



**Raphael de Aquino Ludwig Pereira**

**On the Missing Disinflation Puzzle: A  
Data-Driven Approach**

**Dissertação de Mestrado**

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

Advisor : Prof. Eduardo Zilberman  
Co-advisor: Prof. Marcelo Cunha Medeiros

Rio de Janeiro  
August 2021



**Raphael de Aquino Ludwig Pereira**

**On the Missing Disinflation Puzzle: A  
Data-Driven Approach**

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

**Prof. Eduardo Zilberman**

Advisor

Departamento de Economia – PUC-Rio

**Prof. Marcelo Cunha Medeiros**

Co-advisor

Departamento de Economia – PUC-Rio

**Prof. Márcio Garcia**

Departamento de Economia – PUC-Rio

**Dr. Gabriel Vasconcelos**

Banco BBM

Rio de Janeiro, August the 20th, 2021

All rights reserved.

**Raphael de Aquino Ludwig Pereira**

B.A., Economics, Pontifícia Universidade Católica do Rio de Janeiro, 2017

Bibliographic data

Pereira, Raphael de Aquino Ludwig

On the Missing Disinflation Puzzle: A Data-Driven Approach / Raphael de Aquino Ludwig Pereira; advisor: Eduardo Zilberman; co-advisor: Marcelo Cunha Medeiros. – Rio de Janeiro: PUC-Rio , Departamento de Economia, 2021.

v., 65 f: il. color. ; 30 cm

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

Inclui bibliografia

1. Economia – Teses. 2. Economia – Teses. 3. Inflação;. 4. Machine Learning;. 5. Avaliação de Previsão;. 6. Model Confidence Set;. 7. Seleção de Variáveis;. I. Zilberman, Eduardo. II. Medeiros, Marcelo Cunha. III. Pontifícia Universidade Católica do Rio de Janeiro. Departamento de Economia. IV. Título.

## Acknowledgments

To my advisors Eduardo Zilberman and Marcelo Medeiros. I am very grateful for your patience and attention. You guided and disciplined my research endeavors. Your support and insightful contributions made this project possible.

To my mother Monica and sister Dani. Thank you for always believing in me. Your endless support and unconditional love pushed me forward whenever it was necessary.

To my grandparents Sérgio e Clarita. If not for your support and affection I would never have succeeded. I will never be able to compensate what you did for me.

To my friends, specially Dane and Duda. Thank you for the companionship and fondness through all these years. It is breathtaking how many times you helped me with my adversities and problems.

To PUC-Rio. Where I first started my trajectory in Economics in 2013 and since then fell more and more in love with the place. I'll carry the fond memories of such a happy time forever.

To my father Ruy 'Lula' Ludwig. I would trade everything to have you here with me in this moment. I miss you.

To Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) who financed my Master's studies and allowed me to dedicate myself entirely to Research. The financial support from PUC-Rio is also gratefully acknowledged.

## Abstract

Pereira, Raphael de Aquino Ludwig; Zilberman, Eduardo (Advisor); Medeiros, Marcelo Cunha (Co-Advisor). **On the Missing Disinflation Puzzle: A Data-Driven Approach**. Rio de Janeiro, 2021. 65p. Dissertação de mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

This paper examines the potential explanations for the Missing Disinflation Puzzle (MDP). We construct a data set containing only variables associated with the puzzle, and use of Machine Learning (ML) methods to compute estimates for U.S. Consumer Price Index inflation over the period of interest. These methods can handle large data sets, and perform variable selection. A model selection exercise using Model Confidence Set over pseudo-out-of-sample forecasts is proposed to assess forecasting performance and to analyze the variable selection pattern of these models. We analyze the variable selection performed by the best models and find evidence for explanations associated with different metrics for inflation expectations - in particular those linked to consumers' surveys.

## Keywords

Inflation; Machine Learning; Forecast Evaluation; Model Confidence Set; Variable Selection;

## Resumo

Pereira, Raphael de Aquino Ludwig; Zilberman, Eduardo; Medeiros, Marcelo Cunha. **Sobre o Missing Disinflation Puzzle: Uma Abordagem com Aprendizado de Máquina.** Rio de Janeiro, 2021. 65p. Dissertação de Mestrado – Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro.

O presente trabalho investiga as potenciais explicações para o fenômeno do Missing Disinflation Puzzle. Nós montamos uma base de dados contendo apenas variáveis associadas com o fenômeno, e utilizamos métodos de Machine Learning para calcular estimativas para a inflação do Consumer Price Index durante o período de interesse. Esses métodos podem lidar com bases de dados extensas, e realizar seleção de variáveis. Um exercício de seleção de melhores modelos utilizando a técnica de Model Confidence Set sobre previsões pseudo out-of-sample é proposto. Nós analisamos o padrão de seleção de variáveis entre os melhores modelos selecionados e encontramos evidência a favor das explicações associadas ao uso de diferentes métricas de expectativas de inflação - em especial aquelas ligadas a pesquisas feitas com consumidores.

## Palavras-chave

Inflação; Machine Learning; Avaliação de Previsão; Model Confidence Set; Seleção de Variáveis;

## Table of contents

1	Introduction	<b>10</b>
1.1	Missing Disinflation Puzzle	10
1.2	Contribution	11
1.3	Organization	12
2	Data	<b>13</b>
2.1	Data Description	13
2.2	Database Construction	13
2.2.1	Inflation	14
2.2.2	Expectations	15
2.2.3	Activity	16
2.2.4	Economic Shocks and Financial Variables	18
2.3	Database Specifications	19
3	Methodology	<b>20</b>
3.1	General Framework	20
3.2	Model Confidence Set	22
3.3	Models	23
3.3.1	Autoregressive Models	23
3.3.2	UC-SV	24
3.3.3	Shrinkage Models	24
3.3.4	Random Forest	26
3.3.5	Hybrid Linear-Random Forest Models	28
4	Empirical Results	<b>29</b>
4.1	Reducing the Dimensionality	29
4.2	Best Specifications' Selection	34
4.3	Variable Selection and the Missing Disinflation Puzzle	37
4.3.1	The Sharp Fall of 2008 Q4	41
5	Concluding Remarks	<b>43</b>
A	Data	<b>48</b>
B	Random Forest	<b>52</b>
C	The Sharp Fall of 2008 Q4 - Details	<b>57</b>
D	Computer Codes	<b>65</b>

## List of figures

Figure 1.1	CPI Inflation and U.S. Unemployment	11
Figure 3.1	Forecasting Methods	22
Figure 4.1	Amount of Chosen Models per Family of Specifications	32
Figure 4.2	$\mathcal{M}_{\alpha=0.95}^*$ Specifications Predictive Performance	35
Figure 4.3	$\mathcal{M}_{\alpha=0.95}^*$ Correlation Plot	36
Figure 4.4	$\mathcal{M}_{\alpha=0.95}^*$ Accumulated RMSE Dynamic	37
Figure 4.5	LASSO and adaLASSO Variable Selection per Groups	38
Figure 4.6	adaLASSO Coefficients' Values Evolution over Time	39
Figure 4.7	LASSO Coefficients' Values Evolution over Time	40
Figure B.1	Mean Decrease in Accuracy I	54
Figure B.2	Mean Decrease in Accuracy II	55
Figure B.3	Mean Decrease in Accuracy III	56
Figure C.1	Regressors' Impact on the 2008 Q4 Forecast for the adaLASSO Model	57
Figure C.2	Regressors' Impact on the 2008 Q4 Forecast for the adaLASSO Model	58
Figure C.3	Chosen adaLASSO Specification Predictive Performance With Selected Variables	59
Figure C.4	Chosen adaLASSO Specification Predictive Performance With Selected Variables	60
Figure C.5	Linear Model Predictive Performance With Selected Variables	61
Figure C.6	Models' Specifications According to Number of Variables Utilized I	62
Figure C.7	Models' Specifications According to Number of Variables Utilized II	63
Figure C.8	adaLASSO Specification versus Selected Model	64

## List of tables

Table 3.1	Shrinkage Models	25
Table 4.1	Shrinkage Models: Top 10 Specifications	30
Table 4.2	SCAD-Penalty Models: Top 10 Specifications	30
Table 4.3	Random Forest: Top 10 Specifications	31
Table 4.4	Shrinkage Family: Specifications Selected for $\alpha = 0.95$	33
Table 4.5	SCAD-Penalty Family: Specifications Selected for $\alpha \in \{0.9, 0.95\}$	33
Table 4.6	random Forest Family: Specifications Selected for $\alpha \in \{0.9, 0.95\}$	33
Table 4.7	Specifications in $\mathcal{M}_{\alpha=0.95}^*$	34

# 1 Introduction

## 1.1 Missing Disinflation Puzzle

The Great Recession was one of the most severe financial crisis in the modern era. The period was marked by a general and sharp contraction in economic activity across the world. Inflation behaviour, however, was quite puzzling during this time. According to simple linear models, - as explored in Coibion and Gorodnichenko (2015) - prices should have fallen much more than what actually happened. In reality, inflation was relatively stable. Advanced economies, for example, experienced little decline in inflation during the recession and subsequent years. Many authors, like Hall (2011) and Powell (2019), started calling into question one of the fundamental tenets of macroeconomic theory: the Phillips curve linking the rate of change in prices to the level of economic activity. There was understandably an inability of this framework to explain what soon was called the Missing Disinflation Puzzle (MDP). The central message was relatively clear: the prices' dynamics were not behaving as expected.

To illustrate this phenomenon, we reproduce the exercise made in Coibion and Gorodnichenko (2015). Figure 1.1 illustrates it. We present a scatter plot of quarterly unemployment rates and the deviations of the Consumer Price Index (CPI) inflation from expected inflation. Here, expected inflation is represented as backward-looking expectations in the form of the average of the last 4 observations before time  $t$ . The solid line represents the slope of the average relationship between unemployment and inflation surprises over 1960 Q1 until 2007 Q3.

It is notable that there is a disproportionate amount of observations in the far upper-right part of the graph after 2009. For the authors, this period from the second quarter of 2009 until the end of 2011 represents the MDP, a period in which inflation was well above what might have been expected given the severity of the economic downturn. As shown in the study, this result is robust for different metrics of economic activity and inflation expectations.



Figure 1.1: CPI Inflation and U.S. Unemployment

Economists have suggested a number of possible explanations for the missing disinflation. For example, Bernanke (2010) argue that the credibility of modern central banks has convinced people that neither high inflation nor deflation are likely outcomes (anchored expectations hypothesis). Gordon (2013), in its turn, argue that the "triangle model", which incorporates explicit variables for supply shocks, can successfully track inflation behavior during this period.

Others studies, like the aforementioned Coibion and Gorodnichenko (2015) and Binder (2015), focus on the role of different kinds of survey expectations. Additionally, works in the tradition of Llaudes (2005) and Stock and Watson (2010) highlight the importance of different economic activity variables. Lastly, Del Negro et al. (2015) and Gilchrist et al. (2017) shed light on the role of financial frictions in explaining the missing disinflation.

In short, there is a vast group of possible explanations for the MDP, which sometimes go in different directions. Understanding price dynamics is an important goal for economists, therefore being able to disentangle these possible explanations is key for future analysis of inflation.

## 1.2 Contribution

We find evidence that inflation expectations associated to households - along with a lagged inflation component and the FED inflation projections - are key to predict inflation during the desired period of time. This result is in

line with the conclusions of Coibion and Gorodnichenko (2015), Binder (2017) and Binder (2015).

We construct a data set focused exclusively on variables used in the investigation of the puzzle and exploring it. Our main goal is to disentangle the puzzle and generate evidence for potential explanations explored by different authors. We do that by using several Machine Learning methods in a completely data-driven way.

### **1.3 Organization**

In Section 2, we describe the construction of the data set used in details and review the literature. Section 3 presents the general framework and the methodology used in this work. We also describe the models utilized and the forecasting scheme. In Section 4, we present our empirical results, with forecasts evaluation and the analysis of selected variables' pattern. Section 5 concludes the study.

## 2 Data

### 2.1 Data Description

Our target variable is the U.S. Consumer Price Index inflation (CPI inflation) in quarter-over-quarter format without any transformation. The majority of articles use this metric as default in the Missing Disinflation Puzzle (MDP) investigation. Our data set covers the period 1986 Q3-2011 Q4 (102 observations) and it is in quarterly frequency. The values of monthly series are averaged over the quarter.

The data set of regressors presented in this article consists only of variables used across the literature in the investigation of the MDP. It is split into 4 different groups: **Inflation** (7 variables), **Expectations** (44 variables), **Economic Activity** (23 variables), and **Economic Shocks and Financial Variables** (19 variables). After missing data removal, we remain with 93 variables.

We apply augmented Dickey-Fuller tests to check each series' stationarity in each possible window utilized in the expanding window procedure (Section 3.1) and transform them when necessary. Some explanatory variables also appear in the database described in McCracken and Ng (2016). When possible, we perform the same transformations realized in the aforementioned article.

Appendix A summarizes all information regarding the data set, e.g. transformations performed, time frequency and period covered.

### 2.2 Database Construction

In this section, we go further into details on the reasons behind the choice for each regressor. The main idea here is that the data set must consist only of variables explicitly used in articles which tackle the MDP in some way. We also gather the variables used for robustness checks in these studies.

This section is divided in 4 subsections, each one of them covering a specific variables' group.

### 2.2.1 Inflation

As mentioned before, our target variable is the U.S. CPI inflation which is the same response variable employed in key articles - for example Coibion and Gorodnichenko (2015), Blanchard et al. (2015), and Ball and Mazumder (2011).

We also include in the data set Personal Consumption Expenditures (PCE) - and wage inflation measures explored by the literature. We utilize them as explanatory variables in our models' specifications. Gordon (2013) and Van Zandweghe (2019), e.g., use PCE inflation and its lags in their Phillips curve (PC) models. Coibion and Gorodnichenko (2015) and Binder (2015) also employ this metric in robustness checks for their results.

For wage inflation, we make use of the quarterly change rate in average hourly earnings of manufacturing and construction workers, compensation per hour in the nonfarm business sector, and average hourly earnings of the goods-producing sector. The first three are utilized in Coibion and Gorodnichenko (2015) to generate evidence for the nonexistence of a missing wage disinflation puzzle during the Great Recession. The latter appeared in Heise et al. (2020) as the authors examined puzzling inflation dynamics during and after the Great Recession. These earnings metrics are also explored in Donayre and Panovska (2018).

We estimate a permanent component for the CPI inflation with an Unobserved Components-Stochastic Volatility (UC-SV) model following Stock and Watson (2010) and Medeiros et al. (2019). Ball and Mazumder (2011) and Ball and Mazumder (2014) both highlight the explanatory power of this metric.

Moreover, Stock and Watson (2010) points out that this estimated permanent component of inflation move together very closely with the median 5-year ahead inflation reported in the Survey of Professional Forecasters (SPF) since 2007. They suggest that this metric can also be thought as a measure of long-term inflation expectation, which has been quite stable during the period studied in this article.

Bernanke (2010) emphasizes the role of this stability of longer-term inflation expectations in keeping inflation stable. Del Negro et al. (2015) provide formal support for this view. The authors include data on long-run inflation expectations as an observable variable in the estimation of a Dynamic Stochastic General Equilibrium (DSGE) model. They show that a standard DSGE model with financial frictions successfully predicts a sharp

contraction in economic activity along with a protracted but relatively modest decline in inflation. In this article, however, we employ the estimated inflation's permanent component as an inflation - and not expectation - metric. That said, we can also think of it as a proxy for the SPF 5-year ahead inflation expectations.

Even though many articles entangled with the MDP utilize core inflation measures as the response variable, in this article we decided to use it as an economic shock component (further details in Subsection 2.2.4).

### 2.2.2 Expectations

Coibion and Gorodnichenko (2015) suggest that the missing disinflation during the Great Recession can be explained by the rise of 1 year ahead households' inflation expectations computed by the University of Michigan Survey of Consumers. The authors also argue that this variable serves as a better proxy for firms' inflation expectations than professionals' expectations<sup>1</sup>. This is important, because firms' expectations are the relevant metric for inflation expectations in traditional New-Keynesian (NK) models (see Galí and Gertler (1999)).

Binder (2015) and Binder (2017) further investigate this idea by performing the same regressions of the latter study with slightly different metrics for households' inflation expectations. The first article explores the different subgroups within the Survey of Consumers. The latter introduces a micro-level measure of consumer inflation uncertainty based on behavioral psychology literature on cognition and communication. With this measure, the author estimates expectations of less- and high-uncertain consumers.

Fuhrer et al. (2012) find that US inflation from 1990 to 2010 is well described by expectations-augmented PC models that explores the 1 year ahead median SPF inflation expectation. McLeay and Tenreyro (2019) in turn discusses the optimal identification of the PC using US data, focusing on different time periods and exploring the MDP. When estimating an empirical PC, the authors test different specifications. The majority of them utilize SPF inflation expectations or the combination of SPF inflation expectations and lags of inflation.

