclude all the price quotes to which the ONS attaches a sales indicator. For a price to be marked as being associated with a sale, the ONS requires the latter

---

<sup>7</sup>The base price is typically relied upon in order to normalize price quotes and calculate price indices, or to adjust for changes in the quality and/or quantity of a given good.

<sup>8</sup>Due to a particularly low coverage, Housing, Water, Electricity, Gas and Other Fuels (COICOP 4) and Education (COICOP 10) are excluded from the sample. We also exclude price changes larger than 300%, which we deem as being due to measurement errors. These take place rarely (< 0.01%). Appendix 3.8.1 provides additional details on the construction of the dataset.

<sup>9</sup>The total number of available price quotes denotes a weak downward trend. However, it is important to stress that the composition in terms of categories accounted for by Table 3.1 is roughly stable over time. This implies the presence of no particular trends in the behaviour of product substitutions and sales.

<sup>10</sup>See Chapter 7 of the ONS CPI Manual (ONS, 2014).

<sup>11</sup>In this case the ONS enters a single price for a pint of milk, but the weight attached to this is ‘large’, so as to reflect that all Tesco branches within the region have posted the same price.

<sup>12</sup>In this respect, four levels of sampling are considered for local price collection: locations, outlets within location, items within location-outlet section and individual product varieties. For each geographical region, locations and outlets are based on a probability-proportional-to-size systematic sampling, where size accounts for the number of employees in the retail sector (locations) and the net retail floor space (outlets).

to be available to all potential costumers—so as to exclude quantity discounts and membership deals—and that it only entails a temporary or an end-of-season price reduction.<sup>13</sup> As a second step, we apply a symmetric V-shaped filter, as defined by Nakamura and Steinsson (2010b), on the remaining price quotes. According to the filter, the sale price of item  $i$  at time  $t$ ,  $P_{i,t}^s$ , is identified as follows: i) it is lower than last period’s price (i.e.,  $P_{i,t}^s < P_{i,t-1}$ ) and ii) the next period’s price is equal to last period’s price (i.e.,  $P_{i,t+1} = P_{i,t-1}$ ). A recovery price  $P_{i,t}^r$ , instead, meets the following criteria: i) it is greater than last period’s price (i.e.,  $P_{i,t}^r > P_{i,t-1}$ ) and ii) it is such that  $P_{i,t}^r = P_{i,t-2}$ . Once a price quote has been identified as being a sale or a recovery price, we discard it from the sample.<sup>14</sup>

Item substitutions are a further reason of concern when trying to identify price trajectories, as they require a certain judgement to establish what portion of a price change is due to quality adjustment and which component reflects a pure price adjustment. Product substitutions occur whenever an item in the sample has been discontinued from its outlet, and the ONS identifies a similar replacement item to the price going forward. Therefore, it is reasonable to expect that product turnovers are followed by price changes that either reflect uncaptured quality changes (Bils, 2009), or simply reflect a low-cost opportunity to reset prices that has nothing to do with the underlying sources of price rigidity, as argued by Nakamura and Steinsson (2008). In line with the literature (see, e.g., Berardi et al., 2015, Berger and Vavra, 2017, and Kryvtsov and Vincent, 2017), we interrupt a trajectory whenever it encounters a substitution flag, as indicated by the ONS.

Table 3.1 shows that, after these preliminary steps, we are down to a monthly average of 79,000 price quotes. Finally, we define the price change of item  $i$  at time  $t$  as  $\Delta p_{i,t} = \log(P_{i,t}/P_{i,t-1})$ .<sup>15</sup>

### 3.2.2 Data facts

This section unveils a number of stylized facts about the behaviour of the ONS microdata.<sup>16</sup> The top panels of Figure 3.1 report the time path of the frequency of adjustment and the average magnitude of price changes: decomposing inflation as the product of these statistics carries important information on the relationship

<sup>13</sup>This definition excludes clearance sales of products that have reached the end of their life cycle.

<sup>14</sup>See also Nakamura and Steinsson (2008) and Vavra (2014). As an alternative approach, in place of the price associated with a sale Klenow and Kryvtsov (2008) report the last regular price, until a new regular price is observed.

<sup>15</sup>We also compute price changes as  $\Delta p_{i,t} = 2 \frac{P_{i,t} - P_{i,t-1}}{P_{i,t} + P_{i,t-1}}$ . This definition has the advantage of being bounded and less sensitive to outliers. The results—virtually unchanged with respect to the ones we report—are available from the authors, upon request.

<sup>16</sup>Throughout the paper all statistics derived from the microdata on prices are reported as a 12-month moving average, so as to get rid of the seasonality in the data.

between the distribution of price changes and inflation itself (see, e.g., Gagnon, 2009). As expected, the average price change tends to display a high degree of positive co-movement with CPI inflation, at least until the end of the Great Recession. Thus, in the last part of 2015 the two series are back moving in tandem. As for the frequency of adjustment, it is interesting to notice how this tracks very closely the contraction in the rate of inflation that starts in 2012—going well below its sample average up to that point—while only displaying a weak reversion towards the end of 2015.<sup>17</sup> In the bottom panels of the figure, both statistics are split between positive and negative price changes. The frequency of positive price changes is greater than that associated with negative adjustments throughout the entire sample, while the opposite broadly holds true when comparing the average price changes in either direction. Focusing on the post-recession sample, we appreciate two key aspects: i) the downward trend in the frequency, as observed in the first panel of the figure, is mostly due to the component associated with positive price changes; ii) notwithstanding that the average of positive price changes displays a weak tendency to increase, the (mirror image of the) average of negative price changes denotes a more robust upward trend.<sup>18</sup> Both facts point to a certain degree of asymmetry in price adjustment.

Figure 3.2 plots higher moments of the distribution of price changes.<sup>19</sup> Notably, the standard deviation displays a very large increase in the aftermath of the Great Recession. In fact, as displayed by the top-right panel of the figure, dispersion increases on either side of the median, though negative price changes denote a stronger acceleration in volatility, as compared with positive price changes. In light of this it should be stressed that the fall in CPI inflation occurring in the post-2010 sample is to a large extent a manifestation of the trend in the dispersion of negative price changes—relative to that of positive ones—rather than reflecting a mere shift in the mode of the density. This fact, coupled with the observation of diverging trends in the relative size of average positive/negative price changes, inevitably reflects into the dynamics of the skewness, which fluctuates around a positive mean in the pre-2010 sample, and becomes persistently negative thereafter.

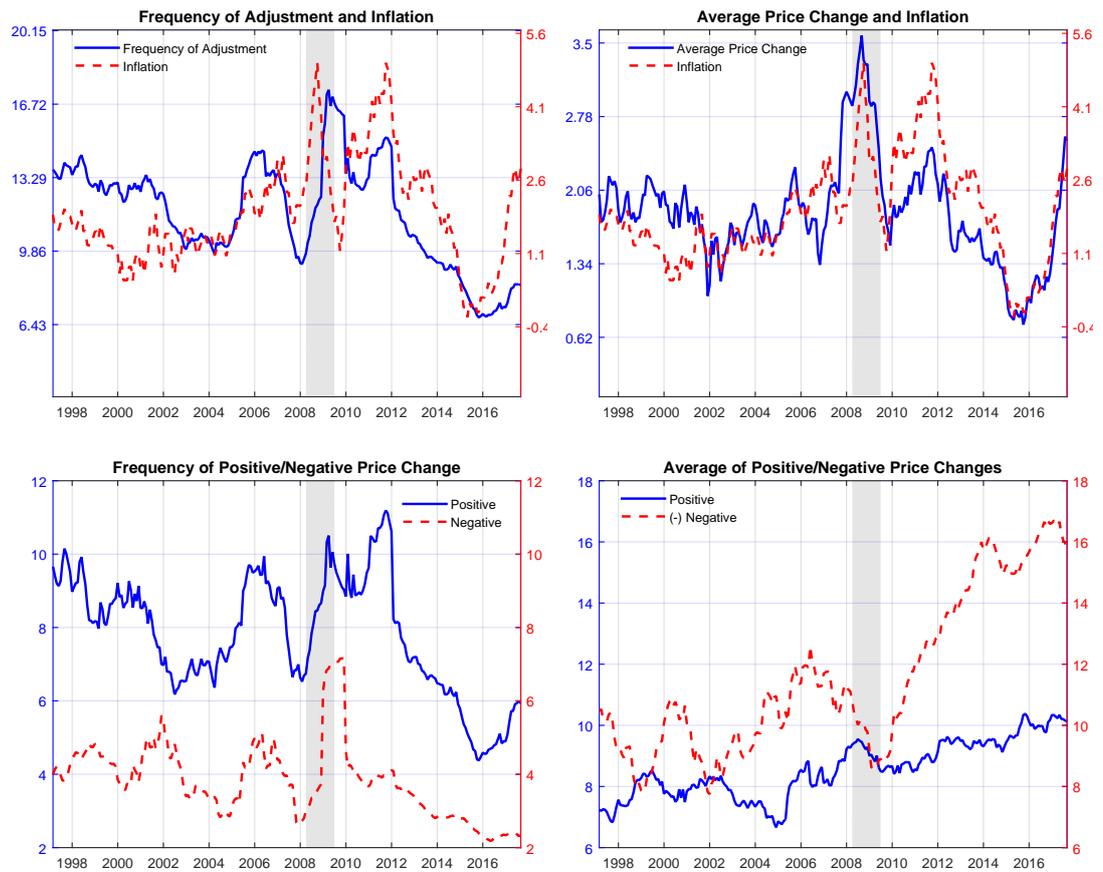
---

<sup>17</sup>The average frequency of price adjustment prior to the fall is broadly in line with the figures reported by previous studies on UK micro price data. To see this, one has to account for the fact that we exclude both utility prices (COICOP 4) and sales. Bunn and Ellis (2012), instead, consider both categories, while Dixon and LeBihan (2012) and Dixon and Tian (2017) include sales, but exclude utility prices.

<sup>18</sup>Figure 3.10 in Appendix 3.8.2 shows that composition effects have no role in generating the facts presented in this subsection. To this end, we compare the moments of the distribution of price changes with their homologues obtained by averaging the corresponding moments of the price quotes for each of the 25 COICOP group categories.

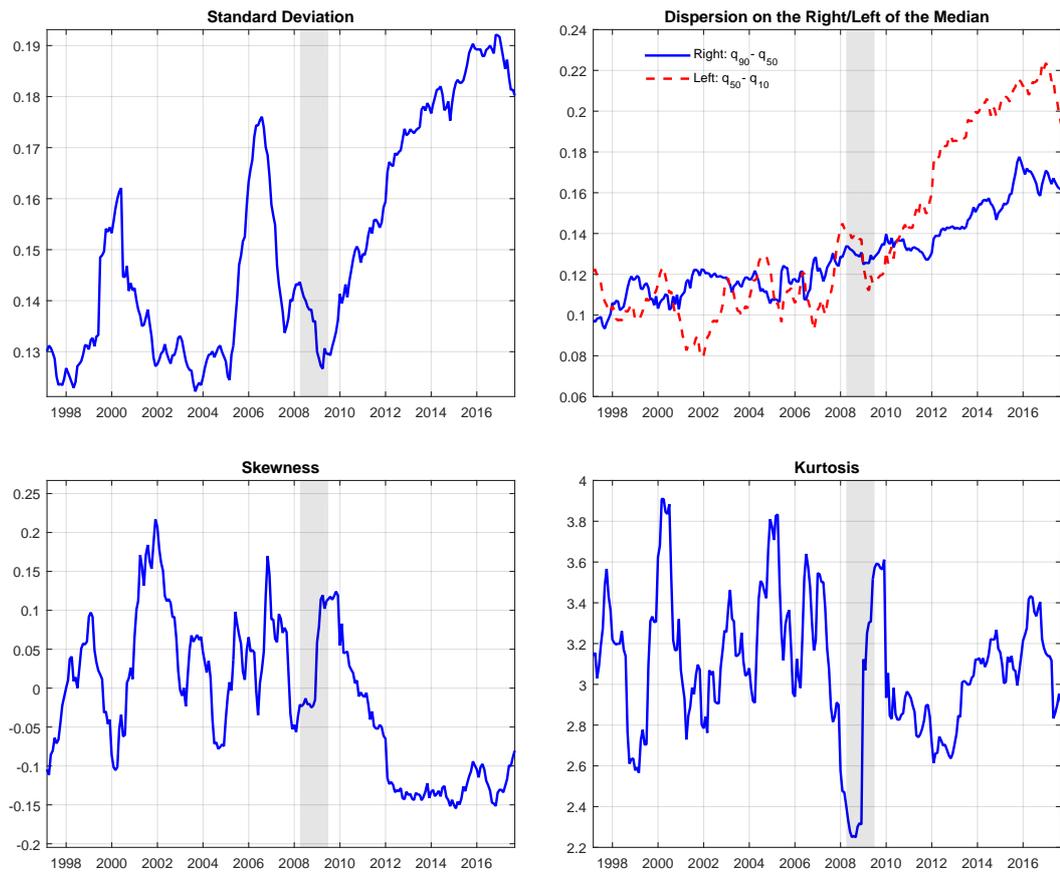
<sup>19</sup>To avoid that zero price changes dominate the distribution, we follow Vavra (2014) and much of the literature in that they consider only non-zero price movements.

Figure 3.1: Frequency of Adjustment and Average Price Changes



Notes: The shaded vertical band indicates the duration of the Great Recession. The inflation rate graphed in the upper panel of the figure is the official CPI inflation rate published by the ONS. In the bottom-right panel we report the the absolute value of average negative price changes.

Figure 3.2: Moments of the Distribution of Price Changes



Notes: Price dispersion on the right (left) side of the median price quote is computed as  $q_{50} - q_{10}$  ( $q_{90} - q_{50}$ ). The skewness and kurtosis of the distribution of price changes are measured as  $\frac{q_{90,t} + q_{10,t} - 2q_{50,t}}{q_{90,t} - q_{10,t}}$  and  $\frac{q_{90,t} - q_{62.5,t} + q_{37.5,t} - q_{10,t}}{q_{75,t} - q_{25,t}}$ , respectively. The shaded vertical band indicates the duration of the Great Recession.

Table 3.2: Correlations of Pricing Moments with Macroeconomic Variables

Full Sample						
	$fr_t$	$\sigma_t^2$	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
$y_t$	-0.569***	0.264***	0.334***	0.422***	-0.363***	-0.322***
$\pi_t$	0.169***	0.000	-0.016	-0.147**	-0.024	-0.281***
$fr_t$	-	0.162**	-0.510***	-0.737***	0.470***	0.286***
Pre-Recession						
	$fr_t$	$\sigma_t^2$	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
$y_t$	0.455***	0.612***	-0.121	-0.092	-0.015	0.171*
$\pi_t$	0.387***	0.213**	-0.416***	-0.410***	0.177*	0.181**
$fr_t$	-	0.569***	-0.120	-0.511***	0.356***	-0.055
Post-Recession						
	$fr_t$	$\sigma_t^2$	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
$y_t$	-0.399***	0.221**	0.137	0.428***	-0.244**	0.291***
$\pi_t$	0.467***	0.077	-0.275***	-0.303***	-0.216**	-0.530***
$fr_t$	-	-0.475***	-0.646***	-0.854***	0.383***	-0.292***

Notes:  $fr_t$  denotes the frequency of adjustment;  $\sigma_t^2$  stands for the volatility of the distribution of price changes;  $q_{n,t}$  measures the  $n$ -th quantile of the distribution of price changes;  $Skew_t$  denotes the skewness of the distribution of price changes and is measured as  $\frac{q_{90,t} + q_{10,t} - 2q_{50,t}}{q_{90,t} - q_{10,t}}$ ;  $Kurt_t$  denotes the kurtosis of the distribution of price changes and is measured as  $\frac{q_{90,t} - q_{62.5,t} + q_{37.5,t} - q_{10,t}}{q_{75,t} - q_{25,t}}$ ;  $y_t$  is a business cycle indicator;  $\pi_t$  indicates aggregate CPI inflation. Aside of the inflation rate, all series are obtained by de-trending their raw counterparts by means of Rotemberg's (1999) version of the HP filter, which sets the smoothing coefficient so as to minimize the correlation between the cycle and the first difference of the trend estimate. \*\*\*/\*\*/\* indicates statistical significance at the 1/5/10% level, respectively.

Table 3.3: Correlations of Pricing Moments with Macroeconomic Variables: the Role of Asymmetry

Full Sample								
	$fr_t^+$	$fr_t^-$	$dp_t^+$	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
$y_t$	-0.330***	-0.636***	0.306***	0.388***	0.188***	0.295***	0.127**	0.417***
$\pi_t$	0.529***	-0.110*	0.031	0.285***	0.253***	-0.370***	-0.366***	-0.203***
Pre-Recession								
	$fr_t^+$	$fr_t^-$	$dp_t^+$	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
$y_t$	0.466***	0.210**	0.213**	0.427***	-0.162*	-0.018	-0.037	-0.059
$\pi_t$	0.154*	0.173*	-0.001	-0.018	0.057	-0.231***	0.045	-0.406***
Post-Recession								
	$fr_t^+$	$fr_t^-$	$dp_t^+$	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
$y_t$	-0.373***	-0.415***	-0.117	-0.489***	-0.696***	0.479***	0.325***	0.415***
$\pi_t$	0.858***	0.556***	-0.171	0.606***	0.535***	-0.760***	-0.702***	-0.619***

Notes:  $fr_t^+/fr_t^-$  stands for the frequency of positive/negative price changes;  $dp_t^+/dp_t^-$  indicates the average size of positive/negative price changes;  $q_{n,t}$  measures the  $n$ -th quantile of the distribution of price changes;  $y_t$  is a (monthly) business cycle indicator;  $\pi_t$  indicates aggregate CPI inflation. Aside of the inflation rate, all series are obtained by de-trending their raw counterparts by means of Rotemberg's (1999) version of the HP filter, which sets the smoothing coefficient so as to minimize the correlation between the cycle and the first difference of the trend estimate. \*\*\*/\*\*/\* indicates statistical significance at the 1/5/10% level, respectively.

Table 3.2 reports the correlation between some of the key moments of the distribution of price changes, CPI inflation and a business cycle indicator.<sup>20</sup> To set aside potential spurious correlation emanating from the low-frequency behaviour of the series under examination, we detrend all of them, aside of the inflation rate.<sup>21</sup> Turning our attention to the frequency of adjustment and the dispersion of price changes, it is important to stress that they also display somewhat different cyclical behaviours. Looking at the entire sample, the frequency moves countercyclically, while dispersion is procyclical. However, the sign of these correlations is only preserved in the post-recession sample, while during the previous decade both statistics have behaved procyclically. Also their pairwise correlation seems to vary substantially across the two subsamples—going from being positive in the first decade to negative thereafter—though measuring dispersion through inter-quantile differences points to a negative correlation.

As for the higher moments of the distribution, the skewness signals a marked countercyclical behaviour. There is an interesting parallel to this in the empirical literature on the dynamics of the cross-sectional distribution of output growth. Using the growth rates of real sales for micro-datasets of individual firms a distinct skewness with a negative correlation to the business cycle has been documented for the US (Higson et al., 2002, Holly et al., 2013), UK (Higson et al., 2004, Holly et al., 2013), and Germany Döpke et al. (2005). A possible extension for future research, would be to investigate if the differences of reaction function to the business cycle for real sales growth given the size of the firm found in the above papers holds for the skewness of price changes as well. Such result would give further evidence for financial frictions such as found in Holly et al. (2013).

The correlation between kurtosis and the cyclical indicator is heavily influenced by the only recession in the time window considered, being negative in the whole sample, while turning positive in the subsamples that exclude the Great Recession.<sup>22</sup> Table 3.3 broadly confirms these tendencies, while showing that the frequency of

---

<sup>20</sup>Appendix 3.8.3 contains more details on the derivation of the monthly coincident indicator of economic activity.

<sup>21</sup>Moreover, when splitting the sample we exclude the period around the Great Recession (2007:M3-2010:M6), so as to avoid that the correlations among the key variables are dominated by the major macroeconomic turmoil in that period. In light of this it is worth stressing that, when interpreting the cyclical properties of the data in the two subsamples, the correlations are likely to be picked up by the behaviour of the series in periods of relatively stronger/weaker expansion, rather than by different cyclical phases.

<sup>22</sup>Villar and Luo (2017) show how different models of price setting may account for different signs of the correlation between inflation and the skewness of price changes. In this respect, menu cost models—which feature the price change distribution becoming less skewed as inflation rises—could well rationalize our data in the second subsample. On the other hand, the Calvo model—which features a positive correlation—could better account for the first subsample. In the remainder of the analysis we will show how such characterization is also supported by the behaviour of the extensive margin of price adjustment—a hallmark of menu cost models—assumes a prominent role in the aftermath of the Great Recession.

negative price changes denotes stronger countercyclicality—as compared with its counterpart computed for positive adjustments—both in the full sample and in the last decade. Concurrently, the procyclicality of the dispersion is a phenomenon that tends to characterize price changes taking place on the left side of the median—mostly in the post-recession sample—while the dispersion of negative price changes varies substantially depending on both the specific subsample and the way dispersion is measured.

To summarize the most consistent patterns of the frequency of adjustment and the dispersion of price changes: after the Great Recession, the former has displayed pronounced countercyclicality, while dispersion has been markedly procyclical throughout the entire sample, with both comovements appearing more marked in the case of negative price changes. Otherwise, the pairwise correlation between these statistics has turned deeply negative after the Great Recession. Notably, this picture stands in contrast with the analysis on US microdata by Vavra (2014), who reports that the cross-sectional standard deviation of price changes is strongly countercyclical and positively comoves with the frequency of adjustment. To rationalize these facts, he employs a stylized menu cost model, showing how shocks to the dispersion of price gaps may play an important role. In the next section we use the same framework to show that changes in the incentives firms face when deciding to change prices can provide us with a rationale for the emergence of negative comovement between the dispersion of price adjustments and their frequency.

### 3.3 Analytical framework

To frame the empirical analysis, we consider the menu cost model popularized by Barro (1972) and Dixit (1991). As illustrated by Vavra (2014), the advantage of this framework is to provide us with a simple analytical setting to keep track of the determinants of the frequency and the dispersion of price changes, as well as the dispersion of price gaps, intended as the difference between the actual price of a given good and its reset price (i.e., the price that would have prevailed in the absence of price-setting frictions).

Firms face a dynamic control problem where  $x$ —the deviation of the current price from the optimal price—is defined as the state variable. A wedge between the state variable and zero entails an out-of-equilibrium cost  $\alpha x^2$ , where  $\alpha$  can be inversely related to market power. When not adjusting,  $x$  follows a Brownian motion  $dx = \phi dW$ , where  $W$  is the increment to the Wiener process. It is possible to change the value of  $x$  by applying an instantly effective control at a lump-sum cost  $\lambda$ . A key identifying assumption for our model is that the cost of adjusting prices is a fixed lump sum cost ( $\lambda$ ), which implies that when adjustment takes place it is optimal for

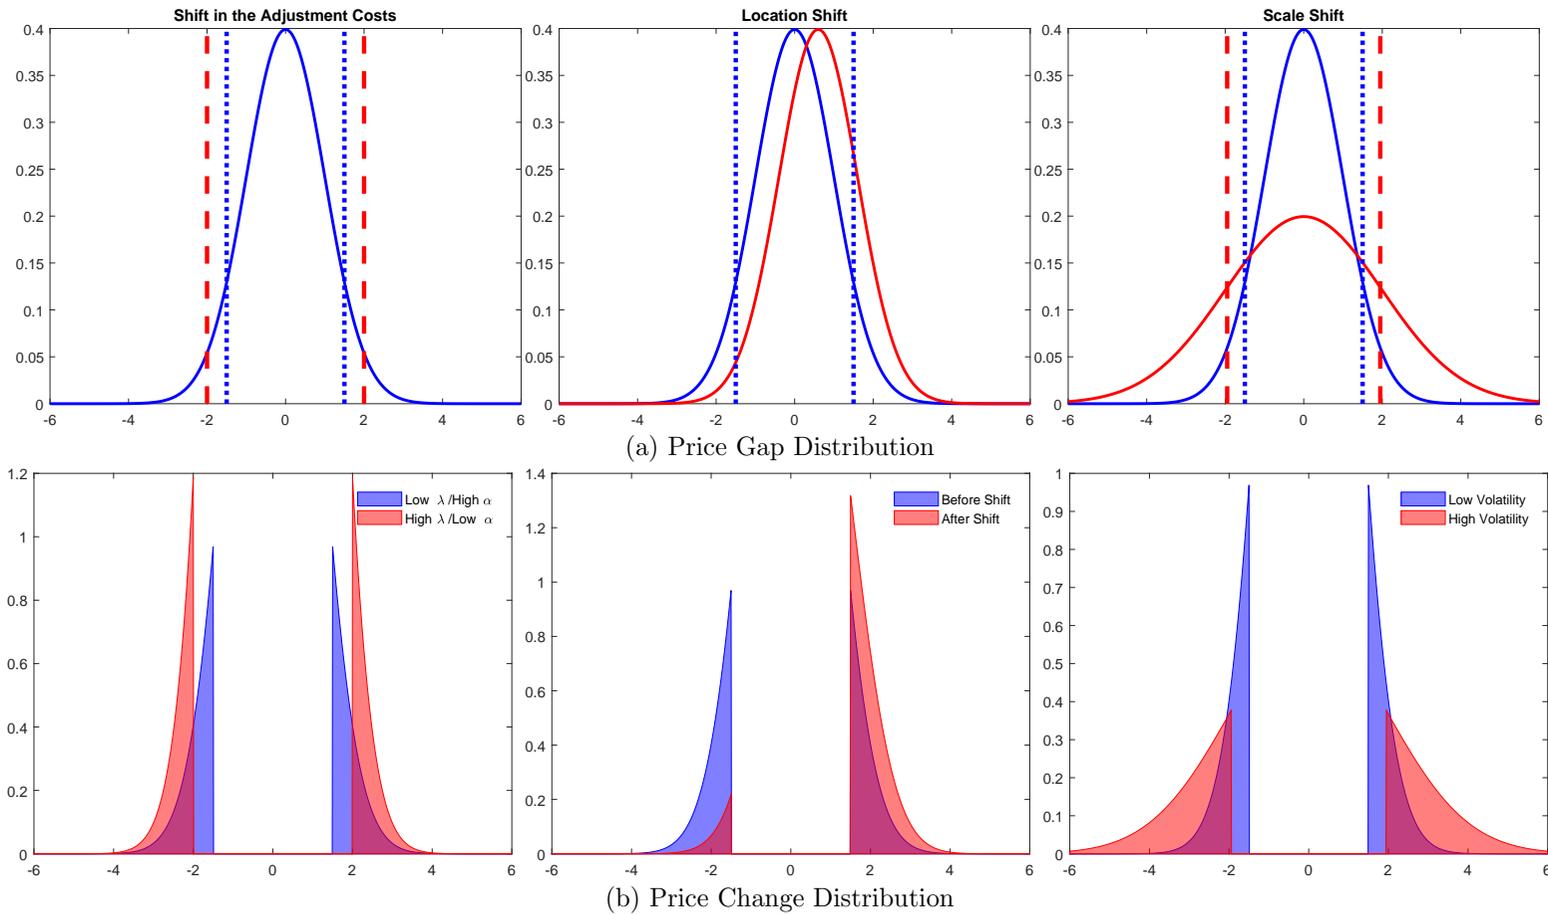
the firms to fully adjust their prices. This is in contrast to the model of Rotemberg (1982) who outlines a quadratic cost function for the adjustment costs (as a function of the price gap), which implies it is optimal for the firms to only partially adjust their prices, when deciding to pay the adjustment cost. From this environment a simple *Ss* rule emerges, according to which the optimal policy is ‘do not adjust’ when  $|x| < \sigma$  and ‘adjust to zero’ when  $|x| \geq \sigma$ , where  $\sigma = (6\lambda\phi^2/\alpha)^{1/4}$  denotes the standard deviation of price changes. Moreover,  $fr = (\alpha/6\lambda)^{1/4} \phi$  is the frequency of adjustment.<sup>23</sup>

To provide an overview of the different determinants of the distribution of price gaps and the associated distribution of price changes, Figure 3.3 considers three possible scenarios: i) a positive shift in the cost of adjustment  $\lambda$  (or, equivalently, a negative shift in  $\alpha$ ) that affects the inaction region, while leaving the distribution of price gaps unaffected; ii) a first-moment shock that causes a shift in the distribution of price gaps, affecting all  $x$ ’s in the same manner; iii) an increase in the dispersion of the distribution of price gaps (i.e., a rise in  $\phi$ ). As for i), a positive change in  $\lambda$  increases the inaction region, translating into a compression in the frequency of adjustment and an increase in the dispersion of price changes. As for ii), the immediate effect of a shift in the distribution of price gaps is to push more firms out of the inaction region, thus inducing an increase in the frequency of adjustment. Importantly, this result does not depend on the specific sign of the shock, as all firms’ desired price changes will be affected in the same way. Thus, all firms pushed out of the inaction region will denote price changes of the same sign, implying a decrease in their dispersion. In fact, Vavra (2014) shows that, while in environments with zero inflation small shocks to  $x$  do not produce any effect on the frequency of adjustment and the dispersion of price changes, in the presence of positive trend inflation the frequency (dispersion) increases (decreases). Finally, a rise in  $\phi$ , as sketched in the last column of the figure, induces both  $fr$  and  $\sigma$  to increase.

---

<sup>23</sup>For analytical details and proofs, see Barro (1972) and Vavra (2014).

