Grab a Bite? Prices in the Food Away From Home Industry During the Covid19 Pandemic*

Diego Solórzano†
Banco de México

[*Work in progress. Please do not circulate without consent.*]

May 2022

Abstract

Due to social distancing measures, among other health related policies, the Food Away From Home (FAFH) industry has been one of the most affected by the Covid19 pandemic. As FAFH prices account for a non-negligible share of CPI baskets, this study examines FAFH price-setting behaviour in Mexico City from two complementary angles. First, using web scraped data from an online food ordering and delivery platform, classified through machine learning techniques, this study shows that in 2020 and 2021 (i) independent and multi-outlet restaurants report similar price trends; (ii) prices of soups and beverages without alcohol, potentially substituted by home-production, exhibit low price growth rates; (iii) in contrast, prices of mains and desserts have been on the rise; (iv) the heterogeneous growth rates across dish categories seem to be explained by the extensive margin; and (v) episodes associated to an escalation in Covid19 cases seem to increase price rigidities. Second, employing shrinkage and non-linear machine learning frameworks able to deal with large number of explanatory variables, this study analyses potential price determinants. Contrary to what is expected, it seems that costs stemming from electricity, permanent workers and prices of beans and rice are more relevant than those from gas, temporal workers and meat prices as inflation determinants for this industry in Mexico City. The results from this paper adds on understanding price-setting during the fast-changing Covid19 pandemic.

JEL: E31, L16.

---

*The views in this paper are exclusively the responsibility of the author and do not necessarily reflect those of Banco de México. I would like to thank Denisse Dueñas and Briseida Galván for excellent research assistance, as well as comments from participants in the *Non-traditional Data, Machine Learning and Natural Language Processing in Macroeconomics Conference* and seminar participants at BoM. This paper has benefited from Banco de Mexico's EconLab initiative. All remaining errors are my own.

†Banco of México, Economic Research Division, Av 5 de Mayo 18, Mexico City, Mexico. E-mail: jsolorzano@banxico.org.mx

1 Introduction

The Food Away From Home (FAFH) industry has been one of the most affected by the Covid19 pandemic. Temporal closures, restricted on-site dining capacity, as well as suppressed demand are among the supply and demand shocks deemed by the industry as it is considered not essential.[1] Understanding how FAFH prices evolve as the pandemic unfolds is important for the conduction of monetary policy as it accounts for a non-negligible weight in the CPI basket. This is particularly true in Emerging Market Economies. For instance, FAFH alone represents around 9% and 11% in the Mexican headline and core CPI, respectively.

This paper studies the FAFH inflation in Mexico City from two complementary angles. First, using web scraped data from an online ordering and delivery platform, this paper examines price dynamics stemming from over 120 million prices from more than 1.7 million dishes offered by around 30,000 restaurants in Mexico City. In particular, the first part of this study focuses on analyzing whether aggregate price levels by type of restaurant (e.g. independent or multi-outlet) and type of dish (e.g. starters or desserts) have shown heterogeneous growth rates during the Covid19 pandemic. Aggregate price dynamics are then decomposed into the extensive and intensive margins of price adjustments, and whether these stylized facts of price changes exhibited diverse patterns around stiff episodes of the Covid19 pandemic. Machine learning techniques on text analysis are deployed in order to deal with the unstructured nature of the dataset, compiled by Banco de México, and it covers since the start of the pandemic, from April 2020 to December 2021.

Second, using shrinkage and non-linear machine learning frameworks, the latter part of this study analyses 86 price determinants likely considered by price-setters surveyed in Mexico City's FAFH CPI and provides a counterfactual of the inflation rate stemming from costs pressures during the Covid19 pandemic. These series reflect costs pressures from three factor markets: wholesale food prices, labor costs and utility bills. The adoption of supervised learning models, as instead of standard regression analysis, is dictated by the large set of likely price determinants, as well as the potential non-linearities each of them carry in the production function of restaurants, and therefore their price-setting decisions. Using the FAFH inflation rate in Mexico City at different horizons as target variable, the machine

[1] Restaurants, fast-food establishments, pizza places, taco shops, among others establishments that provide ready-to eat meals (with or without premises for on-site dining) are considered as part of the FAFH industry. FAFH prices are encompassed in the "Accommodation and Leisure" section in ILO's CPI Handbook.

learning algorithms are trained from 2006 to 2017 in order to assess and learn the relevance of price determinants. Then, based on the state of the determinants, inflation counterfactuals are computed for the 2018-2021 period. Thus, having measured the role of price determinants in pre-pandemic times, this study provides estimates on the FAFH inflation in Mexico City stemming sorely from local input costs and compare them to the observed inflation rate. Hence, as the analysis provides evidence on inflation drivers, it sheds further light in our understanding this phenomena as firms in this industry reopen.

The results from the first part of the study suggest that, first, the growth rate of the aggregate price index computed using web scraped prices from the online platform lead the FAFH CPI in Mexico City since the start of the pandemic. This is in line with the rapid adoption of online ordering and delivery platforms by consumers in 2020. However, the gap between the proposed index and the FAFH CPI in Mexico City has diminished in recent months. Second, multi-outlet and independent restaurants have shared similar price trends during the pandemic. Nonetheless, the aggregate price level of independent restaurants is smoother than multi-outlet eateries. Staggered versus synchronized price adjustments are likely to be behind this result. Third, soups and beverages (without alcohol), which consumers could potentially substitute with home-production, exhibit lower growth rate, on average, than other categories. In contrast, prices of mains and desserts have been on the rise. Forth, by decomposing price changes into extensive and intensive margins, the lower growth rate in the aforementioned dish categories comes from the former. Fifth, using a simple panel data framework, estimates indicate that periods associated with great number of contagions seem to mute price-setting decisions in both margins of adjustment.

With respect to the analysis of price determinants, the results from the second part of this study can be summarized as follows. Electricity fees, the real wage bill and real average wage of permanent workers, as well as wholesale prices of beans, rice and shrimps best describe the FAFH inflation rate in Mexico City at different horizons. Regarding utilities, although in the top 10 of inflation drivers, gas LP and natural gas are not the most important inflation determinants as one might have expected. With respect to labor costs, wage indicators from permanent workers generally outperform wage statistics form temporal workers in predicting the FAFH inflation rate in Mexico City. Lastly, the frequent presence of beans and rice as sides in most meals in the Mexican cuisine could be a potential explanation on their predictive power. In terms of meat, the price of shrimps in wholesale food markets in Mexico

City is a robust explanatory variable across horizons. The price of beef only determines the current inflation rate, suggesting perhaps from a faster cost pass-through.

The results from the former and latter parts of the paper are reconciled through the models' predictions for the 2020 FAFH inflation reported in the latter part. That is, predictions drawn from costs pressures suggest an uptick in the inflation rate in late 2020. However, the FAFH CPI in Mexico City, computed with prices gathered through direct visits to restaurants, remained with little variation throughout 2020 (due to temporal closures and infection risks in on-site dining). Thus, predictions seem to have missed the low inflation rates in 2020. Nonetheless, the experimental index in the first part of the paper, calculated using web scraped data, shows greater pace on price hikes.[2] In fact, predictions made by both Ridge and Random Forest Regressions hinted to a 6% year-on-year FAFH inflation rate by the end of 2020, the experimental index reported a cumulative 5% inflation rate from April to December 2020, while survey data outlined in the CPI reported a 3% year-on-year inflation in December 2020.[3] Hence, despite being different methodologically both price indices (experimental and CPI), the variation of the experimental index is in line with predictions from cost-related variables trained on predicting survey data outcomes. The short time span prevents having a definitive answer but reconciling inflation measures using web scraped data and survey data remains a promising venue of research.[4]

The study of price-setting by type of restaurant (independent or multi-outlet) arises from the literature on firm dynamics showing that firms respond heterogeneously to shocks depending on its characteristics. For instance, in the context of price-setting behaviour, Gilchrist et al. (2017) find that liquidity constrained firms in the US increased prices in 2008, while their unconstrained competitors cut prices. In contrast, I find suggestive evidence that independent and multi-outlet restaurants exhibit similar price trends, partially rationalized by the different nature of shocks. Moreover, some dishes on restaurants' menus are more likely to be substituted by home-production when ordering food delivery. This substitution is not possible when on-site dining is the more prevalent channel of consumption. One might think of soups, beverages or salads in this situation, while mains, desserts and

[2] Explained in great detail in Section 4, the experimental price index is computed by chain linking an index to the products' average variation (unweighted).

[3] I am unable to compute the y-o-y inflation rate of the experimental price index as the data collection of the online ordering and delivery platform started in April 2020.

[4] Flower (2019) and Konny et al. (2019) argue on the benefits and challenges of using web scraped prices when computing price indices.

alcoholic beverages (cocktails) could be more difficult to substitute by home-production. Another type of substitution could come in the form of value of time, potentially associated with the household size. That is, it is better to cook for few than cook for one. For instance, Cortes and Pan (2013) show that outsourcing home production increased female labor participation in Hong Kong. Although in a different context, their findings highlight that cooking time for some type of dishes might make them more prone to substitution, leading to a strategic response from multi-product price-setters, such as restaurants.

Despite the volume, velocity and variety features in the dataset under study in the first part of the paper, web scraped data sources often arrive with little structure for answering relevant economic questions. It is not the exception in this paper. Hence, I use machine learning classifiers for tackling this problem. After examining a battery of classifiers, a multinominal regression proved as the most accurate in classifying dishes in the manually created training set. The model is then deployed in the complete dataset and maps dishes into one of 18 proposed categories. The type of dishes used in this paper are common headers in restaurant menus, like Starters, Soups, Salads, Mains, Tacos, Pizzas, etc.

This paper is related to three strands in the literature. First, although web scraped data is increasingly used for analyzing inflation and its macroeconomic implications, to the author's knowledge this is the first paper studying inflation from the FAFH through the lens of web scraped data. The literature has mainly focused on goods' prices observed at supermarkets or departmental stores.[5] In contrast, this research focuses on prices from an industry in the service sector. Since the FAFH market accounts for a non-negligible weight in the CPI basket, the results from this paper are relevant for policymakers and central banks interested in understanding the price formation process in this industry.