Ball and Mazumder (2014) perceive the SPF-based expectations as a proxy for anchored inflation expectations and use it in their models. Coibion

<sup>1</sup>For a great study about characteristics of firm-expectations using survey-based data, see Coibion et al. (2018b)

and Gorodnichenko (2015), on the other hand, show that this metric cannot account for the MDP<sup>2</sup>.

The majority of the articles revised, however, employ some metric of backward-looking expectations, i.e., expected inflation being determined by past inflation. This metric typically appears as a function of inflation lags among the regressors. Gordon (2013), e. g., try to capture an inertia component of inflation by adding one lag of it in their models.

Following the tradition of Atkeson and Ohanian (2001), Ball and Mazumder (2011) assume that expected inflation is the average of inflation in the past 4 quarters:

$$\pi_t^{\text{BL}} = \frac{1}{4} (\pi_{t-1} + \pi_{t-2} + \pi_{t-3} + \pi_{t-4})$$

where  $\pi$  is inflation in quarterly frequency.

Other works combine different forms of expectations in their models. Doser et al. (2018), among others, utilize a combination of 5 lags of inflation and forward-looking 1 year ahead SPF survey-based inflation expectations.

Lastly, Coibion and Gorodnichenko (2015) use two different survey-based metrics for robustness exercises: Philadelphia FED's GreenBook Data Set CPI inflation expectations, and 1-year ahead Cleveland FED's financial markets inflation expectations calculated as in Haubrich et al. (2012).

We get the median and mean inflation expectation for the Survey of Consumers' expectations. For the SPF and the GreenBook surveys, the median 1-year, 1-quarter, 2-quarters, and 3-quarters ahead are selected. Lastly, the metrics developed by Binder (2015), and the 1-year ahead inflation expectation calculated by the Cleveland FED are also included.

### 2.2.3 Activity

As a measure of economic slack, researchers historically used the unemployment rate or the unemployment gap, i.e., the distance between the current unemployment rate and the estimated natural unemployment rate, to cover the cyclical component of the relationship between inflation and economic activity. Other standard metrics explored in the economic literature are the real GDP and its gap, i. e., the distance between real GDP and the potential output.

<sup>2</sup>For an extensive review of all forms of survey-based inflation expectations computed in the US, and their role in key papers and as policy tool, see Coibion et al. (2018a) and Coibion et al. (2020)

All these measures appear across a wide range of articles studied for this work (Blanchard et al. (2015), Donayre and Panovska (2018), Ball and Mazumder (2011) to cite a few). Llaudes (2005), however, had already shed light on important distinctions in the dynamics of Phillips curve models with short- and long-term unemployment rates (27+ weeks unemployed)<sup>3</sup>.

Since one of the most unique aspects of the US labor market during and after the Great Recession was persistent long run unemployment, many authors perceive in this distinction a potential solution for the MDP. Consequently, it had an impact on the literature with articles like Ball and Mazumder (2014), Coibion and Gorodnichenko (2015) and Gordon (2013) utilizing short-term unemployment and the short-term unemployment gap as the activity variable in their models.

Real GDP, unemployment rate (full, short-term and long-term) and their gap measures are all mapped into the data set. All natural rates and potential output measures are retrieved from the Congressional Budget Office (CBO).

Furthermore, some authors suggest possible transformations using these variables. Stock and Watson (2010) propose the unemployment recession gap (also explored in Ball and Mazumder (2011) and Ball and Mazumder (2014)):  $u_t^{RG} = u_t - \min(u_t, \dots, u_{t-11})$ , where  $u$  is the unemployment rate. The authors also utilize the four quarter change in unemployment,  $(u_t - u_{t-4})$ , and other metrics like total capacity utilization, real GDP, and the Chicago Fed National Activity Index - all of them also subject to the recession gap transformation.

Laurence Ball and Sandeep Mazumder, in turn, investigate in their sequence of articles some other transformations, like the full unemployment gap relative to the level of unemployment (introduced in Debelle and Laxton (1997)) and the unemployment gap averaged over the previous 4 quarters.

In the NK tradition, one can use microfounded models to derive a relationship between inflation and marginal costs. Gali and Gertler (1999), e. g., argue that the real unit labor cost, commonly measured by labor share of income, is the superior forcing variable because of strong contemporaneous correlation with inflation. Doser et al. (2018) follows this path. The authors estimate their models mainly with the full unemployment gap, but also report estimates utilizing the labor share of income.

Van Zandweghe (2019) expand on Del Negro et al. (2015) NK model. In addition to a different way to model firms' behaviour, the authors model marginal costs as the log deviation of real unit labor costs from a linear

<sup>3</sup>See Krueger et al. (2014) for a deep study on the profile of each kind of unemployed

trend plus intercept (instead of simply using real unit labor costs without any transformation).

All the aforementioned measures are contemplated in the database. We retrieve the labor share of income from the U.S. Bureau of Labor Statistics (BLS) and the real unit labor costs from the U.S. Bureau of Economic Analysis (BEA).

#### 2.2.4 Economic Shocks and Financial Variables

Many studies highlight the importance of supply shocks to explain the MDP phenomena. Gordon (2013), e.g., uses the triangle model approach for the Phillips curve - a Phillips curve which depends on 3 elements: inertia, demand and supply shocks - to model the missing disinflation. In their framework, three metrics for shocks appear explicitly in the PC equation: relative price of non-food non-oil imports (defined as the rate of change of the non-food non-oil import minus the rate of change of CPI inflation); change in productivity trend (a Hodrick-Prescott filter trend in labor productivity minus the value of that trend eight quarters earlier); and the relative price food-energy price (defined as the difference between CPI inflation and CPI core).

The subtraction of core inflation metrics from the headline CPI inflation (introduced by Ball and Mankiw (1995)) is a shock measure repeatedly used in this literature. Stock and Watson (2010) and Coibion and Gorodnichenko (2015) underscore the importance of core inflation metrics to filter supply shocks. Ball and Mazumder (2011) and Ball and Mazumder (2014) point out that, instead of core inflation, the median CPI inflation calculated by Cleveland's FED should be used to expurgate supply shocks.

In addition to aforementioned variables, Blanchard et al. (2015) append the import price inflation relative to headline inflation to their models.

Coibion and Gorodnichenko (2015) also call attention to the importance of the link between temporary surge in oil prices, which reached its peak in mid 2008, and inflation expectations. They argue that high energy prices (and food prices to a lesser extent) raised short-term inflation expectations, putting upward pressure on inflation that offset the downward pull from the activity slack. Stock and Watson (2010) comments on the oil-price pass through too. We get both real and nominal prices for the database.

McLeay and Tenreyro (2019) - in their investigations of the optimal Phillips Curve identification and its possible flattening during the Great

Recession - mention supply shocks as one possibility to recover the structural PC. Change in oil prices and housing prices are mentioned by the authors.

Financial market stress surged during the Great Recession, raising borrowing costs and curtailing access to credit for some firms. As a result, authors who tried to solve the MDP exploring a NK framework added different kinds of financial frictions in their models.

Christiano et al. (2015) show that elevate interest rate spreads during the financial crisis put upward pressure on inflation through increased costs via working capital channel, explaining the absence of some kind of disinflation. Furthermore, Gilchrist et al. (2017) note that firms with ample liquidity lowered prices in 2008, while those with limited liquidity raised prices to avoid costly external finance.

Gilchrist and Zakrajsek (2012), in its turn, emphasize the predictive power of corporate bond spreads for inflation and other macroeconomic variables. Doser et al. (2018) test whether these measures of the state of financial markets can explain the MDP.

By trying to cover these financial frictions channels in their VAR model, Bobeica and Jarociński (2019) suggest some variables like the Gilchrist and Zakrajsek (2012) excess bond premium and credit spread measures, medium- and long term interest rate spread, and mortgage bank lending spreads. We construct and gather all these shock metrics for the data set.

### 2.3 Database Specifications

Throughout this work we split the data set into 3 different subsets:

- **Flexible:** Contains all collected variables;
- **Semi-Parsimonious:** The Flexible data set excluding the subdivisions of the University of Michigan households' inflation expectations
- **Parsimonious:** The Semi-Parsimonious data set excluding SPF- and GreenBook inflation expectations not in 1 year ahead format. Furthermore, we exclude from this subgroup some financial variables used in Bobeica and Jarociński (2019), which do not have a explicit link with the literature associated with the MDP.

We use the 3 different data sets in all exercises.

## 3 Methodology

### 3.1 General Framework

We use the following framework developed in Medeiros et al. (2019):

$$y_{t+h} = G_h(\mathbf{x}_t) + u_{t+h} \quad (3-1)$$

where  $y_{t+h}$  is the response variable at time  $t + h$ ;  $\mathbf{x}_t = (x_{1t}, \dots, x_{nt})'$  is a  $n$ -vector of covariates at time  $t$ , containing lags of  $y_t$  and the set of predictors and their lags following the different specifications of the database;  $G_h(\cdot)$  is the mapping between the regressors and future values of the response variable; and  $u_t$  is a zero-mean random error. The target function  $G_h(\mathbf{x}_t)$  assume different mappings at each forecasting horizon  $h$ . These mappings, however, follow the same chosen model type.

The direct forecasting equation is given by

$$\hat{y}_{t+h|t} = \hat{G}_{h,t-R_h+1:t}(\mathbf{x}_t) \quad (3-2)$$

where  $\hat{G}_{h,t-R_h+1:t}$  is the estimated target function based on data from time  $t - R_h + 1$  up to  $t$ , and  $R_h$  is the window size. For each forecasting horizon, we estimate a new model. Therefore, we don't predict any covariate and always use only information available at time  $t$ .

We use this framework to evaluate the 1 step ahead predictions of the US CPI inflation across the Great Recession and following quarters, covering the period of 2007 Q4 to 2011 Q4 (the same time period studied in Coibion and Gorodnichenko (2015)). After each single prediction, we compute forecast error metrics, update the database to the current quarter, and predict the next step. Thereby, the data set is always up-to-date.

In order to access the predictive power of the models, we calculate the squared- and the absolute error for each predictive step. In the end we aggregate them, obtaining the root-mean-square error (RMSE) and the mean-absolute error (MAE) for each run.

$$RMSE = \sqrt{\frac{\sum_{t=1}^T (y_i - \hat{y}_i)^2}{T}} \quad (3-3)$$

$$MAE = \frac{\sum_{t=1}^T y_t - \hat{y}_t}{T} \quad (3-4)$$

We estimate each model using every one of the 3 data set's configurations (Flexible, Parsimonious, and Semi-Parsimonious) individually in 4 different temporal specifications:  $y_t = G_h(\mathbf{x}_{\{t-1\}}, \beta) + u_{t+h}$ ,  $y_t = G_h(\mathbf{x}_{\{t-1, t-2\}}, \beta) + u_{t+h}$ ,  $y_t = G_h(\mathbf{x}_{\{t-1, t-2, t-3\}}, \beta) + u_{t+h}$ , and  $y_t = G_h(\mathbf{x}_{\{t-1, t-2, t-3, t-4\}}, \beta) + u_{t+h}$ . It's particularly important to have these contrasting models' designs, because studied authors specify their models with different lags for the target and independent variables.

The sample used for the forecasts are constructed in two different ways: rolling window and expanding window. Assume that the database consists of  $T$  observations. Define  $t_0 < T$  as the last period used for estimating the model and  $W$  as the window's size. Expanding window procedure follows the below steps:

- for  $t_0 = 1, \dots, T - 1$ 
  - Estimate the model using the  $1, \dots, t_0$  observations;
  - Compute the one-step ahead forecast and update the database

Rolling window procedure, on the other hand, follow these other steps:

- for  $t_0 = 1, \dots, T - 1$ 
  - Estimate the model using the  $t_0 - W + 1, \dots, t_0$  observations;
  - Compute the one-step ahead forecast and update the database

We represent both processes as scheme of Figure 3.1.

The main reason for using the rolling window procedure is that often recent lags have a higher predictive power than older ones. Therefore, a model that is applied on the whole series could belie several characteristics of the whole series, because some variables might have significant changes on their dynamics over longer periods of time. In this exercise, we use a window of 10 years. Nevertheless, we generally find better results with the expanding window's specifications.

In the end, the forecasts are estimated using 2 forecasting window procedures with samples constructed using 3 different data sets' organizations in 4 different temporal arrangements. Additionally, all models always use 3 lags of the dependent variable.

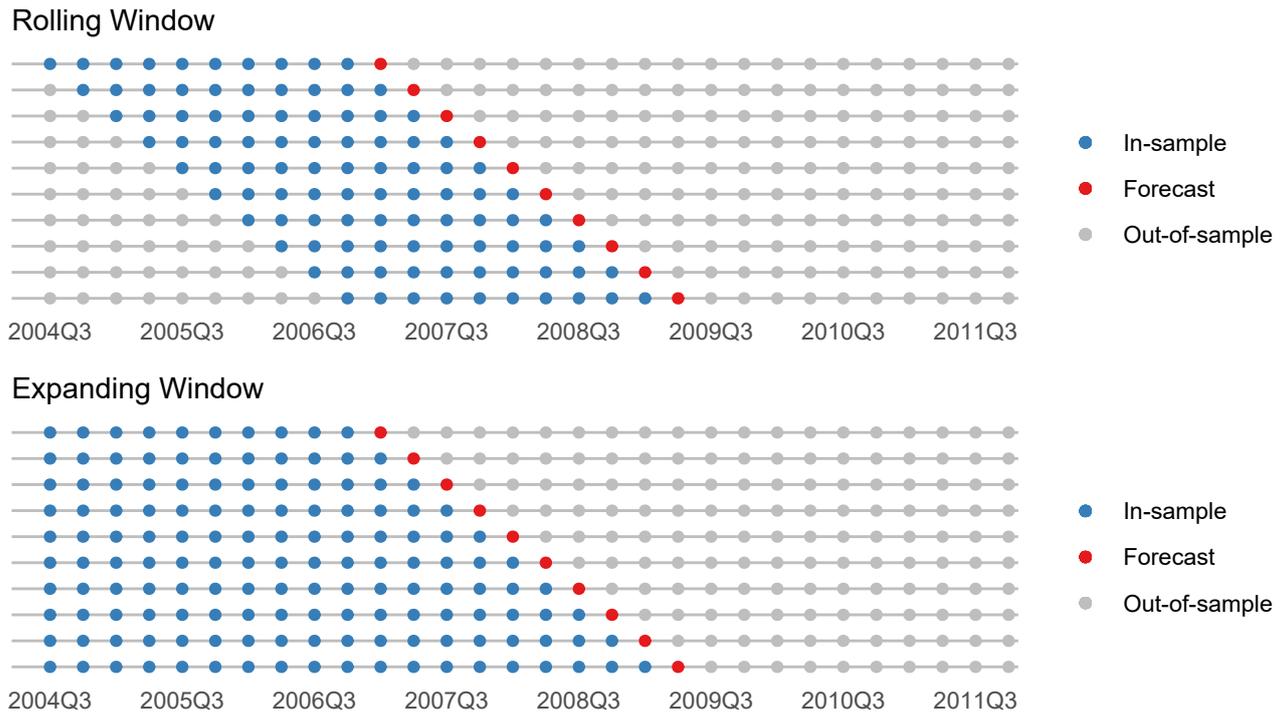


Figure 3.1: Forecasting Methods

### 3.2

#### Model Confidence Set

Following Hansen et al. (2011), we use the Model Confidence Set (MCS) methodology to compare the models and select the best ones. Since we perform multiple forecasting models' comparisons, tests for equal predictive ability like those by West and Cho (1995) would be the natural choice. This method, however, require an estimate of a covariance matrix which is difficult to achieve when the number of models is large.

The MCS procedure does not assume that a particular model is the true model. A MCS is a set of models that is constructed such that it will contain the best model with a given level of confidence  $(1 - \alpha)$ . As  $\alpha$  decreases, the set becomes larger. For large values of  $\alpha$ , that set can consist of one model only. This method is specially useful here, because we do not have a definitive benchmark to beat.

Formally, a MCS is constructed from a collection of competing objects,  $\mathcal{M}_0$ , and a criterion for evaluating these objects empirically. The MCS procedure is based on an equivalence test,  $\delta_{\mathcal{M}}$ , and an elimination rule,  $e_{\mathcal{M}}$ . The equivalence test is applied to the set of objects  $\mathcal{M} = \mathcal{M}_0$ . If  $\delta_{\mathcal{M}}$  is rejected, there is evidence that the models in  $\mathcal{M}$  are not equally 'good' and  $e_{\mathcal{M}}$  is used to eliminate an object with poor sample performance from  $\mathcal{M}$ . This procedure

is repeated until  $\delta_{\mathcal{M}}$  is ‘accepted’, and the MCS is now defined by the set of ‘surviving’ models.