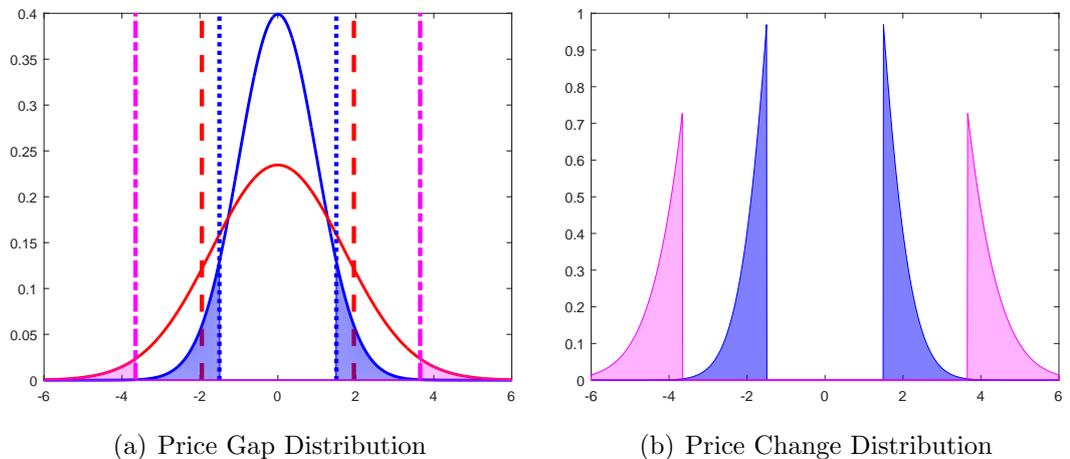
Figure 3.3: Analytical Framework



Note: The first column considers a positive shift in  $\lambda$  (or a negative shift in  $\alpha$ ) that affects the inaction region, while leaving the distribution of price gaps unaffected. The second column considers the effects of a first-moment shock that affects all  $x$ 's in the same direction. The last column depicts the effects of an increase in  $\phi$ . The upper panels report the ex-ante distribution of price gaps and the corresponding bands delimiting the inaction region (dotted-blue lines), together with their ex-post counterparts (dashed-red lines). The bottom panels report the corresponding distributions of price changes.

Vavra (2014) points to second-moment shocks as potential drivers of the positive comovement between the frequency of adjustment and the price-change dispersion in U.S. CPI data. However, in the microdata under examination the comovement between these two statistics is positive only in the first part of the sample, while turning negative in the following decade, when the two series display diverging trending behaviours. In light of this, second-moment shocks might provide a good account of what has happened up to the Great Recession. Moreover, shocks to  $x$  of either sign would determine relative movements in the dispersion of price changes and the frequency of adjustment which do not square with the data, regardless of the time window we consider. In fact, Section 3.4 will show that episodes of major repricing—such as those occurring due to changes in the VAT—do not only reflect into pre-determined price adjustments (i.e., adjustments that are determined ahead of the shock and would materialize into a mere shift of the distribution). As a result, the so-called extensive margin of price flexibility, which accounts for adjustments that are either triggered or cancelled by the VAT change, is shown to play an important role.

Figure 3.4: A combined increase in  $\phi$  and  $\lambda$



Note: We consider a positive shift in  $\lambda$  that affects the inaction region (while leaving the distribution of price gaps unaffected), combined with an increase in  $\phi$ . The left panel reports the transformations occurring to the distribution of price gaps and the corresponding bands delimiting the inaction region: the dotted (blue) line refers to the ex-ante situation, the dashed (red) line denotes the effects of the volatility shift, while the dashed-dotted (magenta) line refers to the effects produced by the increase joint in  $\phi$  and  $\lambda$ . The right panel reports the distribution of price changes, both in the ex-ante situation and in the case of a combined increase in  $\phi$  and  $\lambda$ .

When looking at the post-recession experience, among the free parameters of the model only a persistent increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price may account for the diverging trends

we observe, conditional on the resulting expansion of the inaction region dominating the effects of positive shifts in the dispersion of price gaps. A caveat is in order at this stage: the menu cost model we are employing has been explicitly envisaged to investigate the effects of shocks to  $x$  in the neighbourhood of the steady state. Therefore, within this framework secular movements in the frequency of adjustment and the price-change dispersion—as those observed in the post-recession period—can be thought of as resulting from as a sequence of persistent changes in the volatility and the cost parameters. In this respect, Figure 3.4 considers a situation in which both  $\phi$  and  $\lambda$  increase:<sup>24</sup> the rise in the dispersion of price changes determines an expansion in the inaction region, thus increasing the density outside the adjustment bands and, in turn, the frequency of adjustment. This effect is counteracted by the rise in  $\lambda$ , which widens the inaction region further and restricts the density outside the adjustment bands beyond the initial situation. If the expansion in the inaction region is large enough to overcome the increase in dispersion, we observe negative comovement between the cross-sectional dispersion of prices and the frequency of adjustment, which is consistent with what observed in the post-recession period. To dig deeper into these aspects, the next section introduces an accounting framework that proves to be particularly useful at quantifying the link between changes in the timing of individual price adjustments and macro price flexibility, along with formalizing the distinction between predetermined price adjustments and those which are triggered or cancelled by shocks.

### 3.4 A generalized $Ss$ model

To verify our conjecture, while accounting for the connection between price setting at the micro level and price flexibility at the aggregate level, we use the generalized  $Ss$  model developed by Caballero and Engel (2007). This framework is consistent with lumpy and infrequent price adjustments—which are typically perceived as distinctive traits of price setting—along with encompassing several pricing protocols.<sup>25</sup> Berger and Vavra (2017) also show that such an accounting approach is capable of providing a good fit to the data generated by different structural models (e.g., Golosov and Lucas, 2007 and Nakamura and Steinsson, 2010a). To allow for time variation in different determinants of price adjustment, we estimate the model over each cross section of micro price data, matching different price-setting statistics. More details

---

<sup>24</sup>Once again, a drop in  $\alpha$  would lead to qualitatively similar results.

<sup>25</sup>To mention two extreme examples, the generalized  $Ss$  model can account for both price setting à la Calvo (1983)—where firms are selected to adjust prices at random and price flexibility is fully determined by the frequency of adjustment—as well as for schemes à la Caplin and Spulber (1987) model—where adjusting firms change prices by such large amounts that the aggregate price is fully flexible, regardless of the frequency of adjustment.

on the estimation are reported in Section 3.4.1. In the remainder of this section, instead, we discuss the analytical details of the accounting framework.

Assume that, due to price rigidities, firm  $i$ 's (log of) the actual price may deviate from the (log of) the target or reset price, which is denoted by  $p_{it}^*$ . Thus, we define the price gap as  $x_{it} \equiv p_{it-1} - p_{it}^*$ , implying that a positive (negative) price gap is associated with a falling (increasing) price when the adjustment is actually made. In a simple  $Ss$  model, as the one detailed in the previous section, the price is adjusted when the price gap is large enough, and  $p_{it} = p_{it}^*$  after the adjustment has taken place. Assuming  $l_{it}$  periods since the last price change, the adjustment reflects the cumulated shocks:  $\Delta p_{it} = \sum_{j=0}^{l_{it}} \Delta p_{it-j}^*$ , with  $\Delta p_{it}^* = \mu_t + v_{it}$ , where  $\mu_t$  is a shock to nominal demand and  $v_{it}$  is an idiosyncratic shock.

As discussed by Caballero and Engel (2007), the basic  $Ss$  setting of the previous section can be generalized by assuming *iid* idiosyncratic shocks to the adjustment costs. Thus, by integrating over their possible realizations, we obtain an adjustment hazard  $\Lambda_t(x)$ . This is defined as the (time  $t$ ) probability of adjusting—prior to knowing the current adjustment cost draw—by a firm that would adjust by  $x$  in the absence of adjustment costs (i.e., as if the adjustment cost draw was equal to zero). Caballero and Engel (1993a) prove that the probability of adjusting is non-decreasing in the absolute size of a firm's price gap (i.e., the so-called ‘increasing hazard property’). Denoting with  $f_t(x)$  the cross-sectional distribution of price gaps immediately before an adjustment takes place at time  $t$ , aggregate inflation can be recovered as

$$\pi_t = - \int x \Lambda_t(x) f_t(x) dx. \quad (3.4.1)$$

Notice that the Calvo pricing protocol implies the same hazard across  $x$ 's (i.e.,  $\Lambda_t(x) = \Lambda_t > 0, \forall x$ ).

### 3.4.1 Taking the model to the data

In order to take the model to the data we need to specify generic functional forms for the distribution of price gaps and the hazard function. Specifically, we postulate that the distribution of price gaps at time  $t$ ,  $f_t(x)$ , can be accounted for by the Asymmetric Power Distribution (APD henceforth; see Komunjer, 2007). The probability density function of an APD random variable is defined as

$$f(x) = \begin{cases} \frac{\delta(\varrho, \nu)^{1/\nu}}{\Gamma(1+1/\nu)} \exp \left[ -\frac{\delta(\varrho, \nu)}{\varrho^\nu} \left| \frac{x-\theta}{\phi} \right|^\nu \right] & \text{if } x \leq \theta \\ \frac{\delta(\varrho, \nu)^{1/\nu}}{\Gamma(1+1/\nu)} \exp \left[ -\frac{\delta(\varrho, \nu)}{(1-\varrho)^\nu} \left| \frac{x-\theta}{\phi} \right|^\nu \right] & \text{if } x > \theta \end{cases}, \quad (3.4.2)$$

with  $\delta(\varrho, \nu) = \frac{2\varrho^\nu(1-\varrho)^\nu}{\varrho^\nu + (1-\varrho)^\nu}$ . The parameters  $\theta$  and  $\phi > 0$  capture the location and the scale of the distribution, whereas  $0 < \varrho < 1$  accounts for its degree of asymmetry.

Last, the parameter  $\nu > 0$  measures the degree of tail decay: for  $\infty > \nu \geq 2$  the distribution is characterized by short tails, whereas it features fat tails when  $2 > \nu > 0$ . This functional form nests a number of standard specifications, such as the Normal ( $\nu = 2$ ), the Laplace ( $\nu = 1$ ) and the Uniform ( $\nu \rightarrow \infty$ ). Most importantly, it can capture intermediate cases between the Normal and the Laplace distribution, which is consistent with the steady-state distribution of price changes according to Alvarez et al. (2016).

We then assume that the hazard function can be characterized by an asymmetric quadratic function:

$$\Lambda_t(x) = \min \{ a_t + b_t x^2 \mathbf{I}(x > 0) + c_t x^2 \mathbf{I}(x < 0), 1 \}, \quad (3.4.3)$$

where  $\mathbf{I}(z)$  is an indicator function taking value 1 when condition  $z$  is verified, and zero otherwise. This parsimonious specification nests the Calvo pricing protocol for  $b_t = c_t = 0$ , while potentially allowing for asymmetric costs of adjustment (so as to be able to capture, for instance, downward stickiness, as implied by  $b_t > c_t$ ).<sup>26</sup>

Given the parametric specifications of  $f_t(x)$  and  $\Lambda_t(x)$ , we estimate seven parameters for each cross section of micro price data, so as to match the following moments of the distribution of price changes: mean, median, standard deviation, interquartile range, difference between the 90th and 10th quantile of the distribution, as well as (quantile-based) skewness and kurtosis.<sup>27</sup> We also match the frequency and the average size of prices movements, after distinguishing between positive and negative price changes. Last, we match the observed rate of inflation. The estimates are obtained by simulated minimum distance, using the identity matrix to weight different moments.<sup>28</sup>

### 3.4.2 Making sense of diverging trends in the frequency and dispersion of price changes

The first two panels of Figure 3.5 report the estimated scale parameter of  $f(x)$  and the inaction region associated with two hazard probabilities (5% and 7%). Both statistics increase substantially in the second decade of the sample, thus implying

---

<sup>26</sup>We have checked that the results are robust to plausible variations to this specification. Specifically, using a mixture of two Normal distributions for the price gap and/or the asymmetric inverted normal function for the hazard function delivers results that are qualitatively similar to those reported in the next section.

<sup>27</sup>We match quantilic moments, as the 3rd and 4th moments of the cross-sectional distribution are quite sensitive to outliers.

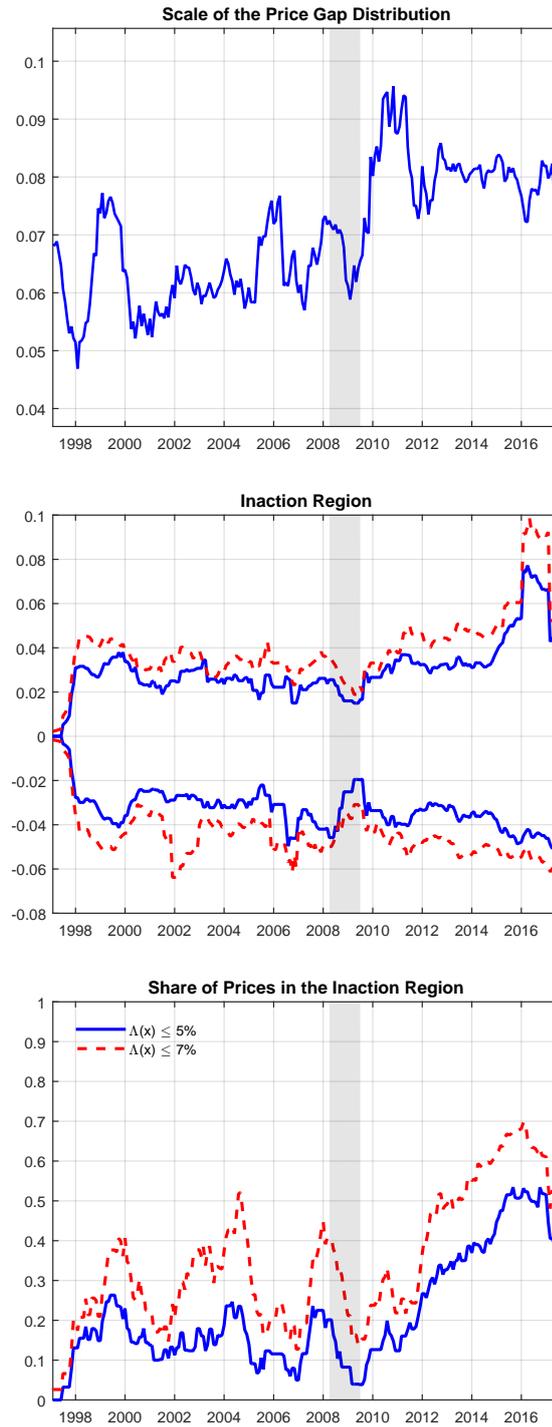
<sup>28</sup>Altonji and Segal (1996) highlight that matching the unweighted distance between moments often performs better in small samples, as compared with using optimal weights. The moments of the simulated distribution are estimated by drawing 100,000 price quotes. We use the Genetic Algorithm to minimize the quadratic distance between data moments and simulated moments, so as avoid ending up in local minima (see, e.g., Dorsey and Mayer, 1995).

that—at least over this period—first-moment shocks do not appear as the main determinant of price adjustment. According to our comparative statics analysis in Section 3.3, a prolonged decline in the frequency of adjustment, coupled with a surge in its dispersion, may be rationalized by an expansion in the inaction region—as prompted by an increase in the fixed cost of adjustment and/or a drop in the cost of deviating from the optimal price, for instance—that overcomes the effects of a positive shift in the dispersion of price gaps. To verify this is indeed the case, the last panel of Figure 3.5 reports the share of prices in the inaction region, obtained as the proportion of prices whose  $\Lambda(x)$  is lower than a given hazard rate. Notably, by the end of the sample about five times as many firms are inactive, as compared with the pre-2010 time window. This stands as indirect evidence that the expansion in the inaction region, as captured by the downward shift in the hazard function, dominates the increase in the dispersion of  $f(x)$ .<sup>29</sup> As we will discuss in the next section, changes in the shape of the distribution of price gaps, coupled with the expansion of the inaction region, imply that non-predetermined price adjustments—which are more likely to occur for large price gaps—have played an increasingly important role in the recent past.

---

<sup>29</sup>Figure 3.13 reports the estimated parameters of the APD, while Figure 3.12 graphs the dynamics of both  $f(x)$  and  $\Lambda(x)$ .

Figure 3.5: Dispersion of Price Gaps and the Inaction Region



Note: The three panels of the figure report the estimated scale parameter of  $f(x)$ , the inaction region (for two different hazard rates), and the corresponding share of prices within the inaction region, respectively. The shaded vertical band indicates the duration of the Great Recession.

### 3.5 Implications for aggregate price adjustment

The estimation of the generalized  $Ss$  model highlights the importance of changes in the distribution of price gaps and the hazard function. To dig deeper into the connection between individual price adjustment and the response of aggregate inflation to nominal demand, Caballero and Engel (2007) highlight that, within their accounting framework, one can derive a measure of aggregate price flexibility that accounts for the impact response of realized inflation to a one-off aggregate nominal shock:

$$\mathcal{F}_t = \lim_{\mu_t \rightarrow 0} \frac{\partial \pi_t}{\partial \mu_t} = \underbrace{\int \Lambda_t(x) f_t(x) dx}_{\text{Intensive Margin}} + \underbrace{\int x \Lambda'_t(x) f_t(x) dx}_{\text{Extensive Margin}}. \quad (3.5.1)$$

Since this flexibility index is simply derived from the accounting identity (3.4.1), its validity as a measure of aggregate flexibility does not require that we take a stand on a specific model of price setting.<sup>30</sup>

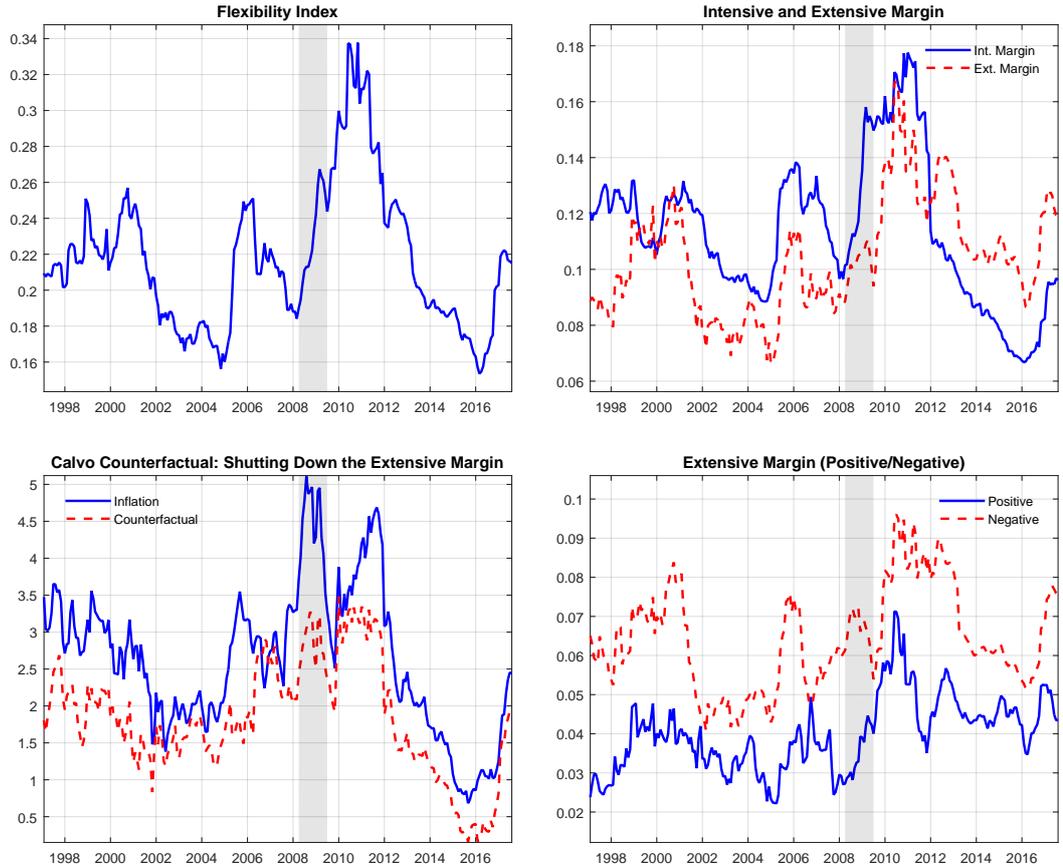
The flexibility index can be naturally decomposed into an intensive and an extensive margin component. On one hand, the intensive margin (*Int*) measures the average frequency of adjustment, and accounts for the part of inflation that reflects price adjustments that would have happened even in the absence of the nominal shock. On the other hand, the extensive margin (*Ext*) accounts for the additional inflation contribution of firms whose decision to adjust is either triggered or cancelled by the nominal shock. Therefore, it comprises both firms who would have kept their price constant and instead change it, as well as firms who would have adjusted their price but choose not to do it. In this respect, it is useful to recall that, being characterized by a constant hazard function, Calvo price setting implicitly assumes that the extensive margin is null.

The top panels of Figure 3.6 report the estimated price flexibility index and its decomposition into the intensive and the extensive margin of price adjustment for the period under investigation. Aggregate price flexibility displays sizeable variation over time, and even more so in the last part of the sample, rising substantially during the Great Recession, and declining thereafter. This is consistent with our analysis of the distribution of price gaps. In fact, after the Great Recession both the intensive and the extensive margin of price adjustment display a contraction, though the fall in the former is much more abrupt, in line with the sustained drop in the frequency of adjustment. As for the extensive margin, the expansion in the inaction region

---

<sup>30</sup>In this respect, Alvarez et al. (2016) show that the steady-state ratio of kurtosis to frequency is a sufficient statistic for monetary non-neutrality in a wide variety of frameworks. However, as highlighted by Berger and Vavra (2017), while their characterization provides us with a measure of cumulative output response, it does not apply to settings that allow for large shocks to the price gap distribution. Despite these fundamental differences, when comparing the two measures obtained from our data, they display a strong negative correlation, as one would expect on theoretical grounds.

Figure 3.6: Price Flexibility and Different Margins of Price Adjustment



Notes: The bottom-left panel reports both the rate of inflation obtained from our sample of ONS price quotes and its counterfactual, obtained by setting the period hazard function to a constant equal to the intensive margin. The shaded vertical band indicates the duration of the Great Recession.

implies that fewer firms are pushed near the adjustment boundaries. Moreover, it should be stressed that, over most of the decline, the extensive margin tends to contribute more to price flexibility, as compared with the intensive one, even after they both revert in 2016. Otherwise, the relative importance of the frequency of adjustment has generally been higher prior to 2012, with few short lived exceptions. To see why we observe such a switch in the relative contribution of the two margins, it is useful to recall Caballero and Engel (2007) and their transformation of (3.5.1):

$$\mathcal{F}_t = \int \Lambda_t(x) f_t(x) [1 + \eta_t(x)] dx \quad (3.5.2)$$

where  $\eta_t(x) = x\Lambda'_t(x)/\Lambda_t(x)$  is the elasticity of the hazard function with respect to the price gap. A downward shift in the hazard function magnifies  $\eta_t(x)$  and, as a

consequence, the importance of the extensive margin relative to the intensive one. This is exactly what happens in the period under examination, as it can be appreciated by inspecting the estimated constant of the hazard function (see Figure 3.14 in Appendix 3.8.6). Alternatively, the same point can be made by approximating the flexibility index as  $F_t \cong Int_t + 2[Int_t - \Lambda_t(0)]$ :<sup>31</sup> from this it is clear how a downward shift in  $a_t$ —which is equivalent to  $\Lambda_t(0)$ —translates into an increase in the importance of the extensive margin relative to the intensive one, *ceteris paribus*.

Table 3.4: Flexibility in Price Adjustment: Correlation with Real Activity and Inflation

Full Sample							
	$\mathcal{F}_t$	$Int_t$	$Ext_t$	$Int_t^+$	$Int_t^-$	$Ext_t^+$	$Ext_t^-$
$y_t$	-0.233***	-0.352***	-0.060	-0.532***	-0.190***	-0.210***	0.044
$\pi_t$	0.380***	0.398***	0.281***	0.005	0.565***	-0.061	0.467***
Pre-Recession							
	$\mathcal{F}_t$	$Int_t$	$Ext_t$	$Int_t^+$	$Int_t^-$	$Ext_t^+$	$Ext_t^-$
$y_t$	0.456***	0.368***	0.412***	0.223**	0.395***	0.403***	0.331***
$\pi_t$	-0.012	0.269***	-0.279***	0.062	0.345***	-0.311***	-0.221**
Post-Recession							
	$\mathcal{F}_t$	$Int_t$	$Ext_t$	$Int_t^+$	$Int_t^-$	$Ext_t^+$	$Ext_t^-$
$y_t$	-0.527***	-0.428***	-0.559***	-0.363***	-0.416***	-0.289***	-0.632***
$\pi_t$	0.678***	0.718***	0.512***	0.372***	0.787***	0.084	0.721***

Notes: The table reports pairwise correlations of output and inflation with the flexibility index, as well as the intensive margin and the extensive margin of price adjustment (together with their counterparts corresponding to positive and negative price gaps). Aside of the inflation rate, all series are obtained by detrending their raw counterparts by means of Rotemberg (1999) version of the HP filter, which sets the smoothing coefficient so as to minimize the correlation between the cycle and the first difference of the trend estimate. \*\*\*/\*\*/\* indicates statistical significance at the 1/5/10% level, respectively.

To gauge the actual contribution of the extensive margin to inflation dynamics, we can take a step further: the bottom-left panel of Figure 3.6 reports both the overall rate of inflation and its counterfactual, obtained by setting the period hazard function to a constant equal to the intensive margin. As pointed out by Gagnon et al. (2013), this is equivalent to calibrating the Calvo model to match the intensive margin of price adjustment by assuming that the probability of price adjustment, while exogenous to the firm, can vary with the state of the economy (i.e.,  $\pi_t^{Calvo} = -f r_t^{Calvo} \int x f_t(x) dx$ , where  $f r_t^{Calvo} = \int \Lambda_t(x) f_t(x) dx$ ). The presence of an increasing hazard function tends to exacerbate the impact of large shocks (Caballero and Engel, 1991). In fact, the extensive margin proves to be rather important in periods of particularly volatile inflation, so that the difference between

<sup>31</sup>For a formal proof, please refer to Caballero and Engel (2007).

the latter and its ‘Calvo counterfactual’ is sizeable. In this respect, it is important to appreciate how movements along the extensive margin may reflect some asymmetries in the adjustment of prices in either direction. To this end, the last panel of the figure reports the extensive margin associated with positive and negative price gaps ( $Ext^+$  and  $Ext^-$ , respectively).<sup>32</sup> Both statistics denote a shift during the last part of the sample, with  $Ext^+$  leading the increase in the wake of the Great Recession, and  $Ext^-$  reflecting the two hikes in the VAT at the beginning of 2010 and 2011. This aspect will be examined in further detail in the next subsection.

From a business cycle perspective, variations in price flexibility do not seem to occur at random: in fact,  $\mathcal{F}_t$  goes from being markedly procyclical in the first part of the sample to invert its cyclicity in the last decade (see Table 3.4). As for the correlation with the rate of inflation, this is generally positive, and more so in the post-recession sample, while it is not statistically different from zero in the previous decade. Analogous changes in the correlation with real activity occur when looking at  $Ext$  and  $Int$  over the two subsamples, while the strong positive correlation with the rate of inflation is even more pronounced in the last decade. It is interesting to notice that, over the full sample, both margins denote a negative correlation with the cyclical indicator, and even more so for  $Ext^+$  and  $Int^+$ .<sup>33</sup> This fact, in conjunction with a correlation with the rate of inflation that is not statistically different from zero, might indicate a certain degree of downward rigidity, given that nominal shocks appear not to be able to stimulate price cuts along both margins. In this respect, the correlation structure of both margins of (negative) price adjustment changes markedly over the two subsamples, indicating that price cuts might have been particularly sticky during the Great Recession.<sup>34</sup>

As a final note on the change in correlation we observe over the two subsamples, it is worth emphasizing how this is consistent with a shift from an environment where the intensive margin dominates the extensive one, to an environment where the extensive margin assumes a prominent role and inflation volatility is particularly marked (see Figure 3.6).

---

<sup>32</sup>To this end, we simply rely on the following decomposition of the extensive margin:  $\int_{-\infty}^{0^-} x\Lambda'_t(x) f_t(x) dx + \int_0^{\infty} x\Lambda'_t(x) f_t(x) dx$ , where  $Ext_t^-$  ( $Ext_t^+$ ) is the first (second) term on the right side of the equality. To see a similar split for the intensive margin, recall that the bottom-left panel Figure 3.1 reports the frequency of positive and negative price changes.

<sup>33</sup>When looking at the two subsamples separately, we notice that such a countercyclicity is a hallmark of the last decade, while comovement is positive in the pre-recession period.

<sup>34</sup>In this respect, Gilchrist et al. (2017) have shown how the interplay between price stickiness and financial frictions faced by firms operating in customer markets might have acted as sources of downward price rigidity during the Great Recession.

### 3.5.1 Price adjustment and the importance of state-dependent pricing: a VAT event study

Examining the relative importance of price adjustment along the extensive margin is of key importance to contrast time-dependent models that are widely employed in quantitative macroeconomic frameworks, with state-dependent models. In this respect Gagnon et al. (2013) suggest that, if the timing of all price changes was predetermined, following a nominal shock we should observe a shift in the gap distribution, with the shape of the distribution being preserved (see, e.g., the middle panel of Figure 3.3). Thus, one can measure the importance of adjustment along the extensive margin by comparing the observed distribution of price changes to a counterfactual distribution that obtains in the absence of the shock. Any evidence that the two distributions differ by more than a shift can be attributed to the extensive margin. To this end, we can usefully exploit episodes of massive repricing triggered by changes in the VAT. These are relatively simple to study, because their timing and size are directly observable.

The recent UK history has been characterized by three episodes of changes in the VAT: a reduction from 17.5% to 15% on December 1, 2008, followed by two hikes: one up to 17.5% on January 1, 2010 and one, further up to 20%, on January 4, 2011. To examine the contribution of VAT changes to the overall degree of price flexibility, Figure 3.7 reports the distribution of price gaps and that of price changes, together with the corresponding hazard function. Moreover, we report their counterfactuals, obtained by averaging the same function, for the same month of the year, in the previous six years.<sup>35</sup>

Looking at the inflation rate in the month corresponding to a VAT change, we notice that shifts in the distribution of price changes are such that many firms seize the opportunity to adjust prices by more than the VAT change, thus implying that inflationary/deflationary pressures from other sources have been released in the process. In support of the view that episodes of massive repricing cannot be seen as mere translations of the distribution of price gaps, we appreciate both a major upward shift and a steepening of the hazard function across all the three episodes of VAT change: in fact, these are associated with a large rise in the frequency of adjustment.