Second, the use survey micro-data from the underlying FAFH component of the CPI. For instance, Hobijn et al. (2006) report that restaurant prices in the euro area increased dramatically after the introduction of the Euro, while EU countries that did not adopt the euro did not observed such increase. While the Covid19 pandemic does not provide a clear focal period for resetting prices as the Euro changeover in Hobijn et al. (2006), I do find firms concentrate otherwise staggered price increases around periods with a downward trend

[5] See, among others, Cavallo (2018); Peña and Prades (2021); Solórzano (2021) for an overview of web scraped prices and its implications on price-statistics typically employed by macroeconomics models. Hull et al. (2017) and Macias et al. (2019) used web scraped consumer prices for improving their inflation forecasting models.

of infection rates. Furthermore, Fougère et al. (2010) study the impact of minimum wage increases in France on price quotes from restaurants encompassed in the french CPI.[6] Consistent with the literature, I find that, among numerous inflation drivers in different factor markets, labor costs remain a key determinant in FAFH inflation.[7]

Third, not directly addressing price responses but rather analyzing the economic implications of the Covid19 pandemic in the FAFH industry, Fetzer (2022) reports that an intervention designed to actively increase demand for on-site dining contributed to subsequent clusters of new infections in the UK.[8] Fetzer (2022) elaborates on the known risks of infection in restaurant settings, leading to low demand for on-site dining. However, this paper studies prices from an online ordering and delivery platform, which experienced high demand in the wake of the Covid19 pandemic.

The paper is organized as follows. Section 2, presents the data characteristics of web scraped prices from the online ordering and delivery platform. Section 3 describes the classification of the unstructured data at hand. Section 4 centers at presenting aggregate price indices used to analyze the FAFH price dynamics during the pandemic. Section 5 outlines the stylized facts on the frequency and size of price adjustments. Section 6 presents price determinants and counterfactuals of the FAFH inflation rate in Mexico City stemming from costs pressures in 2020 and 2021. Section 7 concludes.

2 Data Description

The dataset used in this research comes from daily observations of dishes advertised by restaurants in an online food ordering and delivery platform in Mexico City.[9] The price collection, compiled by Banco de Mexico, is carried out by a robot parsing the platform's website. In broad terms, the robot's price collection consists in gathering data from each and every dish or item displayed on the platform. That is, the robot collects the product's

[6]Card and Krueger (1994) is other influential study on the fast-food industry. Though, the authors mainly focus on employment.

[7]However, the inflation drivers stemming from the labor market used in this paper do not distinguish between minimum wage workers or not.

[8]The policy, "Eat out to help out" (EOHO), subsidised the cost of meals and non-alcoholic drinks by up to 50% across participating restaurants across the UK for meals served on Mondays-Wednesdays (capped to GBP 10 per person). González-Pampillón et al. (2021) also look at the EOHO scheme and find that it induced higher footfall and increased recruitment in the industry.

[9]For confidentiality reasons, the name of the platform cannot be disclosed. Nonetheless, the dataset is available for research purposes through a non-disclosure agreement with Banco de México's EconLab.

identifier, description and price for each dish, as well as the restaurant offering the dish.

The dataset at hand starts in April 1st 2020 and ends in June 30th 2021.[10] Figure 1 provides some descriptive statistics on the price collection task. Panel 1a depicts the number of observations (dishes) reported by the robot throughout the day. Notably, prior April 2021, the robot took around 17 hours for parsing all items in the platform. The collection time has decreased to 12 hours on average since then. The number of items by the minute is greater in recent months as highlighted by the red and orange colors in the figure.[11]

Panel 1b summarizes the number of food items or dishes (observations) and restaurants reported on a daily basis. All in all, the median number of observations and restaurants per day is about 210,000 and 5,000, respectively. One can observe a steady increase in both the number of meals and restaurants between April and June 2020. This period was characterized by the toughest restrictions policies in terms of social distancing and temporal closures of restaurants in Mexico City due to Covid19's first wave. From July 2020 and before February 2021, the numbers remain fairly stable. Finally, there seems to be a new average number of dishes and restaurants since March 2020.

Figure 1: Data Collection

Hourly and Daily Observations

(a) Robot's Collection (b) Daily Observations

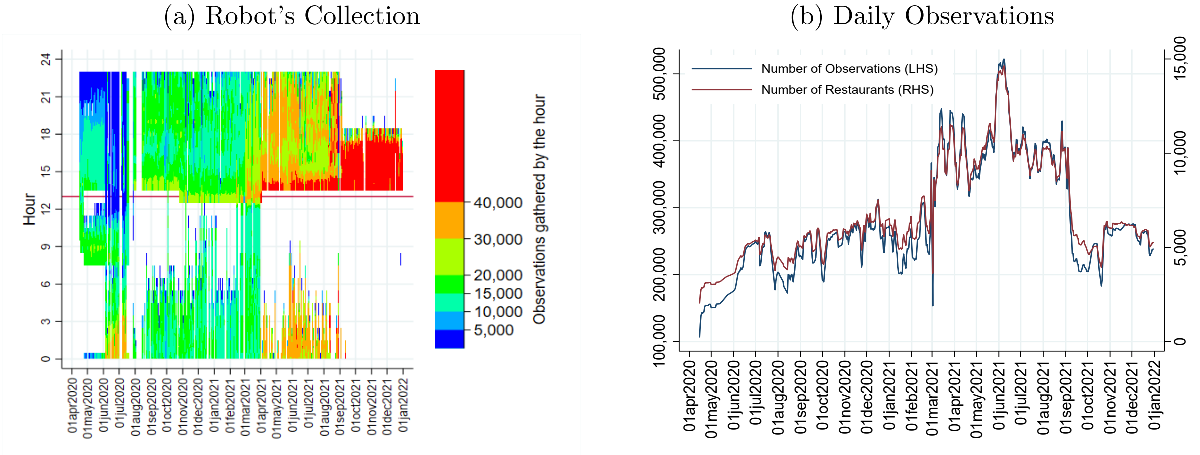

Figure 2 shows the distribution of dishes and restaurants by the number of days in the

[10] The state of emergency in Mexico City started in late March 2020. Hence, it is not possible to have a pre-pandemic benchmark unfortunately. Nonetheless, because of the rapid transition of events, it is believed the dataset encompasses the responses taken by restaurants and the platform in the early stages of the pandemic as it is later described in the paper.

[11] Changes in the offer of meals and restaurants, coupled with modifications on the platform's operation and adjustments in the robot's execution explain the different colors in the graph.

sample. For instance, Panel 2a reports that over 1.7 million different meals have appeared in the sample. On average, a meal is observed for nearly 95 days in the sample. Moreover, Panel 2b summarizes the distribution of the panel of restaurants by the number of days in the dataset. There has been over 33,000 firms in the sample, each of them appearing, on average, about 120 days in the sample.

Figure 2: Histogram on the Panel of Dishes and Restaurants

(a) Dishes

(b) Restaurants

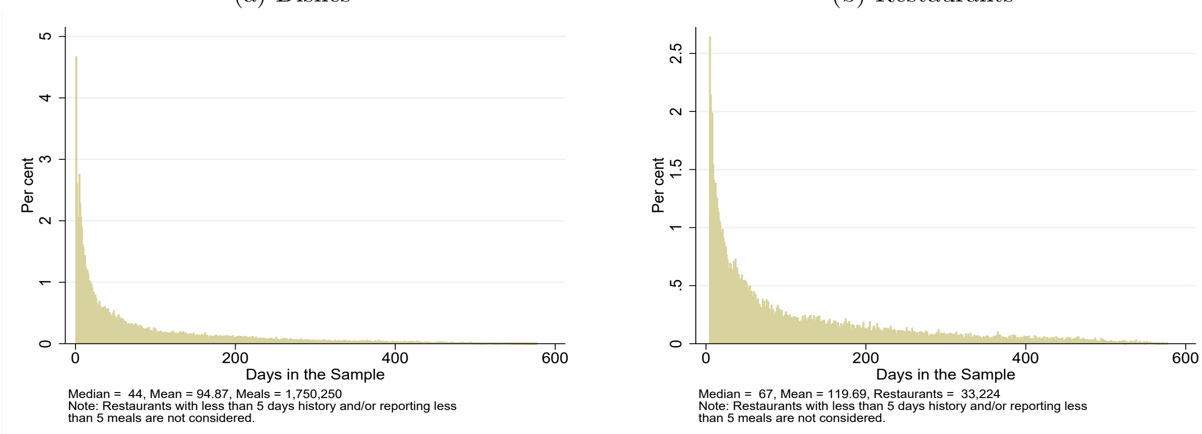

Finally, Figure 3 depicts if a given restaurant identifier (y-axis) is effectively observed on a given day (x-axis). The number of scatters makes it difficult to observe when a single scatter disappear. Nonetheless, when few firms stops appearing in the sample, it is reflected as an horizontal white bar. One of these bar can be observed starting in October 2020, for instance. Also, there are a number of ids which had disappeared from the sample in the last three months.[12]

3 Data Classification

All the benefits of big data sources do not come without a price. Perhaps one of the most cited drawbacks stemming from the use of big data is its unstructured nature. Unstructured data can undermine all the benefits of these granular and fast arriving datasets. To this end, the deployment of machine learning techniques come in hand for this task.

[12]Figure 3 and Panel 2b, on top of some investigative work on the dataset, suggest of some obfuscation strategy i.e. the restaurant's id changes over time. Although this strategy might widens standard errors in the panel data estimates below, the classification and computation of experimental prices indices are less sensitive to this strategy.

Figure 3: Heatmap of Restaurants

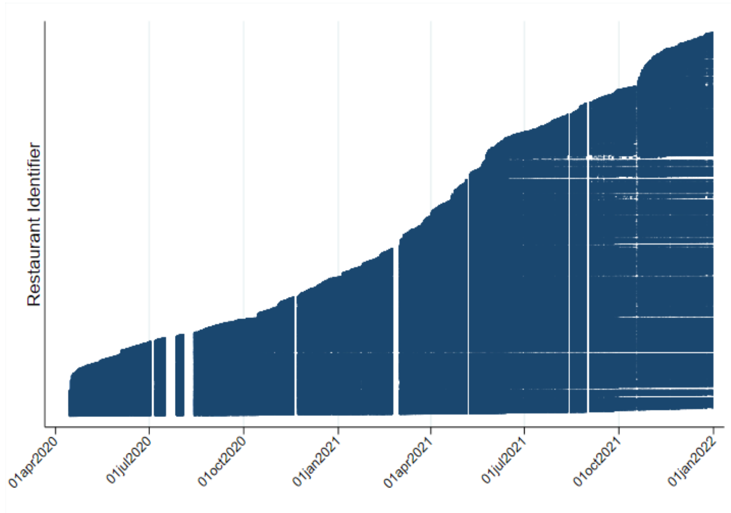

Hence, in this Section I provide details on the classification of the dataset in two dimensions. The first classification divides dishes (observations) into 18 categories. These categories come from common headers in restaurants' menus (e.g. starters, desserts), as well as few subcategories contained in the Mexican CPI (e.g. pizza or grilled chicken). The second classification opens up firms into three types of restaurants: independent, with branches and franchises.