As in Hansen et al. (2011), we use the squared errors of bootstrapped samples as sample losses to create the test statistic:

$$T_{\max, \mathcal{M}} = \max_{i \in \mathcal{M}} t_i \quad (3-5)$$

where  $t_i$  is  $t_i = \frac{\bar{d}_i}{\sqrt{\widehat{\text{var}}(\bar{d}_i)}}$  and  $\bar{d}_i$  is the sample loss of the  $i$ th model relative to the average across models in  $\mathcal{M}$ .

The test statistic  $T_{\max, \mathcal{M}}$  is then used to test the null hypothesis that the expected value of the relative sample loss between models  $i$ th and  $j$ th, and the expected value of the sample loss of the  $i$ th model relative to the average across models in  $\mathcal{M}$  are 0  $\forall i, j \in \mathcal{M}$ .

Thereby, for this test statistic the natural elimination rule is  $e_{\max, \mathcal{M}} \equiv \arg \max_{i \in \mathcal{M}} t_i$ .

For the implementation, we follow the algorithm from Vasconcelos (2018). The confidence set algorithm estimates  $p$ -values for all models using bootstrapped samples and uses  $\alpha$  to select which models are inside the set. The models are removed from the set interactively, thus creating a ranking. We exclude models until the null hypothesis is not rejected.

### 3.3 Models

#### 3.3.1 Autoregressive Models

We split the autoregressive models into 3 groups: (a) AR(1), (b) best ARMA( $p, q$ ), and (c) best ARIMA( $p, d, q$ ). The group (a) procedure consists of using a simple AR(1) model each window to forecast the next step. As for group (b), each window we estimate an ARMA model following the equation (3-6), where  $p \in \{0, 1, 2, 3, 4, 5\}$ ,  $q \in \{0, 1, 2, 3, 4, 5\}$  and  $d = 0$ . We then select the best model using the Bayesian Information Criteria (BIC). The group (c) uses the same strategy as group (b), but with  $d \in \{1, 2\}$ .

$$\left(1 - \sum_{i=1}^p \varphi_i L^i\right) (1 - L)^d X_t = \left(1 + \sum_{i=1}^q \theta_i L^i\right) \varepsilon_t \quad (3-6)$$

### 3.3.2 UC-SV

The Unobserved Components-Stochastic Volatility (UC-SV) model was first introduced by Stock and Watson (2010). It was typically used as a way to model inflation in the US. Here we follow Medeiros et al. (2019) methodology. The UC-SV model is described by the following equations:

$$\begin{aligned}\pi_t &= \tau_t + e^{h_t/2}\varepsilon_t \\ \tau_t &= \tau_{t-1} + u_t \\ h_t &= h_{t-1} + v_t\end{aligned}\tag{3-7}$$

where  $\{\varepsilon_t\}$  is a sequence of independent and normally distributed random variables with zero mean and unit variance,  $\{\varepsilon_t\} \sim \mathcal{N}(0, 1)$ .  $u_t$  and  $v_t$  are also normally distributed with zero mean and variance given by inverse-gamma priors. In this case,  $\tau \sim \mathcal{N}(0, V_\tau)$  and  $h \sim \mathcal{N}(0, V_h)$ , where  $V_\tau = V_h = 0.12$ . The  $h$ -steps-ahead is computed as  $\hat{\pi}_{t+h} = \hat{\tau}_{t|t}$ .<sup>1</sup>

### 3.3.3 Shrinkage Models

The basic idea behind shrinkage models is to penalize the size of parameter estimates in order to shrink those associated to irrelevant variables towards zero. Thereby, these methods are a well-established alternative to deal with a high-dimensional environment and data sets whose number of regressors is larger than the number of observations.

Generally, these models are estimated by solving the following problem:

$$\beta(\lambda_1, \lambda_2) = \left\{ \arg \min_{\beta} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + p(\beta) \right\}\tag{3-8}$$

where  $p$  is a penalty function, which varies according to the model. One less general specification can be written as:

$$\beta(\lambda_1, \lambda_2) = \left\{ \arg \min_{\beta} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda_2 \|\beta\|_2^2 + \lambda_1 \sum_{j=1}^p \hat{w}_j |\beta_j| \right\}\tag{3-9}$$

Most of the models used in this work follow Equation 3-8. Table 3.1 summarises them:

<sup>1</sup>We also introduced a Random Walk model, where  $\hat{\pi}_{t+h} = \pi_t$ , but the results are fairly similar to the AR(1) and UC-SV models. Since neither of them are chosen as best models in any exercise, we left the Random Walk model out of this work.

	$\lambda_1$	$\lambda_2$	$\hat{w}_j$
Ridge	0	$\lambda$	-
LASSO	$\lambda$	0	$1 \forall j$
adaLASSO	$\lambda$	0	$ \boldsymbol{\beta}_{LASSO} + \frac{1}{\sqrt{n}} ^{-1}$
Elastic Net	$\lambda_1$	$\lambda_2$	$1 \forall j$
adaElastic Net	$\lambda_1$	$\lambda_2$	$ \boldsymbol{\beta}_{EN} + \frac{1}{\sqrt{n}} ^{-1}$

Table 3.1: Shrinkage Models

where  $n$  is the number of observations. We add this term to deal with possible zero weights from LASSO estimations. Moreover, the  $p$  function in Elastic Net models can be rewritten as  $\kappa\|\boldsymbol{\beta}\|^2 + (1 - \kappa)\hat{w}_j\|\boldsymbol{\beta}\|_1$ . We fixate  $\kappa = 0.5$ .

The Ridge regression was one of the first techniques capable of dealing with large datasets (Hoerl and Kennard (1970)). It imposes a quadratic penalty to the coefficients and has the appealing feature of an analytical solution. Despite that, the coefficient solution vector consists exclusively of non-zero entries, which is a disadvantage in developing a variable selection analysis.

The Least Absolute Shrinkage and Selection Operator (LASSO) is a method presented originally in Tibshirani (1996). Opposed to Ridge regression, LASSO imposes a penalty on the sum of the coefficients absolute values. Consequently, it shrinks irrelevant variables exactly to zero, allowing it to perform variable selection.

However, Zhao and Yu (2006) and Zou (2006) highlight that LASSO models require strong conditions to achieve consistency, and do not have the oracle property. Thus, Zou (2006) proposes the Adaptive LASSO (adaLASSO) to overcome these problems. The adaLASSO model is a two-step method which uses the first-step LASSO coefficients' estimates to weight the relative importance of the regressors in the second-step as noted in 3.1.

The Elastic Net (EINet) models are generalizations which include both the LASSO and Ridge as particular cases. The penalization used for these models is a simple convex combination of both latter models penalization functions (Zou and Hastie (2005)).

Traditionally, the  $\lambda$  parameters are chosen by cross-validation methods. In our work, on the other hand, we use the BIC to choose the parameter.

Furthermore, we investigate a nonconvex form for the penalty function in the Smoothly Clipped Absolute Deviations (SCAD) penalty (Xie and Huang (2009)). We write the SCAD penalty as the following equation:

$$p(\boldsymbol{\beta}) = \begin{cases} \lambda|\boldsymbol{\beta}| & \text{if } |\boldsymbol{\beta}| \leq \lambda \\ \frac{2a\lambda|\boldsymbol{\beta}| - \boldsymbol{\beta}^2 - \lambda^2}{2(a-1)} & \text{if } \lambda < |\boldsymbol{\beta}| \leq a\lambda \\ \frac{\lambda^2(a+1)}{2} & \text{otherwise} \end{cases} \quad (3-10)$$

for  $a > 2$ . Note that SCAD coincides with the LASSO until  $\beta = \lambda$ . Then it transitions to a quadratic function until  $\beta = \delta\lambda$ . The SCAD penalty retains the penalization rate (and bias) of the LASSO for small coefficients, but continuously relaxes the rate of penalization as the absolute value of the coefficient increases.

Except for the Ridge models, all models in this section introduce sparsity and variable selection, which is a key feature for this work.

### 3.3.4 Random Forest

The Random Forest (RF) methodology was initially proposed by Breiman (2001) as a way to reduce the variance of regression trees. The method is based on the bootstrap aggregation (bagging) of randomly constructed regression trees. Coulombe et al. (2019) shows that a random forest performs well when forecasting inflation. Moreover, Medeiros et al. (2019) also finds that RF models deliver clear predictive gains when forecasting inflation in the US.

Basically, a RF is the average of numerous random regression trees and thus it's able to capture non-linear dynamics between the independent variables and the variable of interest. A regression tree is a non-parametric model based on the recursive binary partitioning of the covariate space  $\mathbb{X}$ .  $y_t$  is the a function of local models, each of them determined in  $K \in \mathbb{N}$  different partitions of  $\mathbb{X}$ .

We can represent the model in a graph of a binary decision tree with  $N \in \mathbb{N}$  parent (or split) nodes and  $K \in \mathbb{N}$  terminal nodes (also called leaves). The graphs grows from a singular root node into  $K$  terminal nodes hence its name. The partitions are defined by a single split in determined regressor  $x_t$ .

Following Garcia et al. (2017), we can represent the RF model formally. The root node is at position 0. At position  $j$ , a parent node generates a left and a right child at positions  $2j + 1$  and  $2j + 2$ , respectively. Every parent node has an associated split variable  $x_{s_j,t} \in \mathbf{x}_t$ , where  $s_j \in \mathbb{S} = 1, 2, \dots, q$ . Furthermore, if we let  $\mathbb{J}$  and  $\mathbb{T}$  be the set of indexes of the parent and terminal nodes, respectively, a tree architecture can be fully determined from  $\mathbb{J}$  and  $\mathbb{T}$ .

The forecasting model based on regression trees can be represented mathematically as:

$$y_t = H_{\mathbb{T}}(\mathbf{x}_t; \boldsymbol{\psi}) + u_t = \sum_{i \in \mathbb{T}} \beta_i B_{\mathbb{J}_i}(\mathbf{x}_t; \boldsymbol{\theta}_i) + u_t \quad (3-11)$$

where

$$B_{\mathbb{J}_i}(\mathbf{x}_t; \boldsymbol{\theta}_i) = \prod_{j \in \mathbb{J}} I(x_{s_j, t}; c_j)^{\frac{n_{i,j}(1+n_{i,j})}{2}} [1 - I(x_{s_j, t}; c_j)]^{(1-n_{i,j})(1+n_{i,j})} \quad (3-12)$$

$$I(x_{s_j, t}; c_j) = \begin{cases} 1 & \text{if } x_{s_j, t} \leq c_j \\ 0 & \text{otherwise} \end{cases} \quad (3-13)$$

$$n_{i,j} = \begin{cases} -1 & \text{if the path to leaf } i \text{ does not include the parent node } j \\ 0 & \text{if the path to leaf } i \text{ includes the right-child node of the parent node } j \\ 1 & \text{if the path to leaf } i \text{ includes the left-child node of the parent node } j. \end{cases} \quad (3-14)$$

Let  $\mathbb{J}_i$  be the subset of  $\mathbb{J}$  that contains the indexes of the parent nodes that form the path to leaf  $i$ , then  $\boldsymbol{\theta}_i$  is the vector that contains all of the parameters  $c_k$  such that  $k \in \mathbb{J}_i$ ,  $i \in \mathbb{T}$ . Note that  $\sum_{j \in \mathbb{J}} B_{\mathbb{J}_i}(\mathbf{x}_t; \boldsymbol{\theta}_j) = 1, \forall \mathbf{x}_t \in \mathbb{R}^{q+1}$ .

A Random Forest is a collection of regression trees, each of which is specified in a bootstrapped sub-sample of the original data set. Suppose that there are  $B$  bootstrapped sub-samples, and denote the estimated regression tree for each of the sub-samples by  $H_{J_b \mathbb{T}_b}(\cdot; \boldsymbol{\psi})$ . The final prediction is defined as:

$$\hat{y}_t = \frac{1}{B} \sum_{b=1}^B H_{J_b \mathbb{T}_b}(\mathbf{x}_t; \boldsymbol{\psi}) \quad (3-15)$$

For each of the bootstrapped sub-samples a regression tree is estimated by recursively repeating the following steps until the terminal node:

1. Randomly select  $m$  out of  $q$  covariates as possible split variables
2. Pick the best variable/split point among the  $m$  candidates
3. Split the node into two child nodes

Random Forests can deal with very large numbers of explanatory variables, and the predicted model is highly nonlinear. Since we are dealing with time series, we use a block bootstrap.

### 3.3.5 Hybrid Linear-Random Forest Models

We follow Medeiros et al. (2019) hybrid adaptations to disentangle the relative importance of variable selection and non-linearity in our Random Forest models. Since we're interested in the variable selection aspect, these models can help us open the black-box associated to Random Forest models. We estimate the RF/OLS and the adaLASSO/RF.

The RF/OLS is estimated using the following steps:

1. For each bootstrap sample  $b = 1, \dots, B$ :
  - Grow a single tree with 20 nodes, and save the  $N \leq k$  split variables
  - Run an OLS on the selected splitting variables
  - Compute the forecast  $\hat{\pi}_{t+h}^b$
2. The final forecast will be  $\hat{\pi}_{t+h} = B^{-1} \sum_{b=1}^B \hat{\pi}_{t+h}^b$

With this model we can check the performance of a simple linear model using variables selected by the Random Forest algorithm. Then we can compare the predictive results of this model with the pure Random Forest models. If the difference between the results is not statistically significant, we understand that non-linearity is not a issue.

The second class of adapted models is the adaLASSO/RF. Here we use the sparse shrinkage model adaLASSO for variable selection and run the Random Forest algorithm on this selected variables. If adaLASSO/RF performs similarly to the Random Forest, we understand that variable selection in RF is less relevant and non-linearity is more important.

In Appendix D, we indicate the computer codes and algorithms used in this work.

## 4 Empirical Results

### 4.1 Reducing the Dimensionality

In the literature of the Missing Disinflation Puzzle (MDP), there is no clear benchmark model. Therefore, the Model Confidence Set (MCS) plays a vital role in this study. The MCS procedure theoretically can deal with a large number of models. However, we first perform some actions to reduce the dimensionality of the set  $\mathcal{M}$  of possible specifications.

First, we investigate if the addition of regressors' lags is capable of generating predictive gains. We split the models into 3 groups: (a) Shrinkage models excluding SCAD; (b) SCAD-penalty models; and (c) Random Forest. For each group, we calculate the one-year ahead predictions in the period corresponding to the MDP as explained in Chapter 3. For each iteration of the forecasting model, we collect the error and squared error, and calculate the RMSE and MAE accumulated up until that point in time. The ranking of top 10 best models according to RMSE criteria for each group is presented in Tables 4.1, 4.2, and 4.3. Each table exhibits the models' specifications, and the RMSE- and MAE descriptive statistics.