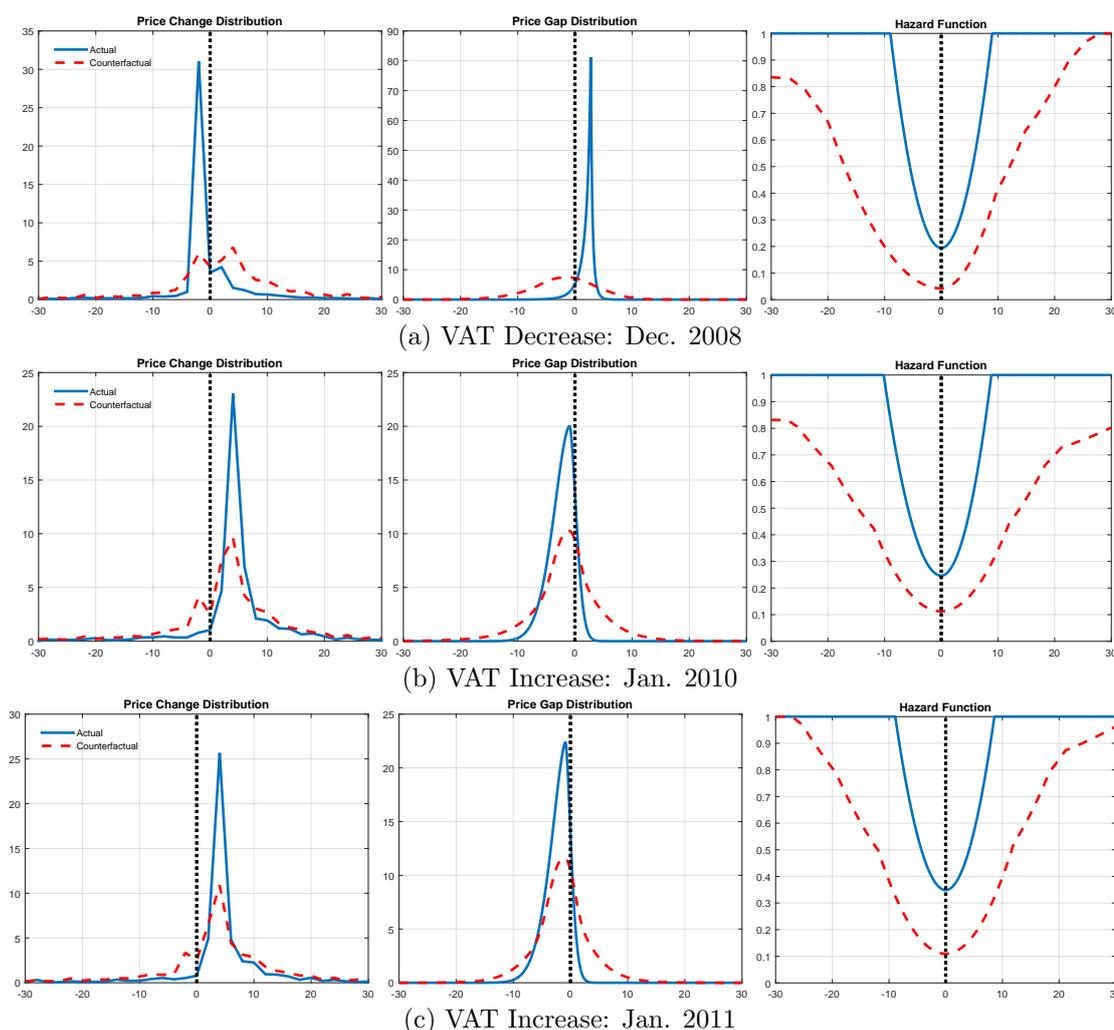
To dig deeper into the role of state-dependent pricing, Table 3.5 reports some statistics in coincidence with the three VAT changes, as well as two alternative scenarios.<sup>36</sup> In the first scenario, we keep the hazard function as that computed in

---

<sup>35</sup>January 2010 has not been included when computing the counterfactual distribution for January 2011, so as to avoid that the second VAT change affects the counterfactual distribution corresponding to the last episode.

<sup>36</sup>More details on the computation of two alternative scenarios are provided in Appendix 3.8.4.

Figure 3.7: Event Study: VAT Changes



Notes: Each line of the figure reports the distribution of price changes, the distribution of price gaps, and the hazard function in the month corresponding to a VAT change. The distribution of price changes is computed by grouping observations into bins of 2% (excluding zeros), and weighting them by their relative importance in the CPI. In all cases, the counterfactuals are computed by averaging the same function, for the same month of the year in the previous 6 years. Three recent episodes of changes in the VAT are considered: a reduction from 17.5% to 15% on December 1, 2008, followed by two hikes, on up to 17.5% on January 1, 2010 and then up to 20% on January 4, 2011.

the counterfactual exercise, but let the price gap distribution vary as a result of the VAT change. Therefore, we abstract from any amplification that could be potentially induced by state-dependent pricing through upward shifts of the hazard function. The second scenario, instead, considers a hypothetical case in which neither the price gap distribution nor the hazard function are affected by the VAT change. From the comparison between actual inflation in the occurrence of a VAT change and its counterfactuals in the alternative scenarios, two observations are worth emphasizing.

Table 3.5: VAT Changes: Actual and Counterfactual Statistics

VAT 1								
	$\pi$	$\mathcal{F}$	$Int$	$Ext$	$Int^+$	$Int^-$	$Ext^+$	$Ext^-$
Actual	-5.941	0.346	0.235	0.111	0.211	0.023	0.105	0.006
Scenario 1	-1.604	0.101	0.060	0.041	0.055	0.005	0.040	0.001
Scenario 2	1.863	0.200	0.096	0.104	0.038	0.058	0.048	0.056
VAT 2								
	$\pi$	$\mathcal{F}$	$Int$	$Ext$	$Int^+$	$Int^-$	$Ext^+$	$Ext^-$
Actual	11.631	0.471	0.322	0.149	0.019	0.304	0.003	0.146
Scenario 1	4.580	0.181	0.135	0.045	0.008	0.127	0.001	0.045
Scenario 2	4.111	0.218	0.148	0.070	0.043	0.105	0.016	0.054
VAT 3								
	$\pi$	$\mathcal{F}$	$Int$	$Ext$	$Int^+$	$Int^-$	$Ext^+$	$Ext^-$
Actual	14.487	0.573	0.428	0.145	0.019	0.409	0.002	0.143
Scenario 1	4.708	0.190	0.136	0.053	0.006	0.130	0.001	0.053
Scenario 2	4.258	0.239	0.154	0.086	0.041	0.113	0.020	0.066

Notes: The table reports the inflation rate, the inflation rate that would have been observed had there not been any extensive margin, the flexibility index, the intensive and extensive margins of price adjustment (as well as their counterparts computed for positive and negative price gaps), all in the month of a VAT change. Three recent episodes of changes in the VAT are considered: a reduction from 17.5% to 15% on December 1, 2008 (indicated by VAT 1), followed by two hikes, on up to 17.5% on January 1, 2010 and then up to 20% on January 4, 2011 (indicated by VAT 2 and VAT 3, respectively). For every episode we contrast the actual numbers with two alternative scenarios. Scenario 1 considers a case in which the VAT change only impacts on the distribution of price gaps, while keeping the hazard function at its counterfactual in Figure 7. Scenario 2, instead, consider an alternative where neither the hazard function nor the price gap distribution change.

First, state-dependent pricing accounts for most of the change in the rate of inflation. Second, in the absence of state-dependent pricing, shifts of the price gap distribution and drops in its dispersion would result in a substantial drop of price flexibility, with both the intensive and the extensive margin decreasing.

Importantly, when comparing the two margins of adjustment, the intensive one is typically twice as large as its counterfactual—indicating that upward shifts in  $\Lambda(0)$  are the dominant feature in the occurrence of changes in the VAT—while movements along the extensive margin appear less evident. However, this conclusion is not warranted after conditioning both margins to positive and negative price changes. In this case, substantial variation takes place along the extensive margin coherent with the sign of the underlying price change. For instance, in the occurrence of the drop in the VAT from 17.5% to 15% (December 2008),  $Ext^+$  is 0.048 in the counterfactual, while actually being more than twice as big. The same order of magnitude can be appreciated when making the same comparison for two VAT hikes (in this case we need to focus on  $Ext^-$ ). Movements in the extensive margin

are a reflection of the interplay between the hazard function and the distribution of price gaps. In this respect, Figure 3.7 shows that all the episodes of VAT change are associated with a close-to-symmetric increase in the steepness of the hazard function, as well as with a shift in the distribution of price gaps in the direction opposite to the sign of a given VAT change. On one hand, this necessarily implies that the extensive margin associated with price gaps coherent with the sign of the adjustment is substantial. On the other hand, the extensive margin associated with price gaps of the opposite sign is very low, as consequence of the hazard function being weighed by a very small probability mass, after the shift in the distribution of price gaps.

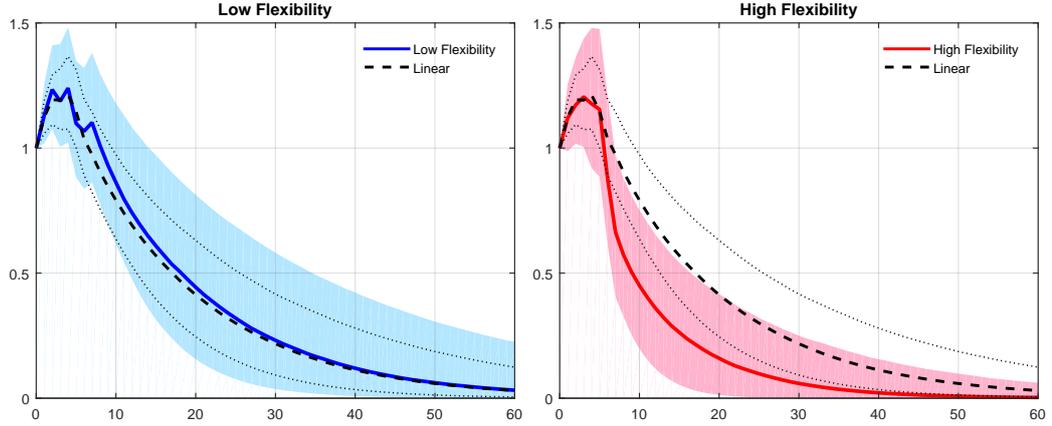
On a slightly different note, it should be stressed that price flexibility reaches its maximum over roughly the same period we observe the most recent VAT price changes. In light of our analysis, this comes as little surprise, given that this type of events typically offer price setters with some windows of opportunity to release at least some of the accumulated (positive/negative) price pressure. However, from a normative perspective the opportunity to enhance coordination between fiscal and monetary policy should be carefully considered. Such a prescription might be particularly relevant in contexts such as the one examined, where the potential real effects of the accommodative monetary-policy stance might have been baffled by VAT changes that were mainly inspired by stimulus- or revenue-based considerations.

### **3.6 Inflation dynamics and state dependence in price flexibility**

The estimation of the generalized  $Ss$  model shows that the pass-through of nominal shocks to inflation is highly variable. We also show that—while not hinging on a specific margin of adjustment—flexibility is higher in connection with positive price changes, while downward price adjustments are typically stickier. These properties bear major implications for evaluating the transmission of shocks to nominal demand. In fact, at the eyes of a hypothetical Central Banker, aggregate price flexibility should be regarded as a key state variable to predict the influence of a given monetary policy stance on prices and quantities.

While aggregate price flexibility only accounts for the response of inflation to a nominal shock, one would expect it to contain valuable information to study state dependence in inflation dynamics. In this section we seek to examine to what extent inflation behaves differently in periods of high and low flexibility. To this end, we employ a regime-switching autoregressive moving average model, where the transition across regimes is a smooth function of the degree of price flexibility.

Figure 3.8: Price Flexibility and Inflation Persistence



Note: This figure reports the responses of inflation to a 1% shock in the STARMA(1,7) model. The left (right) panel graphs the response in the low (high) price flexibility regime. In both cases we also report the the response from a (linear) ARMA(1,7) model. 68% confidence intervals are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003).

The STARMA(p,q) model is a generalization of the smooth transition autoregression model proposed by Granger and Terasvirta (1993).<sup>37</sup> Estimating a traditional ARMA(p,q) for each regime separately entails a certain disadvantage in that we may end up with relatively few observations in a certain regime, which typically renders the estimates unstable and imprecise. By contrast, we can effectively rely upon more information by exploiting variation in the probability of being in a particular regime, so that estimation and inference for each regime are based on a larger set of observations (Auerbach and Gorodnichenko, 2012).<sup>38</sup>

We assume that inflation can be described by the following model:

$$\begin{aligned} \pi_t = & G\left(\tilde{\mathcal{F}}_{t-1}, \gamma\right) \left( \phi_0^H + \sum_{j=1}^p \phi_j^H \pi_{t-j} + \varepsilon_t^H + \sum_{i=1}^q \theta_i^H \varepsilon_{t-i}^H \right) \\ & + \left[ 1 - G\left(\tilde{\mathcal{F}}_{t-1}, \gamma\right) \right] \left( \phi_0^L + \sum_{j=1}^p \phi_j^L \pi_{t-j} + \varepsilon_t^L + \sum_{i=1}^q \theta_i^L \varepsilon_{t-i}^L \right), \end{aligned} \quad (3.6.1)$$

with  $\varepsilon_t^i \sim N(0, \sigma_i^2)$  for  $i = \{L, H\}$ . Moreover, we set  $G\left(\tilde{\mathcal{F}}, \gamma\right) = (1 + e^{-\gamma\tilde{\mathcal{F}}})^{-1}$ ,

<sup>37</sup>In this respect, the STARMA(p,q) model also generalizes the threshold ARMA(p,q) model (DeGooijer, 2017).

<sup>38</sup>Estimating the properties of a given regime by relying on the dynamics of inflation in a different regime would bias our results towards not finding any evidence of non-linearity. In light of this, the asymmetries we will be reporting in the remainder of this section acquire even more statistical relevance.

where  $\tilde{\mathcal{F}}$  denotes the normalized flexibility index and  $\gamma$  is the speed of transition across regimes.<sup>39</sup> We allow for different degrees of inflation persistence across the two regimes, as captured by the regime-specific autoregressive and moving average coefficients, as well as for different volatilities of the innovations. The likelihood of the model can be easily computed by recasting the system in state space (see, e.g., Harvey, 1990). We use Monte Carlo Markov-chain methods developed in Chernozhukov and Hong (2003) for estimation and inference. The parameter estimates, as well as their standard errors, are directly computed from the generated chains.<sup>40</sup>

As we focus on the post-1996 sample, we calibrate the constant terms  $\phi_0^H$  and  $\phi_0^L$  so that in both regimes the long-run inflation forecast is 2%, consistent with the mandate of the Bank of England. Whereas one can potentially estimate the speed of transition between regimes, the identification of  $\gamma$  relies on nonlinear moments. Moreover, in short samples the estimates may be sensitive to a handful of observations. Therefore, we decide to calibrate  $\gamma$  so that roughly 25% of the observations are classified to be in the high-flexibility (low-flexibility) regime, where this is defined by  $G(\tilde{\mathcal{F}}_{t-1}; \gamma) > 0.8$  ( $G(\tilde{\mathcal{F}}_{t-1}; \gamma) < 0.2$ ).<sup>41</sup> Thus, based on the Akaike criterion, we choose  $p = 1$  and  $q = 7$ .<sup>42</sup>

Figure 3.8 reports the impulse-response functions to a 1% shock to inflation in each of the two regimes, and compares them to the response from an equivalent linear model. Inflation is much more persistent in periods characterized by a relatively low price flexibility, with the half-life of the shock being almost twice as large, as compared with periods of high flexibility. In fact, the estimated inflation volatility is 1.44 in the high-flexibility regime and 0.91 in the low-flexibility regime. These results are broadly supportive of the basic insights of the  $S$ s model illustrated in the previous sections, and highlight the importance of keeping track of the degree of price flexibility.

Notably, the impulse-response function from the linear model is consistent with the behaviour of inflation in the low-flexibility regime. A direct implication of this is that neglecting that shocks are propagated at different speeds—depending on the overall degree of price flexibility—would entail an overestimation of their inflationary impact during windows of relatively high flexibility. This should be particularly

---

<sup>39</sup>We employ a backward-looking MA(12) of the flexibility index to get rid of seasonality in the data. Moreover, we lag the index by one month, in order to avoid potential endogeneity with respect to CPI inflation.

<sup>40</sup>See Appendix 3.8.5 for further details.

<sup>41</sup>Figure 3.15 in Appendix 3.8.7 reports the dynamics of  $G(\tilde{\mathcal{F}}_{t-1}; \gamma)$ . Clearly, this specification identifies the 2009-2012 period as being characterized by a high-flexibility regime, whereas the 2002-2005 and 2015-2016 periods are marked by low price flexibility. The qualitative results are robust to variations in  $\gamma$ .

<sup>42</sup>Note that the modified AIC information criterion indicates a STARMA(1,3). Figures 3.16 and 3.17 in Appendix 3.8.7 report the results for this alternative setting. Our key insights are not affected by the exact specification of the STARMA(p,q) model.

Table 3.6: Forecast Errors and Price Flexibility

(a) BoE MPC RPIX/CPI (Absolute) Forecast Errors						
Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F- <i>stat</i>	$\tilde{R}^2$
1	0.093	[0.628]	0.840	[0.092]	0.229	1.69
2	-0.330	[0.279]	2.319	[0.011]	0.045	6.41
3	-0.484	[0.145]	4.117	[0.010]	0.003	13.82
4	-0.344	[0.437]	6.161	[0.003]	0.000	26.45
5	-0.144	[0.811]	5.945	[0.011]	0.000	20.10
6	0.309	[0.603]	4.858	[0.032]	0.003	13.70
7	0.634	[0.236]	4.402	[0.021]	0.006	12.32
8	0.691	[0.182]	3.029	[0.055]	0.063	5.93
(b) Market Participants' (Absolute) Forecast Errors						
Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F- <i>stat</i>	$\tilde{R}^2$
1	0.265	[0.361]	0.826	[0.122]	0.278	1.11
2	-0.383	[0.264]	2.448	[0.010]	0.053	6.12
3	-0.561	[0.150]	4.293	[0.008]	0.004	13.10
4	-0.382	[0.418]	6.398	[0.002]	0.000	25.60
5	-0.103	[0.862]	6.042	[0.009]	0.000	18.74
6	0.453	[0.412]	4.516	[0.049]	0.013	10.48
7	0.903	[0.052]	3.631	[0.052]	0.019	9.47
8	0.883	[0.099]	1.935	[0.221]	0.211	2.19

Notes: The table reports the results of a quadratic spline regression of the absolute forecast errors  $e_{T+h|T}$  (for different forecast horizons,  $h$ ) on a quarterly average of an indicator of the normalized price flexibility index,  $G_t = G(\tilde{\mathcal{F}}_t; \gamma) = (1 + e^{-\gamma\tilde{\mathcal{F}}_t})^{-1}$ , where  $\tilde{\mathcal{F}}$  denotes the normalized flexibility index. The regression takes the form:  $|e_{T+h|T}| = a_0 + a_1G_t + a_2G_t^2 + a_3G_t^2I_{\{G_t > 0.5\}}$ . The upper panel refers to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panel considers market participants' forecast errors. In each panel, the first two pairs of columns report the slope of the relationship evaluated at different levels of the indicator, together the p-value associated with the null hypothesis that the slope is equal to 0 (this is calculated using Newey-West standard errors). The penultimate column (F-*stat*) reports the p-value of the null hypothesis that all the coefficients associated to the flexibility regime are equal to 0. The last column reports the adjusted R-squared, denoted by  $\tilde{R}^2$ .

evident at medium-term forecast horizons, i.e. when the difference between the responses from the linear and the nonlinear model is somewhat larger. This begs the following question: do the Bank of England or other market participants take price flexibility into account when computing their inflation expectations? In the remainder of this section we turn our attention to addressing this issue. In this respect, our premise delivers a key testable implication: if state dependence in price flexibility is accounted for by the forecaster, the resulting inflation forecast errors should be orthogonal to the flexibility regime.

In every quarter, the Inflation Report of the Bank of England publishes (year-on-year) Monetary Policy Committee's inflation forecasts, along with market participants' forecasts. Both types of forecasts refer to the Bank of England's inflation target, which has switched from RPIX inflation to CPI inflation in 2004:Q1. Thus, we construct quarterly (absolute) forecast errors as the (absolute value of the) dif-

ference between the mean forecast<sup>43</sup> and the appropriate forecast target at a given horizon. These are then regressed on a nonlinear function of the flexibility regime indicator,  $G(\tilde{\mathcal{F}}_{t-1}; \gamma)$ : specifically, we use a quadratic spline function with a knot at 0.5. This function is a rather flexible tool, as it allows us to capture a number of potential shapes characterizing the relationship between the flexibility regime and the forecast errors.

Table 3.6 provides a summary of the results from our regression exercise. The fitted function tends to reach a minimum at about  $G(\tilde{\mathcal{F}}_{t-1}; \gamma) = 0.6$ , for all forecast horizons. Thus, we report the slope of the function at values of the indicator equal to 0.3 and 0.9 (so as to consider an equal distance from the minimum point). The last two columns of the table also report the p-value associated with the null that no relationship between the forecast error and the flexibility regime exists, as well as the R-squared (adjusted for the number of regressors), so as to get an idea of the strength of the relationship. The results are consistent with the idea that information about the degree of price flexibility is not fully exploited by the Central Bank or by market participants. In line with Figure 3.8, we find that the relationship tends to be stronger at medium-term horizons, while weakening at both short-term and long-term horizons. Specifically, around a four-quarter horizon, price flexibility accounts for roughly 25% of the variability in the absolute forecast error. The relationship is not statistically significant in periods of low flexibility ( $G = 0.3$ ), whereas it is positive and usually significant when flexibility is relatively high ( $G = 0.9$ ), with the slope displaying larger values at medium-term forecast horizons. These results are roughly the same, no matter which source of forecasts we consider.

The pronounced time variation in price flexibility after the Great Recession helps us to get a better understanding of the concurrent dynamics of the inflation rate. Over this time window inflation peaks twice between 2008 and 2011, while reaching its sample minimum in 2016, partially reflecting sharp movements in the value of the GBP and commodity prices.<sup>44</sup> The Bank of England has generally underestimated the speed and impact of shocks to inflation in the 2008-2011 period. In light of our results, this points to a potential failure in appreciating that price flexibility was itself at its historical peak, possibly as a reflection of the three VAT adjustments taking place over a rather short time window. Conversely, the low-flexibility regime can explain the protracted period of low inflation towards the end of the sample, during which the Bank of England has displayed greater predictive accuracy. This

---

<sup>43</sup>Table 3.11 in Appendix 3.8.7 reports similar results using squared forecast errors. In both cases, the results are virtually unchanged if we use median in place of mean forecasts.

<sup>44</sup>Two main facts are worth noticing with respect to the UK experience in the post-recession sample: i) inflation has been outside the 1%-3% interval for a total of 22 out of 40 quarters, while the same has happened only in 11 quarters during the previous decade; ii) over the same period, inflation has also shot above and below the target, reaching both its maximum (+4.8%) and minimum value (-0.1%) in the overall sample.

regime of low price flexibility has then reversed in the summer of 2016, in coincidence with the sharp movements of the GBP in the aftermath of the Brexit referendum.

### 3.7 Concluding remarks

We document some distinctive patterns in the evolution of the distribution of micro price changes in the UK, and discuss their implications for the transmission of nominal stimulus to output and inflation. By estimating the generalized  $Ss$  model of Caballero and Engel (2007), we are able to report that price flexibility displays pronounced time variation. In fact, over the last decade the capacity of nominal stimulus to generate inflation has decreased substantially. Despite the marked non-linearity in the price response to inflationary shocks—which is crucially dictated by the degree of price flexibility—neither the Bank of England nor professional forecasters appear to account for this type of state dependence when forecasting CPI inflation. In fact, both of them tend to overestimate the impact of inflationary shocks in periods of relatively high price flexibility, especially at medium-term forecast horizons. In light of this, we point to price flexibility as a state variable that both practitioners and policy makers should carefully account for in their forecasting routine. In this respect, we also suggest that time variation in price flexibility should be considered as a key dimension of monetary-policy making. To this end, we observe that changes in price flexibility correlate with departures of CPI inflation from the target, potentially providing the policy-maker with a basis for assessing her state-contingent capability to influence prices and output growth.

A final note on the implications of our results for modelling price setting: by imposing a Calvo price-setting protocol to match the frequency of adjustment one could understate time-variation in price flexibility, which is heavily influenced by the extensive margin of price setting, especially during periods of high volatility in inflation dynamics. In this respect, our work does not just emphasize the importance of time variation in higher moments of the distribution of price changes and their connection with price flexibility—one of the main conclusions of works in this area of research—but also assigns a prominent role to state-dependent price setting in order to understand inflation dynamics, which is what Central Banks and practitioners are ultimately concerned with.

## 3.8 Appendices

### 3.8.1 On the representativeness of the data

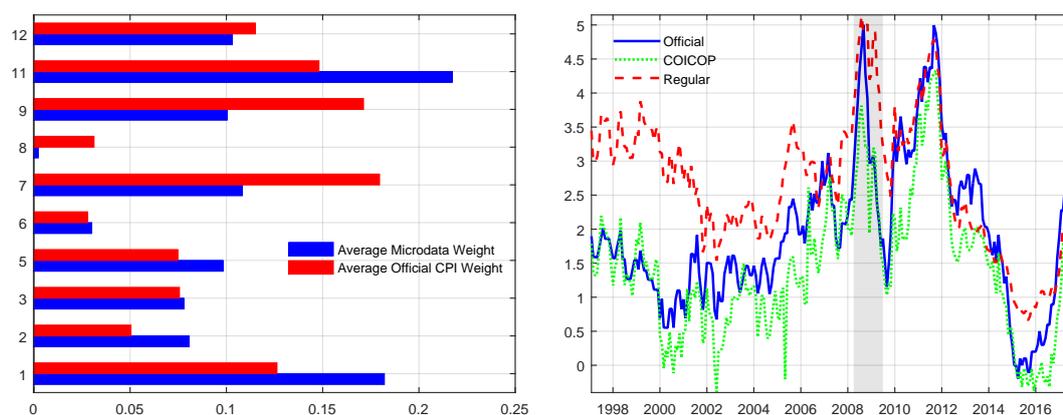
This section provides additional details on the construction of the dataset used in the empirical analysis. The ONS data have a good coverage of all COICOP sectors, with the exception of Housing, Water, Electricity, Gas And Other Fuels (COICOP 4), Communication (COICOP 8) and Education (COICOP 10), whose coverage are less than 15%, 4%, and 3%, respectively. Given the extremely low coverage, we exclude COICOP 4 and 10. We keep COICOP 8, as the available price quotes are clustered in a small subset of items, such as Flower Delivery, Telephone for home use and Phone Accessories.<sup>45</sup>

The left panel of Figure 3.9 contrasts the weights assigned to each of the COICOP sectors to those employed to build the CPI (re-normalized to exclude COICOP 4 and 10). Overall, we observe that using the available price quotes results into relatively larger weights for COICOP 1 and 11, whereas sectors 7 and 9 are underweighed. The right panel of Figure 3.9 reports the official CPI inflation together with the inflation series retrieved from all the available price quotes (labelled *COICOP*) and the inflation obtained once all filters described in Section 3.2 are applied (labelled *Regular*). Unfiltered data track quite closely the official numbers, whereas the ‘regular’ series displays a robust correlation with the official data (roughly 0.7), and shows a positive bias. The latter mainly emerges from the exclusion of sales from the sample.

---

<sup>45</sup>Due to the small number of price quotes in this sector, the results would be little affected by its exclusion from the analysis.

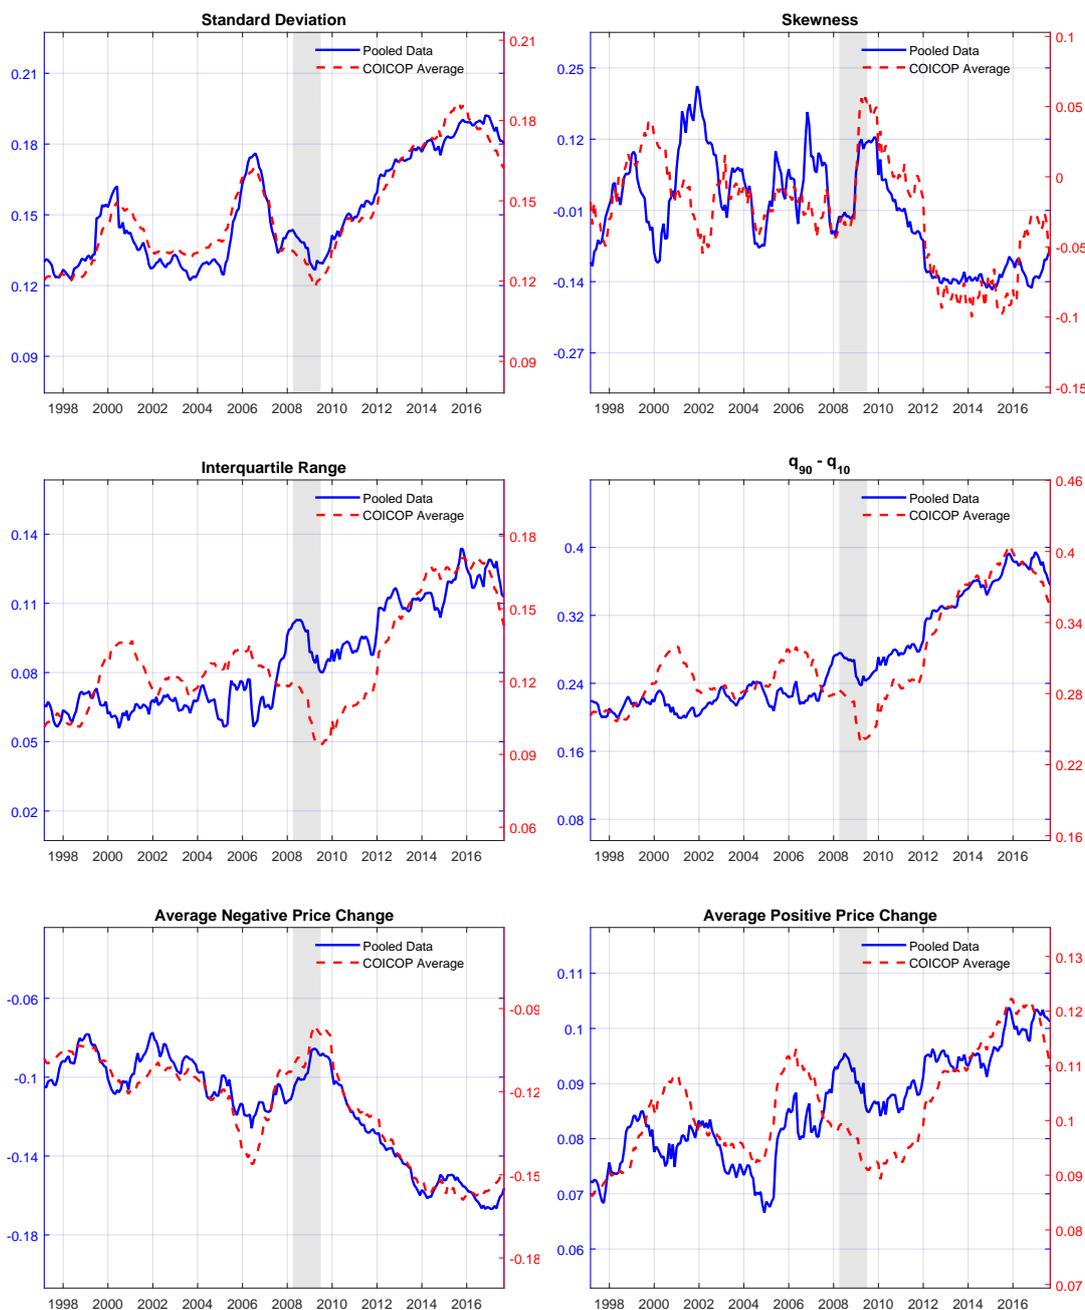
Figure 3.9: Representativeness



Notes: The left panel contrasts the weights assigned to each of the COICOP sectors to those assigned to build the CPI (re-normalized to exclude COICOP 4 and 10). The right panel reports the official CPI inflation, together with the inflation series retrieved from all the available price quotes (labelled *COICOP*) and the inflation obtained once all filters described in Section 3.2 are applied (labelled *Regular*). The COICOP codes are (1) Food And Non-Alcoholic Beverages, (2) Alcoholic Beverages, Tobacco And Narcotics, Clothing And Footwear (3), Furnishings, Household Equipment And Routine Household Maintenance (5), Health (6), Transport (7), Communication (8), Recreation And Culture (9), Hotels, Cafes And Restaurants (11), Miscellaneous Goods And Services (12).