3.1 Dish Classification

Dishes (cross-section dimension of the panel) are classified using machine learning techniques. In what follows, I briefly outline the step taken for this task. A detailed description of the classification process, as well as some forensic statistics on the classifiers performance are left in the Appendix.

3.1.1 Automated Classification

First, this approach requieres the construction of a manually produced training set, under which a number of algorithms are trained. To that end, out of the around 616,000 unique descriptions in the dataset, I manually classify more than 13,000 random dishes based on the descriptions provided by the restaurants. Thus, the manual classification considers a little more than 2% of the dishes in question.

The dishes are classified into 19 categories. The categories are: (1) Starters, (2) Salads,

(3) Soups, (4) Eggs, (5) Mains, (6) Pizzas, (7) Tacos, (8) BBC, (9) Grilled and Roasted Chicken, (10) Desserts, (11) Beverages with Alcohol, (12) Beverages without Alcohol, (13) Meals with Beverages, (14) Meals without Beverages, (15) Group Combos, (16) Dessert Combos, (17) Extras, (18) Others (Non-Food) and (19) Ambiguos.[13] These categories are chosen on the basis of (i) well-recognized headers in many restaurants' menus, (ii) categories with direct mapping to Mexico's CPI categories and (iii) research question at hand.

Second, after applying text cleaning procedures to the dataset, I convert the collection of dish descriptions into a matrix of token (words) counts.[14] Specifically, the matrix contains unigrams (single words) and bigrams (pair of consecutive words) in the descriptions.[15,16] The matrix of unigrams and bigrams has over 32,000 columns, which are then used as explanatory variables by the classifiers.

Third, the classifiers used for this analysis are (i) decision tree, (ii) random forest, (iii) multinomial naive Bayes and (iv) multinominal logistic regression. All classifiers require some form of hyper-parameter selection prior to estimation. To that end, I use k-fold cross validation procedures, which is exposed in great detail in the Appendix.[17]

Forth, after training the classifiers using 80% of the training set, algorithms are deployed over the remaining (unseen) 20% of the manually constructed training set.[18]

As shown in great detail in the Appendix, the *multinomial logistic regression* is picked as the winner across models. It is the one with greatest accuracy (average point estimate), as well with the lowest computational time. Figure 4 depicts the confusion matrix on the prediction of dish labels using the multinomial logistic regression fitted under the complete training

[13]BBC stands for Barbacoa, Birria and Carnitas, which are common taco fillings, and are considered in the Mexican CPI as a specific product category. Meals with/without Beverages consider two or three times meals. For instance, a Meal with Beverage could be a bundle of starter, salad, main and a soft drink.

[14]That is, the columns in the matrix represent each and every single word appearing at least once in the collection of descriptions, the rows of the matrix are the dishes in the dataset, and each matrix cell counts the number of times a word (column) appears in the description (row). See the Appendix for more.

[15]For instance, the unigram representation of "Today is Monday" is ["Today", "is", "Monday"], while the bigram representation is ["Today is", "is Monday"].

[16]Unigrams and bigrams with a frequency less or equal than three in the overall word count of the dataset are neglected. In the Appendix I show there are minor differences in the classifiers' performance when using (i) unigrams only with the same cut-off threshold (frequency less or equal than three) and (ii) unigrams only with no cut-off threshold (universe of words in the dataset).

[17]Importantly, the grid of parameters used for this search is detailed in Table 5 and the optimal set of parameters are reported in Table 6. Also, as category sizes are highly unbalanced in the training set, I outline in the Appendix the steps taken to overcome the over specialisation of classifiers stemming from this feature in the dataset.

[18]Forensic statistics on the performance of the various classifiers over the different matrices of token (words) counts are depicted in Figure 15 in the Appendix.

set. It provides a graphical representation on whether predictions are accurate relative to the true values. Each cell reports the share of each instance such that every row (true labels) adds up to one. Correct predictions lay in the diagonal, values outside the diagonal highlight prediction errors. As shown in Figure 4, most cells on the diagonal report values close to one.

Figure 4: Multinomial Logistic Regression Confusion Matrix
Predictions Over the Entire Training Set

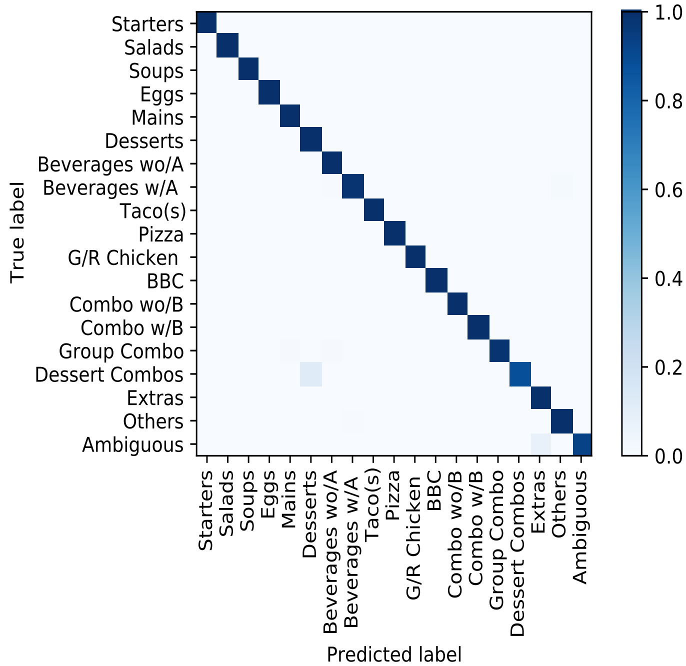

Finally, Table 1 adds on the impact of the machine learning techniques used in this research. The first bloc of columns reports the composition of the manually classified dataset. The second bloc of columns summarizes the outcome labels generated through the logistic regression. Thus, the classification burden of large and fast arriving data is alleviated, while minimizing the classification errors, through the use of machine learning techniques. In turn, this allows us to shed further insights on the highly detailed data at hand.

3.2 Restaurant Classification

The dataset is also classified regarding the nature of the restaurant. As shown by Gilchrist et al. (2017), financially constrainted price-setters increased prices in the 2008 Global Financial Crisis, while their unconstrained counterparts cut prices. Thus, I classify restaurants

Table 1: Dishes' Labels

By Classification Approach

	Course	Classification			
		Manual		Machine Learning	
	Type	Count	Share (%)	Count	Share (%)
1	Starters	640	5.24	16,625	3.27
2	Salads	996	8.15	15,049	2.96
3	Soups	64	0.52	5,604	1.10
4	Eggs	447	3.66	8,819	1.73
5	Mains	3,948	32.31	251,969	49.54
6	Desserts	1,063	8.70	55,010	10.81
7	Beverages wo/Alcohol	2,639	21.60	89,490	17.59
8	Beverages w/Alcohol	208	1.70	9,539	1.88
9	Tacos	429	3.51	23,256	4.57
10	Pizzas	1,444	11.82	20,819	4.09
11	Grilled/Roasted Chicken	15	0.12	345	0.07
12	BBQ, Birria, Carnitas	23	0.19	788	0.15
13	Combo wo/Beverage	55	0.45	355	0.07
14	Combo w/Beverage	100	0.82	7,647	1.50
15	Group Combo	124	1.01	3,171	0.62
16	Dessert Combo	25	0.20	175	0.05
	Total	12,220	100.00	508,661	100.00

Note: Extras, Others and Ambiguous are also considered in the classification exercise but not reported.

into independent, with branches and belonging to a franchise chain.[19] This classification is carried out manually as there is little uncertainty on the classification rules in place.

The composition by type of restaurant is summarized in Figure 5. First, Panel 5a shows that, although restaurant chains have multiple outlets across Mexico City, they constitute a small fraction of restaurants in the sample. In fact, the relative size of independent restaurants has been growing since the pandemic started. Presumably, before the pandemic, restaurant chains were more likely to outsource their online ordering and delivery services to platforms like the one under study. As the pandemic advanced and temporal retail closures were ordered, independent restaurants had no other option than use the online ordering and delivery services.

Panel 5b depicts that restaurant with branches and belonging to a franchise chain offer in general more dishes than independent restaurants. Also, the figure suggests a small trend on the median number of dishes offered by restaurants regardless its type.

[19]Franchise chain are well-known restaurants brands often found on high streets and shopping centers. These restaurants normally has sister-brands and belong to a corporates reporting their balance sheets as they participate in financial markets. Restaurants with branches are those sharing the exact same name (or in some cases the neighbourhood is added to the name e.g. "Taco Shop ABC Reforma" and "Taco Shop ABC Insurgentes"). These groups of restaurants typically operate only in Mexico and are ofter family run. They may or may not participate in financial markets. The remaining restaurants are considered independent.