		<u>Model</u>		<u>RMSE</u>						<u>MAE</u>					
Database		# of Lags	Window	Ranking	Min. Value	Max. Value	Mean Value	SD	Accumulated/ Final	Ranking	Min. Value	Max. Value	Mean Value	SD	Accumulated/ Final
Ridge	Flexible	2	Expanding	1	0.0035	0.0125	0.0084	0.0027	0.0079	3	0.0032	0.0084	0.0058	0.0013	0.0053
Ridge	Flexible	3	Expanding	2	0.0034	0.0130	0.0084	0.0028	0.0079	1	0.0031	0.0087	0.0055	0.0013	0.0046
Ridge	Semi-Parsimonious	3	Expanding	3	0.0033	0.0127	0.0085	0.0029	0.0080	5	0.0029	0.0084	0.0060	0.0016	0.0053
adaLASSO	Flexible	2	Expanding	4	0.0081	0.0140	0.0103	0.0018	0.0081	71	0.0065	0.0133	0.0093	0.002	0.0065
LASSO	Flexible	2	Expanding	5	0.0081	0.0140	0.0103	0.0018	0.0081	72	0.0065	0.0132	0.0093	0.002	0.0065
Ridge	Flexible	1	Rolling	6	0.0029	0.0133	0.0088	0.0034	0.0084	2	0.0022	0.0088	0.0056	0.0018	0.005
Ridge	Parsimonious	3	Expanding	7	0.0031	0.0127	0.0088	0.0031	0.0085	12	0.0027	0.0087	0.0064	0.0018	0.0058
Ridge	Semi-Parsimonious	3	Rolling	8	0.0030	0.0138	0.0091	0.0036	0.0087	3	0.0022	0.0088	0.0058	0.0019	0.0052
LASSO	Flexible	3	Expanding	9	0.0070	0.0134	0.0103	0.0017	0.0087	75	0.0068	0.0123	0.0091	0.0017	0.0068
Ridge	Flexible	3	Rolling	10	0.0034	0.0140	0.0092	0.0034	0.0088	6	0.0027	0.0086	0.0061	0.0018	0.0054

Table 4.1: Shrinkage Models: Top 10 Specifications

		<u>Model</u>		<u>RMSE</u>						<u>MAE</u>					
Database		# of Lags	Window	Ranking	Min. Value	Max. Value	Mean Value	SD	Accumulated/ Final	Ranking	Min. Value	Max. Value	Mean Value	SD	Accumulated/ Final
SCAD	Parsimonious	1	Expanding	1	0.0035	0.0144	0.0091	0.0032	0.0085	1	0.003	0.0092	0.0058	0.0016	0.0051
SCAD	Semi-Parsimonious	1	Expanding	2	0.005	0.0139	0.0094	0.0027	0.0086	2	0.0046	0.0096	0.0066	0.0013	0.0054
SCAD	Flexible	4	Expanding	3	0.0025	0.0155	0.01	0.0042	0.0096	3	0.0025	0.0107	0.0066	0.0024	0.0057
SCAD	Parsimonious	4	Expanding	4	0.0033	0.0155	0.0102	0.004	0.0096	5	0.0032	0.011	0.0069	0.0022	0.0059
SCAD	Semi-Parsimonious	4	Expanding	5	0.0027	0.0159	0.0102	0.0042	0.0097	4	0.0027	0.0109	0.0067	0.0024	0.0058
SCAD	Parsimonious	2	Expanding	6	0.0028	0.0137	0.0099	0.0039	0.0101	6	0.0025	0.0096	0.0071	0.0023	0.0072
SCAD	Flexible	2	Expanding	7	0.0052	0.0139	0.0109	0.0026	0.0104	7	0.0039	0.0101	0.0079	0.0016	0.0075
SCAD	Flexible	3	Expanding	8	0.0052	0.0133	0.0103	0.0025	0.011	11	0.0052	0.0107	0.0084	0.0017	0.0087
SCAD	Flexible	1	Expanding	9	0.0108	0.0188	0.0139	0.0024	0.011	8	0.0076	0.0174	0.011	0.0028	0.0076
SCAD	Parsimonious	2	Rolling	10	0.004	0.0141	0.0106	0.0031	0.0113	12	0.0027	0.0116	0.0086	0.0026	0.0088

Table 4.2: SCAD-Penalty Models: Top 10 Specifications

		<u>Model</u>		<u>RMSE</u>						<u>MAE</u>					
Database		# of Lags	Window	Ranking	Min. Value	Max. Value	Mean Value	SD	Accumulated/ Final	Ranking	Min. Value	Max. Value	Mean Value	SD	Accumulated/ Final
RF	Flexible	1	Expanding	1	0.0036	0.0145	0.0093	0.0033	0.0086	2	0.0032	0.0097	0.0063	0.0017	0.0055
RF	Flexible	1	Rolling	2	0.005	0.0145	0.0097	0.0027	0.0086	7	0.0048	0.0111	0.0072	0.0016	0.0058
RF	Semi-Parsimonious	1	Rolling	3	0.005	0.0144	0.0097	0.0027	0.0087	14	0.0049	0.0111	0.0075	0.0017	0.0061
RF	Parsimonious	1	Rolling	4	0.0054	0.0143	0.0097	0.0026	0.0087	12	0.0053	0.0111	0.0075	0.0016	0.006
RF	Flexible	2	Rolling	5	0.0058	0.0144	0.0098	0.0025	0.0087	17	0.0056	0.0112	0.0078	0.0016	0.0062
RF	Semi-Parsimonious	2	Rolling	6	0.0058	0.0147	0.01	0.0026	0.0088	21	0.0057	0.0115	0.0079	0.0017	0.0063
RF	Parsimonious	1	Expanding	7	0.0032	0.0147	0.0094	0.0035	0.0089	1	0.0027	0.0094	0.0061	0.0018	0.0054
RF	Parsimonious	2	Expanding	8	0.0036	0.0147	0.0095	0.0034	0.0089	9	0.0033	0.0098	0.0067	0.0019	0.0059
RF	Semi-Parsimonious	1	Expanding	9	0.0038	0.0148	0.0096	0.0033	0.0089	6	0.0034	0.0099	0.0066	0.0018	0.0058
RF	Semi-Parsimonious	2	Expanding	10	0.004	0.0147	0.0096	0.0033	0.0089	10	0.0037	0.01	0.007	0.0019	0.006

Table 4.3: Random Forest: Top 10 Specifications

We cannot draw many conclusions from the tables. In Table 4.1 (total of 116 models), the specifications associated with 2 to 3 lags, expanding window, and larger data sets seem to have a better performance than other specifications. However, the variability between the rankings based on RMSE and MAE is noteworthy.

In Table 4.2 (total of 24 models), the variability between the rankings is smaller. Yet, there is no particular pattern in the top 10 specifications, except for the predominance of the expanding window methodology. In Random Forest family models (Table 4.3, total of 24 models) the accumulated RMSE are closer to each other relative to other specifications' results. Shorter regressors' lags seem to perform better too. Ultimately, these tables are more descriptive than conclusive.

In order to access more robust results, we run the MCS procedure for  $\alpha = \{0.8, 0.9, 0.95\}$  in each of the 3 subgroups  $M_0 \in M$  where  $M = \{\text{Shrinkage, SCAD, RF}\}$ . For each combination, we get a set  $\mathcal{M}_{\alpha_0}^{M_0} \in \mathcal{M}_{\alpha}^{M_0}$  of chosen models. Figure 4.1 exhibits the dimensionality of these sets.

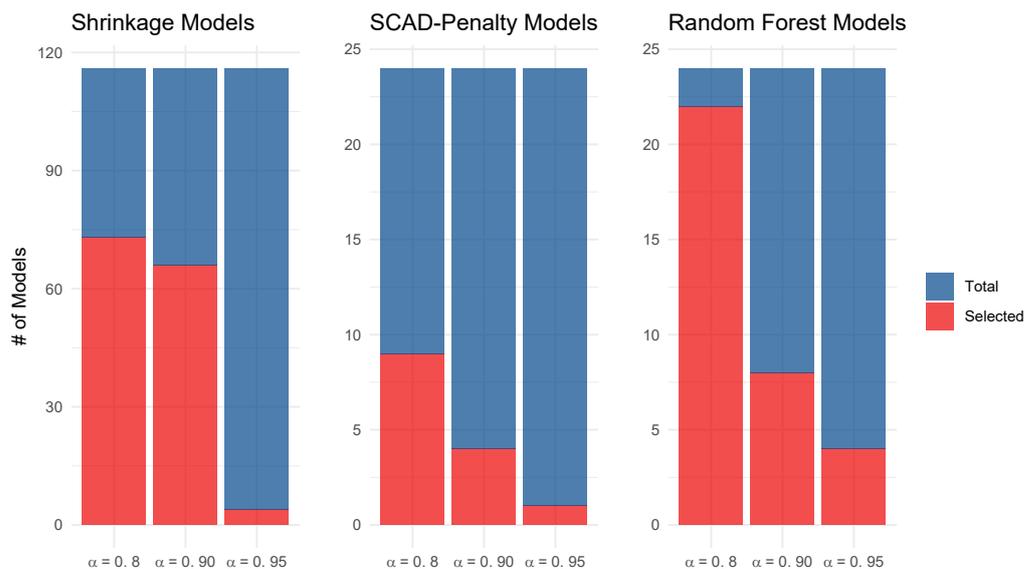


Figure 4.1: Amount of Chosen Models per Family of Specifications

In the Shrinkage models' case, for  $\alpha \in \{0.8, 0.9\}$ , the MCS procedure selects more than 50% of the  $\mathcal{M}^{\text{Shrinkage}}$ . Therefore, we turn our focus to the  $\alpha = 0.95$  case. 4 models are selected:

**Shrinkage Models**

	Database	# of Lags	Window
LASSO	Flexible	2	Expanding
adaLASSO	Flexible	2	Expanding
Ridge	Semi-Parsimonious	3	Expanding
Ridge	Flexible	3	Expanding

Table 4.4: Shrinkage Family: Specifications Selected for  $\alpha = 0.95$ 

All the selected models utilize a expanding window as forecasting procedure. Despite that, there is no clear pattern among them. In the SCAD-penalty- and Random Forest cases, the selection is more parsimonious -  $n(\mathcal{M}_{\alpha=0.9}^{\text{SCAD}}) = 4$  and  $n(\mathcal{M}_{\alpha=0.95}^{\text{SCAD}}) = 1$ ;  $n(\mathcal{M}_{\alpha=0.9}^{\text{RF}}) = 8$  and  $n(\mathcal{M}_{\alpha=0.95}^{\text{RF}}) = 4$ .

**SCAD-penalty Models**

	Database	# of Lags	Window	$\alpha$
SCAD	Parsimonious	1	Expanding	0.9, 0.95
SCAD	Parsimonious	2	Rolling	0.9
SCAD	Semi-Parsimonious	1	Expanding	0.9
SCAD	Flexible	1	Expanding	0.9

Table 4.5: SCAD-Penalty Family: Specifications Selected for  $\alpha \in \{0.9, 0.95\}$ **Random Forest**

	Database	# of Lags	Window	$\alpha$
RF	Flexible	1	Expanding	0.9, 0.95
RF	Flexible	1	Expanding	0.9, 0.95
RF	Semi-Parsimonious	1	Expanding	0.9, 0.95
RF	Flexible	2	Rolling	0.9, 0.95
RF	Semi-Parsimonious	1	Expanding	0.9
RF	Parsimonious	1	Rolling	0.9
RF	Parsimonious	1	Expanding	0.9
RF	Parsimonious	2	Rolling	0.9

Table 4.6: random Forest Family: Specifications Selected for  $\alpha \in \{0.9, 0.95\}$ 

We find that specifications with shorter regressors' lags - except in Shrinkage models - tend to perform better vis-à-vis data sets containing 3 or 4 lags (Table 4.5 and Table 4.6). Therefore, excluding the specifications explicitly contained in sets associated with  $\alpha = 0.95$ , we execute all subsequent exercises

in this work with a data set containing only the first lag of the regressors. Thus, we reduce the dimensionality of  $\mathcal{M}$ .

## 4.2 Best Specifications' Selection

We construct a new set of models  $\mathcal{M}^*$  consisting of: (a) the chosen specifications  $\mathcal{M}_{\alpha=0.95}^M$  in Section 4.1, (b) all models specifications running on data sets with only 1 lag of regressors, and (c) the autoregressive and UC-SV specifications. Then we run the MCS procedure on  $\mathcal{M}^*$ . For  $\alpha = 0.95$ , 9 specifications are selected. Table 4.7 exhibits them and Figure 4.2 presents their predictive performances.

$\mathcal{M}_{\alpha=0.95}^*$			
	Database	# of Lags	Window
RF	Parsimonious	1	Expanding
RF	Parsimonious	1	Rolling
RF	Flexible	1	Rolling
SCAD	Parsimonious	1	Expanding
SCAD	Semi-Parsimonious	1	Expanding
adaLASSO	Flexible	2	Expanding
LASSO	Flexible	2	Expanding
Ridge	Flexible	3	Expanding
Ridge	Semi-Parsimonious	3	Expanding

Table 4.7: Specifications in  $\mathcal{M}_{\alpha=0.95}^*$

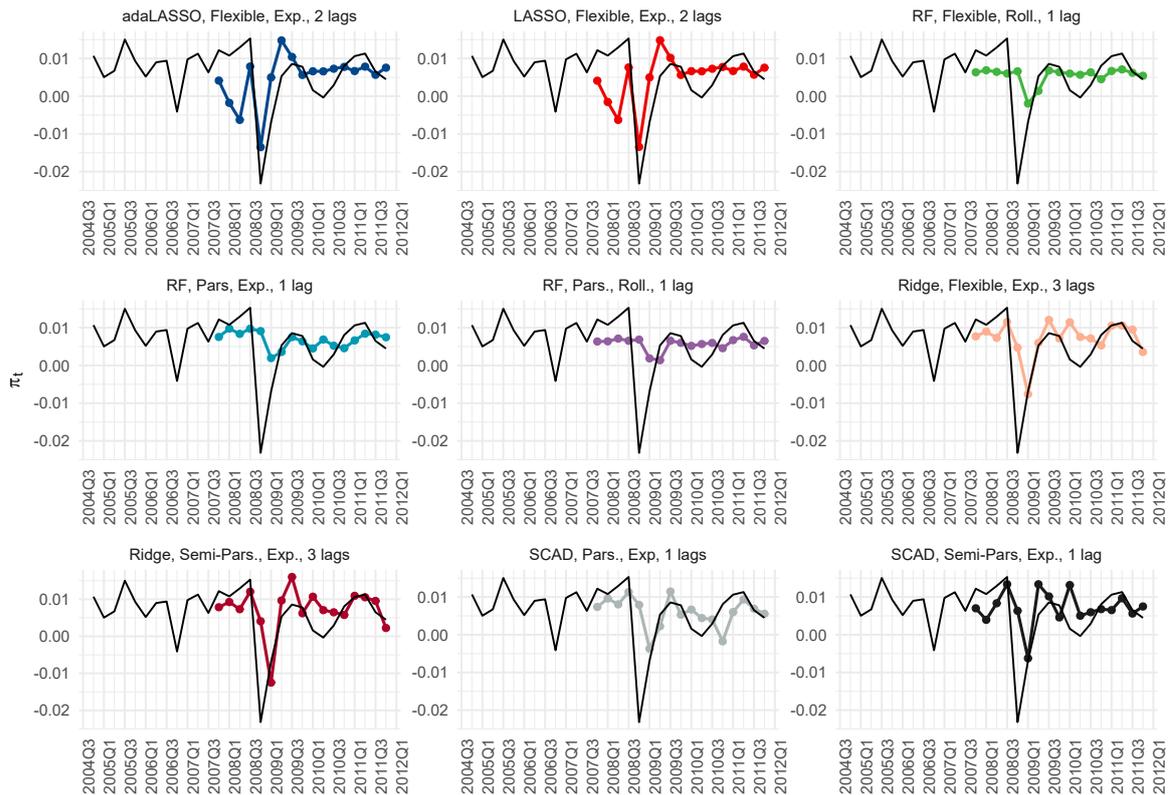


Figure 4.2:  $\mathcal{M}_{\alpha=0.95}^*$  Specifications Predictive Performance

Overall, the different specifications have a similar performance. No model is able to predict the fall between 2009Q4 and 2010Q4, but almost all of them foretell the spikes in 2009 and after the Great Recession. Since the  $\mathcal{M}_{\alpha=0.95}^*$  contains all of them, we interpret these specifications as equally good for the time frame of the MDP. Figures 4.3 and 4.4 exhibit the correlation between the forecasts and the evolution of the RMSE in each model.