### 3.8.2 On the role of aggregation and composition effects

Figure 3.10: Aggregate vs Disaggregated Moments



Notes: The figure compares various moments of the distribution of price changes with their homologues obtained by averaging the corresponding moments of the price quotes obtained for each of the 25 COICOP group categories. The shaded vertical band indicates the duration of the Great Recession.

### 3.8.3 A monthly coincident indicator of economic activity

In Tables 3.2, 3.3 and 3.4 we report the correlation of different variables with respect to a business cycle indicator. The latter is computed as a monthly coincident indicator of GDP growth, where we use monthly information on a number of monthly

macroeconomic indicators of economic activity to infer the underlying movements of GDP at the monthly frequency.<sup>46</sup> Following Mariano and Murasawa (2003), we approximate the (normalized) quarterly growth of real GDP,  $\Delta y_t^q$ , as a moving average of an unobserved month-on-month GDP growth rate,  $\Delta y_t^*$ :

$$\Delta y_t^q = \frac{1}{3}\Delta y_t^* + \frac{2}{3}\Delta y_{t-1}^* + \Delta y_{t-2}^* + \frac{2}{3}\Delta y_{t-3}^* + \frac{1}{3}\Delta y_{t-4}^*.$$

We then assume that  $\Delta y_t^*$  can be decomposed into an aggregate component,  $\alpha_t$ , which is common across a number of other macroeconomic indicators, and an idiosyncratic component,  $\varepsilon_t$ :

$$\Delta y_t^* = \alpha_t + \varepsilon_t.$$

We assume that the idiosyncratic component follows an autoregressive process of order one:

$$\varepsilon_t = \psi\varepsilon_{t-1} + \eta_t.$$

The other macroeconomic indicators are available at a monthly frequency. We specify (the standardized value of) each of them as the sum of two mutually orthogonal components, a common and an idiosyncratic one. The former is captured by the current and lagged values of the aggregate common factor (see, e.g., D'Agostino et al., 2016). Specifically, denoting with  $\Delta x_{it}$  the generic  $i$ -th macroeconomic indicator, we have that

$$\Delta x_{it} = \sum_{j=1}^l \lambda_{ij}\alpha_{t-j} + e_{it},$$

where  $e_{it}$  follows an autoregressive process of order one:

$$e_{it} = \rho_i e_{it-1} + v_{it},$$

where the innovations to the idiosyncratic process are *iid* and uncorrelated across the indicators (i.e.,  $E(v_{it}v_{jt}) = 0, \forall i \neq j$ , and  $E(v_{it}\eta_t) = 0, \forall i$ ).

We let the aggregate factor follow an autoregressive process of order two:

$$\alpha_t = \phi_1\alpha_{t-1} + \phi_2\alpha_{t-2} + u_t.$$

In our specific application, we set  $l = 3$  and all autoregressive processes are restricted to be stationary. The model can be cast in state space. Therefore, the likelihood can be easily computed through the Kalman filter and the factor is retrieved by using the Kalman smoother (see Harvey, 1990).

---

<sup>46</sup>In the following we use a Dynamic Factor Model to estimate a latent index of monthly GDP. An alternative method for estimating such an index is Mitchell et al. (2005), and in future research it would be interesting verify our result using this estimate.

Together with the GDP data, we use following short term (monthly) macroeconomic indicators: (1) the index of manufacturing, (2) the index of services, (3) retail sales (excl. Auto Fuel), (4) Employment and (5) unemployment (claimants count). We use data starting on January 1990: we rely on a sample that is longer than the one employed in our analysis, so as to include two recessionary episodes. The dataset is unbalanced, as some of the indicators are not available from the starting date (and GDP is observed only once in the quarter). This is not an issue, as the Kalman filter can easily deal with an arbitrary pattern of missing observations in the sample.

Table 3.7 reports the fit of the aggregate components for the quarter-on-quarter growth rates of each of the variables being employed. Clearly, the single-factor specification is able to capture a large fraction of the variation in the set of indicators considered here. Figure 3.11 reports quarter-on-quarter variations in the aggregate factor ( $\alpha_t^q = \frac{1}{3}\alpha_t + \frac{2}{3}\alpha_{t-1} + \alpha_{t-2} + \frac{2}{3}\alpha_{t-3} + \frac{1}{3}\alpha_{t-4}$ ), together with the GDP growth. The level of the business cycle indicator is then computed by cumulating the common factor over time, and assuming that trend growth equals the mean of GDP growth over the sample (this is denoted by  $\mu$ ):

$$z_t = \sum_{\tau=1}^t (\hat{\mu} + \hat{\alpha}_\tau),$$

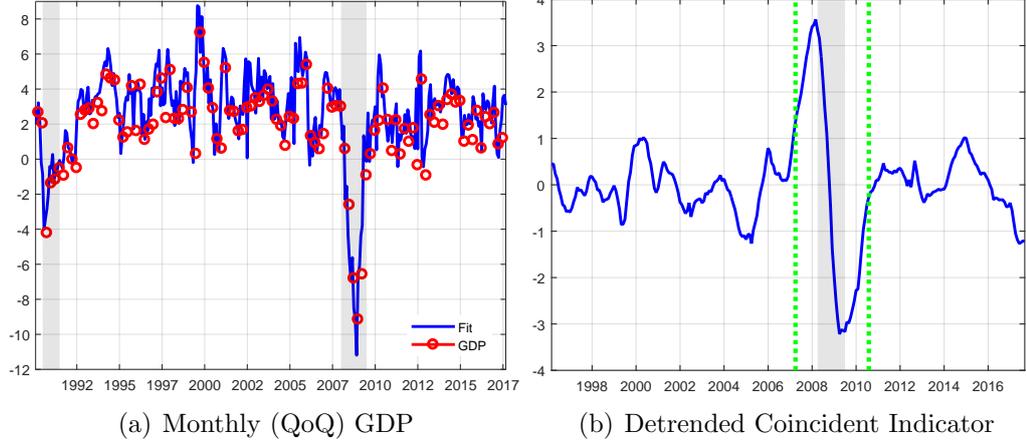
where  $\hat{\alpha}_\tau$  is retrieved by using the Kalman smoother. The business cycle indicator is then computed by applying a simple filter to  $z_t$ . For the baseline results in the paper we use the Rotemberg (1999) version of the HP filter, which chooses the smoothing coefficient of the HP filter so as to minimize the correlation between the cycle and the first difference of the trend estimate.

Table 3.7: Coincident Indicator - Model Fit

	$R^2(\%)$
GDP	87.9
Index of Manufacturing	39.6
Index of Services	82.4
Retail Sales	14.7
Employment	23.3
Unemployment	22.4

Notes: The table reports the fit of the coincident business cycle indicator on the quarter-on-quarter growth rate of the underlying variables.

Figure 3.11: Monthly GDP and Detrended Coincident Indicator



Note: The left panel shows the fit of the (monthly) coincident indicator on the (annualized) quarter-on-quarter growth of real GDP. The right panel reports the detrended GDP using the Rotemberg (1999) version of the HP filter, which sets the smoothing coefficient so as to minimize the correlation between the cycle and the first difference of the trend estimate. The vertical green lines denote the end and the beginning of the subsamples used to exclude the Great Recession from the analysis.

### 3.8.4 Alternative scenarios in the occurrence of a VAT change

Recall that inflation in the occurrence of a VAT change is computed as

$$\pi_t^{VAT} = - \int x \Lambda_t^{VAT}(x) f_t^{VAT}(x) dx,$$

implying that the observed inflation results from both changes in the distribution of price gaps, as well as from shifts in the hazard function. Based on this benchmark, one can envisage two relevant scenarios:

- *Scenario 1*: What rate of inflation would have been observed, had the VAT change only been associated with a change in the price gap distribution, while keeping the incentives of changing prices fixed? To address this question, we compute the following counterfactual rate of inflation

$$\pi_t^{VAT,1} = - \int x \Lambda_t^{No-VAT}(x) f_t^{VAT}(x) dx$$

- *Scenario 2*: What inflation would have been observed in absence of changes in the price gap distribution and the hazard function? This can be retrieved as

$$\pi_t^{VAT,2} = - \int x \Lambda_t^{No-VAT}(x) f_t^{No-VAT}(x) dx$$

The No-VAT counterfactual is computed by averaging the same function, for the

same month of the year in the 6 years before the VAT change.

Comparing  $\pi_t^{VAT,2}$  with the actual rate of inflation highlights the overall effects of the VAT, whereas the comparison between  $\pi_t^{VAT,1}$  and observed inflation quantifies the relevance of the state dependence in price setting (i.e., the fact that incentives to change prices are themselves a function of the underlying environment).

### 3.8.5 Estimation of the STARMA (p,q) model

Recall the smooth transition ARMA model, STARMA(p,q), in Section 3.6:

$$\begin{aligned} \pi_t = & G\left(\tilde{\mathcal{F}}_{t-1}; \gamma\right) \left( \phi_0^H + \sum_{j=1}^p \phi_j^H \pi_{t-j} + \varepsilon_t^H + \sum_{i=1}^q \theta_i^H \varepsilon_{t-i}^H \right) \\ & + \left[ 1 - G\left(\tilde{\mathcal{F}}_{t-1}; \gamma\right) \right] \left( \phi_0^L + \sum_{j=1}^p \phi_j^L \pi_{t-j} + \varepsilon_t^L + \sum_{i=1}^q \theta_i^L \varepsilon_{t-i}^L \right). \end{aligned} \quad (3.8.1)$$

This can be easily cast in state space. Therefore the likelihood can be calculated recursively using the Kalman filter (see Harvey, 1990). Since the model is highly non-linear in the parameters, it is possible to have several local optima and one must try different starting values of the parameters. Furthermore, given the non-linearity of the problem, it may be difficult to construct confidence intervals for parameter estimates, as well as impulse responses. To address these issues, we use a Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003; henceforth CH). This method delivers not only a global optimum but also distributions of parameter estimates.

Denote with  $\theta$  the vector of parameters. We employ the Hastings-Metropolis algorithm to implement CH's estimation method. Specifically, our procedure to construct chains of length  $N$  can be summarized as follows:

- *Step 1:* Draw  $\vartheta^{(n+1)}$ , a candidate vector of parameter values for the chain's  $n + 1$  state, as  $\vartheta^{(n+1)} = \theta^{(n)} + \mathbf{u}_n$  where  $\mathbf{u}_n$  is a vector of *iid* shocks taken from a student-t distribution with zero mean,  $\nu = 5$  degrees of freedom and variance  $\Omega$ .
- *Step 2:* Take the  $n + 1$  state of the chain as

$$\theta^{(n+1)} = \begin{cases} \vartheta^{(n+1)} & \text{with probability } \min \left\{ 1, \frac{L(\vartheta^{(n+1)})}{L(\theta^{(n)})} \right\} \\ \theta^{(n)} & \text{otherwise} \end{cases}$$

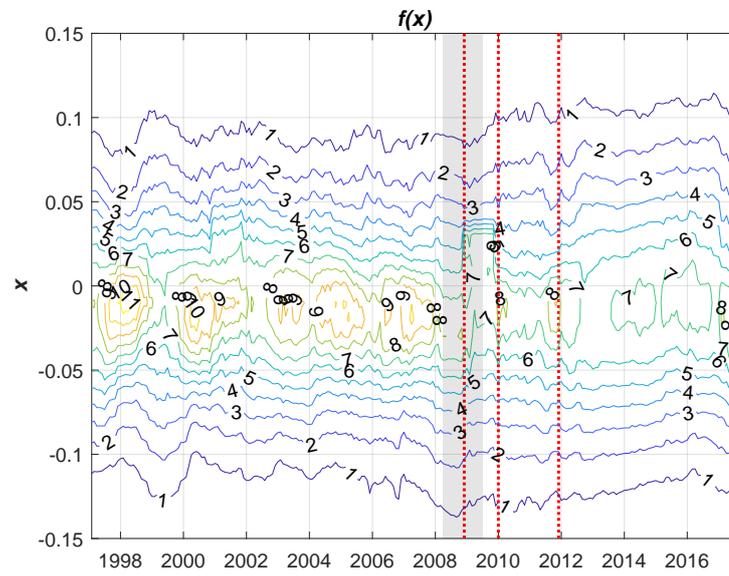
where  $L(\theta)$  denotes the value of the likelihood of the model evaluated at the parameters values  $\theta$ .

Specifically, we use an adaptive step for the value of  $\Omega$ , i.e. this is recalibrated using the accepted draws in the initial part of the chain and then adjusted on the fly to generate 25 – 35% acceptance rates of candidate draws, as proposed in Gelman et al. (2004). We use a total of 50,000 draws, and drop the first 25,000 draws (i.e., the ‘burn-in’ period). We then pick the 1-every-5 accepted draws to mitigate the possible autocorrelations in the draws. We run a series of diagnostics to check the properties of the resulting distributions from the generated chains. We find that the simulated chains converge to stationary distributions and that simulated parameter values are consistent with good identification of parameters.

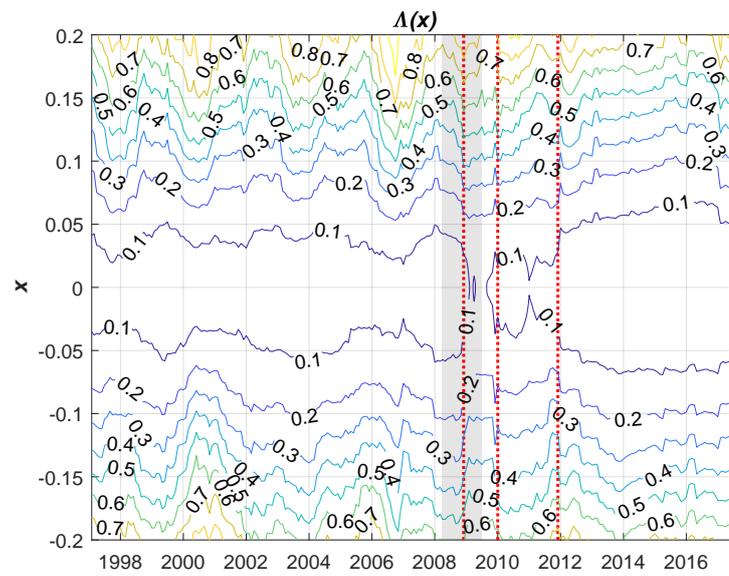
CH show that  $\bar{\theta} = \frac{1}{N} \sum_{i=1}^N \theta^{(i)}$  is a consistent estimate of  $\theta$  under standard regularity assumptions of maximum likelihood estimators. CH also prove that the covariance matrix of the estimate of  $\theta$  is given by the variance of the estimates in the generated chain. Furthermore, we can use the generated chain of parameter values  $\theta^{(i)}$  to construct confidence intervals for the impulse responses.

### 3.8.6 Model estimates

Figure 3.12: Estimated Price Gap Distributions and Hazard Functions



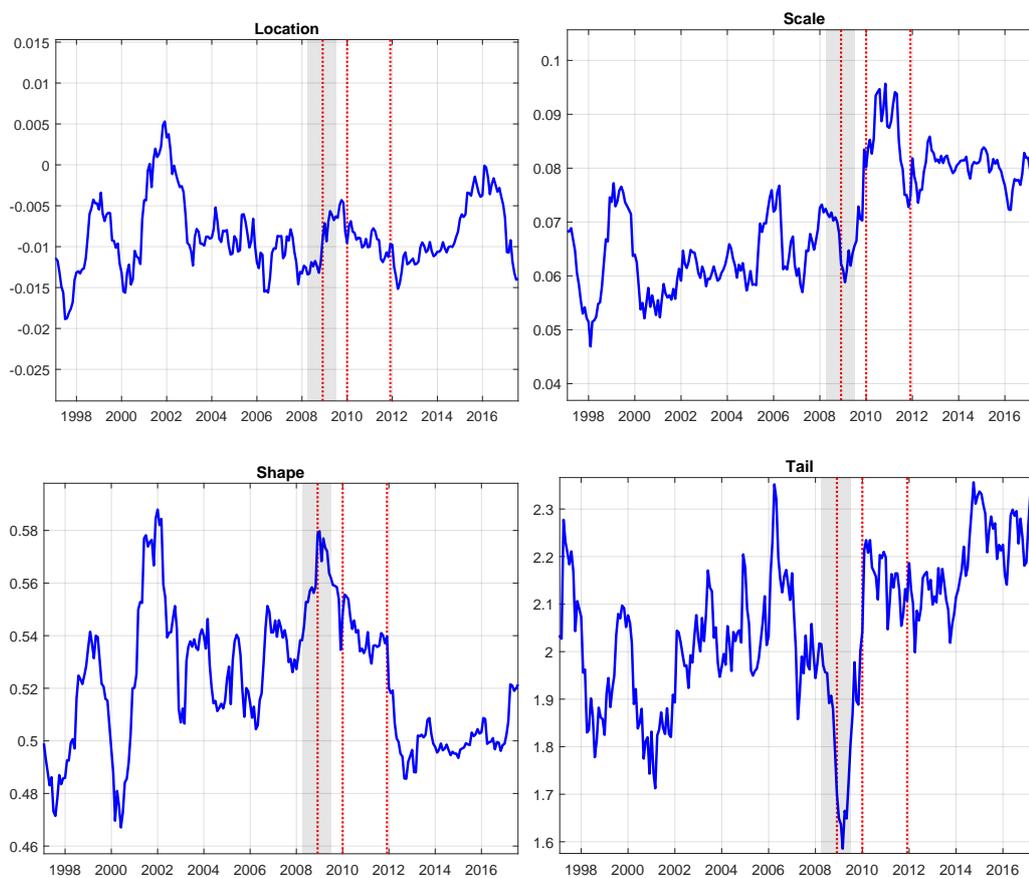
(a) Price Gap Distributions



(b) Hazard Functions

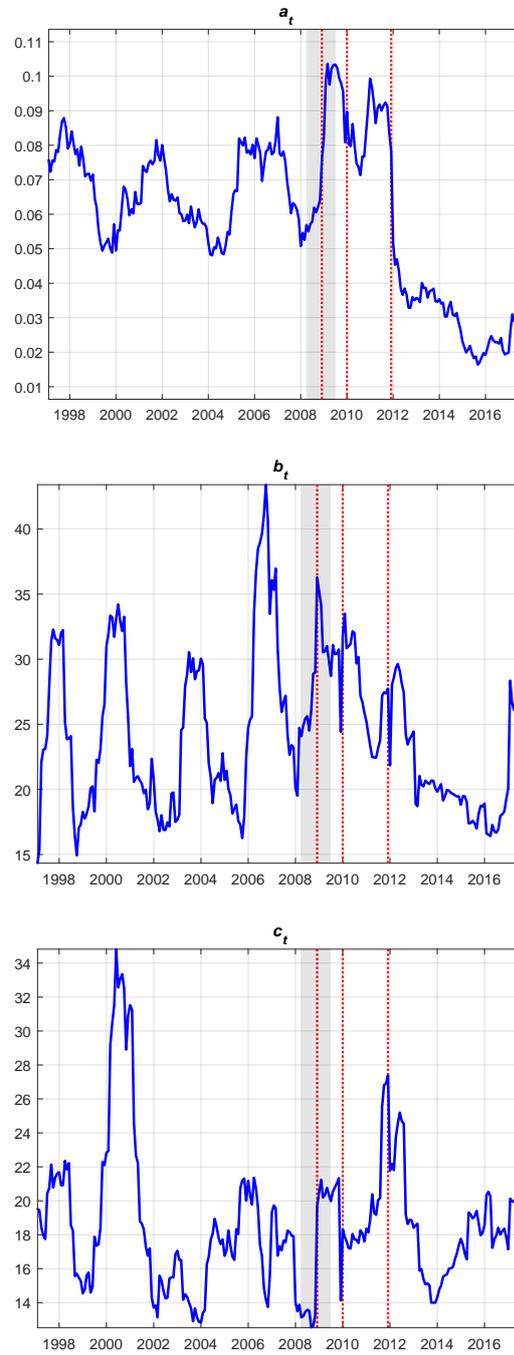
Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

Figure 3.13: Parameters of the Price Gap Distribution



Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

Figure 3.14: Parameters of the Hazard Function



Note: The red lines denote the three VAT changes in the sample. The shaded vertical band indicates the duration of the Great Recession.

### 3.8.7 Additional figures and tables

Table 3.8: Correlations of Pricing Moments with Macroeconomic Variables (Quadratic Trends)

Full Sample						
	$fr_t$	$\sigma_t^2$	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
$y_t$	-0.486***	0.452***	0.143**	0.330***	-0.318***	-0.015
$\pi_t$	0.497***	-0.182***	0.055	-0.265***	-0.065	-0.381***
$fr_t$	-	-0.098	-0.186***	-0.575***	0.267***	0.004
Pre-Recession						
	$fr_t$	$\sigma_t^2$	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
$y_t$	0.381***	0.576***	-0.492***	-0.368***	-0.141	0.406***
$\pi_t$	0.393***	0.206**	-0.420***	-0.539***	0.111	0.169*
$fr_t$	-	0.402***	0.067	-0.484***	0.122	-0.168*
Post-Recession						
	$fr_t$	$\sigma_t^2$	$q_{75,t} - q_{25,t}$	$q_{90,t} - q_{10,t}$	$Skew_t$	$Kurt_t$
$y_t$	-0.733***	0.652***	0.197*	0.578***	-0.172	0.634***
$\pi_t$	0.918***	-0.449***	-0.141	-0.372***	-0.220**	-0.704***
$fr_t$	-	-0.587***	-0.272**	-0.511***	-0.074	-0.619***

Notes:  $fr_t$  denotes the frequency of adjustment;  $\sigma_t^2$  stands for the volatility of the distribution of price changes;  $q_{n,t}$  measures the  $n$ -th quantile of the distribution of price changes;  $Skew_t$  denotes the skewness of the distribution of price changes and is measured as  $\frac{q_{90,t} + q_{10,t} - 2q_{50,t}}{q_{90,t} - q_{10,t}}$ ;  $Kurt_t$  denotes the kurtosis of the distribution of price changes and is measured as  $\frac{q_{90,t} - q_{62.5,t} + q_{37.5,t} - q_{10,t}}{q_{75,t} - q_{25,t}}$ ;  $y_t$  is a business cycle indicator;  $\pi_t$  indicates aggregate CPI inflation. Aside of the inflation rate, all series are detrended with a linear and a quadratic trend. \*\*\*/\*\*/\* indicates statistical significance at the 1/5/10% level, respectively.

Table 3.9: Correlations of Pricing Moments with Macroeconomic Variables: the Role of Asymmetry (Quadratic Trends)

Full Sample								
	$fr_t^+$	$fr_t^-$	$dp_t^+$	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
$y_t$	-0.346***	-0.541***	0.120*	0.572***	-0.078	0.233***	0.058	0.333***
$\pi_t$	0.717***	0.149**	0.135**	-0.105	0.148**	-0.034	-0.411***	-0.127**
Pre-Recession								
	$fr_t^+$	$fr_t^-$	$dp_t^+$	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
$y_t$	0.480***	0.049	-0.141	0.705***	-0.547***	-0.065	-0.399***	-0.120
$\pi_t$	0.733***	0.447***	0.110	-0.039	-0.121	0.242***	-0.295***	-0.179**
Post-Recession								
	$fr_t^+$	$fr_t^-$	$dp_t^+$	$-dp_t^-$	$q_{75,t} - q_{50,t}$	$q_{50,t} - q_{25,t}$	$q_{90,t} - q_{50,t}$	$q_{50,t} - q_{10,t}$
$y_t$	-0.721***	-0.684***	-0.098	0.376***	-0.527***	0.412***	0.525***	0.449***
$\pi_t$	0.891***	0.767***	-0.062	-0.083	0.352***	-0.392***	-0.703***	-0.269**

Notes:  $fr_t^+/fr_t^-$  stands for the frequency of positive/negative price changes;  $dp_t^+/dp_t^-$  indicates the average size of positive/negative price changes;  $q_{n,t}$  measures the  $n$ -th quantile of the distribution of price changes;  $y_t$  is a business cycle indicator;  $\pi_t$  indicates aggregate CPI inflation. Aside of the inflation rate, all series are detrended with a linear and a quadratic trend. \*\*\*/\*\*/\* indicates statistical significance at the 1/5/10% level, respectively.

Table 3.10: Flexibility in Price Adjustment: Correlation with Real Activity and Inflation (Quadratic Trends)

Full Sample							
	$\mathcal{F}_t$	$Int_t$	$Ext_t$	$Int_t^+$	$Int_t^-$	$Ext_t^+$	$Ext_t^-$
$y_t$	-0.492***	-0.502***	-0.398***	-0.615***	-0.396***	-0.462***	-0.297***
$\pi_t$	0.584***	0.620***	0.443***	0.293***	0.725***	0.124*	0.578***
Pre-Recession							
	$\mathcal{F}_t$	$Int_t$	$Ext_t$	$Int_t^+$	$Int_t^-$	$Ext_t^+$	$Ext_t^-$
$y_t$	0.093	0.202**	-0.035	-0.041	0.296***	-0.166*	0.043
$\pi_t$	0.495***	0.705***	0.181**	0.393***	0.772***	-0.033	0.272***
Post-Recession							
	$\mathcal{F}_t$	$Int_t$	$Ext_t$	$Int_t^+$	$Int_t^-$	$Ext_t^+$	$Ext_t^-$
$y_t$	-0.800***	-0.787***	-0.713***	-0.648***	-0.799***	-0.453***	-0.791***
$\pi_t$	0.769***	0.788***	0.645***	0.562***	0.832***	0.304***	0.775***

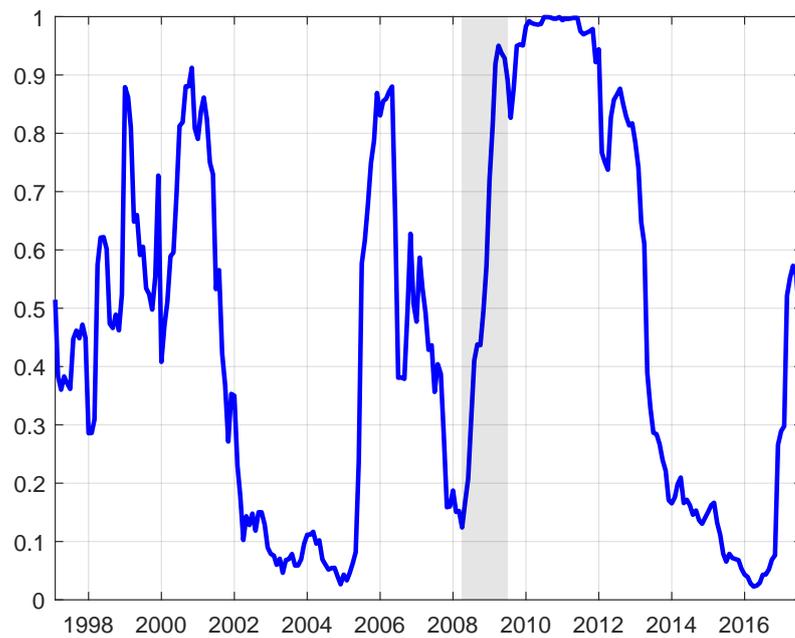
Notes: The table reports pairwise correlations of output and inflation with the flexibility index, as well as the intensive margin and the extensive margin of price adjustment (together with their counterparts corresponding to positive and negative price gaps). Aside of the inflation rate, all series are detrended with a linear and a quadratic trend. \*\*\*/\*\*/\* indicates statistical significance at the 1/5/10% level, respectively.