Figure 5: Dataset Composition by Type of Restaurant

(a) Share of Restaurants (b) Number of Meals

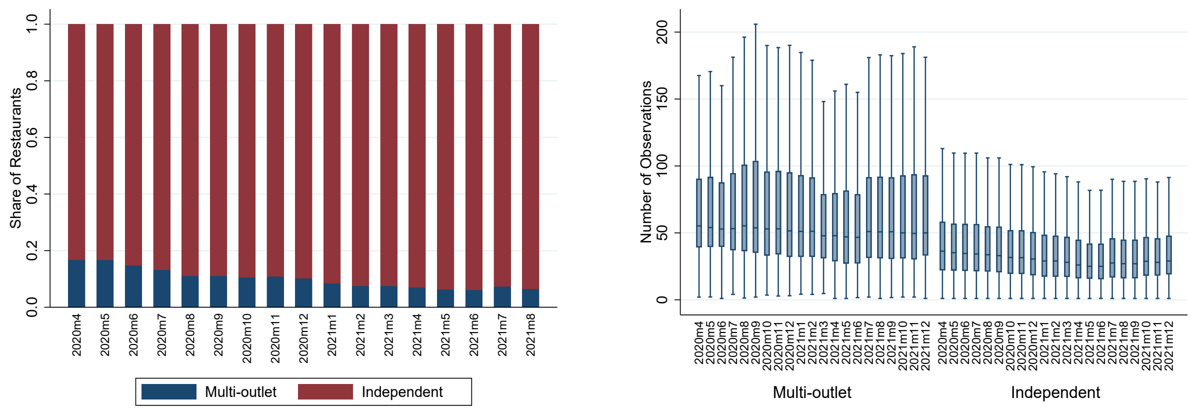

4 Experimental FAFH Price Indices

This Section provides evidence on the evolution of FAFH prices in Mexico City during the pandemic based on the dataset above described. Aggregate price measures, summarizing millions of prices, are computed using two types of price indices. Although they do not follow the orthodox CPI methodology, they share broad similarities as it is explained below.

These aggregate price measures, or experimental prices indices, are reported (i) at the aggregate, (ii) by type of restaurant and (iii) by type of dish. The first one provides a general picture on price dynamics. The second two decompose the heterogeneity in price-setting across dimensions that help understanding the inflation formation process in this industry.

As in any price survey, the treatment of product churn can have important implications on aggregate statistics. This is particularly important in the context of web scraped data as it has been reported that it tends to present greater product churn than survey data.[20] In order to ameliorate this problem, I opt for using dishes (and therefore restaurants) observed at least one day in a given fortnight and appearing in at least 75% of the analyzed fortnights.

4.1 Definition of Price Indices

The first price index is named *"Average Variation Index"* or *AVI*. It is computed as:

[20]See, for instance, Solórzano (2021) and Flower (2019).

$$y_t = \Pi_{i \in \Theta} \left(\frac{p_{i,t}}{p_{i,t-1}} \right)^{\frac{1}{N_t}}$$
$$AVI_t = y_t AVI_{t-1}$$

The term y_t computes the geometric average of prices changes in fortnight t relative to t-1 from dishes observed in at least 75% of fortnights, $i \in \Theta$. The average variation y_t is then chain linked to a Jevons index $Apr2020 = 100$. Note, if a dish is observed in t-1 but not t, it is not considered for the geometric average at time t.[21] Since this index compares a fixed basket of goods between t and t-1, this index is somewhat similar to the methodology followed by most CPIs.[22]

The second index is named *"Average Price Index"* or *API*, which is calculated as:

$$x_t = \Pi_{i \epsilon \Theta} (p_{i,t})^{\frac{1}{N_t}}$$
$$API_t = \frac{x_t}{x_{t-1}} API_{t-1}$$

The term x_t computes the geometric average of prices in fortnight t relative to t-1 from dishes observed in at least 75% of fortnights, $i \in \Theta$. Then, the variation in x_t is chain linked to a Jevons index $Apr2020 = 100$. Hence, API index is a Unit Value Index.[23] Contrary to *AVI*, *API* does consider the entry and exit of goods from one period to the next one (limited by the definition of Θ though).

4.2 FAFH Price Evolution During the Covid19 Pandemic

Figure 6 provides evidence on the dynamics shown by FAFH prices in Mexico City for the sample of restaurants considered in this study. It also overlays the dynamics from the FAFH component of the official CPI for both Mexico City and National (overall).

There are a number of facts to highlight. First, Panel 6a illustrates AVI exhibited a steady increase from April to July 2021, period in which it built a gap with respect to Mexico City's official FAFH CPI. In fact, Mexico City's FAFH CPI lagged behind for nearly a year, from June 2020 to April 2021. Second, Mexico City's FAFH CPI exhibited changes in the positive

[21]By not considering the item in the average nor imputing a zero variation, this approach is equivalent as if the average variation was imputed to dishes not observed in fortnight t. In fact, imputing the average variation of observed goods on missing goods is a common approach used in price surveys by NSOs.

[22]AVI does not follow a fixed basket of goods in all periods as the CPI. It encompasses limited entries and exits of products according to the definition of set Θ.

[23]Recent studies using Unit Value Indices in the context of price data are Diewert (2020); Diewert and Fox (2020); **?**; among others.

trend at times of uplifts of restrictions in terms of the pandemic. That is early December 2020 and July 2021. Third, Mexico City's CPI closed significantly the gap AVI in July 2021.

Panel 6a also illustrates AVI by type of restaurant. According to AVI, both types of restaurants follow a very similar trend. Nonetheless, it seems like price-setters with branches tend to move first than independent restaurants. As Hobijn et al. (2006) document for restaurants in the euro area during the currency exchange over, the bumpy price adjustment form restaurants with branches could be explained by the synchronization of this type of restaurants, while the AVI for independent eateries could be a reflection of the staggered process among them.

Figure 6: Experimental Price Indices

Overall and By Type of Restaurants

(a) Average Variation Index (b) Average Price Index

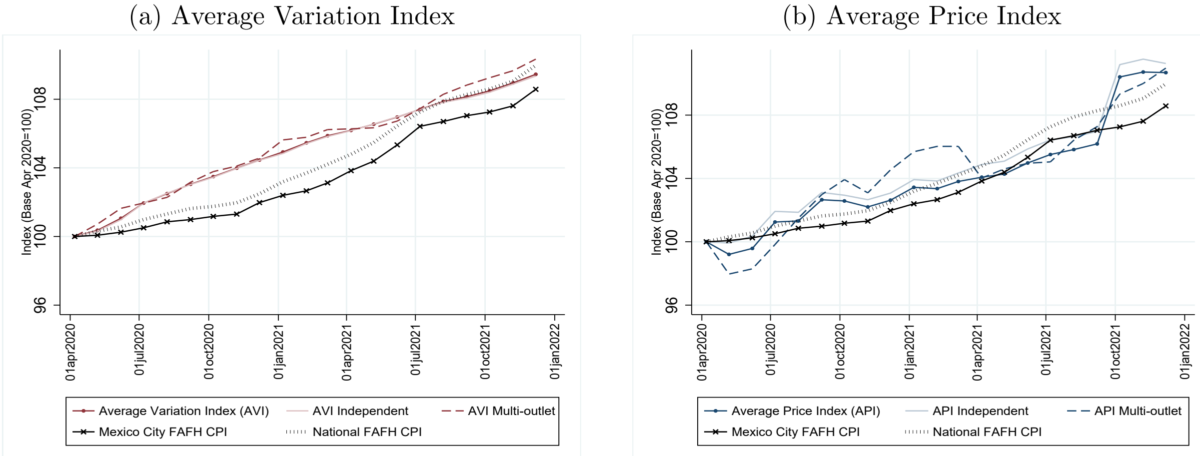

Table 2: Cross-Correlation of Monthly Inflation

From Experimental Price Indices and CPI

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Average Price Index (API)	1.000							
(2) API Independent	0.984***	1.000						
(3) API Multi-outlet	0.471***	0.330**	1.000					
(4) Average Variation Index (AVI)	0.174	0.137	0.255*	1.000				
(5) AVI Independent	0.151	0.120	0.193	0.969***	1.000			
(6) AVI Multi-outlet	0.110	0.076	0.305*	0.256*	0.010	1.000		
(7) Mexico City FAFH CPI	-0.014	-0.022	-0.106	-0.054	-0.091	0.120	1.000	
(8) National FAFH CPI	-0.070	-0.074	-0.052	-0.116	-0.159	0.159	0.748***	1.000

*** p<0.01, ** p<0.05, * p<0.1

On the other hand, API has shown greater divergence between indices. Since October 2020, restaurant with branches has exhibited greater average price than its independent

Table 3: Cross-Correlation of Monthly Inflation in 2021

From Experimental Price Indices and CPI

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Average Price Index (API)	1.000							
(2) API Independent	0.992***	1.000						
(3) API Multi-outlet	0.269	0.161	1.000					
(4) Average Variation Index (AVI)	-0.298	-0.305	-0.004	1.000				
(5) AVI Independent	-0.287	-0.278	-0.170	0.912***	1.000			
(6) AVI Multi-outlet	-0.074	-0.115	0.408**	0.371*	-0.041	1.000		
(7) Mexico City FAFH CPI	0.000	0.012	-0.193	0.210	0.057	0.316	1.000	
(8) National FAFH CPI	-0.276	-0.278	-0.047	0.351*	0.132	0.511**	0.724***	1.000

*** p<0.01, ** p<0.05, * p<0.1

counterparts. The temporal inclusion of pricy items could be behind this differential. See Panel 6b. Though, the gap between types of restaurants has been closing since April 2021.

Figure 7 and Figure 8 summarize the evolution of prices according to the dish classification proposed in Section 3. AVI in Panel 7a suggests that Soups has systematically reported the lowest average price variation among dishes. Beverages without Alcohol and Tacos are categories also reporting lower price increases relative to the start of the pandemic. However, the difference of these dishes with respect to the rest is considerable smaller than for the Soups case. In contrast, Beverages with Alcohol and Desserts have been leading the level level since the start of the pandemic.

If one allows for a more flexible stance in terms of entry and exit of goods, as API does, Panel 8a shows that Starters, Eggs and Salads have increased their price more relative to other categories since the start of the pandemic. Similarly to Panel 7a, Soups and Beverages without Alcohol are among with the lowest increases.

5 Stylized Facts from FAFH Prices in the Pandemic

This Section presents quantitative evidence of the frequency and size of price adjustments. These price statistics have been widely study using survey data as they are informative in the calibration of New Keynesian models. Price moments are computed following standard procedures in the literature (e.g. Bils and Klenow (2004) and Nakamura and Steinsson (2008)).