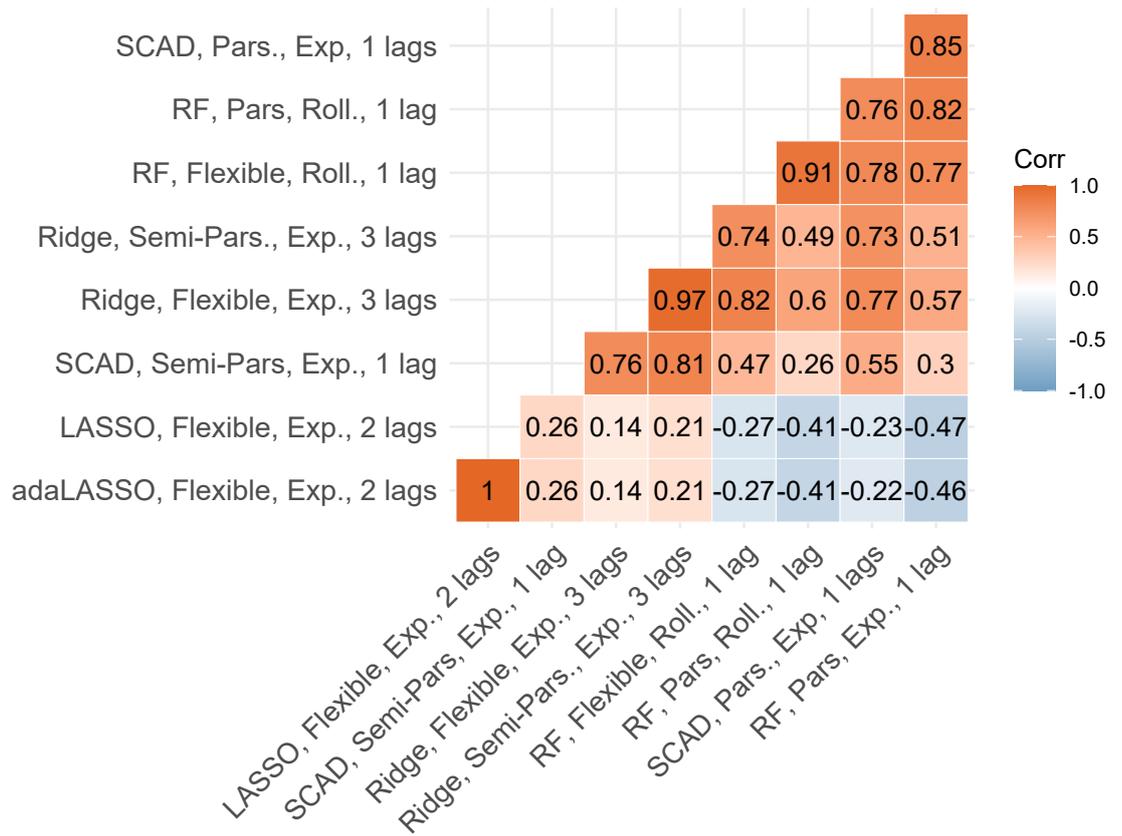


Figure 4.3:  $\mathcal{M}_{\alpha=0.95}^*$  Correlation Plot

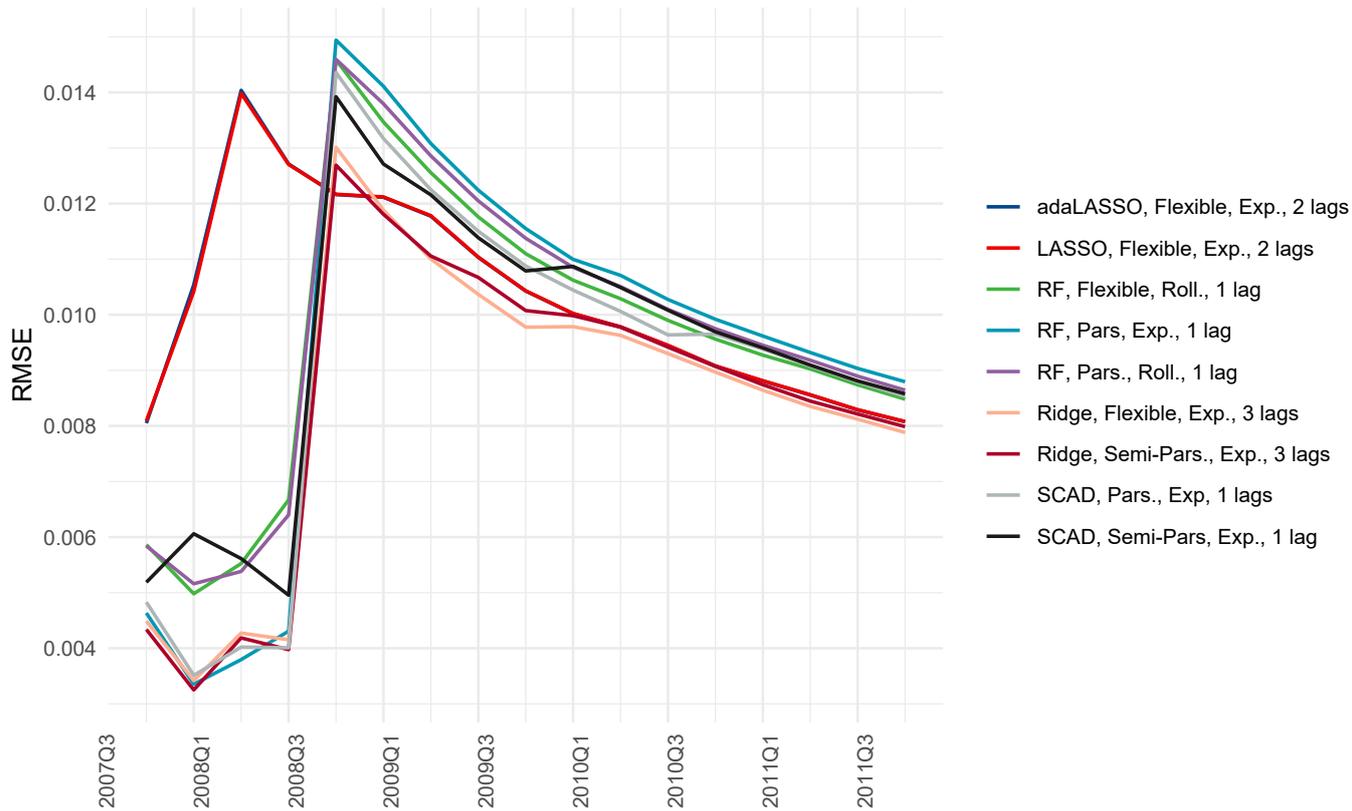


Figure 4.4:  $\mathcal{M}_{\alpha=0.95}^*$  Accumulated RMSE Dynamic

The presence of a LASSO- and a adaLASSO specifications in  $\mathcal{M}_{\alpha=0.95}^*$  is particularly interesting. Both models are linear and perform variable selection. Furthermore, both of them are the only specifications which predict the sharp fall of inflation in 2008 Q4. In Sections 4.3 we further investigate these results.

The Random Forest specifications are also compelling. Using the hybrid models we try to disentangle the relative importance of variable selection and non-linearity in these models. In Appendix B we shed some light in this discussion.

### 4.3 Variable Selection and the Missing Disinflation Puzzle

First, we analyze the variable selection performed by the LASSO- and the adaLASSO models. Figure 4.3 demonstrates the amount of variables selected in each forecasting step and in which groups they are contained.

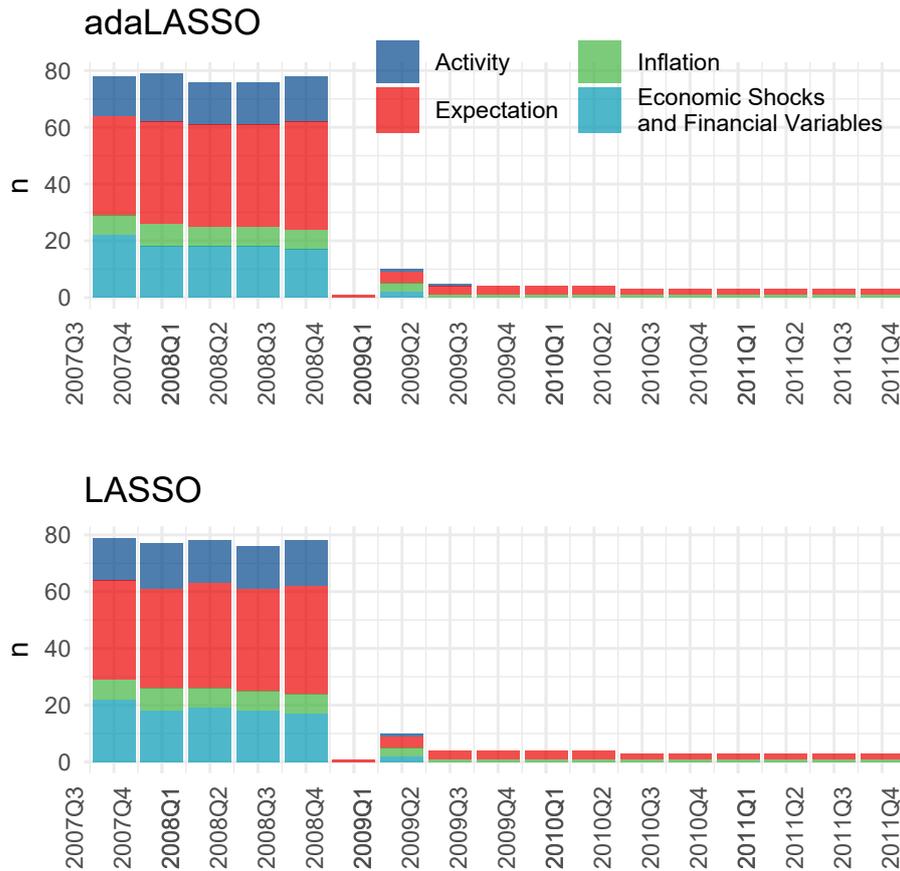


Figure 4.5: LASSO and adaLASSO Variable Selection per Groups

For all periods the intercept is calculated, but  $n$  does not include it. In both cases the story is quite the same. Up until 2008Q4 more than 75 variables are selected. Then just 1 in 2009Q1, 10 in 2009Q2, 5 in 2009 Q3 (4 for LASSO), 4 between 2009 Q4 and 2010 Q2, and 3 up until the end.

From 2009 Q3 until the end of the sample 3 variables are always selected in both specifications: the first lag of PCE headline inflation (`pce_headline_1d`), GreenBook 1 quarter ahead inflation expectation (`gb_fcpi_1q_1d`), and the Michigan Survey of Consumer household's inflation expectation for those between 18 and 34 years old (`px1_mean_a1834_1d`). Moreover, between 2009 Q3 and 2010 Q2 the inflation expectations of households with low uncertainty calculated as in Binder (2015) are also always selected (`mu_1_1d`). Figures 4.6 and 4.7 expose the evolution of the coefficients connected to these variables. We multiply the coefficients by 1000 for better visualization.

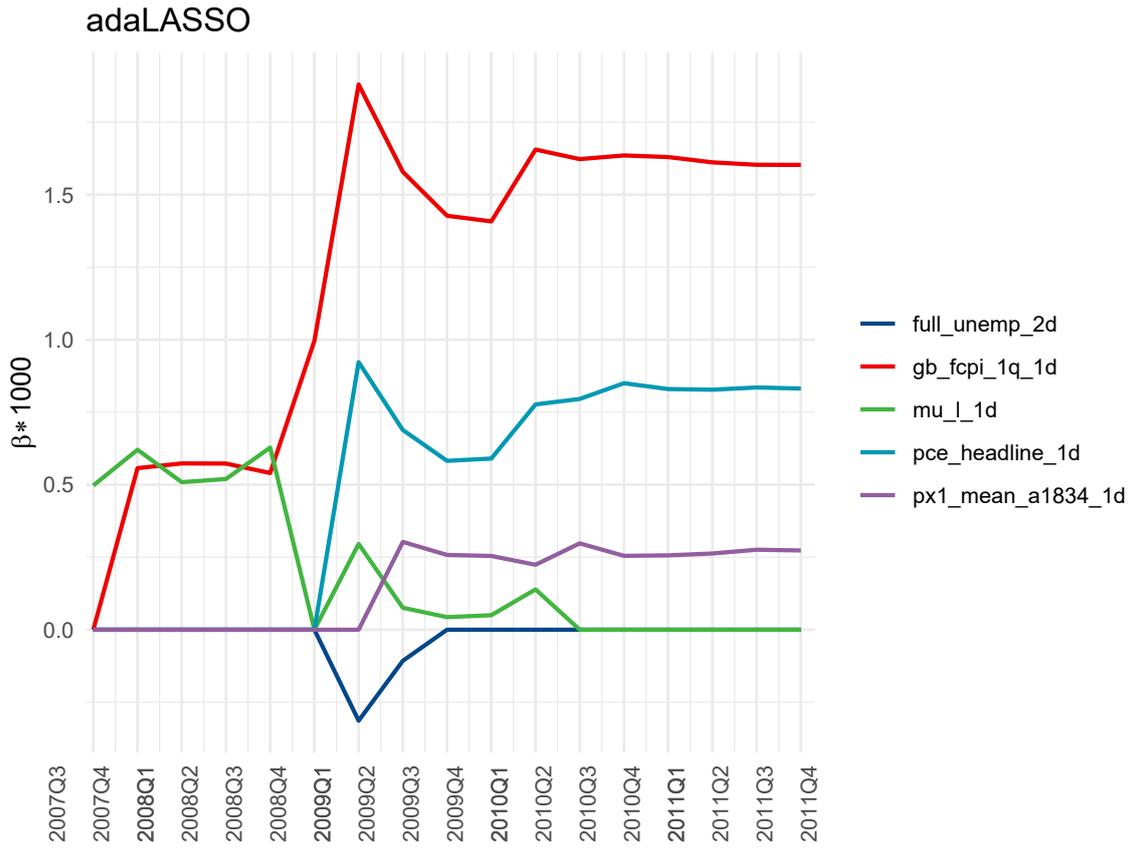


Figure 4.6: adaLASSO Coefficients' Values Evolution over Time

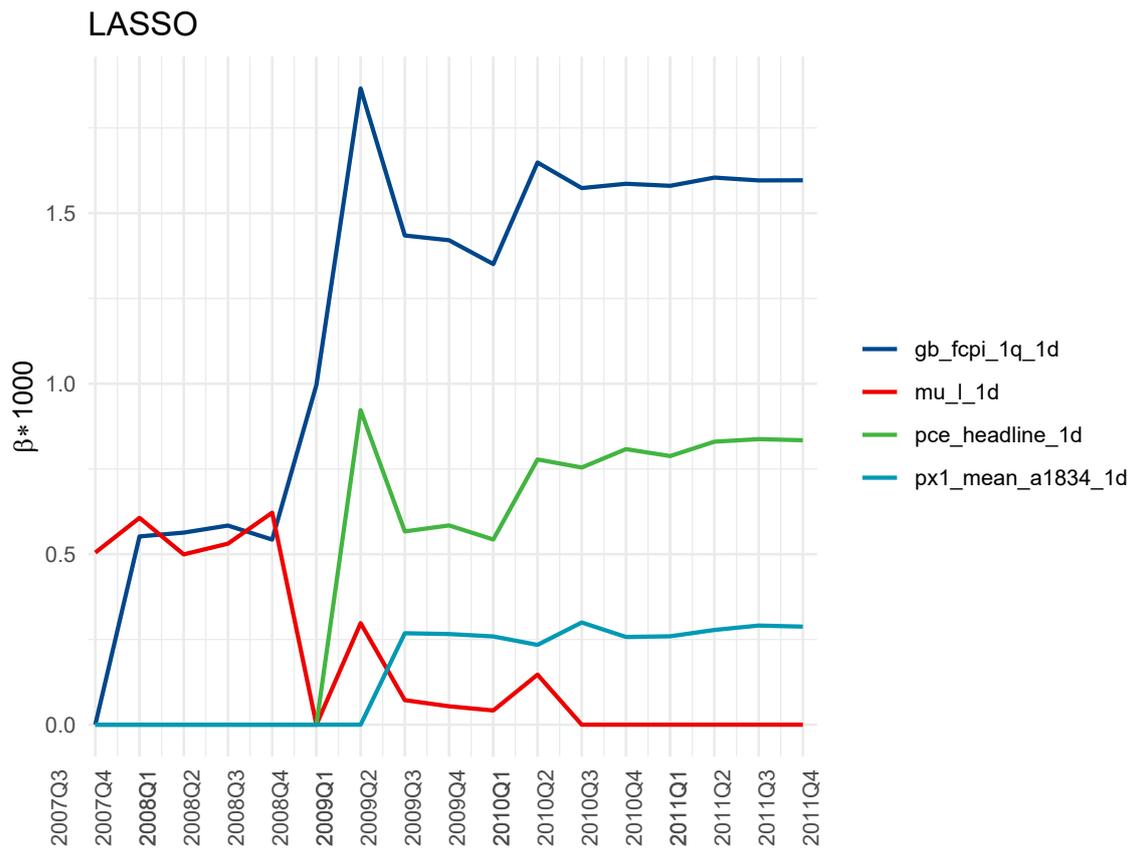


Figure 4.7: LASSO Coefficients' Values Evolution over Time

The presence of `px1_mean_a1834_1d` and `mu_l_1d` - both metrics derived from the University of Michigan Survey of Consumers - generate evidence for the explanation proposed in Coibion and Gorodnichenko (2015) for the puzzle. The authors suggest that firms' inflation expectations are best proxied by household expectations. They show that an expectations-augmented Phillips curve, using household inflation expectations as measured by the Michigan Survey of Consumers, can account for the absence of strong disinflationary pressures since 2009. For them, the rise in inflationary expectations between 2009 and 2011 account for why the inflation did not fall as much as one might have predicted.

These results also generate evidence for Binder (2015). In their article, the authors explain that even among households, there is substantial heterogeneity of expectations, and the average household forecast may not be the best proxy. They proceed to show that the inflation expectations associated with less-uncertain consumers provide a better proxy for firms' expectations than those associated with high-uncertain consumers and the SPF expectations. Lastly, they use this variable to solve the MDP.

In Binder (2017), however, the author find evidence that the inflation expectations associated with college-educated, high-income, working-age (35 to 54 years old), and male consumers drive inflation dynamics more than the expectations of other demographic groups or of professional forecasters. In our work, only expectations associated with agents between 18 and 34 years old are relevant.

The absence of a Economic Shock or Activity variables, and the presence of a lagged inflation component and expectations associated with professionals (GreenBook) are also noteworthy results.