Table 3.11: Forecast Errors and Price Flexibility: Robustness (MSE)

(a) BoE MPC RPIX/CPI (Squared) Forecast Errors						
Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F-stat	$\tilde{R}^2$
1	0.082	[0.668]	0.544	[0.222]	0.559	-1.21
2	-0.304	[0.512]	3.073	[0.009]	0.150	2.98
3	-0.553	[0.333]	8.413	[0.011]	0.005	12.45
4	-0.440	[0.617]	15.556	[0.014]	0.000	25.80
5	0.018	[0.989]	17.463	[0.023]	0.000	22.44
6	0.818	[0.540]	14.810	[0.054]	0.001	16.12
7	1.564	[0.212]	11.514	[0.091]	0.009	11.29
8	2.145	[0.135]	6.578	[0.285]	0.117	4.02
(b) Market Participants' (Squared) Forecast Errors						
Horizon	Slope at $G = 0.3$		Slope at $G = 0.9$		F-stat	$\tilde{R}^2$
1	0.713	[0.291]	0.426	[0.497]	0.363	0.25
2	-0.396	[0.464]	3.491	[0.007]	0.123	3.65
3	-0.763	[0.287]	9.235	[0.008]	0.007	11.63
4	-0.608	[0.517]	16.589	[0.010]	0.000	24.46
5	-0.063	[0.960]	18.043	[0.016]	0.000	20.81
6	0.923	[0.465]	14.287	[0.045]	0.005	13.17
7	1.789	[0.129]	9.562	[0.099]	0.043	7.16
8	2.315	[0.091]	3.916	[0.431]	0.390	0.02

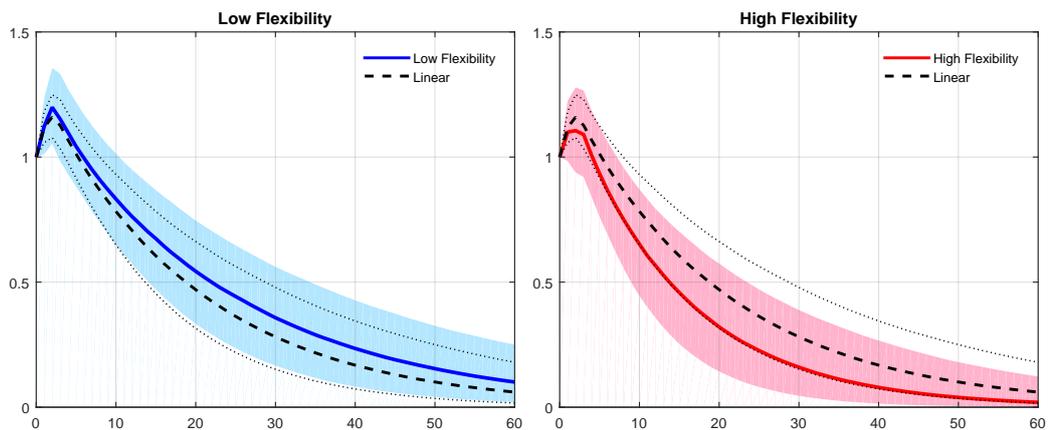
Notes: The table reports the results of a quadratic spline regression of the squared forecast errors  $e_{T+h|T}$  (for different forecast horizons,  $h$ ) on a quarterly average of an indicator of the normalized price flexibility index,  $G_t = G(\tilde{\mathcal{F}}_t; \gamma) = (1 + e^{-\gamma\tilde{\mathcal{F}}_t})^{-1}$ , where  $\tilde{\mathcal{F}}$  denotes the normalized flexibility index. The regression takes the form:  $e_{T+h|T}^2 = a_0 + a_1G_t + a_2G_t^2 + a_3G_t^2I_{\{G_t > 0.5\}}$ . The upper panel refers to the Bank of England MPC's RPIX/CPI forecast errors, while the bottom panel considers market participants' forecast errors. In each panel, the first two pairs of columns report the slope of the relationship evaluated at different levels of the indicator, together the p-value associated with the null hypothesis that the slope is equal to 0 (this is calculated using Newey-West standard errors). The penultimate column (F-stat) reports the p-value of the null hypothesis that all the coefficients associated to the flexibility regime are equal to 0. The last column reports the adjusted R-squared, denoted by  $\tilde{R}^2$ .

Figure 3.15: Probability of a High-flexibility Regime



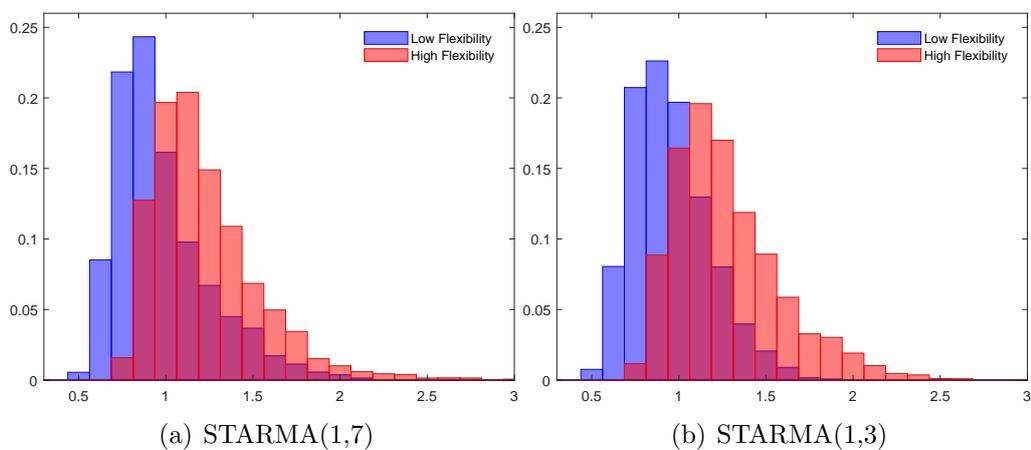
Note: The figure reports the probability of ending up in a high-flexibility regime, obtained in accordance with the STARMA(1,7) model presented in Section 3.6. The shaded vertical band indicates the duration of the Great Recession.

Figure 3.16: Price Flexibility and Inflation Persistence



Note: Figure 3.16 reports the responses of inflation to a 1% shock in the STARMA(1,3) model. The left (right) panel graphs the response in the low (high) price flexibility regime. In both cases we also report the the response from a (linear) ARMA(1,3) model. 68% confidence intervals are built based on the Markov Chain Monte Carlo (MCMC) method developed in Chernozhukov and Hong (2003).

Figure 3.17: Price Flexibility and Inflation Volatility



Notes: Each panel reports the distribution of the estimated inflation volatility in the two regimes. The left panel refers to the STARMA(1,7), while the right panel refers to the STARMA(1,3).

## Chapter 4

# Adding up the UK's consumer price inflation data: Quality change, missing observations and the importance of sales

### Abstract

In this paper I investigate the relative importance of missing observations, product replacements and sales in the measurement of the UK's consumer price inflation rate. Using an inflation accounting framework based on a weighted geometric aggregation formula for the price index, I show how to decompose the inflation rate into four components of regular matched inflation, missing observations, product replacements, and sales as well recovery prices. Using published micro-price data I apply this framework to three different item categories of “women's top, long sleeved not blouse”, “semi-skimmed milk”, and “men's haircut” to show the relative importance for different service and item categories. I find a significant role for non-regular price changes in the clothing item of women's top due to strong seasonal fashion cycles and sales, which is in contrast to the two other items for which the inflation rate is found to be mostly driven by direct price changes with negligible role for the other components. Lastly to check the calculations I calculate a price index for each of the items, assuming no quality changes, which gives evidence of a historical mismeasurement of the clothing category.

### 4.1 Introduction

In this paper I investigate the importance of missing observations, product replacements and sales/recoveries in the measurement of the UK's Consumer Price Index (CPI) inflation rate as published by the Office for National Statistics (ONS). To achieve this I set out a novel decomposition of the inflation rate using a geometric

aggregation formula that allows me to linearly decompose the inflation rate into its different components.

I use published micro price quotes to apply framework to the historical decomposition of the UK's inflation rate. However, as the price trajectories are not uniquely identified, I will only be able to investigate the quality adjustments for product replacements that have been deemed directly comparable by the ONS. Hence, part of the missing observations component will be the excluded component of all product replacements with different quality characteristics, as judged by the ONS' price collectors.

As an empirical example I apply the framework to the two good categories of "women's top, long sleeved not blouse", and "semi-skimmed milk", as well as the service item of "men's haircut". The reason I choose the clothing item of women's top is to highlight the impact of the methodological changes made in January 2010 (and onwards) for how the category of clothing (COICOP division 3) is sampled and what is deemed comparable replacements. Prior to the change, when a clothing item reached the end of its product cycle, and a new replacement had to be found, the new product was deemed to be not comparable to the previous item, and hence the base prices were changed to adjust for the difference in quality. However, given that the main exit of products occurs in February, which is also the price recovery month from the January's sales, this method had the unfortunate side-effect of introducing a permanent deflationary effect from these sales into the measurement of clothing inflation. Hence, to correct this issue from January 2010 and onwards, the ONS allowed replacement of clothing items with similar characteristics to be deemed comparable, and hence eliminated the downward drift in the clothing index.<sup>1</sup> However, a systematic investigation as to whether this change in methodology of allowing new products to be comparable to the old ones, has not taken place to the knowledge of this author. To highlight some of these issues we therefore explore an item within this category as our empirical example. One challenge for our framework is that clothing is a category that displays strong seasonal patterns and for which it is known that consumers does indeed prefer newer products to older ones. Hence a clear interpretation of our results is complicated. Nevertheless, in our empirical investigation we find that the main driver of the inflation rate within the category, is as expected, from the product replacements.

A related work to this chapter is the presentation to the UK's Advisory Panel on Consumer Prices (APCP) by de Vincent-Humphreys (2017) investigating the impact of the formula effect within the item category "Women's vests/strappy tops". This is

---

<sup>1</sup>As an unintended consequence of the methodology change, the dispersion of the ratio between prices and their base prices increased, resulting in a measurement error from the RPI that uses the Carli formula to calculate the lowest level of aggregation, discovered in reports such as Diewert (2012).

a difficult presentation to read as the author does not always clarify his calculations and graphs.<sup>2</sup> He outlines an framework for simulating a counterfactual inflation rate based on randomly setting sampled replacement price relatives to being non-comparable (by creating artificial base prices for them). The problem with this methodology is however, that although it highlights some critical aspects of the difference between the Retail Price Index and Consumer Price Index<sup>3</sup> in terms of outlining a mechanism driving the dispersion of the price relatives (and thereby the formula effect), the presentation does not deal with the actual issue of how to account for fashion cycles in the measurement of the inflation rate for clothing items and how measure its impact internally on the current Consumer Price Index. In contrast this chapter focuses on quantifying the impact of the individual components for the measured inflation rate, which allows for a quantification of the the sales and recovery components of fashion cycles. Another related literature is found in the report by Diewert (2012), who highlighted the issue around measuring fashion items. Diewert suggests an alternative approach to deal with this issue of fashion items, in the form of a stochastic elementary aggregation formula, that explicitly takes into account the strong seasonal patterns for categories such as clothing.<sup>4</sup> As a side note, an (inversely) related question to price measurement is the question of productive and real growth rate, which in recent years have received increasing attention with the developed world seeing a productivity slowdown as documented in the U.S. by Gordon (2016a,b), and for the UK by Bean (2016).

International studies into measurement errors in consumer price inflation include the Boskin Commission in the US (Boskin et al., 1996, 1998), that highlighted important measurement biases such as outlet substitution for consumers changing shopping preferences to cheaper outlets, which results in an upward bias to the measured inflation rate.<sup>5</sup> Other potential biases that were also highlighted by the Boskin commission includes the delayed introduction of new products such as mobile phones, as studied by Hausman (1999, 2003), and insufficient quality adjustments for existing technology as improvement is made, such as for example in the case of lightning technology as studied by Nordhaus (1996, 1998), leading to overestimation of the true inflation rate.

In the following sections I will proceed by first describing the basics of the index

---

<sup>2</sup>As well as missing the point that not all RPI elementary aggregates are calculated using the arithmetic Jevons formula.

<sup>3</sup>A topic already extensively discussed in other places such as Diewert (2012) and Johnson (2015).

<sup>4</sup>Other notable reports for the UK are the Johnson (2015) investigation the Consumer Price Index, with an emphasis on the lack of good owner-occupied housing in measurements, and the Bean (2016) Review that focused on the quality of the national accounts measurements, which implicitly is the mirror image of any price measurements.

<sup>5</sup>See also Gordon (2000, 2006) for a look back at the results achieved by the commission.

construction and the aggregation methodology in section 4.2. I will then outline the accounting framework for the inflation rate and the use it to decompose the measurement into a component of sales, missing observations, products replacement and pure price changes in section 4.3. As a check on my results in section 4.4 I outline a pure price index based only on the sampled prices and assuming no quality changes. Finally, I will conclude in section 4.5. In appendix 4.6.1 I show the validity of my framework, by comparing how well it matches up the officially published item indices from the ONS, and in appendix 4.7, I show how to extend the framework to the higher levels of aggregation.

## 4.2 The Basic Building Blocks: Index Calculations

### 4.2.1 Prices, Base Prices, and Identification of Price Trajectories

Any investigation of potential measurement errors in the UK's consumer price inflation rate requires me to first outline how to construct a price index. The key building blocks for calculating any price index is a sample of prices. To acquire this, on either the second or third Tuesday (or Wednesday) of every month and the days around it, the UK's Office for National Statistics (ONS) collects approximately 108,000 price quotes for a representative basket of consumption products, consisting of around 1,100 items, from shops all around the United Kingdom.<sup>6</sup> This monthly dataset of price quotes, dating back to February 1996, has been made publicly available since September 2011 (ONS, 2011). A further unspecified amount of price quotes for 180 items, or about 40% of the consumer price index by weight, are sampled centrally by the ONS, but are unfortunately not being made publicly available due to confidentiality agreements. The collected prices are in discrete values of pounds sterling with increments of £0.01 – one penny, the lowest denominator of the pound – which is also the lowest possible value of any price quote.<sup>7</sup> From the perspective of the ONS, the price quotes are uniquely identified, by the collection date, the location, the outlet, and the product that was sampled. The location is defined as a large shopping district within one of the UK's regions, e.g. Oxford Street within the region of London. As for the outlet, it is the shop from which the price was sampled,

---

<sup>6</sup>The collection day is referred to as '*Index Day*' by the ONS, with the exact date being kept a secret and varied to avoid any price distortions due to bank holiday weekends, as well as possible price manipulation by retail providers.

<sup>7</sup>I exclude the few exceptions of a zero price, recorded in instances of transition between public and private provision of a product, where the market price for the public provision is zero from the perspective of the consumer.

such as the UK retailer John Lewis. Lastly, the product is the item for which a price quote was collected, with an example being “Women’s top, long sleeved, not blouse” (id 510223).<sup>8</sup> I will refer to an individual price quote as  $P_{m,y,i}$ , where  $i$  is the price trajectory (theoretically defined by the location, outlet, and product), and  $m, y$  are respectively the month and year the price was sampled. To validate the data, the ONS have various procedures in place for both the local collectors of the prices quotes and at the ONS’ central office to exclude any implausible prices and recording errors. In this paper I will only be working with the validated price quotes.<sup>9</sup>

To standardise the prices for aggregation, and to adjust for any possible quality or quantity changes to a product sampled for a given price trajectory, the ONS use what they refer to as base prices. A base price is either the previous January’s price of the trajectory, if observed, or an estimate of it in the case of product replacements (adjusted for any differences in quality or quantity). I will refer to the base price as  $P_{m,y,i}^0$  which, if the previous January’s price for trajectory  $i$  was observed, is defined as:

$$P_{m,y,i}^0 = \begin{cases} P_{1,y-1,i} & m = 1 \\ P_{1,y,i} & \forall m > 1 \end{cases}. \quad (4.2.1)$$

If the January price is not observed, such as is the case when a product replacement takes place, the ONS provides their estimate of it in the dataset. Given that the base prices in some cases consist of estimated values and not only observed January prices, they are defined to be a positive non-zero real number.<sup>10</sup>

Unfortunately, due to confidentiality agreements with the retailers, in the publicly available dataset, the ONS are unable to publish the location from where a price quote was sampled and instead only provides the region. Furthermore, the outlets are only identified by an anonymised integer shop code which is only unique given the unpublished location of the outlet, and is not necessarily identical for the same outlet across different locations. As a result, some of the shop codes for different outlets based in different locations within the same region are identical. I am therefore unable to directly identify a unique price trajectory for all the price quotes. To reconstruct the price trajectories, I use two additional pieces of information from the dataset. Firstly, the ONS provides a shop type indicator for whether the shop, from which a price quote was sampled, was a chain with 10 or more stores

---

<sup>8</sup>The item id uniquely nests the higher aggregation categories of COICOP classes, groups, and divisions to which the sampled product belongs to.

<sup>9</sup>See ONS (2014, p.34-38) for further details on the validation procedures.

<sup>10</sup>It is worth noting however, that the majority of the base prices are discrete, as they are essentially observed January price quotes. Only a small fraction of the price quotes have estimated base prices.

nationally (shop type 1) or a smaller independent firm (shop type 2).<sup>11</sup> Within the same region I can use this shop type to distinguish some of the non-unique shop codes from each other. Secondly, I condition the identification of the trajectories on the base prices for the price quotes, which primarily change in February (with the exception of non-comparable product replacements where the base price is changed to adjust for quality or quantity differences), and which therefore have a direct link to the previous January price quote. Using these two additional characteristics together with the region, shop code, and product information, for each month from February to January the following year I link up the price quotes recursively year by year. I then link these yearly trajectories together across the years, by connecting the base prices to the previous period’s January prices. Applying this method, I can infer a unique trajectory for the majority of all the price quotes.<sup>12</sup> The main drawback is that when the ONS makes changes to a base price due to quality (or quantity) adjustment to a sampled product, this results in a break of the estimated trajectory. The true trajectory will be split into two as the base prices will differ. Hence, my method gives an upper bound on the estimated amount of trajectories, with the actual number observed by the ONS being smaller.<sup>13</sup>

The dataset runs from February 1996 till November 2017, for a total of 262 months with 28.4 million observations divided over 1.2 million estimated trajectories. I will refer to the total amount of estimated trajectories as  $N_{\bar{m},\bar{y}}$ , where  $\bar{m}, \bar{y}$  refers to November 2017 (the last sampling date in our dataset), and the total set of possible price trajectories are given by  $\{1, 2, \dots, N_{\bar{m},\bar{y}}\}$  which is the full set of price trajectories, including products, locations, and outlets not part of the sampling basket at time  $m, y$ . Given that every February the ONS introduces new outlets, items, and locations to be sampled into the dataset, the number of total trajectories,  $N_{\bar{m},\bar{y}}$ , will grow with time.

Table 4.1: Price trajectories

	Trajectory 1		Trajectory 2	
	$P_{t,1}$	$P_{t,1}^0$	$P_{t,2}$	$P_{t,2}^0$
Feb. 2013	£25	£28	£26	£22
Mar. 2013	£22	£28	£26	£22

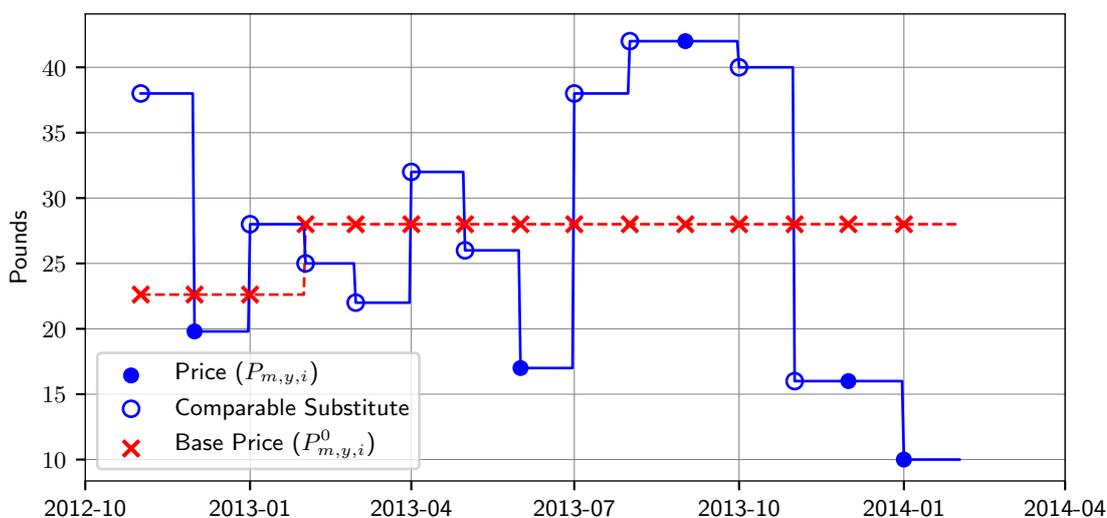
Notes: Example of price trajectories for “Women’s top, long sleeved, not blouse” (id 510223), collected from a chain-store (shop type 1), with shop code 3 within the region of London (region 2).

<sup>11</sup>In the dataset we also observe a few recording errors of shop type 0.

<sup>12</sup>The exception is a small subset of prices of around 0.36% by weight each month, which I am unable to uniquely identify, and hence I exclude them from the analysis.

<sup>13</sup>In future work, we hope to link these breaks together so we are able to investigate the quality and quantity adjustments made by the ONS in further detail as well.

Figure 4.1: A Price Trajectory



Notes: Example of a price trajectory for “Women’s top, long sleeved, not blouse” (id 510223), collected from a chain-store (shop type 1), with shop code 3 within the region of London (region 2). The price is the sampled price quote, comparable substitute indicates that a new product is being sampled for the trajectory, and the base prices are the previous January’s price (or an estimate of it) as defined in equation 4.2.1. In the above we notice that the large changes to the price trajectory are often associated with a product substitution.

As an example from the dataset, if I look at the “Women’s top, long sleeved, not blouse” (id 510223), sampled from the region of London, at a chain store (shop type 1) with shop code 3, I observe two different prices of £25 and £26 in February 2013. One could speculate that I am observing one price from the retailer John Lewis at Oxford Street as the first (unobserved) outlet and location, and the other could be the retailer Debenhams in the London shopping centre Westfields.<sup>14</sup> However, as I observe neither a unique identifier for the outlet (I only have the shop code of 3 and shop type 1 for both of the price quotes), nor the location, I am unable to differentiate the two price trajectories from each other. Once I take into account the individual base prices of the quotes, as shown in table 4.1, I am able to make a distinction between the two price trajectories, even though I still do not know the exact outlet or location from where they were sampled. As a result, when in March 2013 I observe a new price of £22 with base price £28, and another price of £26 with a base price of £22, I can comfortably say that the price drop to £22 is from price trajectory 1 of table 4.1, with the previous sampled price of £25 in February 2013. Figure 4.1 shows the full history of price quotes and base prices of trajectory 1 from table 4.1.

<sup>14</sup>It could also be the same outlet, e.g. John Lewis, in the two different locations, although this need not be the case, as the shop codes are only ordered within the location.

## 4.2.2 Elementary Aggregates and Missing Observations

The first step in the index calculations by the ONS, is to calculate a stratified index. To aggregate the price quotes, I follow the Jevon’s method used by the ONS in taking the ratio of the prices to the base prices (I will refer to this as an element), and then aggregating them into the stratified index, referred to as an elementary aggregate index, using a weighted geometric formula.<sup>15</sup> The stratified index is classified by the item id and a stratum cell.<sup>16</sup> The stratum cell depends on the type of stratification used for the item id to reflect different pricing patterns across different regions and shop types (independent or chains). As a result the stratum cell is either given by the region, the shop type, or both, as well as a few being not stratified at all (for these few items the elementary aggregate index is equivalent to the item index that I will define shortly). I define the full set of item id and stratum cell as being given by  $\{1, \dots, H_{\bar{m}, \bar{y}}\}$ , where  $H_{\bar{m}, \bar{y}}$  is the maximum amount of combinations as seen from the last period in the dataset (November 2017). As the regions and shop types are fixed, and most items do not change stratification method through the sample,  $H_{\bar{m}, \bar{y}}$  only increases when new items are introduced into the sampling basket in February each year.

The elementary aggregate index is calculated as:

$$Z_{m,y,h}^{ea} = \prod_{i \in \mathcal{I}_{m,y,h}^{ea}} \left( \frac{P_{m,y,i}}{P_{m,y,i}^0} \right)^{w_{m,y,i}^{ea}}. \quad (4.2.2)$$

In the above,  $Z_{m,y,h}^{ea}$  is the elementary aggregate which is defined to be a non-negative real number by construction (and the definition of the prices, base prices and elementary aggregate weights),  $h$  refers to the stratum id defined by the item and stratum cell (e.g.  $h \in \{1, \dots, H_{\bar{m}, \bar{y}}\}$ ) and as before  $m, y$  are respectively the month and year.  $\mathcal{I}_{m,y,h}^{ea}$  is the information set of sampled price quotes within stratum  $h$ , and defines the price trajectories which was available for sampling and validated by the ONS for the period  $m, y$ . The elementary aggregate information set is a subset of the total set of price trajectories;  $\mathcal{I}_{m,y,h}^{ea} \subset \{1, \dots, N_{\bar{m}, \bar{y}}\}$ . There exists a unique mapping from the price trajectory ( $i$ ) to the stratum classification ( $h \leftarrow h(i)$ ), but

---

<sup>15</sup>In the price index literature there are three different formulae for calculating elementary aggregate indices: the geometric formula I am using called the Jevon’s formula; the ratio of averages called Dutot (which for a small set of elementary aggregates the ONS uses); and lastly an average ratio formula called the Carli formula (which due to its bias is no longer used, except for the calculation of the Retail Price Index). According to O’Neil et al. (2017, p. 247), in the CPI 63% of the elementary aggregate indices are calculated using the Jevons formula, 5% using the Dutot, and 32% using weights (unfortunately the authors does not specify the formula used with these weights), and with the Carli formula not used at all following European regulation.

<sup>16</sup>In the dataset, there are periods of missing stratum cells, especially for Northern Ireland which seems to be a data error. In my work, I have assumed that the unpublished cells have inflation dynamics identical to the published ones, and normalise the weights to take this into account.

as within each set of stratum cell and item id there can be multiple price trajectories, the opposite is not the case.

The elementary aggregate weights ( $w_{m,y,i}^{ea}$ ) are used to reflect that some of the prices sampled are from stores that have a uniform pricing policy across different locations, regions or indeed nationally. Hence instead of re-sampling the ‘same’ price from different locations, as a cost saving measure the ONS collects a single observation and uses a shop replication factor (shop weight) to reproduce that price quote and base price in the index. Further, missing price quotes are dealt with through variations in the elementary aggregate weights. This implies that the missing observations are estimated as a weighted average of the price quotes available for sampling, a topic I will return to in section 4.3.2. The elementary aggregate weight is given by:

$$w_{m,y,i}^{ea} = \begin{cases} \frac{f_{m,y,i}}{\sum_{j \in \mathcal{I}_{m,y,h(i)}^{ea}} f_{m,y,j}} & \forall i \in \mathcal{I}_{m,y,h(i)}^{ea} \\ 0 & \forall i \notin \mathcal{I}_{m,y,h(i)}^{ea} \end{cases}, \quad (4.2.3)$$

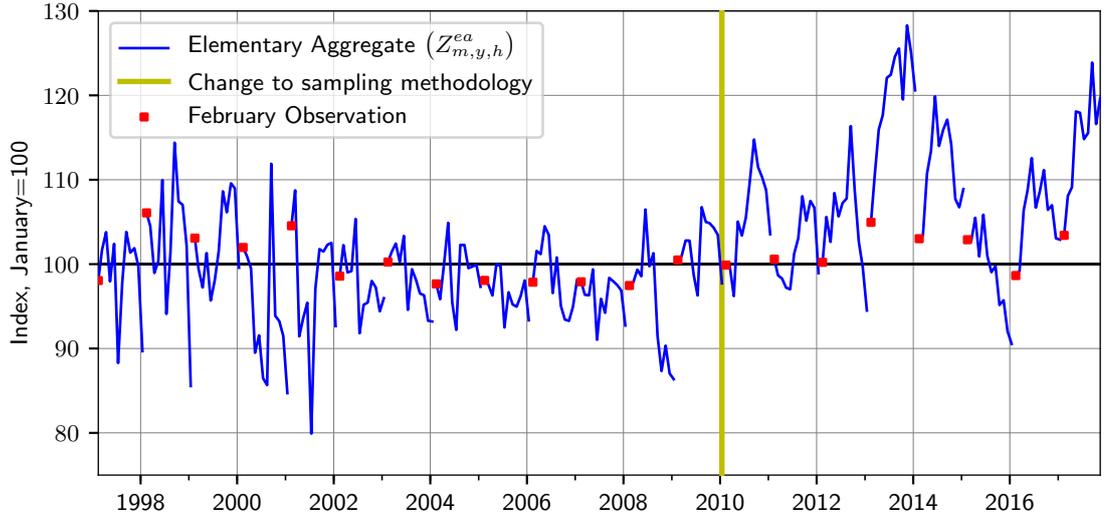
in the above  $w_{m,y,i}^{ea}$  is the elementary aggregate weight which is defined as  $w_{m,y,i}^{ea} \in [0, 1]$ , with the property that it sums to unity ( $\sum_{i \in \mathcal{I}_{m,y,h}^{ea}} w_{m,y,i}^{ea} = 1 \forall h$ ). Further,  $f_{m,y,i}$  is the shop replication factor defined as a positive integer (with the majority taking the value of 1 which would give equal weights if all are one). However, for some of the price trajectories from large chains with a uniform pricing policy, the ONS replicates a single sampled price quote, by setting the shop replication factor to an integer value equal to the amount of shops to be represented by the quote.<sup>17</sup>

As an example of an elementary aggregate index, I again look at “Women’s top, long sleeved not blouse” in figure 4.2. I observe that the elementary aggregates are always relative to the previous January’s price (the base price), and hence the series displays a pattern of constantly rescaling back towards one every February, as the base price changes to the previous January’s value. As can also be observed in figure 4.2 from January 2010 and onwards, the ONS changed their sampling methodology for the clothing category. As a result there is a structural change in the series. Previously, when a clothing item had been on sale and was then taken off the market, a new similar clothing item was deemed to be non-comparable and an adjustment was made to the base price. Hence the index captured the January sales, but never the recovery in February when new collections were introduced. With the new methodology the ONS has allowed for a broader acceptance of comparable products within the clothing category to eliminate this (artificial) deflationary pressure on

---

<sup>17</sup>Before January 2006, the replication factors were updated every January and used till the following December. After 2006 the methodology was changed such that the 2006 January update of the replication factors was used till January 2007. From then on they have been updated every February and used until the subsequent January (in the following year).