Using daily data, I fit a linear probability model of the form:

$$P\left(y_{i,j,t}=1|x\right)=\beta_1 x_{dishtype}+\theta_j+\theta_t+\varepsilon_{i,j,t} \tag{1}$$

Figure 7: Average Variation Index

By Dish Type

(a) Categories 1

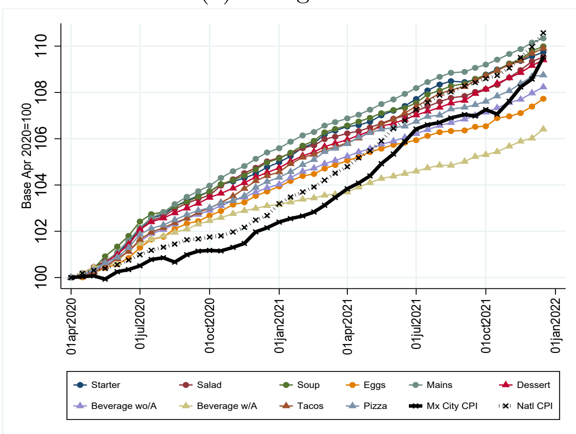

(b) Categories 2

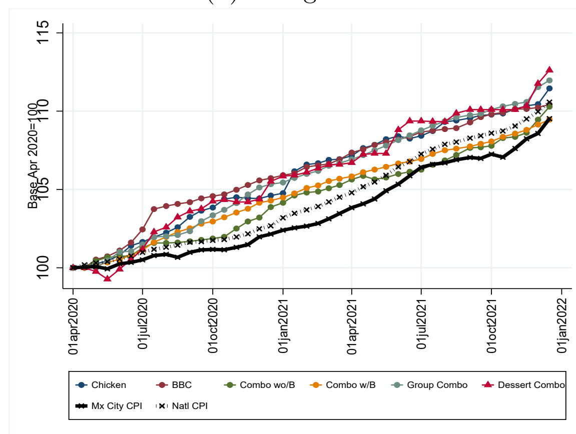

Figure 8: Average Price Index

By Dish Type

(a) Categories 1

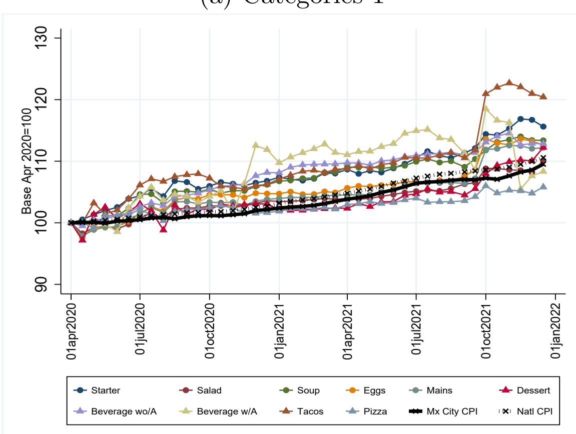

(b) Categories 2

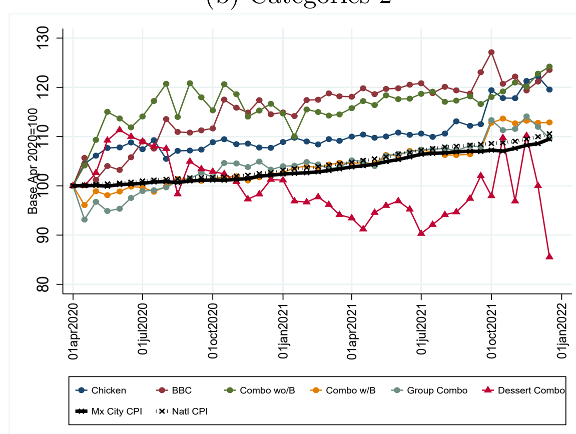

where $y_{i,j,t} = 1$ is a dummy variable if the price of product i at restaurant j on day t changed with respect to day $t-1$, $\Delta p_{i,j,t} \neq 0$, or zero otherwise. θ_j and θ_t represent restaurant and time fixed effects, respectively. Additionally, I decompose price hikes and price drops by running the same model using $y^{Hikes}_{i,j,t} = 1$ if $\Delta p_{i,j,t} > 0$ and zero otherwise; as well as $y^{Drops}_{i,j,t} = 1$ if $\Delta p_{i,j,t} < 0$ and zero otherwise.

A second equation analyzes the size of price adjustments, given a price change:

$$\mid \Delta p_{i,j,t} \mid = \beta_1 x_{dishtype} + \theta_j + \theta_t + \varepsilon_{i,j,t} \quad (2)$$

where $\mid \Delta p_{i,j,t} \mid$ is the absolute value of (log) price changes.[24] Similarly to the linear probability model, two further models are estimated for price hikes ($\Delta p_{i,j,t} > 0$) and price drops ($\Delta p_{i,j,t} < 0$).

Figure 9 summarizes the stylized facts of price setting across dish types.[25] First, Panel 9a highlights that Mains, Tacos, Pizzas, BBC, Combo with Beverages and Group Combos adjusted more frequently their prices than other categories in 2020 and 2021, mainly driven by more frequent price hikes, while price drops remain fairly similar across categories. In contrast, categories with less frequent price changes are Desserts, Beverages without and with Alcohol.

Second, Panel 9b reports the average size of price adjustment, given a price change, by type of dish. Beverages without and with Alcohol are the categories exhibiting greater prices changes, due primary to positive price changes. Given a price change, Mains, Pizza, BBC and Group Combo are those with smaller price changes.

Thus, Figure 9 provides insights on the price levels described in the previous Section and shown in Figure 7. For instance, categories with the greatest AVI increase from April 2020 to December 2021, like Mains, BBC and Group Combos, exhibit more frequent price changes (extensive margin) than other categories; but when their prices change, they do so by a smaller margin (intensive margin) than other categories. Pizzas also show similar price-setting dynamics- more frequent but by smaller amounts, although its price level does not stand out from the rest. In contrast, categories with the least AVI increase in the same period of time, like Beverages without and with Alcohol, report less frequent price changes while their size of adjustment is greater when they do change.

All in all, the heterogeneous pricing behavior along the extensive and intensive margins across their different type of products suggest that price-setters in the FAFH industry do not follow a single pricing rule. Instead, they are able to juggle the margins of price adjustments for their different products.

In the Appendix, I provide further robustness checks on the heterogeneity of price adjustments. First, the above regressions control for time fixed effects that control for time-specific

[24] The log price change is defined as $\Delta p_{i,j,t} = ln(P_{i,j,t}) - ln(P_{i,j,t-1})$, where $P_{i,j,t}$ is the nominal price of product i as observed from the OOD platform. One way to avoid price hikes and price drops cancelling out in the regression coefficient is taking the absolute value of prices changes. This is a common practice in studies analyzing nominal rigidities, like Bils and Klenow (2004), Nakamura and Steinsson (2008), Dhyne et al. (2009), among others.

[25] Regression estimates are reported in Table 10 and Table 11 in the Appendix.

Figure 9: Stylized Facts of Price Adjustments

Representative Dishes
By Sign of Price Adjustment

(a) Frequency of Changes

(b) Size of Adjustments

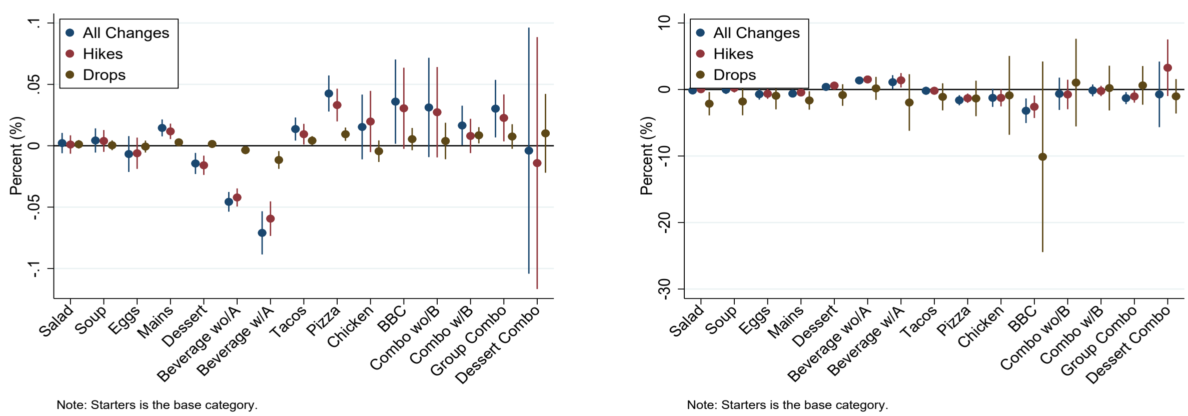

events. However, one could think that seasonal patterns might better fit this type of data (e.g. pay-day effect around the start/end of the month and/or weekend effects). The results suggest that using seasonal fixed effects instead of time fixed effects change very little the coefficients under analysis.[26] Second, I show that qualitatively the conclusion of heterogeneous price-setting holds when all products in dataset are analyzed and not only the set of representative products (defined by the Θ set in Subsection 4.1).[27]

The above results can further decomposed in the context of the Covid19 pandemic. That is, depending on the stage of the Covid19 pandemic, I estimate the following expressions:

$$P\left(y_{i,j,t}=1|x\right)=\beta_1 x_{dishtype}\times Pandemic_n+\theta_j+\theta_t+\varepsilon_{i,j,t}$$

$$\mid \Delta p_{i,j,t} \mid = \beta_1 x_{dishtype}\times Pandemic_n+\theta_j+\theta_t+\varepsilon_{i,j,t}$$

where $Pandemic_n$ is a categorical variable signaling four stages of the pandemic: 1st wave

[26]See Figure 20a in the Appendix. Seasonal controls are day of the week, calendar day, month and year.

[27]In the Appendix, Figure 21a for the comparison between estimates using observations in Θ and all observations in the dataset. Moreover, regression estimates when using all observations while decomposing by sign of adjustment are reported in Table 12 and Table 13 and depicted in Figure 22a.

(1/2),[28] 1st wave (2/2),[29] 2nd wave[30] and 3rd wave.[31] The intuition behind this four stages obeys Mexico City's timeline in terms of Covid19 cases and social distancing measures: (i) the first and toughest lockdown measures in Mexico City, including temporal closures of on-site dining and the rapid increase of users in OOD platforms, (ii) followed by the relaxation of lockdown measures, including the reopening of restaurants for on-site dining with restricted capacity, (iii) new surge of cases, decrease on the on-site dining capacity and limited vaccination rollout and (iv) relaxation of social distancing measures, vaccination rollout and Delta wave.