### 4.3.1 The Sharp Fall of 2008 Q4

The largest disinflation in our sample happens in 2008 Q4 and both models aforementioned are the only ones in  $\mathcal{M}_{\alpha=0.95}^*$  able to predict it. We attempt to disentangle the variables' choice in this period using 3 different data mining processes over the adaLASSO specification model (the results are the same for the LASSO specification). None of them is able to produce conclusive results.

For the rest of this section we describe briefly the processes and their outcomes. More details are in Appendix C.

First, we identify the variables which have the biggest impact on the 1 step ahead 2008 Q4 forecast. We remove these variables from the data set, re-estimate the model, and repeat this exercise iteratively. As we remove the nominal change in the price of oil at  $t - 2$ , and the real change in the price of oil at  $t - 2$  - the two variables with the biggest impact on the forecast - the model becomes unable to predict the sharp fall in 2008 Q4.

Then, we run the same 1-step-ahead forecast procedure described in Section 3.1 using simple linear regressions with a data set containing only both metrics along with other selected variables detailed in Appendix C. These models are not capable of predicting the sharp fall in 2008 Q4.

The second method follows the steps below:

1. We collect the 78 regressors whose coefficients are not zero selected by the adaLASSO specification;
2. We rank them decreasingly according to the absolute impact on the  $\hat{y}_{2008Q4}$  estimate;
3. We construct data sets containing the regressor  $x_I$  and all others  $x_i \leq x_I$ , where  $i, I \in \mathbb{I}$  and  $\mathbb{I}$  is the set of ranking positions;

4. We repeat the forecasting procedure described in Section 3.1 using simple linear regressions;

A total of 78 models are estimated. We pick the models which have the best predictive performance for 2008 Q4 and collect the variables contained in these specifications. Then, we analyze these chosen regressors using MCS.

The third and last method consists of:

1. We run a simple regression of CPI headline on the change in oil nominal prices in  $t - 2$  and collect the one-step-ahead forecast for 2008 Q4;
2. We create all possible data sets consisting of the combination of the change in oil nominal prices in  $t - 2$  and another variable, and collect the one-step-ahead forecast for 2008 Q4;
3. We select the specification which minimizes the predictive error and repeat the process with a new specification containing the change in oil nominal prices in  $t - 2$ , the chosen variable and all other possible variables;
4. We repeat this process iteratively until the last specifications yields the same prediction as the adaLASSO model for 2008 Q4.

None of these methods generates conclusive results. All of them in the end yields models' specifications, which are not able to isolate a particular effect or explanation for the adaLASSO model being able to predict the fall in 2008 Q4.

## 5 Concluding Remarks

This paper contributes to the Missing Disinflation Puzzle (MDP) literature by constructing a data set focused exclusively on variables used in the investigation of the puzzle, and by using several Machine Learning methods to generate evidence for potential explanations explored in the literature. We aim to assess if the use of these methods combined with the data set lead to any pattern of variable selection.

The explanations and econometric approaches for the MDP are not consensual, therefore there is no clear benchmark to beat. To deal with this, we perform several Machine Learning methods to generate 1 step ahead predictions across the period of 2007 Q4 - 2011 Q4, and use the fully data-driven Model Confidence Set (MCS) procedure to select the best models. In the first step, we reduce the dimensionality of possible models' specifications, and, after that, we run the MCS procedure on the whole set of remaining models.

The Machine Learning methods are from different types, such as autoregressive, non-linear and shrinkage models, and we present statistical evidence that two models of the shrinkage family (LASSO, ELNet, Ridge, and adaptives) feature between the best ones. These models are linear and perform variable selection which is our main objective. Furthermore, they are the only specifications among the best ones capable of predicting the sharp decline in 2008 Q4 inflation.

Moreover, we investigate the variable selection in the sharp disinflation in 2008 Q4 through data mining processes, and try to disentangle the trade-off between variable selection and non-linearity in Random Forest models without conclusive results.

Lastly, in the variable selection analysis, the variables selected after 2008 Q4 are consistent across all remaining period. We find evidence that inflation expectations associated to households are key to predict inflation in line with the explanations used for the MDP in Coibion and Gorodnichenko (2015) and Binder (2015).

## Bibliography

- Atkeson, A. and Ohanian, L. E. (2001). Are Phillips curves useful for forecasting inflation? *Quarterly Review*, 25(Win):2–11.
- Ball, L. and Mankiw, N. G. (1995). Relative-price changes as aggregate supply shocks. *The Quarterly Journal of Economics*, 110(1):161–193.
- Ball, L. and Mazumder, S. (2011). Inflation dynamics and the great recession. *Brookings Papers on Economic Activity*, 42(1):337–405.
- Ball, L. and Mazumder, S. (2014). A Phillips Curve with Anchored Expectations and Short-Term Unemployment. (20715).
- Bernanke, B. (2010). Monetary policy and the housing bubble: a speech at the annual meeting of the american economic association, atlanta, georgia, january 3, 2010. Speech 499, Board of Governors of the Federal Reserve System (U.S.).
- Binder, C. (2015). Whose expectations augment the phillips curve? *Economics Letters*, 136.
- Binder, C. (2017). Measuring uncertainty based on rounding: New method and application to inflation expectations. *Journal of Monetary Economics*, 90(C):1–12.
- Blanchard, O., Cerutti, E., and Summers, L. H. (2015). Inflation and Activity: Two Explorations and Their Monetary Policy Implications. (WP15-19).
- Bobeica, E. and Jarociński, M. (2019). Missing Disinflation and Missing Inflation: A VAR Perspective. *International Journal of Central Banking*, 15(1):199–232.
- Breiman, L. (2001). Random forests. *Machine learning*, 45(1):5–32.
- Christiano, L. J., Eichenbaum, M. S., and Trabandt, M. (2015). Understanding the Great Recession. *American Economic Journal: Macroeconomics*, 7(1):110–167.
- Coibion, O. and Gorodnichenko, Y. (2015). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.

- Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018a). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, 56(4):1447–91.
- Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018b). How do firms form their expectations? new survey evidence. *American Economic Review*, 108(9):2671–2713.
- Coibion, O., Gorodnichenko, Y., Kumar, S., and Pedemonte, M. (2020). Inflation expectations as a policy tool? *Journal of International Economics*, 124(C).
- Coulombe, P. G., Leroux, M., Stevanovic, D., Surprenant, S., et al. (2019). How is machine learning useful for macroeconomic forecasting? Technical report, CIRANO.
- Debelle, G. and Laxton, D. (1997). Is the Phillips Curve Really a Curve? Some Evidence for Canada, the United Kingdom, and the United States. *IMF Staff Papers*, 44(2):249–282.
- Del Negro, M., Giannoni, M. P., and Schorfheide, F. (2015). Inflation in the great recession and new keynesian models. *American Economic Journal: Macroeconomics*, 7(1):168–96.
- Donayre, L. and Panovska, I. (2018). U.s. wage growth and nonlinearities: The roles of inflation and unemployment. *Economic Modelling*, 68:273–292.
- Doser, A., Nunes, R., Rao, N., and Sheremirov, V. (2018). Inflation Expectations and Nonlinearities in the Phillips Curve. (1018).
- Fuhrer, J. C., Olivei, G. P., and Tootell, G. M. B. (2012). Inflation Dynamics When Inflation Is Near Zero. *Journal of Money, Credit and Banking*, 44:83–122.
- Gali, J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, 44(2):195–222.
- Garcia, M. G., Medeiros, M. C., and Vasconcelos, G. F. (2017). Real-time inflation forecasting with high-dimensional models: The case of brazil. *International Journal of Forecasting*, 33(3):679–693.
- Gilchrist, S., Schoenle, R., Sim, J., and Egon, Z. (2017). Inflation Dynamics during the Financial Crisis. *American Economic Review*, 107(3):785–823.
- Gilchrist, S. and Zakrajsek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, 102(4):1692–1720.

- Gordon, R. J. (2013). The phillips curve is alive and well: Inflation and the nairu during the slow recovery. (19390).
- Hall, R. E. (2011). The long slump. *American Economic Review*, 101(2):431 – 69.
- Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2):453–497.
- Haubrich, J., Pennacchi, G., and Ritchken, P. (2012). Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *Review of Financial Studies*, 25(5):1588–1629.
- Heise, S., Karahan, F., and Şahin, A. (2020). The Missing Inflation Puzzle: The Role of the Wage-Price Pass-Through. (27663).
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: applications to nonorthogonal problems. *Technometrics*, 12(1):69–82.
- Krueger, A. B., Cramer, J., and Cho, D. (2014). Are the Long-Term Unemployed on the Margins of the Labor Market? *Brookings Papers on Economic Activity*, 45(1 (Spring)):229–299.
- Llaudes, R. (2005). The Phillips curve and long-term unemployment. (441).
- McCracken, M. W. and Ng, S. (2016). Fred-md: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589.
- McLeay, M. and Tenreyro, S. (2019). Optimal inflation and the identification of the phillips curve. NBER Working Papers 25892, National Bureau of Economic Research, Inc.
- Medeiros, M. C., Vasconcelos, G. F., Veiga, Á., and Zilberman, E. (2019). Forecasting inflation in a data-rich environment: the benefits of machine learning methods. *Journal of Business & Economic Statistics*, pages 1–22.
- Powell, J. (2019). Testimony before the senate banking committee, washington d.c.
- Stock, J. H. and Watson, M. W. (2010). Modeling inflation after the crisis. *Proceedings - Economic Policy Symposium - Jackson Hole*, pages 173–220.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.

- Van Zandweghe, W. (2019). The Phillips Curve and the Missing Disinflation from the Great Recession. *Economic Review*, (Q II):5–31.
- Vasconcelos, G. (2018). Forecasting in high-dimension: Inflation and other economic variables. *PUC-Rio*.
- West, K. D. and Cho, D. (1995). The predictive ability of several models of exchange rate volatility. *Journal of Econometrics*, 69(2):367–391.
- Xie, H. and Huang, J. (2009). SCAD-penalized regression in high-dimensional partially linear models. *The Annals of Statistics*, 37(2):673 – 696.
- Zhao, P. and Yu, B. (2006). On model selection consistency of lasso. *Journal of Machine learning research*, 7(Nov):2541–2563.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429.
- Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320.

## A Data

Here we present the data set. The column `tcode` denotes the following data transformation for a series  $\mathbf{x}$ : (1) no transformation; (2)  $\Delta\mathbf{x}_t$ ; (3)  $\Delta^2\mathbf{x}_t$ ; (4)  $\log\mathbf{x}_t$ ; (5)  $\Delta\log\mathbf{x}_t$ ; (6)  $\Delta^2\log\mathbf{x}_t$ ; (7)  $\Delta(\frac{\mathbf{x}_t}{\mathbf{x}_{t-1}} - 1)$ . The transformations made follow McCracken and Ng (2016) when possible and aim to make the variables stationary in all possible windows associated with the expanding window procedure used in this work. The data sources are in the complementary material.

Inflation						
id	tcode	Code	Description	Frequency	Period	
1	5	cpi_headline	Consumer Price Index for All Urban Consumers: All Items in U.S. City Average	Quarterly	1947-2020	
2	5	pce_headline	Personal Consumption Expenditures: Chain-type Price Index	Quarterly	1947-2020	
3	5	comp_nonfarm	Nonfarm Business Sector: Compensation Per Hour, Index	Quarterly	1947-2020	
4	5	earn_const	Average Hourly Earnings of Production and Nonsupervisory Employees, Construction, Dollars per Hour, Quarterly, Seasonally Adjusted	Quarterly	1947-2020	
5	5	earn_goods	Average Hourly Earnings of Production and Nonsupervisory Employees, Goods-Producing, Dollars per Hour	Quarterly	1947-2020	
6	5	earn_manufac	Average Hourly Earnings of Production and Nonsupervisory Employees, Manufacturing, Dollars per Hour	Quarterly	1939-2020	
7	1	tau	Inflation Permanent Component Estimated as Medeiros et al (2019)	Quarterly	1983-2011	

Expectations						
id	tcode	Code	Description	Frequency	Period	
8	2	mu_h	Mean CPI Inflation Expectations of Highly-Uncertain Consumers	Monthly	1978-2020	
9	2	mu_l	Mean CPI Inflation Expectations of Less-Uncertain Consumers	Monthly	1978-2020	
10	2	cle_fcpi_1y	FED Cleveland CPI Inflation Expectations; 1 year ahead	Monthly	1982-2020	
11	2	gb_fcpi_1q	CPI Inflation Projections from the Greenbooks of the Federal Reserve Board of Governors; 1 quarter ahead	Monthly	1979-2020	
12	2	gb_fcpi_2q	CPI Inflation Projections from the Greenbooks of the Federal Reserve Board of Governors; 2 quarters ahead	Monthly	1979-2020	
13	2	gb_fcpi_3q	CPI Inflation Projections from the Greenbooks of the Federal Reserve Board of Governors; 3 quarters ahead	Monthly	1979-2020	
14	2	gb_fcpi_1y	CPI Inflation Projections from the Greenbooks of the Federal Reserve Board of Governors; 1 year ahead	Monthly	1979-2020	
15	2	msc_mean_fcpi_1y	University of Michigan Survey of Consumers Mean CPI Inflation Expectations; 1 year ahead	Quarterly	1970-2020	
16	2	msc_med_fcpi_1y	University of Michigan Survey of Consumers Median CPI Inflation Expectations; 1 year ahead	Quarterly	1978-2020	
17	2	spf_fcpi_1y	Survey of Professional Forecasters Median Inflation Expectations; 1 year ahead	Quarterly	1970-2020	
18	2	spf_fcpi_1q	Survey of Professional Forecasters Median Inflation Expectations; 1 quarter ahead	Quarterly	1981-2020	
19	2	spf_fcpi_2q	Survey of Professional Forecasters Median Inflation Expectations; 2 quarters ahead	Quarterly	1981-2020	
20	2	spf_fcpi_3q	Survey of Professional Forecasters Median Inflation Expectations; 3 quarters ahead	Quarterly	1981-2020	
21	2	pxl_mean_a1834	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Age Between 18 and 34; 1 year ahead	Quarterly	1978-2020	
22	2	pxl_mean_a3544	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Age Between 35 and 44 1 year ahead	Quarterly	1978-2020	
23	2	pxl_mean_a4554	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Age Between 45 and 54 1 year ahead	Quarterly	1978-2020	
24	2	pxl_mean_a5564	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Age Between 55 and 64, 1 year ahead	Quarterly	1978-2020	
25	2	pxl_mean_a6597	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Age Between 65 and 97; 1 year ahead	Quarterly	1978-2020	
26	2	pxl_mean_m	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Male; 1 year ahead	Quarterly	1978-2020	
27	2	pxl_mean_f	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Female; 1 year ahead	Quarterly	1978-2020	
28	2	pxl_mean_y14	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Bottom 25% Income Group; 1 year ahead	Quarterly	1978-2020	
29	2	pxl_mean_y24	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Second 25% Income Group; 1 year ahead	Quarterly	1978-2020	
30	2	pxl_mean_y34	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Third 25% Income Group; 1 year ahead	Quarterly	1978-2020	
31	2	pxl_mean_y44	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Top 25% Income Group; 1 year ahead	Quarterly	1978-2020	
32	2	pxl_mean_ehs	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, High School or Less; 1 year ahead	Quarterly	1978-2020	
33	2	pxl_mean_esc	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, High School Degree; 1 year ahead	Quarterly	1978-2020	
34	2	pxl_mean_ecd	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, College Degree; 1 year ahead	Quarterly	1978-2020	
35	2	pxl_mean_egs	University of Michigan Survey of Consumers Mean CPI Inflation Expectations, Graduate Studies; 1 year ahead	Quarterly	1978-2020	
36	2	pxl_med_a1834	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Age Between 18 and 34; 1 year ahead	Quarterly	1978-2020	
37	2	pxl_med_a3544	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Age Between 35 and 44 1 year ahead	Quarterly	1978-2020	
38	2	pxl_med_a4554	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Age Between 45 and 54 1 year ahead	Quarterly	1978-2020	
39	2	pxl_med_a5564	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Age Between 55 and 64, 1 year ahead	Quarterly	1978-2020	
40	2	pxl_med_a6597	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Age Between 65 and 97; 1 year ahead	Quarterly	1978-2020	