Figure 4.2: Elementary Aggregate Index



Notes: Example of Elementary Aggregate index for the item “Women’s top, long sleeved, not blouse” (id 510223), stratified by shop type, with the above example being from the stratum cell of multiple outlets (shop type 1). The index consists of an average of 199 price quotes per period, and the average shop replication factor for these price quotes is 1.23.

the price index.<sup>18</sup>

### 4.2.3 Item Indices and Chain-linking

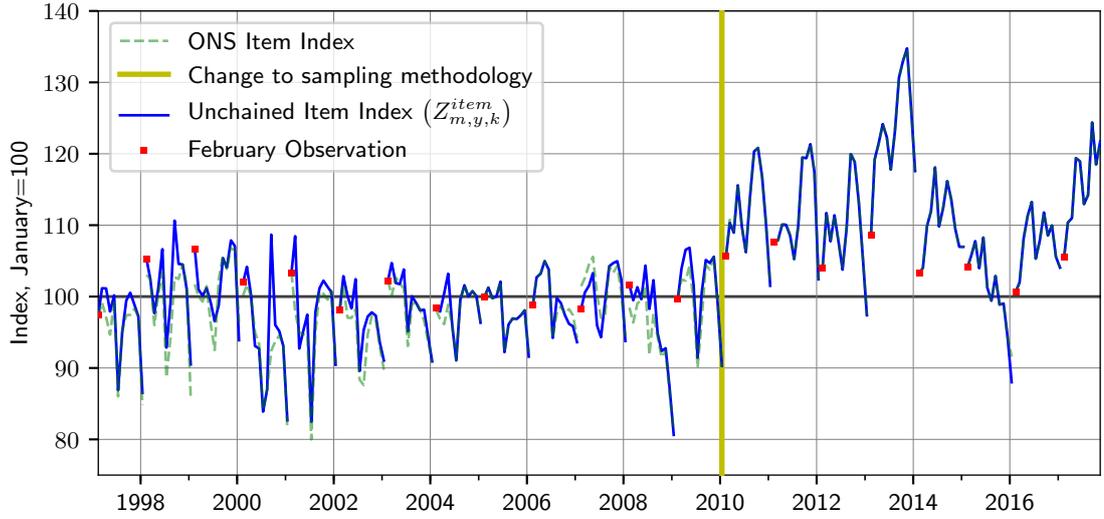
The next step in the aggregation is to calculate the item indices. To aggregate to the item indices, the ONS uses what they call stratum weights. The stratum weights reflect the relative importance of differences in purchasing patterns across the stratum, such as regions or shop types (in contrast to the shop weights, which represents market shares). The stratum weights reflects the ONS’ best available information about these differences in purchasing patterns, and are provided to me in the dataset. However, in contrast to the ONS, I will be calculating the index using a weighted geometric average formula. The ONS uses an arithmetic Lowe index in their current methodology to calculate the item indices and the COICOP aggregates as will be briefly discussed in section 4.7 (Clews et al., 2014, p.40). To achieve an analytical expression for the potential bias, I will proceed with the index construction using a weighted geometric mean formula.<sup>19</sup> As a result my item index is calculated as:

$$Z_{m,y,k}^{item} = \prod_{h \in \mathcal{I}_{m,y,k}^{item}} (Z_{m,y,h}^{ea})^{w_{m,y,h}^{strat}}, \quad (4.2.4)$$

<sup>18</sup>For more information see ONS (2010).

<sup>19</sup>It is well established in the literature that this will result in a lower inflation rate by the nature of a geometric average against an arithmetic one.

Figure 4.3: Item Index (unchained)



Notes: item index “Womens top, long sleeved, not blouse” (id 510223). Notice that the index is unchained, and hence resets back towards 100 every February. The index is the weighted geometric average of the chains given by shop type 1 and an average weight of 72.7% and the independent shops (shop type 1), with an average weight of 27.3%.

where  $Z_{m,y,k}^{item}$  is the (unchained) item index for item id  $k$ , defined as a non-negative real number. The item information set is given by  $\mathcal{I}_{m,y,k}^{item} \subset \{1, \dots, H_{\bar{m},\bar{y}}\}$ , remembering that  $H_{\bar{m},\bar{y}}$  defines the total set of stratum cell and item id combinations. The stratum weight,  $w_{m,y,h}^{strat}$ , is defined as  $w_{m,y,h}^{strat} \in [0, 1]$ , and  $\sum_{h \in \mathcal{I}_{m,y,k}} w_{m,y,h}^{strat} = 1 \forall k$ .<sup>20</sup> The stratum weights are updated every February and used until January the proceeding year.

As can be seen from figure 4.3 the index is still based on the previous January’s prices, which means the index resets every February. The ONS only publishes the item indices as an unchained index, and hence in the above figure 4.3 I also plot the officially published item index series as given by the ONS (the green stippled lines). As can be observed, my item index closely tracks the index as published by the ONS, with some divergence prior to 2010. I also observe that the February index after the methodology change, now includes estimates of the recovery prices from the January sales, and hence have a significant positive inflationary pressure, where previously there was no estimate of the recovery prices, and hence sales had a permanent effect on the estimated index.

I follow the ONS in using a double chain linking of the index (ONS, 2014), to link the index across the years, and to allow for the introduction of new price

<sup>20</sup>For a small sample of the indices, not all stratification are publicly available from the ONS. In these cases I renormalise the weights, and assume the dynamics of the unpublished stratum, is identical to the published stratum.

trajectories into the basket every February (as well as exits) while still keeping a consistent sample.<sup>21</sup> To calculate this, I first define an index based on the current years' January value:<sup>22</sup>

$$J_{m,y,k}^{item} = \begin{cases} 1 & m = 1 \\ Z_{m,y,k}^{item} & \forall m > 1 \end{cases}, \quad (4.2.5)$$

and recall, from equation 4.2.2, that the index  $Z_{m,y,k}^{item}$  is already calculated relative to the current years January's price for any month other than January. Using the two indices from equation 4.2.4 and equation 4.2.5, the chain-linked item index can be calculated as the following recursion:

$$CPI_{m,y,k}^{item} = CPI_{12,y-1,k}^{item} \times \left( \frac{Z_{1,y,k}^{item}}{Z_{12,y-1,k}^{item}} \right) \times J_{m,y,k}^{item}. \quad (4.2.6)$$

The starting condition of the recursion is defined as the arithmetic average of the base year  $y_0$ , which is currently 2015, set equal to the index value of 100 (e.g.  $\frac{1}{12} \sum_{m=1}^{12} CPI_{m,y_0,k}^{item} = 100$ ). If I combine this arithmetic average together with the property that the index  $J_{m,y,k}^{item}$  is calculated as the percentage change since January ( $CPI_{m,y_0,k}^{item} = CPI_{1,y_0,k}^{item} J_{m,y_0,k}^{item}$ ), I can write out the starting condition for January of the base year ( $y_0$ ) as:

$$CPI_{1,y_0,k}^{item} = \frac{100}{\frac{1}{12} \sum_{m=1}^{12} J_{m,y_0,k}^{item}}. \quad (4.2.7)$$

By chain-linking the indices they no-longer follow the zig-zagged pattern, and are lined into one continuous index, as shown in Figure 4.4.

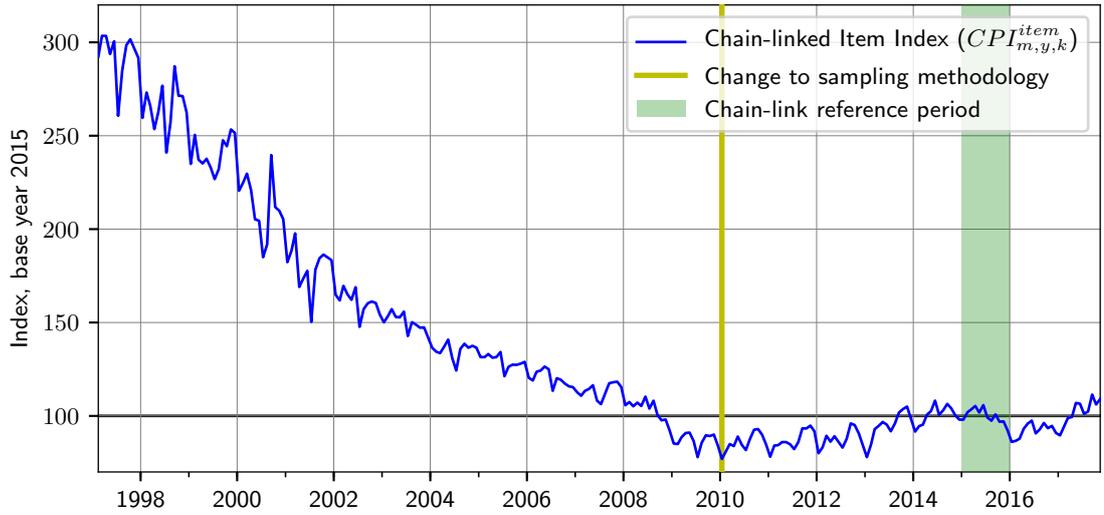
The chain-linking enables the introduction (and exit) of price trajectories into the index, as well as allowing the weights to be updated without impacting the estimated index. In figure 4.4 a clear change can be observed in the trend inflation in 2010. Whereas before 2010 the lack of recovery prices from sales meant the average month-on-month (log) inflation rate of the index was  $-0.81\%$ . After the changes in January 2010 the average monthly inflation rate of the index has been  $0.27\%$ . I speculate that the changing inflation dynamics are mostly due to the changes to the sampling methodology, resulting in a move from a strongly deflationary regime to mildly inflationary. However, the same period also coincides with the largest economic and financial crisis in living memory, and, hence, any speculation about causality merits a more careful investigation.

---

<sup>21</sup>This methodology also allows the ONS to introduce item indices for the higher aggregates of the COICOP classifications.

<sup>22</sup>This is in contrast to the previous observed January as given by the base prices (see equation 4.2.1).

Figure 4.4: Chain-Linked Item Index



Notes: Chain-linked item index “Women’s top, long sleeved, not blouse” (id 510223) with base year 2015. The index implies that the price of a women’s top in 1997 was three times as expensive as today, for the same quality product.

## 4.3 Decomposing the inflation rate

### 4.3.1 Inflation Accounting

Now that I have outlined the framework for how to calculate an alternative version of the item indices in the Consumer Price Index,<sup>23</sup> let me turn applying this framework to decompose the inflation rate into its constituent components of missing observations, product replacements, sales/recovery prices and regular price price changes. To derive a measure get this decomposition I first need to calculate the inflation rate from the chain-linked item index. Taking the log difference of the chain-linked item index from equation 4.2.6, I get the following expression:

$$\Delta \ln CPI_{m,y,k}^{item} = \begin{cases} \ln Z_{1,y,k}^{item} - \ln Z_{12,y-1,k}^{item} & \forall m = 1 \\ \ln Z_{2,y,k}^{item} & \forall m = 2 \\ \ln Z_{m,y,k}^{item} - \ln Z_{m-1,y,k}^{item} & \forall m \geq 3 \end{cases} \quad (4.3.1)$$

Recalling (and combine) the definition of the elementary aggregate index ( $Z_{m,y,h}^{ea}$ ) from equation 4.2.2 and the item index ( $Z_{m,y,k}^{item}$ ) in equation 4.2.4, I can write out the log of the item index in terms of the underlying price quotes, base prices, and

<sup>23</sup>Appendix 4.6.1 shows that the framework is comparable to the current method used by the ONS and appendix 4.7 shows how to extend the framework to headline inflation.

weights:

$$\ln Z_{m,y,k}^{item} = \sum_{h \in \mathcal{I}_{m,y,k}^{item}} \sum_{i \in \mathcal{I}_{m,y,h}^{ea}} w_{m,y,h}^{strat} w_{m,y,i}^{ea} (\ln P_{m,y,i} - \ln P_{m,y,i}^0). \quad (4.3.2)$$

I define the composite item weight  $W_{t,i}^{item} = w_{t,h(i)}^{strat} w_{t,i}^{ea}$ , remembering that there exists a unique mapping from the price trajectories to the stratification id:  $h \leftarrow h(i)$ . If we further define the scaler (date) index of  $t$ , which has the mapping to the month and year of  $t \leftrightarrow \{m, y\}$ , and with the change in the index defined as:

$$\Delta t \leftrightarrow \Delta\{m, y\} = \begin{cases} \Delta y & m = 1 \\ \Delta m & \forall m > 1 \end{cases}, \quad (4.3.3)$$

I can then write out the inflation rate of equation 4.3.1 as:

$$\begin{aligned} \Delta \ln CPI_{t,k}^{item} &= \sum_{h \in \mathcal{I}_{t,k}^{item}} \sum_{i \in \mathcal{I}_{t,h}^{ea}} W_{t,i}^{item} (\ln P_{t,i} - \ln P_{t,i}^0) \\ &\quad - d_t^{feb} \sum_{h \in \mathcal{I}_{t-1,k}^{item}} \sum_{i \in \mathcal{I}_{t-1,h}^{ea}} W_{t-1,i}^{item} (\ln P_{t-1,i} - \ln P_{t-1,i}^0) \end{aligned} \quad (4.3.4)$$

where  $d_t^{feb}$  is an indicator variable that takes the value of zero in the month of February and one in all other months.<sup>24</sup>

To check that the methodology and data source adds up to the inflation rate implied by the official numbers published the ONS, in appendix 4.6.1 I show that for the category of Women's top although not a perfect fit throughout the sample, the two series have very similar characteristics, and magnitudes. After 2010 the two series are almost identical, and hence I conclude that my calculated inflation rate and the decomposition to follow will have similar magnitudes for the series published by the ONS.

To simplify the notation I can write out the item index of equation 4.3.2 in matrix notation as:

$$\ln Z_{t,k}^{item} = \mathbf{W}_{t,k}^{item} \begin{pmatrix} \mathbf{p}_t & - \mathbf{p}_t^0 \\ N_{T,1} & N_{T,1} \end{pmatrix}, \quad (4.3.5)$$

In the above  $N_T$  is total set of price trajectories as defined earlier ( $N_T = N_{\bar{m}, \bar{y}}$  and  $T \leftrightarrow \{\bar{m}, \bar{y}\}$ ). The log price ( $\mathbf{p}_t$ ), log base price ( $\mathbf{p}_t^0$ ), and weight vector ( $\mathbf{W}_{t,k}^{item}$ ) are

---

<sup>24</sup>The variable is defined as  $d_t^{feb} = \mathbb{1}_{\{m \neq 2\}}$ , where  $\mathbb{1}_{\{\cdot\}}$  is the indicator function.

given by:

$$\begin{aligned}\mathbf{p}_t &= \left[ \ln P_{t,1} \quad \dots \quad \ln P_{t,N_T} \right]^\top, \\ \mathbf{p}_t^0 &= \left[ \ln P_{t,1}^0 \quad \dots \quad \ln P_{t,N_T}^0 \right]^\top, \\ \mathbf{W}_{t,k}^{item} &= \left[ W_{t,1}^{item} \quad \dots \quad W_{t,N_T}^{item} \right].\end{aligned}$$

Using the matrix notation for the log of the item index from equation 4.3.5, and the inflation rate calculation from equation 4.3.1, I can write the inflation rate as:

$$\Delta \ln CPI_{t,k}^{item} = \mathbf{W}_{t,k}^{item} (\mathbf{p}_t - \mathbf{p}_t^0) - d_t^{feb} \mathbf{W}_{t-1,k}^{item} (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0). \quad (4.3.6)$$

The above equation 4.3.6 will be my baseline accounting framework for quantifying decomposition of the missing price quotes, product replacements<sup>25</sup> and sales.

The aggregate item weight ( $\mathbf{W}_{t,k}^{item}$ ) can be written as the following multiplication of the stratum weight matrix and the elementary aggregate weight matrix:

$$\mathbf{W}_{t,k}^{item} = \mathbf{w}_{t,k}^{strat} \mathbf{w}_t^{ea} \quad (4.3.7)$$

where  $N_T$  is the total amount of trajectories and  $H_T$  is the total amount of item and stratum cell combinations (recall that  $T \leftrightarrow \{\bar{m}, \bar{y}\}$ ). The elementary aggregate matrix ( $\mathbf{w}_t^{ea}$ ) is defined as:

$$\mathbf{w}_t^{ea} = \begin{bmatrix} w_{t,11}^{ea} & \dots & w_{t,1N_T}^{ea} \\ \vdots & \ddots & \vdots \\ w_{t,H_T1}^{ea} & \dots & w_{t,H_T N_T}^{ea} \end{bmatrix}, \quad (4.3.8)$$

with the individual elements of row  $h$  and column  $i$  in the matrix of the elementary aggregate weight given by:

$$[\mathbf{w}_t^{ea}]_{h,i} = \begin{cases} w_{t,i}^{ea} & \forall i \in \mathcal{I}_{t,h}^{ea} \\ 0 & \forall i \notin \mathcal{I}_{t,h}^{ea} \end{cases} \quad (4.3.9)$$

where a 0 weight is given if a price quote either does not feature in the aggregation of that path or is missing (e.g. is not available for sampling) for the period  $t$ . Recall from equation 4.2.3 that if an observation is missing for a price trajectory because it is not part of the current basket, it was not available for collection, or it is not part of the index, then its elementary aggregate weight is set to zero.

---

<sup>25</sup>For directly comparable replacements. Remember that for non-comparable product replacements my current methodology breaks the trajectory in two, as the base prices changes, and hence will feature into my estimates of the missing observations.

As for the stratum weight row vector, it is given by:

$$\mathbf{w}_{t,k}^{strat} = \begin{bmatrix} w_{t,1}^{strat} & \dots & w_{t,H_T}^{strat} \end{bmatrix}, \quad (4.3.10)$$

and with the individual element of column  $h$  of the vector given by:<sup>26</sup>

$$[\mathbf{w}_{t,k}^{strat}]_{1,h} = \begin{cases} w_{t,h}^{strat} & \forall h \in \mathcal{I}_{t,k}^{item} \\ 0 & \forall h \notin \mathcal{I}_{t,k}^{item} \end{cases}. \quad (4.3.11)$$

Both of the above weight matrices are very sparse, and mostly consisting of zeroes. The hierarchical structure of the aggregation procedure means that there is a unique path from an individual price quote ( $i$ ), that only features in one branch of the aggregation to the item index and inflation rate. This implies that for all of the weight matrices, each column will only have a maximum of one row that is non-zero and the rest will be zeroes. As a consequence it is possible to order the elementary aggregate weight ( $w_t^{ea}$ ) in such a way that it is block diagonal, given that moving from elementary aggregates (right) to the stratum indices (left) has a smaller dimension than the one preceding ( $H_T \ll N_T$ ).

### 4.3.2 Missing Price Quotes

There are two kinds of missing price quotes in the dataset; permanently missing because a product has reached the end of its life-cycle, and temporarily missing price quotes. A permanently missing price quote means that the price line drops out from any further collection and is the termination of that price trajectory. In contrast a temporarily missing quote means a price trajectory was temporarily unavailable for sampling (such as being out of stock when the collectors visited the outlet).<sup>27</sup>

To set the scene, let me first illustrate how the missing observations are currently dealt with using the methodology of the ONS. In table 4.2, I have constructed a theoretical example of three price trajectories where in period 2 trajectory 3 is not available for sampling. As a consequence, the methodology implies that the period 2 ratio between price and base price of trajectory 3 is an estimate constructed as a weighted geometric average of the other two trajectories. I assume trajectory 3 is only temporary missing, which could be due to being out of stock, implying that there is an underlying (but unobserved) price in period 2. Cases of product replacements will be dealt with later. For the first period, all price quotes are

<sup>26</sup>The notation  $[x]_{h,i}$  implies the element on the  $h$ -row and  $i$ -column of the matrix  $x$ .

<sup>27</sup>The missing item is, however, expected to be back on the shelves within the next month or two. If a price trajectory has a missing price for more than two months, that trajectory has to be dropped according to Eurostat's HICP regulation. Temporarily missing price quotes are marked with an indicator "T", and those permanently missing with an "M" in the dataset.

Table 4.2: A numerical example of missing observations

	Formula	Period ( $t$ )		
		1	2	3
Element				
Trajectory 1	$P_{t,1}/P_{t,1}^0$	1.00	1.02	1.03
Trajectory 2	$P_{t,2}/P_{t,2}^0$	1.02	1.01	1.02
Trajectory 3	$P_{t,3}/P_{t,3}^0$	0.98		1.03
Weight				
Trajectory 1	$w_{t,1}^{ea}$	1/3	1/2	1/3
Trajectory 2	$w_{t,2}^{ea}$	1/3	1/2	1/3
Trajectory 3	$w_{t,3}^{ea}$	1/3	0	1/3
Elementary Aggregate				
Information set	$\mathcal{I}_{t,h}^{ea}$	{1, 2, 3}	{1, 2}	{1, 2, 3}
Index	$Z_{t,h}^{ea} = \prod_{i \in \mathcal{I}_{t,h}^{ea}} (P_{t,i}/P_{t,i}^0)^{w_{t,i}^{ea}}$	1.00	1.015	1.027
Inflation	$Z_{t,h}^{ea}/Z_{t-1,h}^{ea} - 1$		1.51%	1.15%
Matched				
Information set	$\mathcal{I}_{t,h}^*$	{1, 2}	{1, 2}	{1, 2}
Index	$Z_{t,h}^{ea*} = \prod_{i \in \mathcal{I}_{t,h}^*} (P_{t,i}/P_{t,i}^0)^{w_{t,i}^{ea*}}$	1.01	1.015	1.025
Weight	$\sum_{i \in \mathcal{I}_{t,h}^{ea*}} w_{t,i}^{ea}$	2/3	2/3	2/3
Inflation	$\pi_t^* = Z_{t,h}^{ea*}/Z_{t-1,h}^{ea*} - 1$		0.50%	0.98%
Missing				
Information set	$\hat{\mathcal{I}}_{t,h}^{ea}$	{3}	{}	{3}
Index	$\hat{Z}_{t,h}^{ea} = \prod_{i \in \hat{\mathcal{I}}_{t,h}^{ea}} (P_{t,i}/P_{t,i}^0)^{\hat{w}_{t,i}^{ea}}$	0.98		1.03
Estimate of Index			1.015	
Weight	$\sum_{i \in \hat{\mathcal{I}}_{t,h}^{ea}} w_{t,i}^{ea}$	1/3	1/3	1/3
Inflation	$\pi_t^M = \hat{Z}_{t,h}^{ea}/\hat{Z}_{t-1,h}^{ea} - 1$		3.51%	1.47%
Difference (inflation)	$\pi_t^M - \pi_t^*$		3.01	0.49

Notes: Example of a missing price trajectory. To simplify exposition I assume all base prices are equal to 1 in the above theoretical example, and hence we leave them out of the calculations. Further, all shop replication factors are set to 1, and hence I am assuming equal weighting.

observed, and hence the elementary aggregate index is calculated based on all of the three trajectories. In the second period, trajectory 3 is unavailable for sampling, such as being out of stock when the price collectors visits the store. Hence the calculation of the elementary aggregate index ( $Z_{t,h}^{ea}$ ), is only based on the first two trajectories, and the weights are re-weighting to take this into account. Hence when I calculate the inflation rate between period 1 and 2 (1.51%), I can decompose this into the part of matched trajectories between the two periods (trajectory 1 and 2), which have an inflation rate of 0.5% (and weights of 2/3 in the overall index), and between the estimated index of trajectory 3 (equivalent to the index calculated from our two observed trajectories), and the previous index value of trajectory 3, with an implied inflation rate of 3.51% (with weights of 1/3 in the total). Lastly, in period 3

I again observe all the trajectories, and hence the overall index is based on the full information set. This implies that when I calculate the inflation between period 1 and 3, which both includes all the trajectories, the estimates used for trajectory 3 in period 2 will cancel out, and hence I am left with only observable quantities in this longer-run effect. The unobserved price of trajectory 3 in period 2 will only feature in terms of how it affects the price level of the trajectory when it is re-observed. As a consequence theoretically, if all trajectories are re-observed, any estimates for missing observations should cancel out in the long-run and only have transitory effects. However, since not all trajectories are re-observed as they might exit the basket, and due to the interaction with quality changes, I do observe some long-run impacts as observed from the clothing example previously.

To pick out the missing observations from the vector of log price quotes and base prices, I use a selection matrix  $\mathbf{M}_t$  which is a diagonal matrix with ones or zeros on the diagonal depending on whether the price quote is missing (ones) or not (zero), and zeros on all off-diagonal elements. A more formal definition of this selection matrix is as follows:

$$\mathbf{M}_t = \underset{N_T \cdot N_T}{diag} \left( \left[ d_t^{feb} \times \mathbb{1}_{(i \notin \mathcal{I}_t)} \forall i = 1, \dots, N_T \right] \right), \quad (4.3.12)$$

where  $\mathbb{1}_{(i \notin \mathcal{I}_t)}$  is an indicator function that takes the value of one if the price quote  $i$  is missing, or zero otherwise.  $\mathcal{I}_t$  is the full information set of available price quotes in period  $t$ , and is a subset of the full set of trajectories ( $\mathcal{I}_t \subset \{1, \dots, N_T\}$ ). Recall that there is a non-linear relationship between the shop weights and the missing observations, as seen from equation 4.2.3. This is equivalent to the methodology used by the ONS to drop missing price quotes from the index (ONS, 2014). In their methodology, the ONS deals with missing observations through the shop replication factors ( $f_{t,i}$ ) by setting them equal to zero for missing observations. The latter has the same effect as estimating the unobserved missing price from the geometric average of the prices and base prices composing the elementary aggregate index. This means that when a price trajectory is missing, there will be a recalculation of all the elementary aggregate weights within that stratum cell as a consequence. Further, because of the base prices, we always have the previous periods price in February hence  $d_t^{feb}$ , which we adjust for in the selection matrix, by always setting the diagonal equal to zeroes in February.

Combining equation 4.3.12 with equation 4.2.3, it straightforwardly follows by virtue of the elementary aggregate weight ( $\mathbf{w}_t^{ea}$ ) being zero for missing observations that:

$$\mathbf{W}_{t,k}^{item} \mathbf{M}_t = \mathbf{0} \Leftrightarrow \mathbf{W}_{t,k}^{item} (\mathbf{I}_{N_T} - \mathbf{M}_t) = \mathbf{W}_{t,k}^{item} \quad \forall t. \quad (4.3.13)$$

Using the above expression in equation 4.3.13, I can write out the Geometric Lowe

index, from equation 4.3.5, for the current month  $t$  as:

$$\ln Z_{t,k}^{item} = \mathbf{W}_{t,k}^{item}(\mathbf{I}_{N_T} - \mathbf{M}_{t-1})(\mathbf{p}_t - \mathbf{p}_t^0) + \mathbf{W}_{t,k}^{item}\mathbf{M}_{t-1}(\mathbf{p}_t - \mathbf{p}_t^0), \quad (4.3.14)$$

and the previous months index as:

$$\ln Z_{t-1,k}^{item} = \mathbf{W}_{t-1,k}^{item}(\mathbf{I}_{N_T} - \mathbf{M}_t)(\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) + \mathbf{W}_{t-1,k}^{item}\mathbf{M}_t(\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0). \quad (4.3.15)$$

Hence using the above equations, I can re-write equation 4.3.6 as:

$$\begin{aligned} \Delta \ln CPI_{t,k}^{item} &= \mathbf{W}_{t,k}^{item}(\mathbf{I}_{N_T} - \mathbf{M}_{t-1})(\mathbf{p}_t - \mathbf{p}_t^0) \\ &\quad - d_t^{feb} \mathbf{W}_{t-1,k}^{item}(\mathbf{I}_{N_T} - \mathbf{M}_t)(\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \\ &\quad + \mathbf{W}_{t,k}^{item}\mathbf{M}_{t-1}(\mathbf{p}_t - \mathbf{p}_t^0) \\ &\quad - d_t^{feb} \mathbf{W}_{t-1,k}^{item}\mathbf{M}_t(\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0). \end{aligned} \quad (4.3.16)$$

In the above equation 4.3.16, the two last terms deal with the missing observations in terms of calculating the inflation rate of the missing price quotes as the weighted log mean of the observed ones (and their log base prices). Notice also by the nature of the chain-linking methodology, there are no missing observations in February.

If I look back at the elementary aggregate weight from equation 4.2.3, I can redefine it in terms of only the price trajectories that was available for sampling in both the current and previous periods:

$$w_{t,i}^{ea*} = \begin{cases} \frac{f_{t,i}}{\sum_{j \in (\mathcal{I}_{t,h(i)}^{ea} \cap \mathcal{I}_{t-1,h(i)}^{ea})} f_{t,j}} & \forall i \in (\mathcal{I}_{t,h(i)}^{ea} \cap \mathcal{I}_{t-1,h(i)}^{ea}) \\ 0 & \forall i \notin (\mathcal{I}_{t,h(i)}^{ea} \cap \mathcal{I}_{t-1,h(i)}^{ea}) \end{cases} \quad (4.3.17)$$

where the information set,  $\mathcal{I}_{t,h(i)}^{ea} \cap \mathcal{I}_{t-1,h(i)}^{ea}$ , is the trajectories that was available for sampling in both periods. I then define the matrix form of the matched elementary aggregate weight  $\mathbf{w}_t^{ea*}$ , as having the row  $h$ , and column  $i$  elements of:

$$[\mathbf{w}_t^{ea*}]_{h,i} = \begin{cases} w_{t,i}^{ea*} & \forall i \in (\mathcal{I}_{t,h}^{ea} \cap \mathcal{I}_{t-1,h}^{ea}) \\ 0 & \forall i \notin (\mathcal{I}_{t,h}^{ea} \cap \mathcal{I}_{t-1,h}^{ea}) \end{cases} \quad (4.3.18)$$

Using the above expression, I can define  $\mathbf{W}_{t,k}^{item*} = \mathbf{w}_t^{strat} \mathbf{w}_t^{ea*}$ , and using it with

equation 4.3.16, I can rewrite the decomposition of the inflation rate as:

$$\begin{aligned}
\Delta \ln CPI_{t,k}^{item} = & \mathbf{W}_{t,k}^{item*} \left[ (\mathbf{p}_t - \mathbf{p}_t^0) - d_t^{feb} (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \right] \\
& + (\mathbf{W}_{t,k}^{item} - \mathbf{W}_{t,k}^{item*}) (\mathbf{I}_{N_T} - \mathbf{M}_{t-1}) (\mathbf{p}_t - \mathbf{p}_t^0) \\
& - d_t^{feb} \mathbf{W}_{t-1,k}^{item} \mathbf{M}_t (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \\
& + \mathbf{W}_{t,k}^{item} \mathbf{M}_{t-1} (\mathbf{p}_t - \mathbf{p}_t^0) \\
& - d_t^{feb} (\mathbf{W}_{t-1,k}^{item} - \mathbf{W}_{t,k}^{item*}) (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0).
\end{aligned} \tag{4.3.19}$$

Recalling for the above that all columns of the  $\mathbf{W}_{t,k}^{item*}$  matrix where either the current or previous trajectory were unavailable for sampling is zero, hence I drop the selection matrices for the first expression (as it is implicitly given by the matched weights;  $\mathbf{W}_t^{item*}$ ).