Regarding the frequency of price changes, Figure 10 shows that for the frequency of price changes, point estimates of the first half of the 1st wave are smaller than those of the second half of the 1st wave. In the majority of cases, point estimates in the first half of the 1st wave remain closer to zero than those of the second half, which are located to the right of zero. For the 2nd wave, the frequency of price changes seems similar to those observed in the first half of the 1st wave. Results from the 3rd wave, in contrast, look like those of the second half of the 1st wave.

A similar pattern can be observed for the size of price changes. On the one hand, estimates from the first half of the 1st wave resemble those of the 2nd wave. They are closer to zero and with narrower confidence intervals. On the other hand, coefficients from the second half of the 1st wave look similar to those of the 3rd wave. They are normally greater than those in first half of the 1st wave or in the second wave, and have wider confidence intervals.

Thus, although the Covid19 pandemic had many direct and indirect economic effects in the period under study, Figure 10 seems to suggest that acute periods of infections (first half of 1st wave and 2nd wave) muted price-setting decisions in both margins. Then, the

[28]It includes the first and toughest lockdown measures adopted throughout the pandemic, including temporal restaurant closures for on-site dining. This is from April 2020 (start of the dataset) until June 2020.

[29]This period considers from July 2020 to September 2020. During this period there was a decrease on the number of cases after the peak of the first wave, as well as some relaxation on social distancing measures, including the reopening of restaurants for on-site dining with limited capacity.

[30]It runs from October 2020 to March 2021, inclusive. This period was characterized by a great number of cases, partially driven by the Day of the Death and Christmas Bank Holidays. Although for most of this period vaccination was unavailable, it is worth noticing that in Mexico vaccination started in February 2021 for the 60+ years old. According to WHO statistics, by late March 2021, about 2% of Mexico City population was fully vaccinated.

[31]This period encompasses April 2021 to December 2021 (end of dataset). It includes the vaccination rollout as well as the summer "Delta Wave". By late December 2021, about 65% of Mexico City population was fully vaccinated according to WHO statistics. The "Omicron Wave" only started to gain momentum by late December 2021 and therefore it is not considered as part of the analysis.

relaxation of social distancing measures (second half of 1st wave) and the vaccination rollout coupled with the "Delta Wave" (3rd wave) were perceived by restaurants as an opportunity to reset their prices.

Taking a closer look at some of the categories described above, the more frequent price changes of Mains, Tacos, Pizzas and Group Combos primarily took place in the second half of 1st wave and 3rd wave indeed. Notably, it is in the other two pandemic stages (first half of 1st wave and 2nd wave) when prices of Beverages without and with Alcohol were particularly more rigid than the remaining categories. With respect to the size of adjustments, it is also in the second half of 1st wave and 3rd wave when one observes larger price changes in Mains, Tacos, Pizzas and Group Combos, as well as in Beverages without and with Alcohol.

Hence, big data sources, coupled with machine learning techniques, allow us to shed further light on the price-setting decisions followed by multi-product agents, such as restaurants. The stylized facts presented in this Section can further be used to validate menu-cost models for multi-product firms as those proposed by Alvarez and Lippi (2014), Nakamura and Steinsson (2008), Hobijn et al. (2006), among others.

6 Determinants of FAFH Inflation

This section evaluates the likely price determinants in the FAFH industry. To that end, I move away from the web scraped dataset due to data restrictions, and focus instead in studying the drivers of the year-on-year FAFH inflation in Mexico City as published by *INEGI*.

In particular, using a large set of determinants (e.g. food prices, labor costs and utility bills), coupled with a number of machine learning frameworks, I study what explains the current, as well as the up to one year ahead, year-on-year FAFH inflation. That is, I retrieve the most important drivers behind the models that best fit the current, as well as future (observed), FAFH inflation in Mexico City.

The numerous determinants, as well as the uncertainty on the weight each of them carry in the production function of the restaurants included in the CPI survey, makes it a suitable task for machine learning tools. The forecasting tools for this analysis are shrinkage methods and non-linear machine learning models.

As it is explained in great detail below, estimates are drawn employing pre-Covid19 time periods. The various demand and supply shocks stemming from the Covid19 pandemic, as

Figure 10: Stylized Facts of Price Changes at Different Stages of the Pandemic

Representative Dishes

Price Changes Regardless Sign of Adjustment

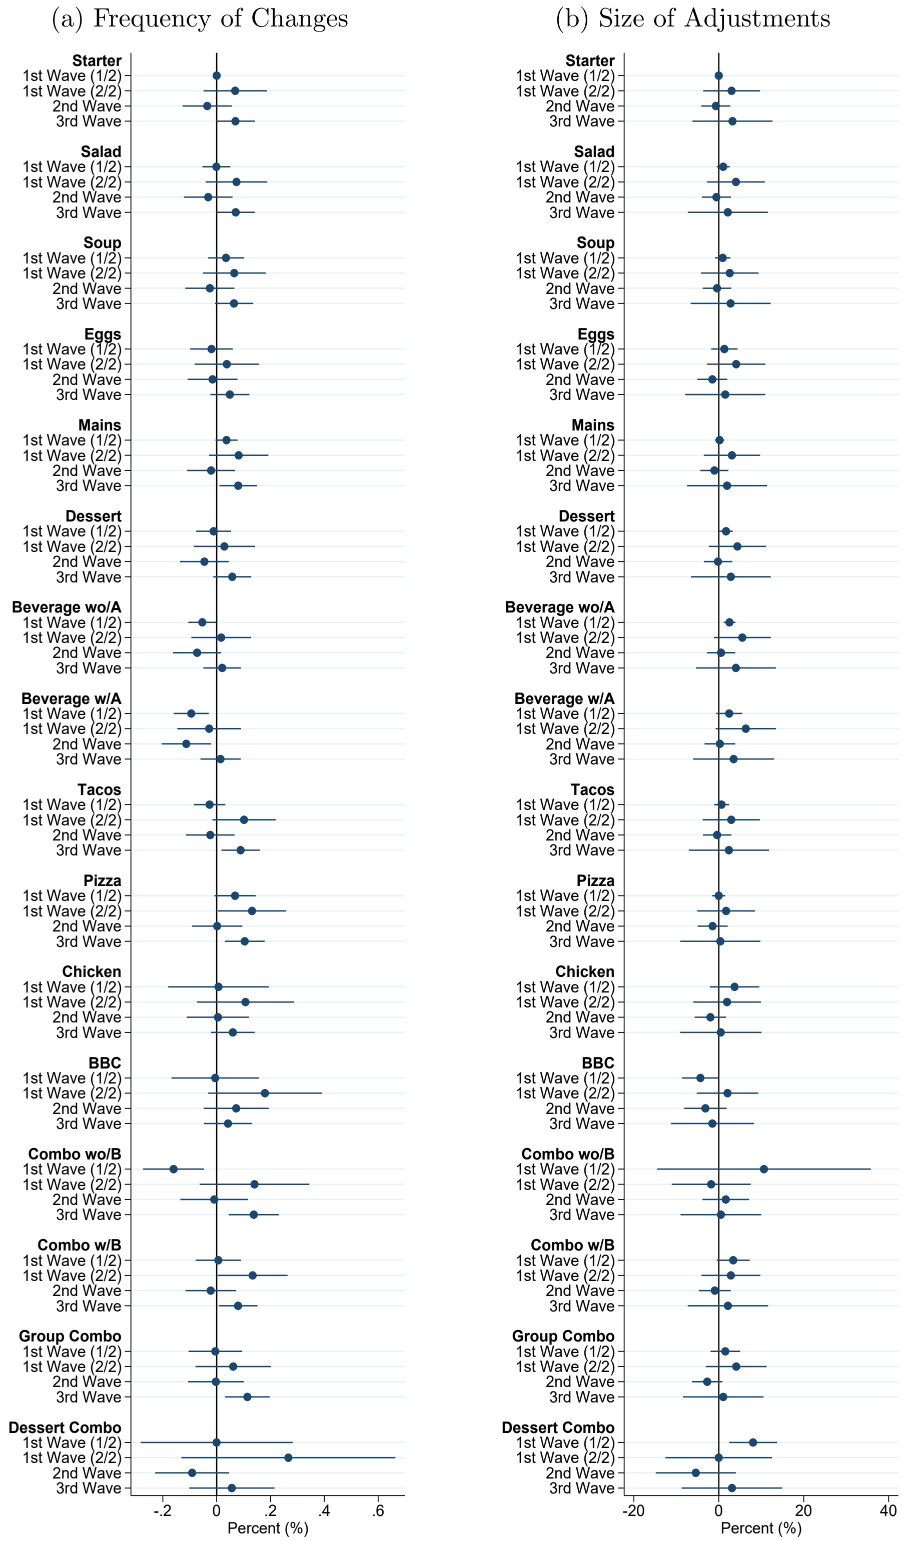

it was discussed in the previous Section, might have altered the weights price-setters factor in determinants in their pricing decisions. This is particular for restaurants included in the CPI survey, which are surveyed at their premises and mainly offer on-site dining.

Thus, I neglect observations gathered throughout the Covid19 pandemic for the benchmark models assessing the role of the different determinants under study. However, as a bypass product, once models are trained using data up to 2018, they are deployed as counterfactuals of the inflation rate in the FAFH industry stemming from costs pressures during the Covid19 pandemic.

6.1 Data and Methodology

The determinants under study reflect supply pressures from three different factor markets in Mexico City: wholesale food prices, labor-related costs and utilities.

First, I use microdata from warehouses and wholesale food markets in Mexico City and Mexico City's Metropolitan Area (MCMA). The data is published by the Ministry of Commerce and it provides average prices observed in these markets by food category.[32] Examples of these food categories are apples, tomatoes, beef, chicken, lemon, onion, etc. These series are further disaggregated geographically: one set of series for supply centers within Mexico City and another set of series for warehouses located in MCMA. There are 73 different series available and included in the analysis.

Second, labor costs data comes from publicly available data published by the Mexican Social Security Institute.[33] I use a number of labor costs indicators (number of workers, wage bill, average wage per worker), specifications (e.g. nominal and real wages) and type of workers (temporal and permanent) associated with the FAFH industry in Mexico City. There are 10 labor-related costs encompassed in the analysis.

Third, utilities like gas LP gas, natural gas and electricity in Mexico City are also incorporated in the analysis. To that end, I use Mexico City's CPI series.