41	2	pxl_med_m	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Male; 1 year ahead	Quarterly	1978-2020
42	2	pxl_med_f	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Female; 1 year ahead	Quarterly	1978-2020
43	2	pxl_med_y14	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Bottom 25% Income Group; 1 year ahead	Quarterly	1978-2020
44	2	pxl_med_y24	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Second 25% Income Group; 1 year ahead	Quarterly	1978-2020
45	2	pxl_med_y34	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Third 25% Income Group; 1 year ahead	Quarterly	1978-2020
46	2	pxl_med_y44	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Top 25% Income Group; 1 year ahead	Quarterly	1978-2020
47	2	pxl_med_ehs	University of Michigan Survey of Consumers Median CPI Inflation Expectations, High School or Less; 1 year ahead	Quarterly	1978-2020
48	2	pxl_med_esc	University of Michigan Survey of Consumers Median CPI Inflation Expectations, High School Degree; 1 year ahead	Quarterly	1978-2020
49	2	pxl_med_ecd	University of Michigan Survey of Consumers Median CPI Inflation Expectations, College Degree; 1 year ahead	Quarterly	1978-2020
50	2	pxl_med_egs	University of Michigan Survey of Consumers Median CPI Inflation Expectations, Graduate Studies; 1 year ahead	Quarterly	1978-2020
51	2	expec_bl_cpi	Backward Looking CPI Inflation Expectations as Average of Last 4 Quarters	Quarterly	1978-2020

**Economic Activity**

id	tcode	Code	Description	Frequency	Period
52	6	tcu	Capacity Utilization: Total Index, Percent of Capacity	Quarterly	1967-2020
53	6	gdp_real	Real Gross Domestic Product, Billions of Chained 2012 Dollars	Quarterly	1947-2020
54	6	gdp_gap	gdp_real - Real Potential Gross Domestic Product, Billions of Chained 2012 Dollars Estimated by U.S. Congressional Budget Office	Quarterly	1949-2020
55	5	cfnai	Chicago Fed National Activity Index	Monthly	1967-2020
56	6	st_unemp	Short-Term Unemployment Rate by Bureau of Labor Statistics	Monthly	1960-2020
57	6	full_unemp	Full Unemployment Rate by Bureau of Labor Statistics	Monthly	1960-2020
58	6	full_unemp_st_gap	full_unemp - Natural Rate of Unemployment (Short-Term) Estimated By U.S. Congressional Budget Office	Quarterly	1960-2020
59	6	full_unemp_lt_gap	full_unemp - Natural Rate of Unemployment Estimated By U.S. Congressional Budget Office	Quarterly	1960-2020
60	6	st_unemp_gap	st_unemp - Natural Rate of Unemployment (Short-Term) Estimated By U.S. Congressional Budget Office	Quarterly	1960-2020
61	6	full_unemp_t4	full_unemp - full_unemp 4 Quarters Before	Quarterly	1961-2020
62	6	bz_average_unemp_gap	st_unemp_gap but using the mean of last 4 quarters in each metric	Quarterly	1961-2020
63	6	rgap_full_unemp	The Recession Gap Transformation created by Stock Watson (2010) applied to full_unemp	Quarterly	1961-2020
64	6	rgap_st_unemp	The Recession Gap Transformation created by Stock Watson (2010) applied to st_unemp	Quarterly	1961-2020
65	6	full_unemp_st_DLgap	full_unemp_st_gap relative to the level of full Unemployment as in Debelle-Laxton (1997) and Ball, Mazumder (2011)	Quarterly	1961-2020
66	6	full_unemp_lt_DLgap	full_unemp_lt_gap relative to the level of full Unemployment as in Debelle-Laxton (1997) and Ball, Mazumder (2011)	Quarterly	1961-2020
67	6	st_unemp_DLgap	st_unemp_gap relative to the level of full Unemployment as in Debelle-Laxton (1997) and Ball, Mazumder (2011)	Quarterly	1961-2020
68	5	rgap_gdp_real	The Recession Gap Transformation created by Stock Watson (2010) applied to gdp_real	Quarterly	1948-2020
69	5	rgap_tcu	The Recession Gap Transformation created by Stock Watson (2010) applied to tcu	Quarterly	1968-2020
70	6	rgap_cfnai	The Recession Gap Transformation created by Stock Watson (2010) applied to cfnai	Monthly	1968-2020
71	5	real_ulc	Nonfarm Business Sector Real Unit Labor Costs from the U.S. Bureau of Economic Analysis	Quarterly	1955-2020
72	5	labor_share	Nonfarm Business Sector: Labor Share, Index 2012=100	Quarterly	1967-2020
73	5	ulc	Nonfarm Business Sector: Unit Labor Cost, Percent Change From Quarter One Year Ago	Quarterly	1948-2020
74	5	real_ulc_logdev	Residuals of a Regression using Log of real_ulc on a linear trend as in Van Zandweghe (2019)	Quarterly	1955-2020

**Economic Shocks and Financial Variables**

id	tcode	Code	Description	Frequency	Period
75	1	oil_eia_nchange	Nominal Change in Oil Prices as Estimated by U.S. Energy Information Administration	Quarterly	1974-2020
76	1	oil_eia_rchange	Real Change in Oil Prices as Estimated by U.S. Energy Information Administration	Quarterly	1974-2020
77	2	gordonprod	Productivity Measure Calculated by Gordon (2013)	Quarterly	1949-2020

78	6	wb_food	World Bank Food Prices Index	Quarterly	1960-2020
79	5	import_deflator	Imports of goods (implicit price deflator)	Quarterly	1947-2020
80	2	MS_import_prices_rel	Imports of goods (implicit price deflator) Relative to GDP Deflator	Quarterly	1947-2020
81	1	head_core	cpi_headline transformed - Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average Inflation	Quarterly	1957-2020
82	1	head_med	cpi_headline transformed - Median CPI Inflation as calculated by the FED Cleveland	Quarterly	1983-2020
83	1	import_head	import_deflator transformed - cpi_headline transformed	Quarterly	1947-2020
84	5	comm_price_index	Producer Price Index by Commodity: All Commodities, Index 1982=100	Quarterly	1913-2020
85	3	corp_bond_spread	Spread Between Moody's Seasoned Aaa Corporate Bond Yield and Moody's Seasoned Baa Corporate Bond Yield, Percent, Quarterly	Quarterly	1986-2020
86	3	GZ_credit_spread	Gilchrist, Zakrajšek (2012) Credit Spread Index	Monthly	1973-2020
87	2	GZ_ebp	Gilchrist, Zakrajšek (2012) Excess Bond Premium Estimates	Monthly	1973-2020
88	2	MORTGAGE30US	30-Year Fixed Rate Mortgage Average in the United States, Percent	Quarterly	1971-2020
89	2	DGS10	10-Year Treasury Constant Maturity Rate, Percent, Quarterly	Quarterly	1976-2020
90	2	T10YFFM	Spread Between 10-Year Treasury Constant Maturity Rate and Fed Funds Rate, Percent	Quarterly	1976-2020
91	2	DGS2	2-Year Treasury Constant Maturity Rate, Percent, Quarterly	Quarterly	1976-2020
92	2	T2YFFM	Spread Between 2-Year Treasury Constant Maturity Rate and Fed Funds Rate, Percent	Quarterly	1976-2020
93	2	MORTGAGESPREAD	MORTGAGE30US - Effective Federal Funds Rate, Percent	Quarterly	1971-2020

## B Random Forest

We repeat the same 1 step ahead forecasting exercise with the RF/OLS and adaLASSO/RF models described in Section 3.3.5.

With the RF/OLS we can check the performance of the linear model using variables selected by the Random Forest algorithm. We compare the predictive results of this model with the pure Random Forest model using the Model Confidence Set (MCS) procedure. If the difference between the results is not statistically significant, we understand that non-linearity is not a issue.

The second class of adapted models is the adaLASSO/RF. If adaLASSO/RF performs similarly to the Random Forest - using the same MCS procedure to test it - we understand that variable selection in RF is less relevant and non-linearity is more important.

After the forecast estimations, we get a set of models  $\mathcal{M}^{\text{Hybrid}}$  with 18 elements - 6 corresponding to each class of models (pure RF, RF/OLS, and adaLASSO/RF). We use these predictions to calculate the squared errors and run the MCS procedure.

We find that  $n(\mathcal{M}_{\alpha=0.8}^{\text{Hybrid}}) = 14$ ,  $n(\mathcal{M}_{\alpha=0.9}^{\text{Hybrid}}) = 9$  and  $n(\mathcal{M}_{\alpha=0.95}^{\text{Hybrid}}) = 1$ . The singular model selected when  $\alpha = 0.95$  is a RF/OLS specification. With  $\alpha = 0.9$ , models from the 3 groups are selected.

Therefore, there's evidence that the RF/OLS is the best model, which implies that variable selection is more relevant than non-linearity.

As a last exercise, we try to access the variables' importance for the Random Forest models selected in  $\mathcal{M}_{\alpha=0.95}^*$  (see Table 4.7 for details) using the mean decrease in accuracy.

Basically, after training the forest we take one variable, scramble it in a way that preserves the distribution of the variable, and measure the predictive accuracy of the resulting tree. For each variable, we repeat this process in all trees. The whole procedure is then replicated for all variables.

This loss in accuracy for each variable is measured by the mean decrease in accuracy. The more the accuracy falls, the more important the variable is for the regression tree. Since the trees are trained on highly correlated variables,

they tend to be resistant to permutations in single variables. Thereby, we focus on the ranking of the variable.

In Figures B.1, B.2, and B.3 we exhibit the variables in descending order of importance as calculated from the mean decrease in accuracy. We report the results for the 3 Random Forest specifications in  $\mathcal{M}_{\alpha=0.95}^*$  for the desired time frame (2007 Q4 - 2011 Q4). We choose to present only the top 8 of 93 possible variables.

There are similar results from the adaLASSO variable selection performed in Section 4.3. When possible, the variables `px1_mean_a1834_1d`, `gb_fcpi_1q_1d`, or `mu_1_1d` are present in the top 8 rankings. However, a deeper analysis is not trivial, since there is no clear outlier in mean decrease accuracy importance for any model at any quarter.







## C The Sharp Fall of 2008 Q4 - Details

First, we identify the variables which have the biggest impact on the 1 step ahead 2008 Q4 forecast using the LASSO and adaLASSO models selected in Section 4.3. The results for both models are qualitatively the same, thereby we present only the tables and figures relative to the adaLASSO specification.

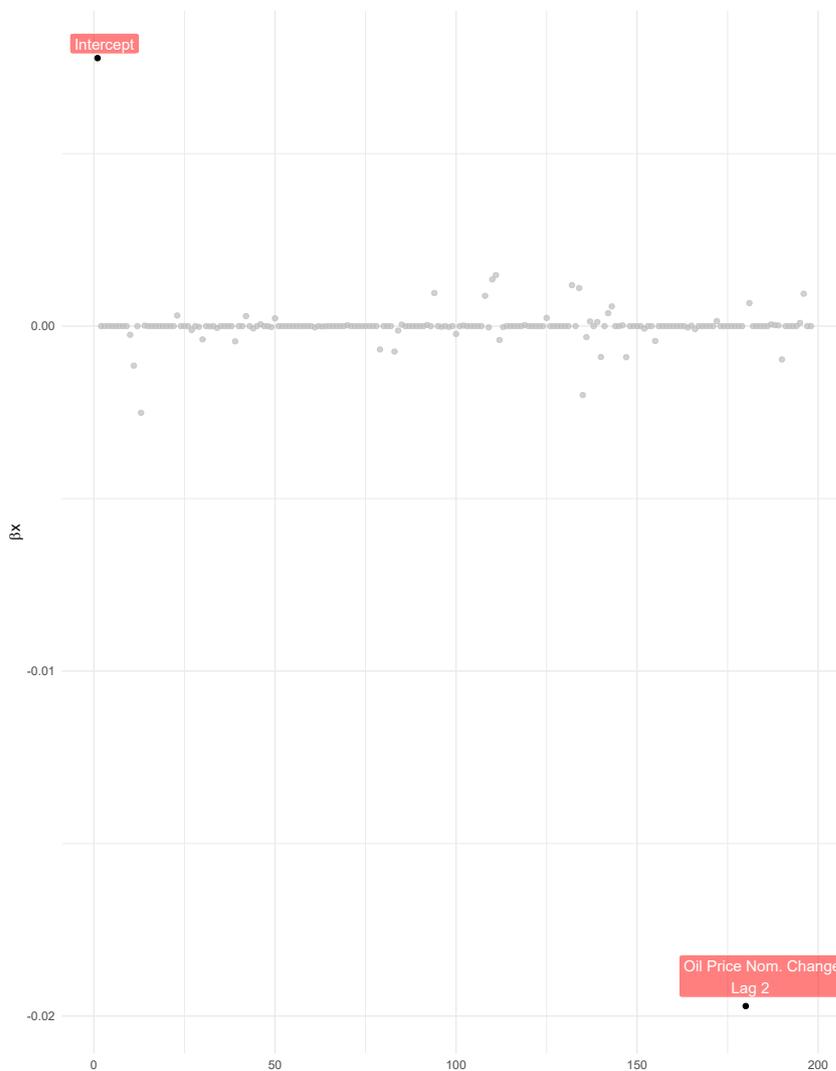


Figure C.1: Regressors' Impact on the 2008 Q4 Forecast for the adaLASSO Model

Figure C.1 presents the result of this exercise. The lag 2 of nominal

change in oil prices is a clear outlier. By removing it and repeating the same exercise, we achieve a similar result with the lag 2 of real change in oil prices having the biggest impact on the forecast. Both variables have a correlation of 0.9837.

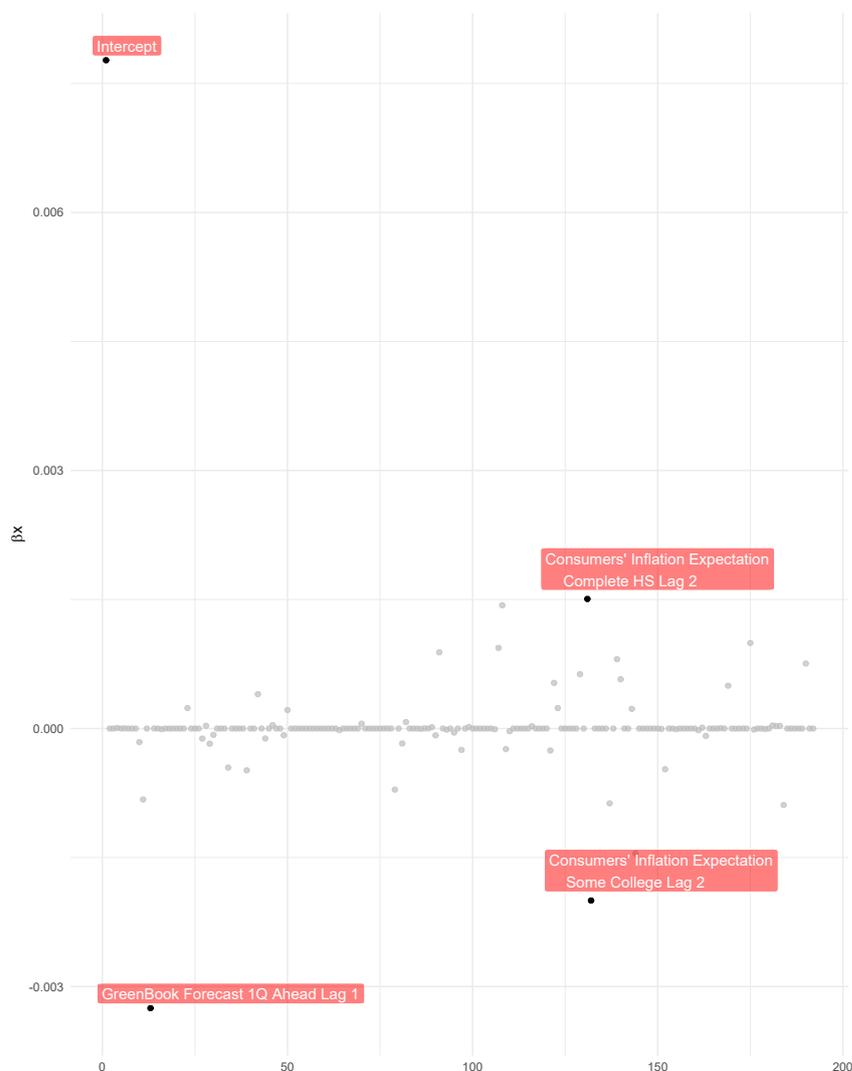


Figure C.2: Regressors' Impact on the 2008 Q4 Forecast for the adaLASSO Model



Figure C.3: Chosen adaLASSO Specification Predictive Performance With Selected Variables

We remove both of them from the data set and perform the forecasting procedure of Section 4. Figure C.2 show that after removing the oil prices there is no clear outlier anymore. Figure C.3 exhibits the forecasting of the adaLASSO model without the selected variables. As we remove the nominal change in the price of oil at  $t - 2$ , and the real change in the price of oil at  $t - 2$ , the model becomes unable to predict the sharp fall in 2008 Q4.

Moreover, we remove in addition to oil prices the `gb_fcpi_1q_1d` and collect the forecasts for the studied time period. Figure C.4 presents the result.