The re-weighting terms of the missing observations takes place at the elementary aggregate level and hence using equation 4.2.3:<sup>28</sup>

$$\Delta \mathbf{w}_t^{ea} \underset{H_T \cdot N_T \cdot N_T \cdot 1}{\iota_{N_T}} = \underset{H_T \cdot 1}{\mathbf{0}} \quad \forall t \tag{4.3.20}$$

The above weight changes are the rebalancing required to adjust for the missing price quotes in the sample. The implicit assumption for the change in the weights due to missing observations is that of an outlet substitution, where consumers change their place of shopping if a product becomes temporarily unavailable in a shop. Crucially, however, this substitution effect does not take into account the product substitution that happens in reality for a large number of items when they become temporarily unavailable in an outlet.

### 4.3.3 Replacement products and quality change

The ONS has two different concepts of product substitutions: comparable and non-comparable. Comparable product substitutions are instances when a product is no longer available for sampling but an alternative is found that is deemed to be comparable in quality and quantity. Hence no adjustment is made to the base prices for that trajectory. In contrast for non-comparable product substitutions either the quality or quantity is deemed to have changed, and hence the ONS adjust the base prices accordingly to offset this real change for the trajectory.

For a subset of price quotes, they are flagged as either being a comparable product replacement (a substitution) or a non-comparable replacement for a product that is no longer available for sampling. In the dataset I use the selection matrix  $\mathbf{R}_t$  for

---

<sup>28</sup>Follows trivially given that  $\underset{H_T \cdot N_t}{w_t^{ea}} \iota_{N_T} = \iota_{H_T} \forall t$ , where  $\iota_N$  is a unit vector is dimension  $N$ .

Table 4.3: Example of Quality Change

	Formula	Period ( $t$ )		
		1	2	3
<b>Trajectory 1</b>				
Price	$P_{t,1}$	1.01	1.2	1.3
Base Price	$P_{t,1}^0$	1.00	1.2	1.2
Weight	$w_{t,1}^{ea}$	1/2	1/2	1/2
Price Change	$P_{t,1}/P_{t-1,1} - 1$		18.8%	8.3%
Quality Adjustment	$P_{t,1}^0/P_{t-1,1}^0 - 1$		20%	0%
Inflation	$\left(\frac{P_{t,1}}{P_{t-1,1}}\right) / \left(\frac{P_{t,1}^0}{P_{t-1,1}^0}\right) - 1$		-1.0%	8.3%
<b>Trajectory 2</b>				
Price	$P_{t,2}$	1.01	1.02	1.03
Base Price	$P_{t,2}^0$	1	1	1
Weight	$w_{t,2}^{ea}$	1/2	1/2	1/2
Price Change	$P_{t,2}/P_{t-1,2} - 1$		0.99%	0.98%
Quality Adjustment	$P_{t,2}^0/P_{t-1,2}^0 - 1$		0%	0%
Inflation	$\left(\frac{P_{t,2}}{P_{t-1,2}}\right) / \left(\frac{P_{t,2}^0}{P_{t-1,2}^0}\right) - 1$		0.99%	0.98%
<b>Aggregation</b>				
Index	$Z_t^{ea} = \left(\frac{P_{t,1}}{P_{t,1}^0}\right)^{w_{t,1}^{ea}} \left(\frac{P_{t,2}}{P_{t,2}^0}\right)^{w_{t,2}^{ea}}$	1.01	1.01	1.06
Inflation	$Z_t^{ea}/Z_{t-1}^{ea} - 1$		0.0%	4.6%

Notes: Example of a quality change to a price trajectory.

both a substitution or a replacement. I define  $\mathbf{R}_t$  similarly to the selection matrix for missing prices as a diagonal matrix, with one on the diagonal if a price quote is marked as a substitute or replacement product, and zero otherwise.

$$\mathbf{R}_t = \underset{N_T \cdot N_T}{diag} \left( \left[ d_{t,i}^R \forall i = 1, \dots, N_T \right] \right), \quad (4.3.21)$$

where  $d_{t,i}^R$  is an indicator function that takes the value of one, if the price quote  $i$  is marked with an indicator as being either a non-comparable replacement good, a comparable substitution or a weight change, and zero otherwise.<sup>29</sup>

Recall the first expression from the missing prices accounting framework in equation 4.3.16, where I am looking at only price trajectories available for sampling in

<sup>29</sup>A replacement is a replacement product that is non-comparable to the previous sampled product (and is marked with the indicator codes N or Z). In contrast a substitute product is comparable to the previous sampled product and is marked with the indicator codes C or X. The indicator implicitly nests further disaggregate indicators for each of these categories in terms of:

$$d_{t,i}^R = d_{t,i}^{RC} + d_t^{RN} + d_t^{RW}, \quad (4.3.22)$$

where  $d_{t,i}^{RC}$  is an indicator variable for a comparable substitution in the trajectory,  $d_{t,i}^{RN}$  is a non-comparable product replacement, and  $d_{t,i}^{RW}$  is a weight change. Given we only observe one flag for each data point only one of the indicators can be non-zero at any time and hence  $d_{t,i}^R \in \{0, 1\} \forall t \forall i$ .

both the previous and current period:

$$\mathbf{W}_{t,k}^{item*} \left[ (\mathbf{p}_t - \mathbf{p}_t^0) - d_t^{feb}(\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \right]$$

This means that I can pick out the changes related to a replacement, substitution or weight change of a product as:

$$\mathbf{W}_{t,k}^{item*} \mathbf{R}_t \left[ (\mathbf{p}_t - \mathbf{p}_t^0) - d_t^{feb}(\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \right] \quad (4.3.23)$$

where the first element is defined by the fact that for all non-replacement and non-substitute price quotes the base prices only change in the month of February and hence:

$$(\mathbf{I}_{N_T} - \mathbf{R}_t) d_t^{feb} \Delta \mathbf{p}_t^0 = 0. \quad (4.3.24)$$

For the substitutions which are comparable, the base price will be identical and hence  $\Delta p_{t,i}^0 = 0$ . The replacement price lines (non-comparable) are made comparable by adjusting the base prices. Any changes to quality (or quantity) are adjusted for by having different base prices, which is done through the vector:  $\mathbf{R}_t \Delta \mathbf{p}_t^0$ , having non-zero elements.

The ONS uses three different methods for dealing with changes to the sampled items. Firstly, if possible, a direct comparison for which the base price of the item previously sampled is carried forward. Secondly, a direct comparison of the base price by either quantity adjustments (if clear observable changes) or hedonic regressions to adjust for more complex specifications of the item base price.<sup>30</sup> The last method used is imputations of the base price if no information is available to quantify the difference. The imputations are done by assuming that the price change of the item from the base month to the current month is equal to the average change within the elementary aggregate it belongs to, and hence a base price is imputed from this information.<sup>31</sup>

#### 4.3.4 Sales and Recovery Prices

Lastly, we set out the decomposition for when price is marked by the ONS as either being a sales price or a recovery price. To do this we use the selection matrix  $\mathbf{S}_t$ , using the sales and recovery flag as given by the ONS in the dataset. The matrix

---

<sup>30</sup>The ONS only does quality adjustment for a subset of the item indices. PCs (since 2003), laptops (since 2005), Tablet PCs (since 2013), Digital Cameras (since 2004), Smartphones (since 2011) and mobile phones (since 2007) (Wells and Restieaux, 2014). All the price data for which these indices are adjusted are centrally collected by the ONS, and therefore not published in the public dataset (Wells and Restieaux, 2014, Annex A).

<sup>31</sup>See ONS (2014, p.49–55) for further details. Further, as already mentioned in January 2010 the ONS changed their sampling methodology for clothing, such that more product substitutions were deemed comparable, whereas previously the majority was deemed non-comparable.

follow previous notation in being a diagonal matrix with ones if a price is indicated to be either a sale or recovery from a sale price and zero otherwise. More formally we set out the matrix as

$$\mathbf{S}_t = \underset{N_T \cdot N_T}{diag} \left( \left[ d_{t,i}^S \forall i = 1, \dots, N_T \right] \right), \quad (4.3.25)$$

where  $d_{t,i}^S$  is an indicator function which takes the value of 1 for a sale/recovery price in this period or a sale price in the previous period and zero otherwise.

### 4.3.5 Applying the decomposition

Define the combined selection matrix of  $\mathbb{R}_t$  as

$$\mathbb{R}_t = (\mathbf{I}_{N_T} - \mathbf{M}_{t-1}) (\mathbf{I}_{N_T} - \mathbf{R}_t) (\mathbf{I}_{N_T} - \mathbf{S}_t). \quad (4.3.26)$$

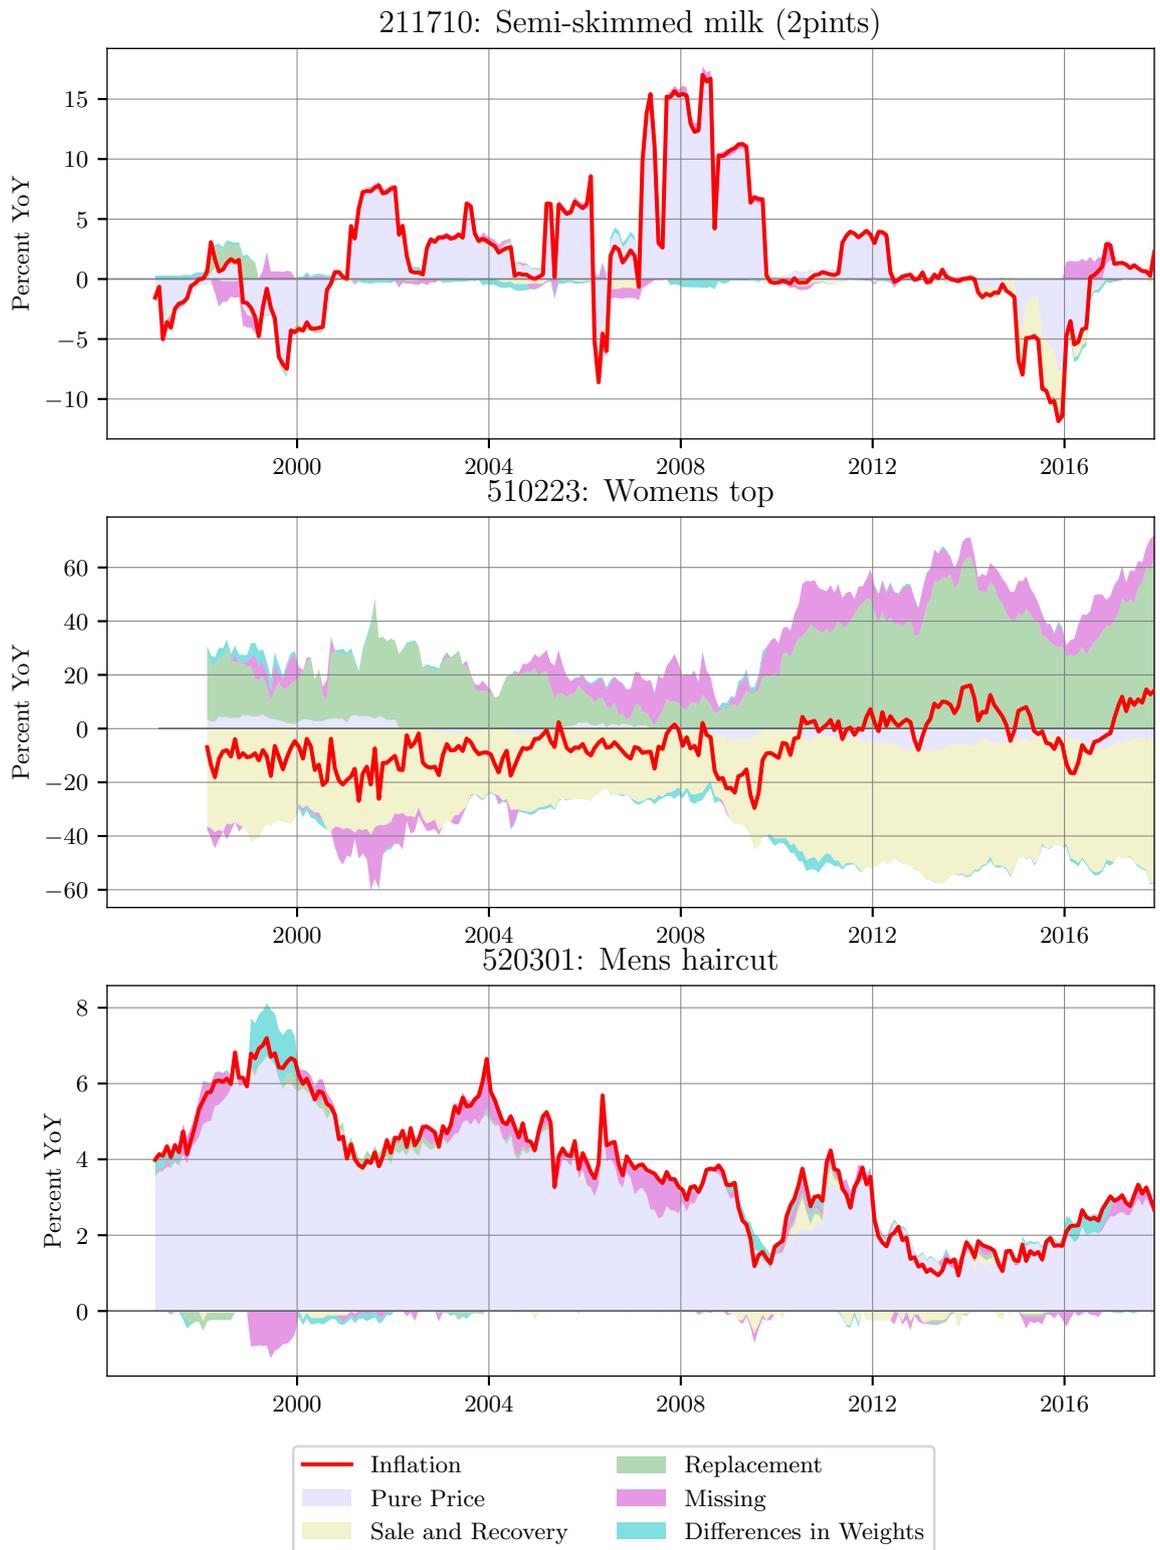
Using this notation we can then combine all our elements into the decomposition of the inflation rate as

$$\begin{aligned} \Delta \ln CPI_{t,k}^{item} = & \mathbf{W}_{t,k}^{item} \mathbb{R}_t \left[ (\mathbf{p}_t - \mathbf{p}_t^0) - d_t^{feb} (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \right] \\ & + \mathbf{W}_{t,k}^{item} (\mathbf{I}_{N_T} - \mathbf{M}_{t-1}) (\mathbf{I}_{N_T} - \mathbf{R}_t) \mathbf{S}_t (\mathbf{p}_t - \mathbf{p}_t^0) \\ & - d_t^{feb} \mathbf{W}_{t-1,k}^{item} (\mathbf{I}_{N_T} - \mathbf{M}_t) (\mathbf{I}_{N_T} - \mathbf{R}_t) \mathbf{S}_t (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \\ & + \mathbf{W}_{t,k}^{item} (\mathbf{I}_{N_T} - \mathbf{M}_{t-1}) \mathbf{R}_t (\mathbf{p}_t - \mathbf{p}_t^0) \\ & - d_t^{feb} \mathbf{W}_{t-1,k}^{item} (\mathbf{I}_{N_T} - \mathbf{M}_t) \mathbf{R}_t (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \\ & - d_t^{feb} (\mathbf{W}_{t-1,k}^{item} - \mathbf{W}_{t,k}^{item}) (\mathbf{I}_{N_T} - \mathbf{M}_t) \mathbb{R}_t (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0) \\ & + \mathbf{W}_{t,k}^{item} \mathbf{M}_{t-1} (\mathbf{p}_t - \mathbf{p}_t^0) - d_t^{feb} \mathbf{W}_{t-1,k}^{item} \mathbf{M}_t (\mathbf{p}_{t-1} - \mathbf{p}_{t-1}^0). \end{aligned} \quad (4.3.27)$$

The first part of the decomposition can be seen as the regular price change (or pure price change) resulting from directly observable changes in prices. Any deviations from this process being the main drive of the inflation rate as is the case for clothing categories raises question the measurement of the inflation rates. The second components is the sale and recovery component, which given the transitory nature of a sale (by construction) should be mean zero over a period such as a year. Thirdly we the impact of product replacements and lastly the missing observations and how they are offset by changes to the weights.

I apply this framework to decomposing the inflation rate for the three different item categories we are investigating in figure 4.5. In the figure 4.5 we clearly see that for both men's haircut and semi-skimmed milk the primary drive of the inflation rate is the regular "pure" price changes. Although for a few periods there is a minor impact of other effects, such as during 2015 where sales and recovery price

Figure 4.5: Decomposition of the inflation rate



Notes: item inflation “Women’s top, long sleeved, not blouse” (id 510223), “Men’s haircut” (id 520301), “2 pint of skimmed milk” (id 211710).

dynamics plays some role for semi-skimmed milk inflation, it is an overall minor impact. This compares to the category of women’s top, where we see the opposite pattern of regular price dynamics only having a minor impact on the inflation rate, which is mostly driven by the opposing forces of deflationary sales prices and inflationary replacement items, which could be interpreted as pattern of a fashion cycle. Although we would expect sales and recovery to have no impact on the inflation rate on an annual base, the measurement of a clothing fashion cycle result in sales having a permanent deflationary impact, that is offset (partially in the first of the sample) by replacement products.

## 4.4 An alternative pure price index: No quality change

As an alternative measure of the index construction, we construct a price index purely based on the prices and the weights as provided by the ONS, abstracting from any quality change through changes to the base prices.

The geometric price index is constructed directly as

$$y_t = \sum_{i \in \mathcal{I}_t} W_{t,i}^{item} \ln P_{t,i} \quad (4.4.1)$$

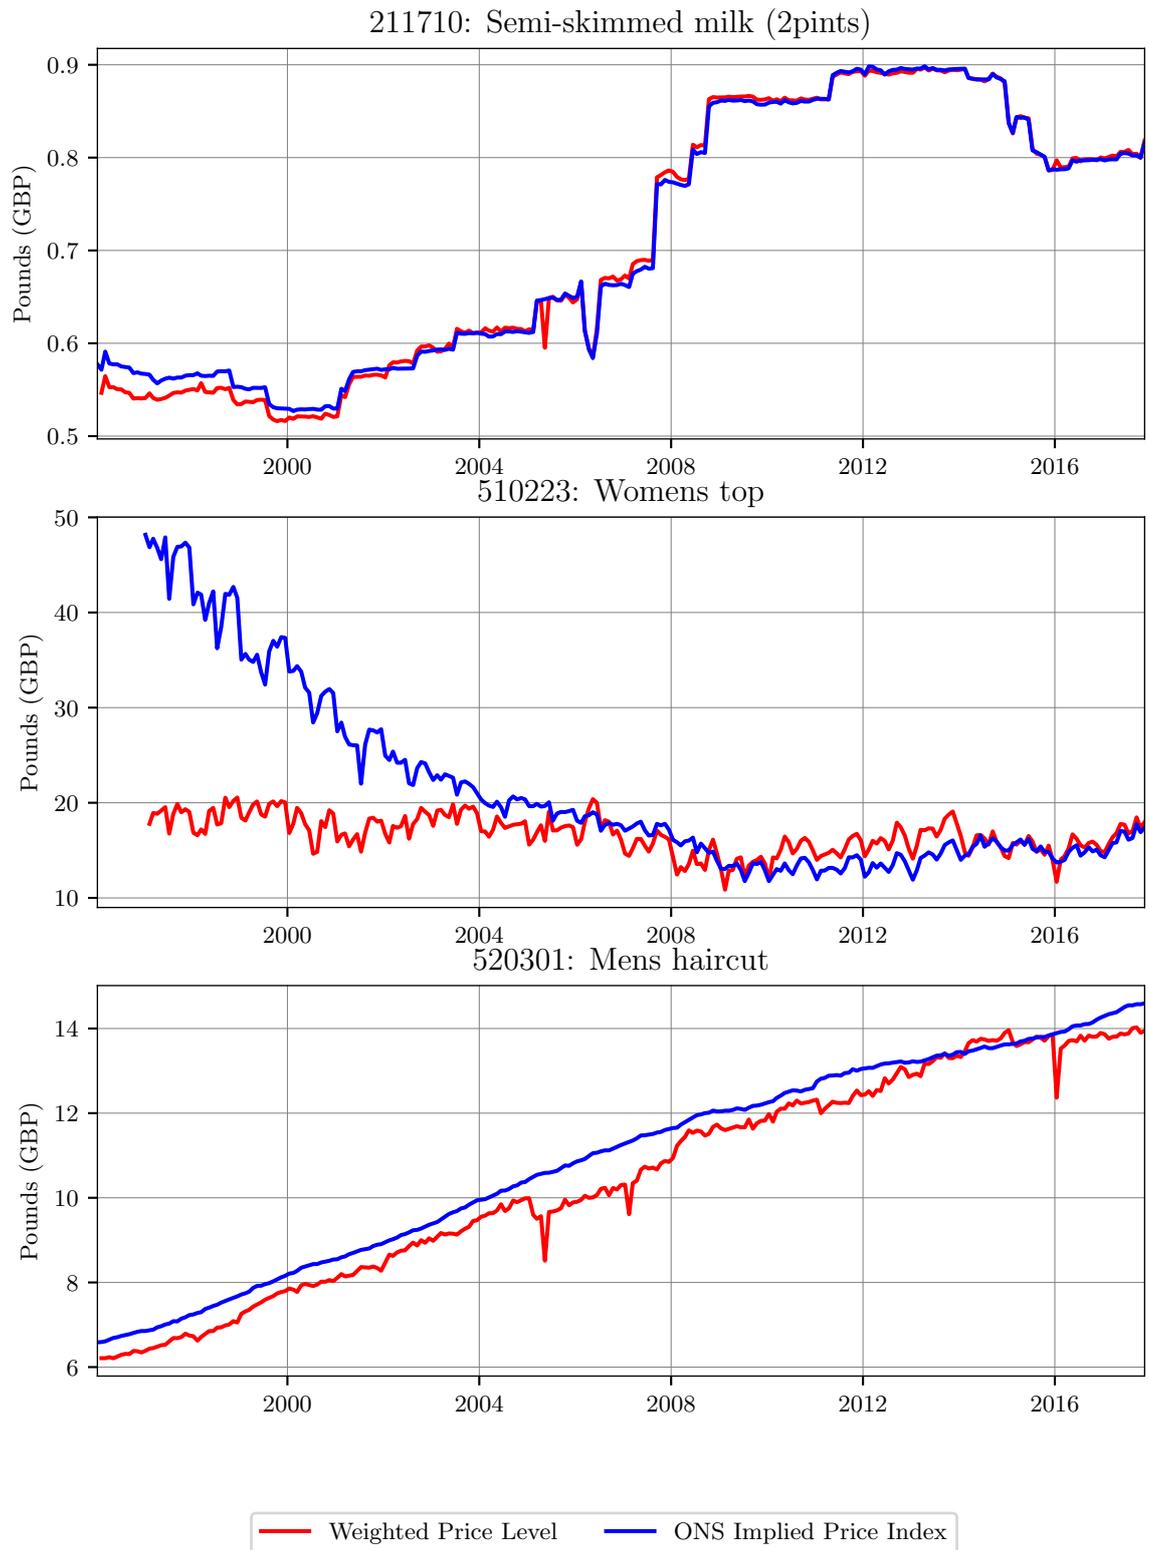
which does not take into account any adjustment of the base prices. The results are shown in figure 4.6 which is the geometric weighted mean of the price level. The first observation we make is that for both semi-skimmed milk and men’s haircut the pure price index is in line with the implied price index as given by the ONS, which confirms our previous results that the primary driver for these items are pure price changes. In contrast women’s top clearly shows a very different behaviour in the early sample before the methodological changes made in 2010, which gives evidence of a mismeasurement in the historical values.

To show the underlying changes in the distribution of sampled prices, figure 4.7 plots the (log) distributions using the weights of the prices. For women’s top we saw a very large increase in the dispersion of the sampled prices after 2010, when the changes are made for how the category is sampled.

## 4.5 Conclusion

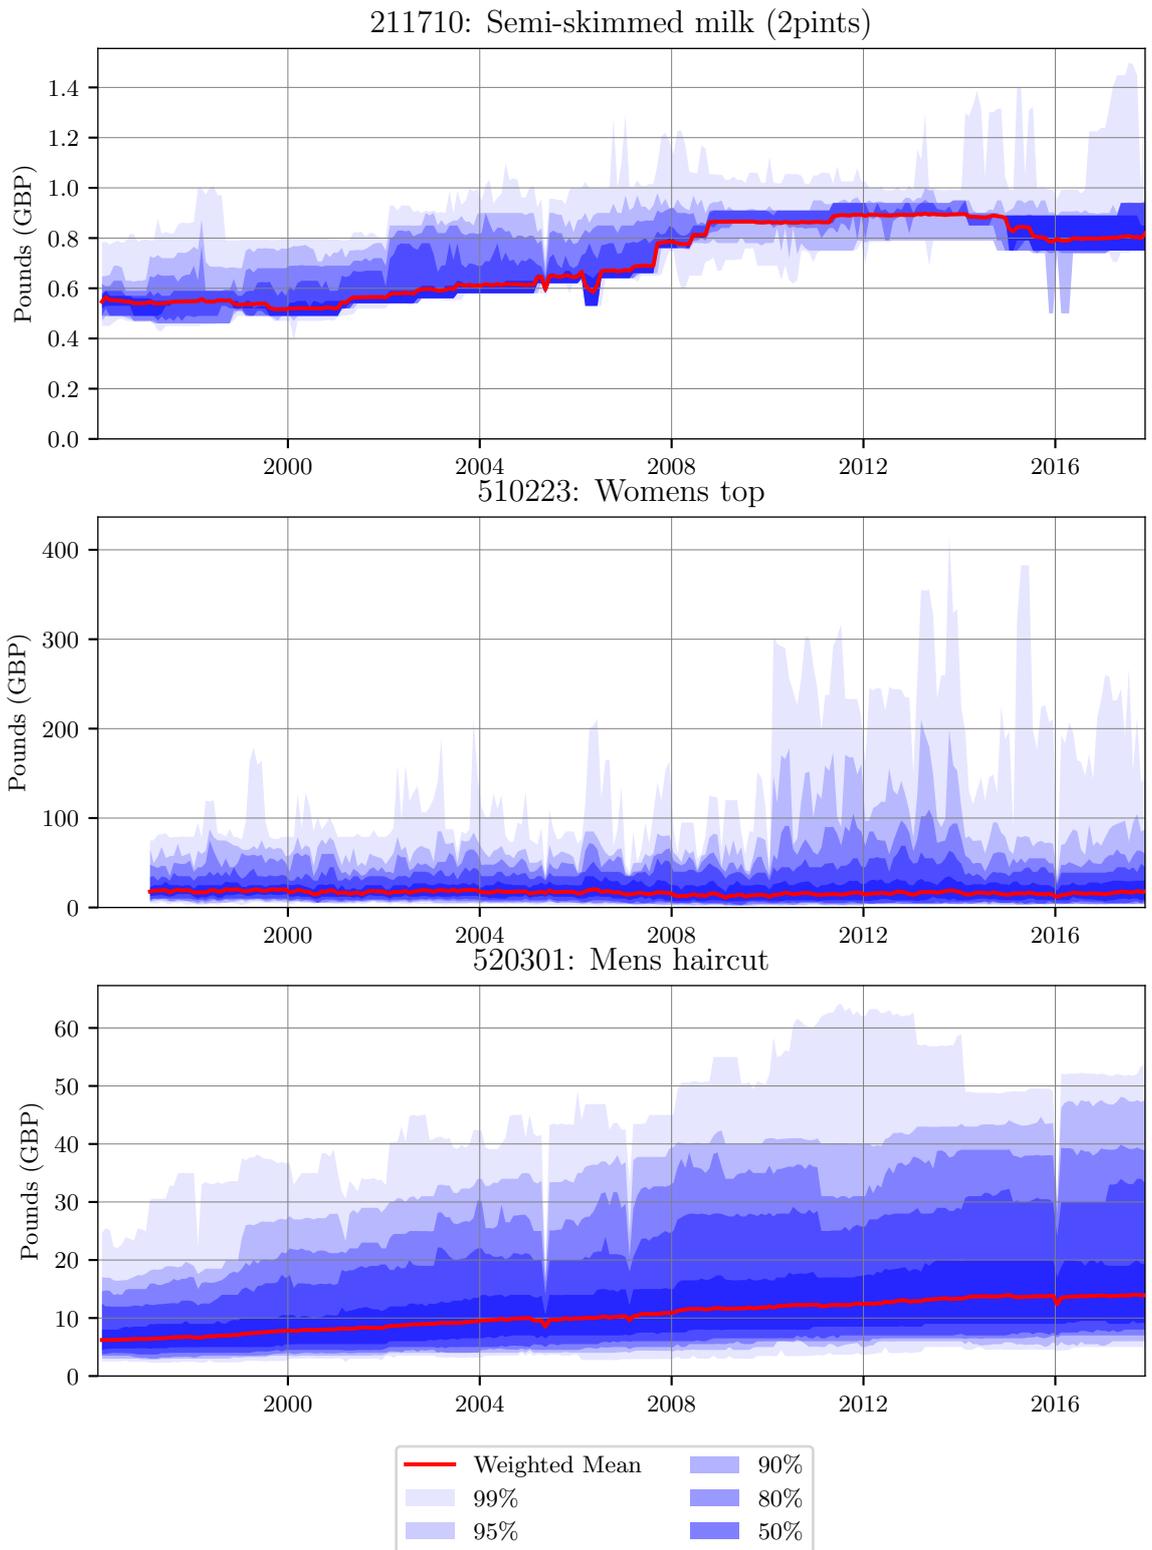
In this paper I have shown that missing observations, product replacements and sales are important features for the measurement UK’s inflation rate in clothing items.

Figure 4.6: Weighted index of prices



Notes: item inflation “2 pint skimmed milk” (id 211710) “Women’s top, long sleeved, not blouse” (id 510223), and “Men’s haircut” (id 520301).

Figure 4.7: Distribution of sampled prices



Notes: item inflation “2 pint skimmed milk” (id 211710) “Women’s top, long sleeved, not blouse” (id 510223), and “Men’s haircut” (id 520301).

I postulate that the reason the sales and recoveries are not mean zero, is due to the way price trajectories are often missing in their recovery phase, and hence a deeper investigation of interaction with missing price quotes is needed to established whether there is indeed a mismeasurement of the sales and recoveries prices. For the product replacements, although there is a clear difference on average between the inflation rate of the matched price inflation and the replacements inflation, the fact that the category good I am investigating is a fashion item makes it difficult to interpret this results as a measurement bias.