All in all, 86 series are at hand for studying the determinants in the FAFH inflation in Mexico City.[34] All series start in 2005. Furthermore, in compliance with the frequency in

[32]In Spanish, *Secretaría de Economía* publishes wholesale food prices via its *Sistema Nacional de Información de Mercados* or *SNIIM*.

[33]In Spanish, *Instituto Mexicano del Seguro Social*.

[34]See Table 14 in the Appendix for the complete list of explanatory variables and some descriptive statistics.

which FAFH inflation is measured by INEGI, and in order to maximize the number of available observations, estimates are calculated at fortnightly frequencies (twice a month). Thus, series from wholesale food prices and labor-related costs are linearly interpolated as they are released on a monthly basis, while utilities are already available on a fortnightly frequency.

The recognition of inflation drivers at different time horizons is carried out as follows. First, I train five frameworks. They are (i) Elastic Net Regression, (ii) Lasso Regression, (iii) Ridge Regression, (iv) Random Forest Regression and (v) Support Vector Machine Regression.[35] Second, I select the framework that best fit the data in terms of forecasting accuracy from 2006 to 2018. Specifically, I test for statistical difference in forecasting performance using Diebold-Mariano tests. Thirdly, I take a closer look at the predictors carrying greater weight in the best performing model. These three steps are repeated for each of the five horizons at which results are reported: current, as well as the 3-, 6-, 9- and 12-month ahead inflation horizon.

For the estimation set-up, I follow Joseph et al. (2021) from the Bank of England. Joseph et al. (2021) study the subindices in the CPI that are relevant for forecasting UK inflation at different horizons. The strategy in this research is similar in terms of assessing what are the most likely determinants of future FAFH inflation in Mexico City. The specification in mind is of the form:

$$\Delta_{24} y_{t+j} = g(\Delta_{24} X_t, \beta^0) + \varepsilon_t \tag{3}$$

where y_{t+j} is the fortnightly FAFH CPI in Mexico City, and thus $\Delta_{24} y_{t+j}$ is the (lead) annual inflation rate at $j = [0, 6, 12, 18, 24]$ fortnights ahead i.e. contemporaneous, as well as three, six, nine and twelve months ahead. X_t is the vector of 86 determinants previously described, which enters the model in annual growth rates as well. These determinants run from 2006 to 2017. Note, however, as the target variable $\Delta_{24} y_{t+j}$ leads determinants X_t in Equation 3, the different models may or may not use data from 2006 or 2018.[36] β^0 is the set of hyper-parameters shaping the form of $g(.)$, more on the hyper-parameter selection below.

I deviate from Joseph et al. (2021) methodology by not including lagged inflation ($\Delta_{24} y_{t-1}$) as an explanatory viable. The main reason behind this decision is to leave the inflation de-

[35] Araujo and Gaglianone (2020), Joseph et al. (2021), and references therein, provide brief descriptions of these models.

[36] For instance, frameworks fitting the current inflation rate ($j = 0$) employ the contemporaneous target variable $\Delta_{24} y_t$ and, thus, run from 2006 to 2017 while neglecting 2018 inflation dynamics. In contrast, models fitting the 12-month ahead inflation ($j = 24$) use a leading target variable $\Delta_{24} y_{t+24}$, therefore they neglect the 2006 FAFH inflation rate, while considering 2007-2018 period.

terminants speak for themselves and not relying on the persistency of the inflation rate. The decision becomes even more relevant for the 2018-2021 counterfactual exercise, described in great detail below, as it would reflect merely costs pressures while neglecting other type of shocks due to the Covid19 pandemic.

With respect to the estimation and deployment of machine learning models, I follow standard procedures. The series are transformed to year-on-year log differences and standardized. Then, the hyper-parameters for each model are selected through cross-validation taking special care of the time series context at hand. Specifically, the time span of the training period is divided into equidistant (time-wise) folds. I use 12 folds in this exercise, equivalent a fold every 12 months. In the first iteration, the model is trained using the first fold and then evaluated in the subsequent folds. In the second iteration, the training period is extended with the second fold and evaluated in the remaining folds once again. Hence, the estimation setting follows an expanding window approach.[37] The grid of hyper-parameters are summarized in Table 15 in the Appendix.

As a bypass product, having fitted numerous frameworks using determinants from 2006 to 2017, I provide counterfactuals for the year-on-year Mexico City's FAFH inflation rate throughout the 2018-2021 period. The intuition behind leaving out these last four years from the training period is mainly due to the Covid19 pandemic: the 2018-2019 could be seen as an out-of-sample validation period; while the predictions for 2020-2021 period could be seen as an inflation counterfactual stemming from costs-related pressures considered in the analysis.

6.2 Results

Estimates suggest that Random Forest Regressions (RFR) exhibit the lowest Root Mean Squared Error (RMSE) across most time horizons, as shown in Figure 11. Furthermore, Ridge and Support Vector Machine Regressions (SVMR) are the runner ups with fairly stable RMSE around the different horizons, while Elastic Net and Lasso are generally the least accurate models.

This is confirmed by the Diebold-Mariano tests reported in Table 4. Each entry in Table 4 indicates the model with better accuracy between the column-model and row-model in

[37] This approach is different to the traditional K-fold cross-validation in the machine learning literature. K-fold cross-validation would lead to bias results in the forecasting exercise as it assumes that observations are independent and identically distributed. See Joseph et al. (2021) and Coulombe et al. (2020) for more on time-series cross-validation.

Figure 11: Models' Root Mean Squared Error

By Horizon of Prediction[a]

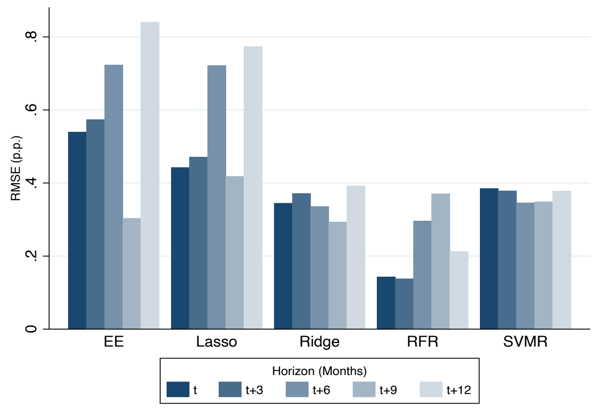

[a]Note: The hyper-parameters for each model and horizon are selected through cross-validation following an expanding window approach. As the training period considers 12 years of data, I use 12 folds. The grid of parameters explored in the cross-validation is reported in Table 15 in the Appendix, while Table 16 summarizes the hyper-parameters that fit best the data according to the cross-valation. RMSE is computed using the models' predictions in the training period i.e. from 2006 to 2017. Bear in mind that, as described in Footnote 36, the target variable leads the determinants. Thus, depending on the models' horizon, some frameworks may or may not fit Mexico City's FAFH annual inflation from 2006 or 2018 in their training period. Also in the Appendix, Table 17 reports the point estimates behind Figure 11. Also in the Appendix, Figure 25 provides the models' fitted values over the training period.

question if statistically significant differences in their predictions are found according to a Diebold-Mariano test. The cell is left empty if the Diebold-Mariano test does not suggest any statistically significant difference between the column- and row-models' predictions.

By looking at the RFR rows and columns, it seems that its predictions are more accurate than its competitors' predictions for all but the nine-month ahead Mexico City's FAFH annual inflation rate. For the nine-month ahead inflation rate, Ridge outperforms all its competitors, including RFR. Finally, Table 4 shows that SVMR is in general more accurate than EE and Lasso, but underperforms relative to Ridge and RFR.

Moving into the determinants, I analyze the determinants of the FAFH inflation through the lens of the RFR as it is the model that best fits the inflation rate in all but one horizon. Figure 12 provides a heatmap of the most important variables explaining the FAFH inflation by horizon according to the RFR framework.[38]

It highlights that the annual growth rate of the wage bill and mean wage, both in real

[38]In particular, Figure 12 plots the top 10 features by horizon. There might be more than 10 colored boxes by horizon because a feature might be in the top 10 for a given horizon but it might not be in the top 10 in another horizon. In the Appendix, Figure 26 reports the complete set of variables under analysis.

Table 4: Diebold-Mariano's Predictive Accuracy by Horizon[a]

Horizon t

	EE	Lasso	Ridge	RFR	SVMR
EE	- o -				
Lasso	Lasso**	- o -			
Ridge	Ridge***	Ridge***	- o -		
RFR	RFR***	RFR***	RFR***	- o -	
SVMR	SVMR***	SVMR**	Ridge*	RFR***	- o -

Horizon t + 3 months

	EE	Lasso	Ridge	RFR	SVMR
EE	- o -				
Lasso	Lasso***	- o -			
Ridge	Ridge***	Ridge***	- o -		
RFR	RFR***	RFR***	RFR***	- o -	
SVMR	SVMR***	SVMR***		RFR***	- o -

Horizon t + 6 months

	EE	Lasso	Ridge	RFR	SVMR
EE	- o -				
Lasso		- o -			
Ridge	Ridge***	Ridge***	- o -		
RFR	RFR***	RFR***	RFR*	- o -	
SVMR	SVMR***	SVMR***		RFR**	- o -

Horizon t + 9 months

	EE	Lasso	Ridge	RFR	SVMR
EE	- o -				
Lasso	EE***	- o -			
Ridge	Ridge**	Ridge***	- o -		
RFR	EE***	RFR***	Ridge***	- o -	
SVMR	EE***	SVMR***	Ridge***		- o -

Horizon t + 12 months

	EE	Lasso	Ridge	RFR	SVMR
EE	- o -				
Lasso		- o -			
Ridge	Ridge***	Ridge***	- o -		
RFR	RFR***	RFR***	RFR***	- o -	
SVMR	SVMR***	SVMR***		RFR***	- o -

[a]Note: Each cell reports the model with better accuracy between the column-model and row-model in question if there is a statistically significant difference between their predictions. Empty spaces in the lower diagonal imply no statistically significant differences. Accuracy over the training period, 2006-2017. *, **, *** represent $p < 0.10$, $p < 0.05$ and $p < 0.01$, respectively, in the Diebold-Mariano test.