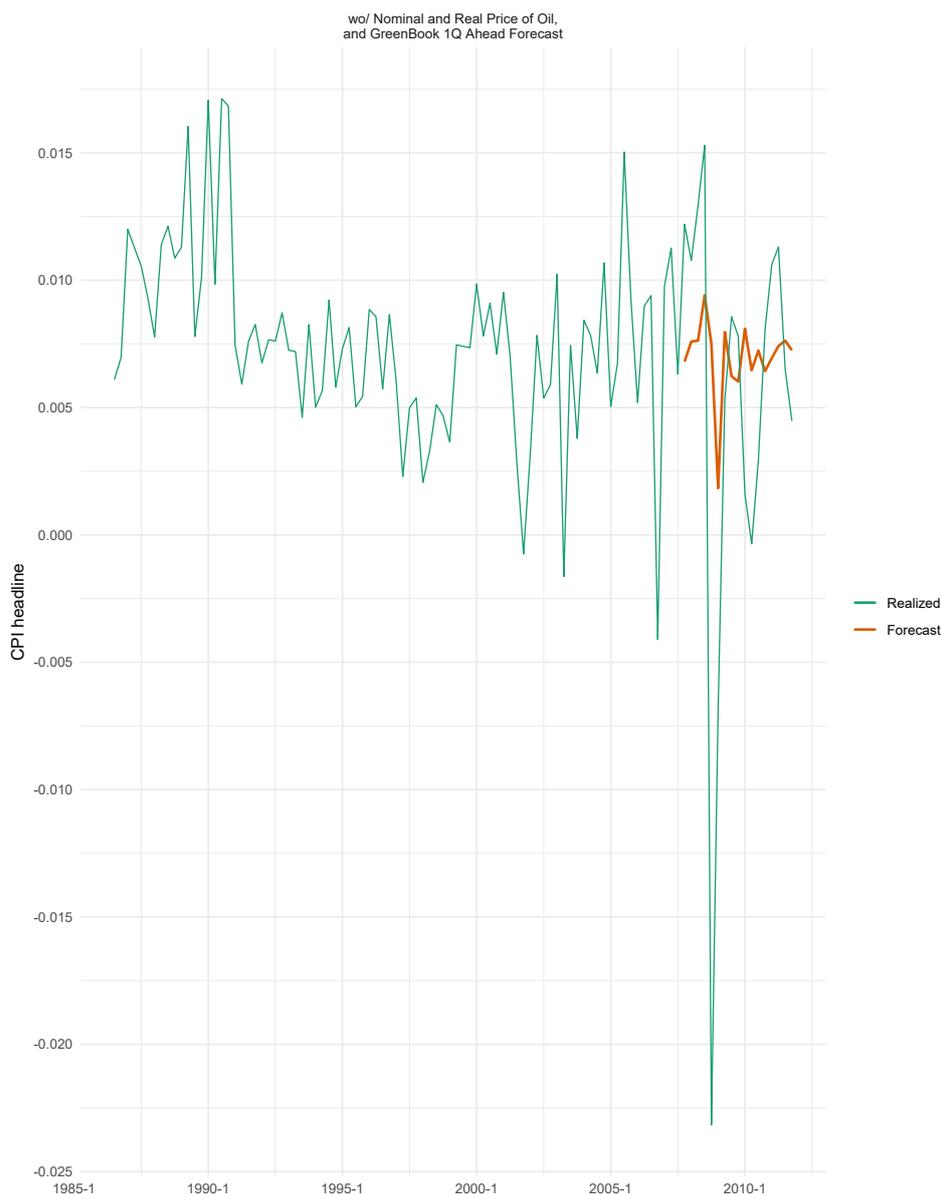


Figure C.4: Chosen adaLASSO Specification Predictive Performance With Selected Variables

Lastly, we perform the aforementioned forecasting exercise with a simple linear model using the change in nominal oil prices and GreenBook 1 quarter ahead inflation expectations. Notwithstanding their impacts on the adaLASSO specification forecast for 2008 Q4, the linear model is unable to predict the sharp fall in 2008 Q4.

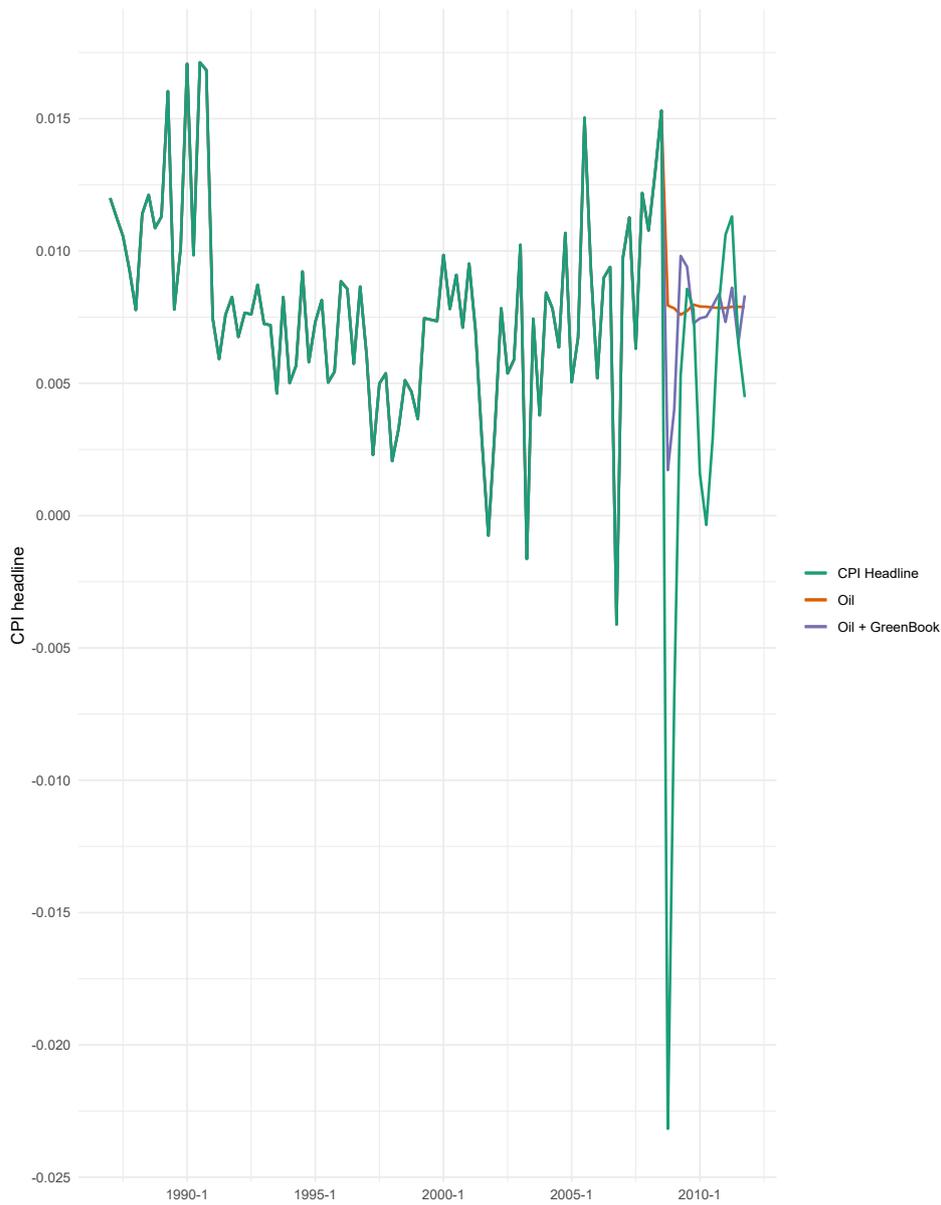


Figure C.5: Linear Model Predictive Performance With Selected Variables

The results for the first method are not conclusive. Next, we present the results of the second method described in Section 4.3.1 :



Figure C.6: Models' Specifications According to Number of Variables Utilized I

From the total of 78 estimated models, we exhibit the 3 specifications which have the best predictive performance for 2008 Q4 in Figure C.6. The specification which minimizes the squared error uses 18 variables. The lag 2 changes in nominal- and real oil prices, and 1 quarter ahead GreenBook inflation expectations are contained among them.

We also run the MCS procedure on the 78 created models' squared errors and collect the specifications associated with  $\alpha = 0.95$  and their forecasts. At this  $\alpha$  value, only the specification associated with 8 variables is selected. As shown in Figure C.7, no model is able to predict the magnitude of 2008 Q4 disinflation.

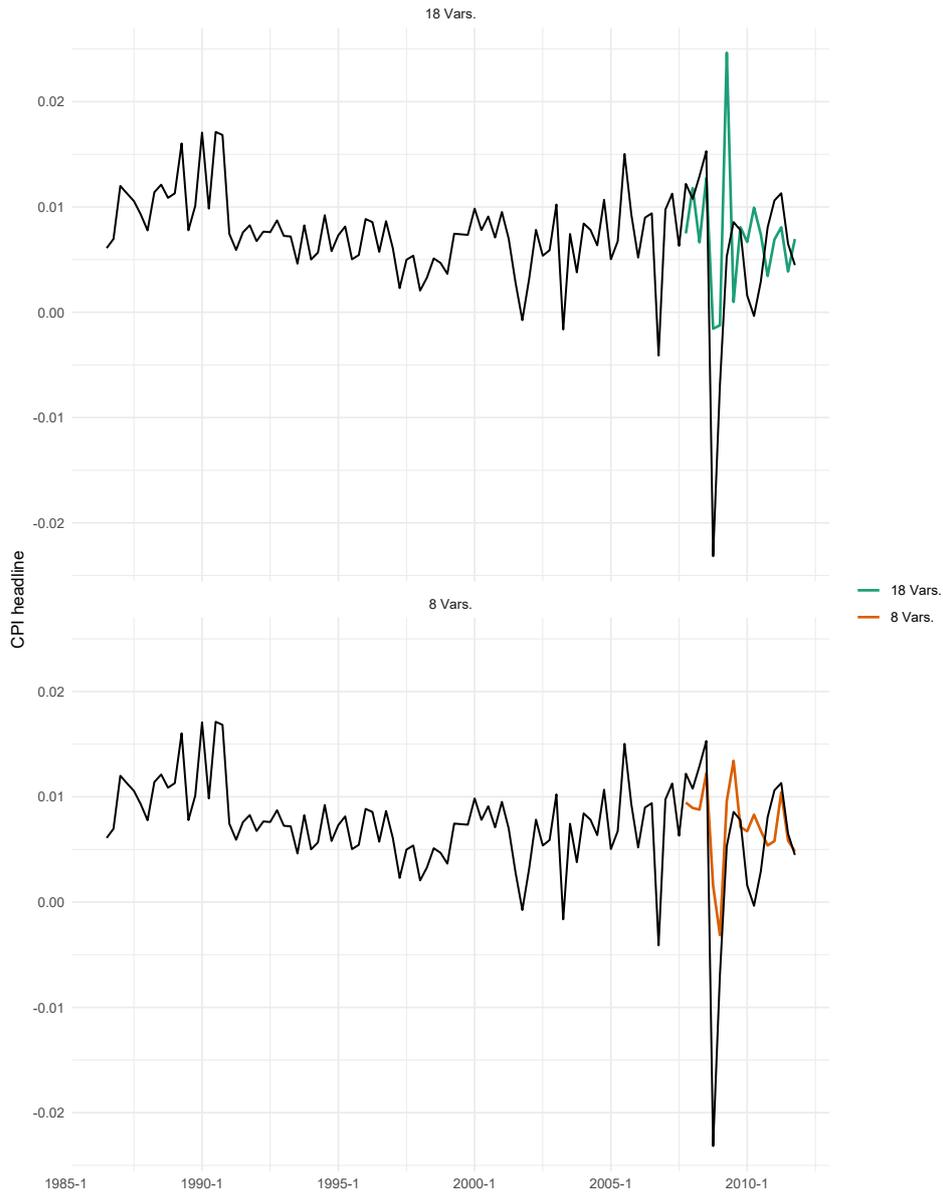


Figure C.7: Models' Specifications According to Number of Variables Utilized II

Lastly, we show the results of the third proposed procedure in Section 4.3.1. The selected model uses 16 variables for its estimation and its forecast is represented in Figure C.8. The GreenBook 1 quarter ahead inflation expectations and oil prices figure among the selected variables. but in general the variables used in this method and in the second method are not the same.

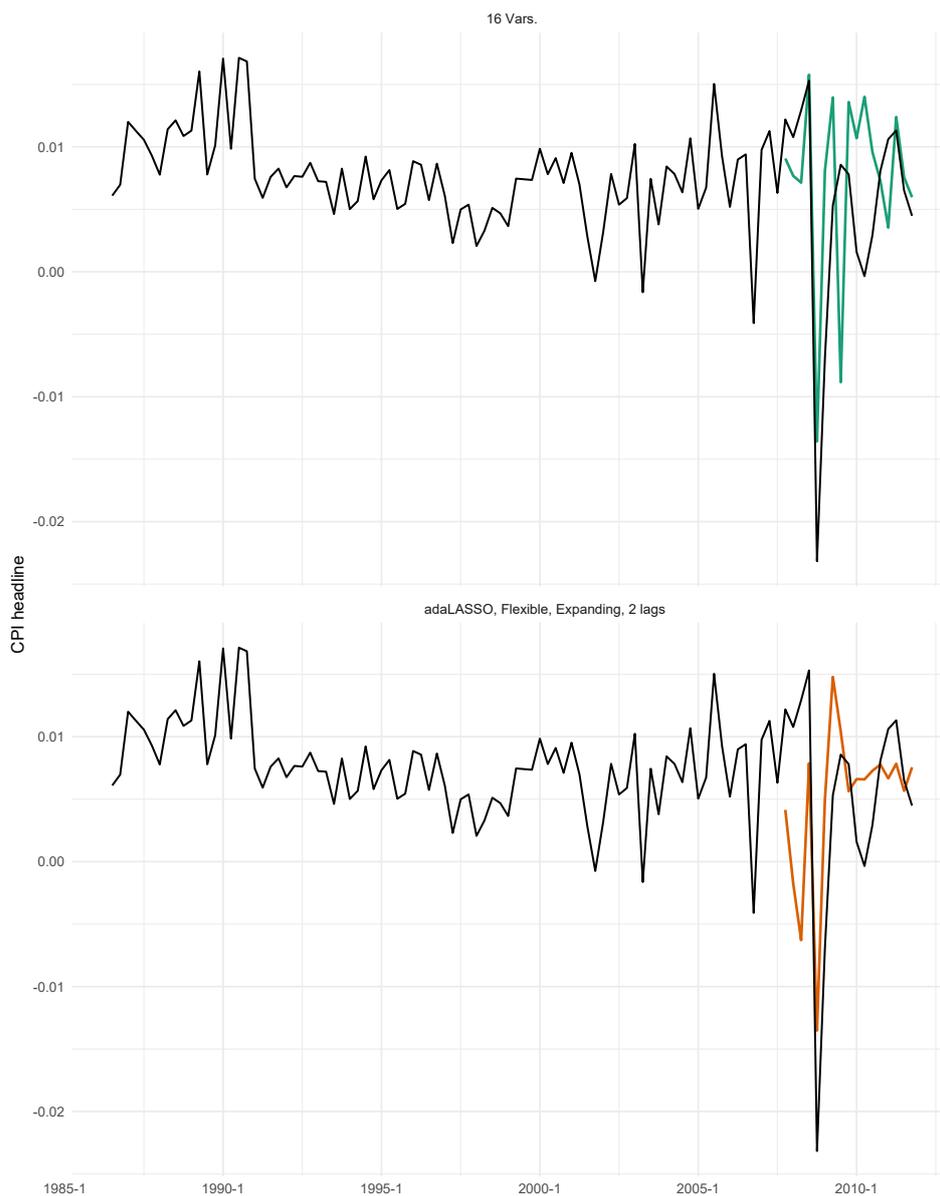


Figure C.8: adaLASSO Specification versus Selected Model

None of these data mining procedures generates conclusive results. All of them in the end yields models' specifications, which are not able to isolate a particular effect or explanation for the adaLASSO model being able to predict the fall in 2008 Q4.

## D Computer Codes

All models are estimated in R using standard and well-established packages. For the shrinkage models we use the `glmnet` and `HDEconometrics` packages. For the RF models, we use the `ranger` package. For the SCAD models, the `ncvreg` package. We follow the algorithm of Medeiros et al. (2019) for the Model Confidence Set and UC-SV.

For all methods within the LASSO family, the penalty parameter is selected using the Bayesian Information Criteria (BIC). As noted previously, we set  $\kappa = 0.5$  in the Elastic Net's family models. For the RF models, each individual tree is grown until there are only 5 observations in each leaf. The number of bootstrap samples we use is 500.