In contrast for the item categories of men’s haircut and semi-skimmed milk I find less importance for the non direct price change components in the measurement of their inflation rates.

I see this work as a first step in evaluating more systematically follow the importance of quality adjustment, the estimates of the missing observations, and sales. The framework outlined should be a cost-efficient method to implement as a check on the different measures. Further, the framework in this chapter is easily scaleable and able to be implemented across all of the COICOP categories and aggregation indices. However, given the nature of clothing being a fashion item and the complications this pose for my results, there is still more work to be done to fully understand the implications of what underlying technical elements drive the measurement of the UK’s inflation rate.

An additional benefit of the framework outlined in this chapter is that is links the measurement literature of the price index, to the method applied in the macroeconomic literature on micro price data dynamics such as seen in chapter 3 of this thesis. It allows a quantification of the assumptions used when including or excluding certain prices from the analysis in terms of their impact of the overall inflation rate.

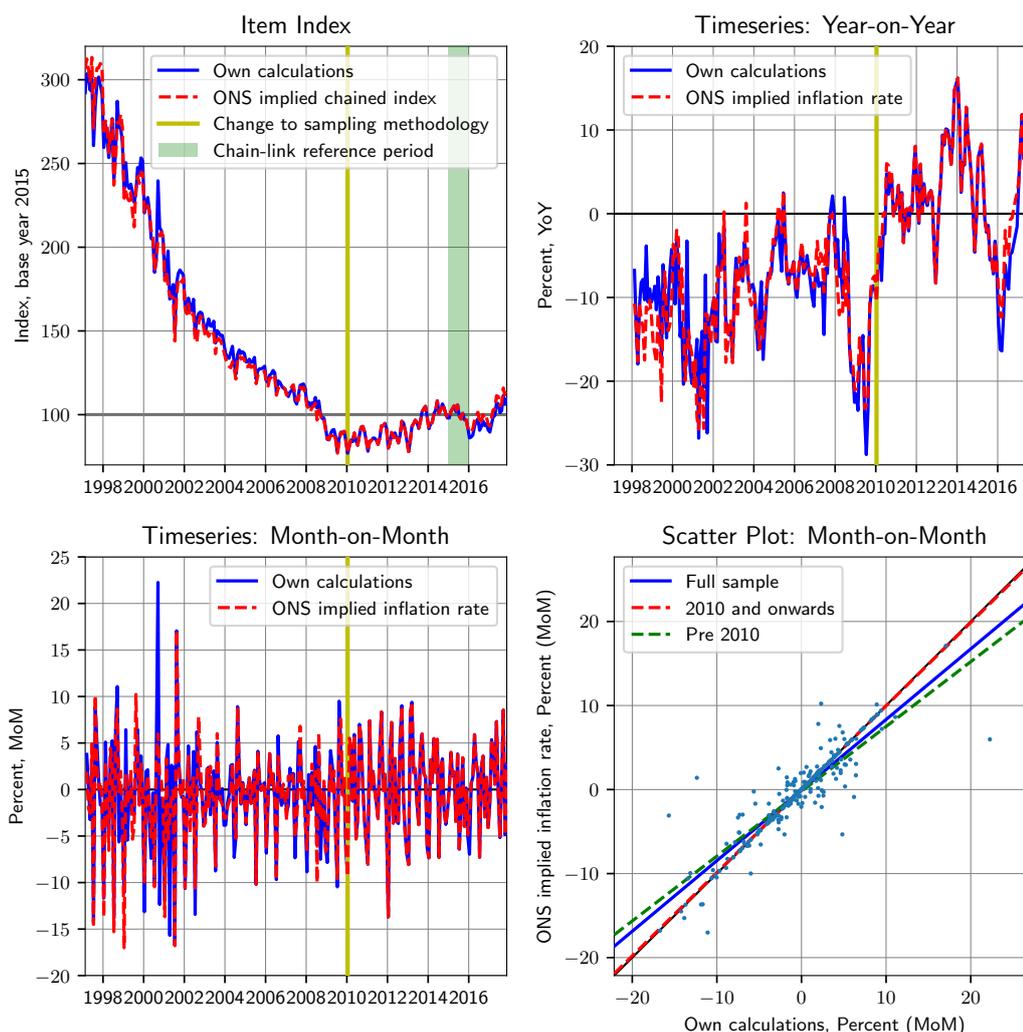
## **4.6 Appendix**

### **4.6.1 Comparison with the official data**

If I compare the implied inflation rate from the officially published data from the ONS for the “Women’s top, long sleeved, not blouse” item category I get a reasonably good fit as can be seen in figure 4.8. After the methodology changes in 2010 the officially published series from the ONS and my own calculated index have very similar properties except for one episode in January 2016, where there is a divergence between the two chain linked indices (with permanent effects thereafter for the level of the index).

As can be seen in the above table 4.4 for all the periods the constant term is

Figure 4.8: Fit of the inflation rate



Notes: item “Women’s top, long sleeved, not blouse” (id 510223).

not significantly different from zero but has a qualitative meaningful sign in terms of matching the property that a geometric mean series such as our index will have a lower volatility compared with an arithmetic counterpart such as used by the ONS. Hence, for the period when the the item category experience deflationary pressures (prior to 2010), the constant is negative, as the deflation from the geometric index is less negative than from the arithmetic index used by the ONS (and vice verse for the following period of inflationary pressures). After 2010 I cannot reject the hypothesis of a unity intercept and zero constant, which is also reflected in the very high  $r^2$ , and close fit of the series in figure 4.8. However, for the period preceding the methodological changes in 2010 (and as a consequence for the full sample), I can reject the hypothesis of a unitary intercept, as the intercept is significantly below 1 for this period. Hence, there is some difference between our calculated inflation rate and the inflation rate as implied from the item index published by the ONS.

Table 4.4: Regression fit of own calculations to the ONS implied inflation rate

	Full sample	Pre 2010	2010 and onwards
Constant	-0.07 (0.14)	-0.21 (0.22)	0.04 (0.04)
Intercept	0.84 (0.03)	0.77 (0.04)	0.99 (0.01)
$r^2$	0.81	0.73	0.99

Notes: Regression fit of item index of “Women’s top, long sleeved, not blouse” (id 510223) for the model  $\pi_t^{ons} = a + b\pi_t^{own} + e_t$ , where  $\pi_t^{ons}$  is the ONS implied inflation rate and  $\pi_t^{own}$  is our own calculations. Standard deviations of estimated coefficients in brackets.

Table 4.5: Moments of the series

	Full Sample	Pre 2010	2010 and onwards
ONS: Mean (St.Dev.)			
MoM	-0.40 (5.07)	-0.84 (5.20)	0.31 (4.80)
YoY	-5.27 (8.95)	-10.50 (6.13)	2.59 (6.41)
Own: Mean (St.Dev.)			
MoM	-0.40 (5.43)	-0.81 (5.74)	0.27 (4.82)
YoY	-5.31 (8.81)	-10.12 (6.14)	1.92 (7.12)
Correlations			
MoM	0.90	0.85	1.00
YoY	0.94	0.81	0.98

Notes: Properties of the inflation series for “Women’s top, long sleeved, not blouse” (id 510223).

## 4.7 Higher Aggregates: COICOP Classifications

### 4.7.1 COICOP Class

For future work it would be of interest to aggregate the potential sources of bias into higher levels of aggregation. In the following section I will sketch out the methodology to translate the index aggregations to the higher COICOP levels and how this maps into the framework of inflation decompositions. The first higher aggregate level is to use the item weights to aggregate from the item indices to the lowest COICOP level of “Class” indices. These weights are updated every February (together with the basket change), and used until the subsequent January.<sup>32</sup>

$$Z_{m,y,c}^{class} = \prod_{k \in \mathcal{I}_{m,y,c}^{class}} (Z_{m,y,k}^{item})^{w_{m,y,k}^{item}} \quad (4.7.1)$$

<sup>32</sup>For a small subset of the price quotes, historically some changes to the item weights took place throughout the year to adjust for seasonal variation in the consumption basket. However in recent years the ONS have stopped this procedure, with the argument that globalisation has eliminated such seasonal variation.

Again, I define a January based index as:

$$J_{m,y,c}^{class} = \begin{cases} 1 & m = 1 \\ Z_{m,y,c}^{class} & \forall m > 1 \end{cases} \quad (4.7.2)$$

Combining the two equations, and using the same double chain-linking methodology as earlier, I can calculate the COICOP class indices:

$$CPI_{m,y,c}^{class} = CPI_{12,y,c}^{class} \times \frac{Z_{1,y,c}^{class}}{Z_{12,y-1,c}^{class}} \times J_{m,y,c}^{class} \quad (4.7.3)$$

Except for an additional item weight, the methodology of the inflation decomposition follows readily from the above equation. A minor modification is that the item weights historically have had some seasonal variation which one would ideally want to split out as an individual component by itself.

#### 4.7.2 Higher COICOP Indices: Headline, Divisions and Groups

The COICOP weights are used to aggregate from the COICOP class indices to the higher categories of the COICOP classification (Groups, Divisions and Headline). The data for the COICOP weights are mainly from the Household Expenditure data of the National Accounts (Johnson, 2015, p.151). To follow Eurostats requirement, these weights are updated every January and used until December the same year.<sup>33</sup> The first step is to define the unchained headline Consumer Price Index as:

$$Z_{m,y} = \prod_{c \in \mathcal{T}_{m,y}^{cpi}} (Z_{m,y,c}^{class})^{w_{m,y,c}^{class}} \quad (4.7.4)$$

However, as the COICOP class weights are updated every January, to follow the methodology of Eurostat, I now need to introduce one additional feature to the double chain linking to adjust for this non-February weight changes. To do this, I first define the January based index as:

$$J_{m,y} = \begin{cases} 1 & m = 1 \\ Z_{m,y} & \forall m > 1 \end{cases} \quad (4.7.5)$$

and then introduce the double chain linking for the COICOP headline index level

---

<sup>33</sup>However, from 2017 and onwards, the COICOP weights are price updated in February, to bring them in line with the chain-linking procedure.



# Chapter 5

## Concluding Remarks

In this thesis I have demonstrated the benefits of using larger disaggregated datasets relative to the smaller aggregated sets usually used within macroeconomics. In chapter 2, using a large dataset of photo-interpreted points, we found evidence for an over restrictive land use policy across 27 European countries which is keeping the supply of residential housing below the social optimum. The political decisions for how to allocate land use permits seems to give priority specifically to house owners, evidencing what is often referred to as Not-In-My-Back-Yard behaviour. It thereby restricts the amount of land used for residential purposes to below the social optimum, as given from a Pigovian equilibrium of negative externalities. It is difficult to find a readily available and comparable measure of land use within countries, not to mention across different countries. Hence, our analysis was only made possible by using this large non-standard dataset of points which allowed us to derive a comparable measure of residential land across the European countries.

In chapter 3 on inflation dynamics in the UK, we use a data set of 27.5 million individual price quotes and find that the pricing mechanism of most UK firms displays secular movements within the period investigated. The secular trends in the distribution are especially seen in changes to the dispersion and frequency of reset price inflation. Both are movements that in the aggregate can cancel each other out, making it unobservable in the aggregate series. Indeed for some of the periods where we observe increased dispersion, a fall in the frequency of adjustments results in the dispersion between the aggregate division inflation rates fell. Further, we found evidence that VAT episodes have a significant effect in raising the flexibility of firms, making monetary policy less effective during these episodes. When price flexibility is high, any monetary policy shocks are passed through quicker by price changes. This highlight the importance of coordination between fiscal and monetary policy. Lastly, we found that both the Bank of England and professional forecasters can improve their inflation forecasting accuracy by taking into account the state of price flexibility in the economy.

In chapter 4 I explored the construction methods for the UK's Consumer Price Index, to investigate if evidence can be found towards any bias in the measurements of the inflation rate. I found that the inflation components of both missing observations and quality changes are significant drivers of the inflation rate, for a particular narrow category of goods as measured for "Women's Top, long sleeved not blouse". However, although a naive interpretation of this result is of a bias in the measurement, fashion cycles makes this ambiguous and an alternative interpretation is that it reflects a taste for newer products (fashion), rather than an error in measurement. For both chapter 3 and 4 by looking at the dataset in details at a highly disaggregated level allowed new results to be uncovered. Especially the investigation into the composition of the index makes us appreciate dynamics in the series otherwise not observed.

In conclusion, I find for all three chapters, that although for most work of a macroeconomist the aggregate series provides a good summary of the key dynamics of interest, there are underlying trends and dynamics that can only be captured from large disaggregated datasets. Further, non-standard datasets, such as the LUCAS points used in chapter 2, help us answer more relevant questions for policy and allows us to tackle some of the challenges facing society today. I would encourage more engagement with less regular datasets to try and expand what we can measure and how we measure it. Lastly, as always with any empirical discipline, being critical of your data and how it is constructed is always key. As my last chapter on inflation shows, although we have come a long way, there are still questions about whether we can find better ways to measure issues such as fashion products, and capture a more holistic view of the distribution for the underlying series of aggregated data.

# Chapter 6

## Bibliography

- ALTONJI, J. G. AND L. M. SEGAL (1996): “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business & Economic Statistics*, 14, 353–366. 28
- ALVAREZ, F., H. LEBIHAN, AND F. LIPPI (2016): “The Real Effect of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach,” *American Economic Review*, 106, 2817–2851. 3.1, 4, 3.4.1, 30
- ALVAREZ, F. AND F. LIPPI (2014): “Price Setting With Menu Cost for Multi-product Firms,” *Econometrica*, 82, 89–135. 4
- ALVAREZ, L. J., E. DHYNE, M. HOEBERICHTS, C. KWAPIL, H. L. BIHAN, P. LÜNNEMANN, F. MARTINS, R. SABBATINI, H. STAHL, P. VERMEULEN, AND J. VILMUNEN (2006): “Sticky Prices in the Euro Area: A Summary of New Micro-Evidence,” *Journal of the European Economic Association*, 4, 575–584. 4
- ANGRIST, J. D. AND J.-S. PISCHKE (2009): *Mostly Harmless Econometrics – An Empiricist’s Companion*, 41 William Street, Princeton, New Jersey 08540, US: Princeton University Press, first edition ed. 2.4.2
- AUERBACH, A. J. AND Y. GORODNICHENKO (2012): “Measuring the Output Responses to Fiscal Policy,” *American Economic Journal: Economic Policy*, 4, 1–27. 3.6
- BANK FOR INTERNATIONAL SETTLEMENTS, T. (2015): “Residential Property Statistics,” <http://www.bis.org/statistics/pp.htm>, accessed: 05-11-2015. 2
- BARRO, R. J. (1972): “A Theory of Monopolistic Price Adjustment,” *Review of Economic Studies*, 39, 17–26. 3.1, 3.3, 23
- BEAN, C. (2016): *Independent Review of UK Economic Statistics*, Independent Commision. 4.1, 4
- BERARDI, N., E. GAUTIER, AND H. L. BIHAN (2015): “More Facts about Prices: France Before and During the Great Recession,” *Journal of Money, Credit and Banking*, 47, 1465–1502. 4, 3.2.1
- BERGER, D. AND J. VAVRA (2017): “Dynamics of the U.S. Price Distribution,” *European Economic Review*, Forthcoming. 3.1, 2, 3.2.1, 3.4, 30

- BILS, M. (2009): “Do Higher Prices for New Goods Reflect Quality Growth or Inflation?” *The Quarterly Journal of Economics*, 124, 947–985. 3.2.1
- BILS, M. AND P. J. KLENOW (2004): “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, 947–985. 4
- BOSKIN, M., E. R. DULBERGER, R. J. GORDON, Z. GRILICHES, AND D. JORGENSEN (1996): *Final Report of the Advisory Commission to Study the Consumer Price Index*, Committee on Finance, United States Senate, One Hundred Fourth Congress, Second Session. 4.1
- (1998): “Consumer Prices, the Consumer Price Index, and the Cost of Living,” *Journal of Economic Perspectives*, 12, 3–26. 4.1
- BRYNJOLFSSON, E. AND A. MCAFEE (2014): *The Second Machine Age: Work, Progress, and Brilliant Technologies*, Castle House, 75/76 Wells Street, London W1T 3QT: W. W. Norton & Company, ebook ed. 1
- BUNN, P. AND C. ELLIS (2012): “Examining The Behaviour Of Individual UK Consumer Prices,” *Economic Journal*, 122, 35–55. 3.1, 17
- CABALLERO, R. J. AND E. M. R. A. ENGEL (1991): “Dynamic (S, s) Economies,” *Econometrica*, 59, 1659–1686. 3.5
- (1993a): “Microeconomic Adjustment Hazards and Aggregate Dynamics,” *The Quarterly Journal of Economics*, 108, 359–383. 3.4
- (1993b): “Microeconomic rigidities and aggregate price dynamics,” *European Economic Review*, 37, 697–711. 2
- (2007): “Price stickiness in Ss models: New interpretations of old results,” *Journal of Monetary Economics*, 54, 100–121. 3.1, 3.4, 3.5, 3.5, 31, 3.7
- CALVO, G. A. (1983): “Staggered prices in a utility-maximizing framework,” *Journal of Monetary Economics*, 12, 383–398. 25
- CAPLIN, A. S. AND D. F. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *The Quarterly Journal of Economics*, 102, 703–725. 25
- CARVALHO, C. AND O. KRYVTSOV (2017): “Price Selection,” Manuscript. 3.1
- CHERNOZHUKOV, V. AND H. HONG (2003): “An MCMC approach to classical estimation,” *Journal of Econometrics*, 115, 293–346. 3.8, 3.6, 3.8.5, 3.16
- CHESHIRE, P. AND S. SHEPPARD (2002): “The welfare economics of land use planning,” *Journal of Urban Economics*, 52, 242–269. 2.1.2
- CLEWS, G., R. SANDERSON, AND J. RALPH (2014): “Calculating a Retrospective Superlative Consumer Price Index for the UK,” *ONS*. 4.2.3
- COCHRAN, W. G. (1977): *Sampling Techniques*, Reading, Massachusetts: Addison-Wesley, third edition ed. 2.7.3

- COSTAIN, J. AND A. NAKOV (2011): “Distributional dynamics under smoothly state-dependent pricing,” *Journal of Monetary Economics*, 58, 646–665. 4
- D’AGOSTINO, A., D. GIANNONE, M. LENZA, AND M. MODUGNO (2016): “Nowcasting Business Cycles: A Bayesian Approach to Dynamic Heterogeneous Factor Models,” *Advances in Econometrics*, in: *Dynamic Factor Models*, 35, 569–594. 3.8.3
- DE VINCENT-HUMPHREYS, R. (2017): “RPI and CPI: A tale of two formulae,” Presentation, Advisory Panel on Consumer Prices (APCP) – Technical. 4.1
- DEGOOIJER, J. G. (2017): *Elements of Nonlinear Time Series Analysis and Forecasting*, Springer. 37
- DIEWERT, W. E. (2012): “Consumer Price Statistics in the UK,” Website link (last accessed 6. May 2017). 1, 4.1, 3
- DIXIT, A. (1991): “Analytical Approximations in Models of Hysteresis,” *Review of Economic Studies*, 58, 141–151. 3.3
- DIXON, H. D. AND H. LEBIHAN (2012): “Generalised Taylor and Generalised Calvo Price and Wage Setting: Micro-evidence with Macro Implications,” *Economic Journal*, 122, 532–554. 17
- DIXON, H. D., K. B. LUNTEL, AND K. TIAN (2014): “The impact of the 2008 crisis on UK prices: what we can learn from the CPI microdata,” Cardiff Economics Working Papers E2014/7, Cardiff University, Cardiff Business School, Economics Section. 3.1
- DIXON, H. D. AND K. TIAN (2017): “What We can Learn About the Behaviour of Firms from the Average Monthly Frequency of Price-Changes: An Application to the UK CPI Data,” *Oxford Bulletin of Economics and Statistics*, 79, 907–932. 17
- DÖPKE, J., M. FUNKE, S. HOLLY, AND S. WEBER (2005): “The cross-sectional dynamics of German business cycles: a bird’s eye view,” *Economic Studies* 23/2005, Deutsche Bundesbank. 3.2.2
- DORSEY, R. E. AND W. J. MAYER (1995): “Genetic Algorithms for Estimation Problems with Multiple Optima, Nondifferentiability, and Other Irregular Features,” *Journal of Business & Economic Statistics*, 13, 53–66. 28
- DOTSEY, M. AND R. G. KING (2005): “Implications of state-dependent pricing for dynamic macroeconomic models,” *Journal of Monetary Economics*, 52, 213–242. 4
- EICHENBAUM, M., N. JAIMOVICH, AND S. REBELO (2011): “Reference Prices, Costs, and Nominal Rigidities,” *American Economic Review*, 101, 234–262. 3.2.1
- EUROSTAT (2012): “Land cover/use statistics,” <http://www.ec.europa.eu/eurostat/web/lucas/overview/>, accessed: 05-11-2015. 2.1

- (2013): *LUCAS 2012 (Land Use / Cover Area Frame Survey) Technical Reference Document: C-3 Land use and Land Cover Classification (revised)*, European Commission. Eurostat. Directorate E: Sectoral and Regional statistics Unit E-4. 2.2.1, 2.7.3
- FEDERAL RESERVE BANK OF DALLAS, T. (2015): “International House Price Database,” <http://www.dallasfed.org/institute/houseprice/>, accessed: 05-11-2015. 2
- GAGNON, E. (2009): “Price Setting during Low and High Inflation: Evidence from Mexico,” *The Quarterly Journal of Economics*, 124, 1221–1263. 4, 3.2.2
- GAGNON, E., D. LOPEZ-SALIDO, AND N. VINCENT (2013): “Individual Price Adjustment along the Extensive Margin,” *NBER Macroeconomics Annual*, 27, 235–281. 3.1, 3.5, 3.5.1
- GELMAN, A., D. DUNSON, D. RUBIN, H. S. STERN, AND J. B. CARLIN (2004): *Bayesian Data Analysis*, Chapman & Hall/CRC Texts in Statistical Science. 3.8.5
- GERTLER, M. AND J. LEAHY (2008): “A Phillips Curve with an Ss Foundation,” *Journal of Political Economy*, 116, 533–572. 4
- GIANNONE, D., L. REICHLIN, AND D. SMALL (2008): “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of Monetary Economics*, 55, 665–676. 1
- GILCHRIST, S., R. SCHOENLE, J. SIM, AND E. ZAKRAJŠEK (2017): “Inflation Dynamics during the Financial Crisis,” *American Economic Review*, 107, 785–823. 34
- GOLOSOV, M. AND R. E. LUCAS (2007): “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 115, 171–199. 3.4
- GORDON, R. J. (2000): “The Boskin Commission Report and its aftermath,” *NBER Working Paper Series*, 7795. 5
- (2006): “The Boskin Commission Report: A Restrospective One Decade Later,” *NBER Working Paper Series*, 12311. 5
- (2016a): “Perspectives on The Rise and Fall of American Growth,” *American Economic Review*, 106, 1–7. 4.1
- (2016b): *The Rse and Fall of American Growth*, 6 Oxford Street, Woodstock, Oxfordshire OX20 1TW, UK: Princeton University Press, first edition ed. 4.1
- GRANGER, C. W. J. AND T. TERASVIRTA (1993): *Modelling Non-Linear Economic Relationships*, Oxford University Press. 3.6
- HARVEY, A. C. (1990): *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press. 3.6, 3.8.3, 3.8.5
- HAUSMAN, J. (1999): “Cellular Telephone, New Products, and the CPI,” *Journal of Business & Economic Statistics*, 17, 188–194. 4.1

- (2003): “Sources of Bias and Solutions to Bias in the Consumer Price Index,” *Journal of Economic Perspective*, 17, 23–44. 4.1
- HIGSON, C., S. HOLLY, AND P. KATTUMAN (2002): “The cross-sectional dynamics of US business cycle: 1950-1999,” *Journal of Economic Dynamics & Control*, 26, 1599–155. 3.2.2
- HIGSON, C., S. HOLLY, P. KATTUMAN, AND S. PLATIS (2004): “The Business Cycle, Macroeconomic Shocks and the Cross-Section: The Growth of UK Quoted Companies,” *Economica*, 71, 299–318. 3.2.2
- HOLLY, S., I. PETRELLA, AND E. SANTORO (2013): “Aggregate fluctuations and the cross-sectional dynamics of firm growth,” *Journal of the Royal Statistical Society: Series A*, 176, 459–479. 3.2.2
- INTERNATIONAL MONETARY FUND, T. (2015): “Global House Price Index,” <http://www.imf.org/external/research/housing/>, accessed: 05-11-2015. 2
- JENKINS, C. AND S. PATEL (2017): “Consumer price inflation, updating weights: 2017,” *ONS*. 34
- JOHNSON, P. (2015): *UK Consumer Price Statistics: A Review*, Independent Commission. 3, 4, 4.7.2
- KARADI, P. AND A. REIFF (2014): “Menu Costs, Aggregate Fluctuations, and Large Shocks,” Discussion Papers 10138, CEPR. 3.1, 4
- KEHOE, P. AND V. MIDRIGAN (2015): “Prices are sticky after all,” *Journal of Monetary Economics*, 75, 35–53. 3.2.1
- KLENOW, P. J. AND O. KRYVTSOV (2008): “State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?” *The Quarterly Journal of Economics*, 123, 863–904. 4, 14
- KOMUNJER, I. (2007): “Asymmetric power distribution: Theory and applications to risk measurement,” *Journal of Applied Econometrics*, 22, 891–921. 3.4.1
- KRYVTSOV, O. AND N. VINCENT (2017): “The Cyclicalities of Sales and Aggregate Price Flexibility,” Manuscript. 3.2.1
- LUCAS, R. E. (2001): “Externalities and Cities,” *Review of Economic Dynamics*, 4, 245–274. 2.3.4
- MARIANO, R. S. AND Y. MURASAWA (2003): “A new coincident index of business cycles based on monthly and quarterly series,” *Journal of Applied Econometrics*, 18, 427–443. 3.8.3
- MIDRIGAN, V. (2011): “Menu Costs, Multiproduct Firms, and Aggregate Fluctuations,” *Econometrica*, 79, 1139–1180. 3.1, 4
- MILES, D. (2012): “Population Density, House Prices and Mortgage Design,” *Scottish Journal of Political Economy*, 59, 444–466. 2.1.2

- MITCHELL, J., R. J. SMITH, M. R. WEALE, S. WRIGHT, AND E. L. SALAZAR (2005): “An Indicator of Monthly GDP and an Early Estimate of Quarterly GDP Growth,” *The Economic Journal*, 115, F108–F129. 46
- NAKAMURA, A. O., E. NAKAMURA, AND L. I. NAKAMURA (2011): “Price dynamics, retail chains and inflation measurement,” *Journal of Econometrics*, 161, 47–55. 4
- NAKAMURA, E. AND J. STEINSSON (2008): “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *The Quarterly Journal of Economics*, 123, 1415–1464. 4, 3.2.1, 14
- (2010a): “Monetary Non-neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 125, 961–1013. 3.4
- (2010b): “More Facts About Prices – Supplement to “Five Facts About Prices: A Reevaluation of Menu Cost Models”,” Online Supplement. 3.2.1
- NAKAMURA, E., J. STEINSSON, P. SUN, AND D. VILLAR (2018): “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation,” *Quarterly Journal of Economics*, Forthcoming. 4
- NORDHAUS, W. D. (1996): “Do Real-Output and Real-Wage Measures Capture Reality? The History of Lightning Suggests Not,” in *The Economics of New Goods*, ed. by T. F. Bresnahan and R. J. Gordon, University of Chicago Press, chap. 1, 27–70. 4.1
- (1998): “Quality Change in Price Indexes,” *Journal of Economic Perspectives*, 12, 59–68. 4.1
- OECD (2015): “Focus on house prices,” <http://www.oecd.org/eco/outlook/focusonhouseprices.htm>, accessed: 01-12-2015. 2
- O’NEIL, R., J. RALPH, AND P. A. SMITH (2017): *Inflation – History and Measurement*, Palgrave Macmillan. 15
- ONS (2010): “CPI and RPI: increased impact of the formula effect in 2010,” Website link (last accessed 15. July 2018). 18
- (2011): “Responding to the Open Data Agenda: An ONS case study,” Website link (last accessed 6. May 2017). 4.2.1
- (2014): *Consumer Price Indices: Technical Manual*, Office for National Statistics, 2014 ed. 10, 9, 4.2.3, 4.3.2, 31
- PAYNE, C. (2016): “Assessing the Impact of methodological improvements on the Consumer Prices Index,” *ONS*. 34
- PESARAN, H. H., R. P. SMITH, AND K. S. IM (1996): “Dynamic Linear Models for Heterogenous Panels,” in *The Economics of Panel Data: A Handbook of Theory with Applications*, ed. by L. Mátyás and P. Sevestra, Dordrecht: Kluwer. 2.2, 2.4.1, ??

- PESARAN, M. H. (2015): *Time Series and Panel Data Econometrics*, Great Clarendon Street, Oxford, OX2 6DP, United Kingdom: Oxford University Press, first edition ed. 2.4.1
- ROTEMBERG, J. J. (1982): “Sticky Prices in the United States,” *Journal of Political Economy*, 90, 1187–1211. 3.3
- (1999): “A Heuristic Method for Extracting Smooth Trends from Economic Time Series,” NBER Working Papers 7439, National Bureau of Economic Research, Inc. 3.4, 3.8.3, 3.11
- STOCK, J. H. AND M. W. WATSON (2002): “Macroeconomic Forecasting Using Diffusion Indexes,” *Journal of Business & Economic Statistics*, 20, 147–162. 1
- VARIAN, H. R. (2014): “Big Data: New Tricks for Econometrics,” *Journal of Economic Perspectives*, 28, 3–28. 1
- VAVRA, J. (2014): “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly Journal of Economics*, 129, 215–258. 3.1, 14, 19, 3.2.2, 3.3, 23, 3.3
- VEGA, S. H. AND J. P. ELHORST (2013): “On spatial econometrics models, spillover effect and W,” Paper presented at 53rd ERSa conference, Palermo. 2.4.1
- VILLAR, D. V. AND S. LUO (2017): “The Skewness of the Price Change Distribution: A New Touchstone for Sticky Price Models,” Finance and Economics Discussion Series 2017-028, Board of Governors of the Federal Reserve System (U.S.). 22
- WELLS, J. AND A. RESTIEAUX (2014): “Review of Hedonic Quality Adjustment in UK Consumer Price Statistics and Internationally,” *ONS*. 30