terms, of permanent workers, as well as the price of beef in MAMC best describe the current FAFH inflation rate in Mexico City.[39] The three month ahead inflation rate is mainly driven by the price of beans in MAMC.[40] The annual growth in the price of beans in MAMC and the price of shrimps in Mexico City seems to best explain the six month ahead inflation rate.[41] Moreover, the annual price variation of beans and rice in both MAMC and Mexico City, the price of shrimps and avocados in Mexico City, as well as the year-on-year growth of the nominal mean wage of permanent workers are relevant features determining the nine month ahead FAFH inflation rate. Lastly, the nominal mean wage of permanent workers also describes well the 12-month ahead inflation rate, in addition to the annual price growth of rice in MAMC.[42]

By looking within each of the factor markets, few salient facts emerge. First, within the determinants associated to utility prices, it seems that electricity is relatively more important than LP gas and natural gas as it is relevant in four out of the five time horizons explored. Although the three utilities are important enough to be considered in Figure 12, they are not the most relevant features in the RFR as one might have expected. Second, regarding the covariates related to labor costs, the mean real wage of permanent workers remains relevant in most horizons; while the real wage bill and the mean nominal wage of permanent workers are relevant in explaining the current and 12-month ahead inflation, respectively. Despite being considered as a buffer between tight and slack conditions in labor markets, it looks like costs-pressures stemming from temporal workers are not as relevant as those from permanent workers. Third, prices of beans and rice, both in MAMC and Mexico City, are among the most important determinants, not only across the wholesale food market but among all determinants under study. Their frequent presence as a side dish in most meals in the Mexican cuisine could be a potential explanation. In terms of meat, the price of shrimps

[39] Other determinants driving the contemporaneous inflation rate include the annual price variation in the price LP gas and electricity; the wage bill in real and nominal terms of temporal workers and the number of temporal workers; as well as the annual price growth of pineapple and dried chilli in the MAMC; the price of chicken, dried chilli, shrimps and rice in Mexico City.

[40] Other relevant covariates are price variations of LP and natural gas; the real wage bill, mean real wage and number of permanent workers; as well as the annual price growth of pineapple, pear and dried chilli in the MAMC; the price of shrimps, carrots, fish (others), legumes and poblano chilli in Mexico.

[41] Other explanatory variables at this horizon are the variation in the price of electricity; number of permanent workers and the wage bill, both in real and nominal terms, for temporal workers; as well as the annual growth rate of rice in MAMC; price of rice and avocados in Mexico City.

[42] Other key determinants for the 12-month ahead inflation rate are the annual price change of electricity and natural gas; the annual variation of the mean real wage of permanent workers and the mean wage, in both nominal and real terms, of temporal workers; as well as the year-on-year price growth of chicken, fish (others), shrimps and rice in Mexico City.

in Mexico City is a robust determinant across horizons, whilst the price of beef determines the current inflation rate only, suggesting perhaps a faster cost pass-through.

As described in the Introduction, the target variable used for this analysis comes from the FAFH component of the CPI survey in Mexico City. Specifically, it is the index computed summarizing observations from a number of restaurants, which in turn prepare many different types of food, not to mention price collectors only gather prices of few items per restaurant.[43] Thus, frameworks that are able to deal with numerous covariates come in hand for assessing the potential role they have in describing the aggregate index.[44]

Figure 12: Relevance of Features

Features on the vertical axis and horizons in the horizontal axis.[a]
Selected features

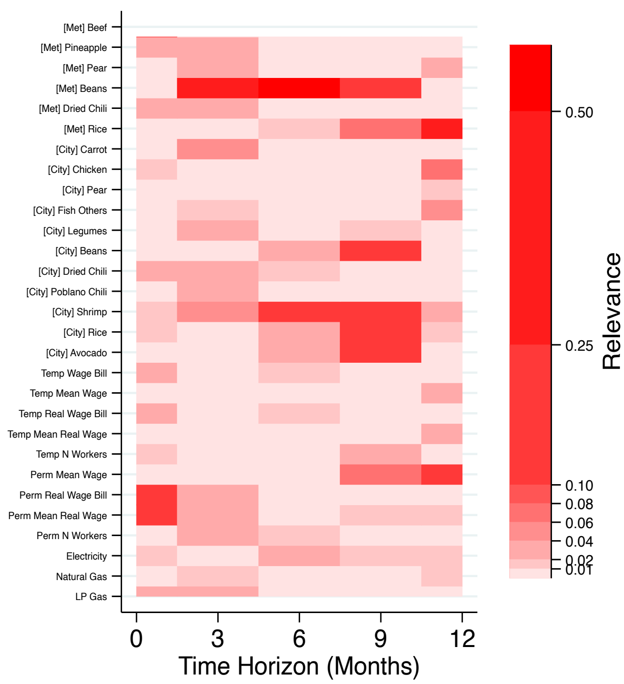

[a]Features' relevance by horizon adds up to 1.

Finally, and as a bypass product, the shrinkage and non-linear machine learning models trained above are able to provide a counterfactual, given the state of the determinants, on Mexico City's FAFH inflation during the Covid19 pandemic. Needless to say, on the one

[43]Price collectors gather prices from about three to five items on the menu, which can be seen as a set menu for a single person e.g. starter, main, dessert plus beverage.

[44]Requesting access to CPI microdata in order to validate if, for instance, a number of observations consider beans or seafood/shrimps is left for future work.

hand, the state of the determinants, specially those from wholesale food markets were affected in 2020 by excess and depress demand of household and restaurant consumption, respectively. On the other hand, temporal but compulsory on-site consumption might have also drastically change price-setting patterns. Hence, while results should be taken with caution due to the different shocks at the time, they remain illustrative of how pre-pandemic price-setters might have reacted should they have faced such dynamics from the factor markets.

In order to validate the models that best predict the target variable out of the training period, I repeat the Diebold-Mariano tests based on the models' predictions throughout 2018-2019 i.e. pre-Covid19 pandemic. As a benchmark of the out-of-sample performance of the shrinkage and non-linear machine learning models, I also include in the Diebold-Mariano tests a comparisons with a simple autoregressive model.[45]

As shown in the Appendix, RFR is no longer the model with the lowest RMSE across horizons.[46] In turn, Diebold-Mariano tests do not suggest any particular model systematically outperforming its competitors. The AR model seems to perform best when predicting the current, as well as the three and six-month ahead inflation. However, for the nine and twelve-month ahead horizon there is no statistically significant winner in terms of accuracy. The liberalization of fuel prices in Mexico in 2017, reshaping in 2018 and 2019 not only price-setting decisions of restaurants but also decisions from suppliers in the factor markets under analysis, could be a potential explanation behind this result. Unfortunately, the short time span between the 2017 fuel prices liberalization and the start of the pandemic leave little room to retrain the models.

Figure 13 depicts the RFR and Ridge predictions before and during the Covid19 outbreak.[47] It seems that, after the start of the Covid19 pandemic in March 2020 and based on the state of the FAFH determinants, both models suggest there could have been an increase of the FAFH inflation rate in Mexico City. These predictions sharply contrast with observed

[45]The AR model is computed following the same approach as the other frameworks in the analysis. That is, I fit an AR(1) from 2006 to 2017, and generate out-of-sample predictions from 2018 to 2021. Results of the Diebold-Mariano tests are also computed using data from 2018 and 2019 i.e. pre-Covid19.

[46]For brevity these results are left in the Appendix. Figure 27 illustrates the RMSE computed using predictions from 2018 and 2019 by model and horizon, while Table 18 summarizes the Diebold-Mariano tests for the same period.

[47]These are the models that performed better, in terms of the lowest point estimate of the RMSE, during the training period. It seems that, before 2020, both models partially followed the annual FAFH inflation rate in Mexico City. For instance, during the first half of 2018, RFR overpredicted the FAFH inflation. Then, in the remaining of 2018 and first half of 2019, RFR underpredicted observed inflation, while overpredicting again in late 2019. A similar pattern is observed for Ridge in 2018 and 2019.

inflation. One can rationalize the deviation between predicted and observed inflation by the temporal but compulsory restaurant closures, low demand for on-site dining, as well as the households' food hoarding. Nonetheless, predictions put forward the idea of an inflation rate between 6% and 7% in annual terms for Mexico City by early 2021. In fact, the experimental price index reported in Section 4 and depicted in Figure 6a provide a cumulative 5% inflation rate from April to December 2020.[48] The CPI only reports such inflation levels by late 2021, period in which the models suggest the year-on-year inflation rate could be around 5%.

As the Covid19 pandemic continues to unfold (e.g. the Omicron-wave only hit Mexico around December 2021), it remains to be seen if the frameworks trained for this study recover their validity. Though, the results illustrate that price-setters in this industry in Mexico City might have faced cost-related pressures, which were not immediately pass-through to prices, highlighting once again the precence of price stickiness.

Figure 13: Models' Predictions

By Horizon of Prediction

(a) Ridge

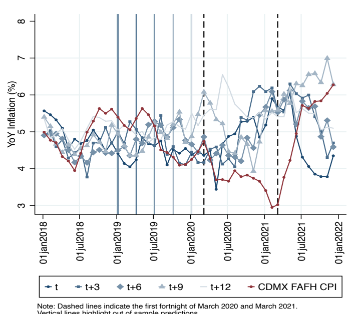

(b) RFR

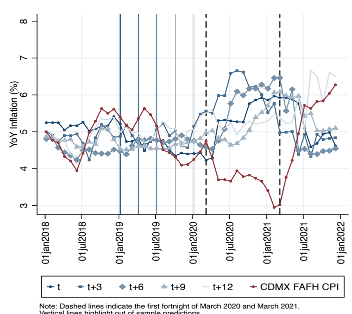

7 Conclusions

The Food Away From Home (FAFH) industry has been one of the most affected by the Covid19 pandemic. Temporal closures, restricted on-site dining capacity, as well as suppressed demand are among the supply and demand shocks deemed by the industry. Understanding how FAFH prices evolve as the pandemic unfolds is important for the conduction of monetary policy as it accounts for a non-negligible weight in the CPI basket.

This paper studies the FAFH inflation in Mexico City from two complementary angles. First, using web scraped data from an online ordering and delivery platform, this paper

[48] I am unable to compute the y-o-y inflation rate of the experimental price index as the data collection of the online ordering and delivery platform started in April 2020.

examines price dynamics stemming from over 120 million prices from more than 1.7 million dishes offered by around 30,000 restaurants in Mexico City. Second, using machine learning algorithms, the latter part of this study analyses 86 price determinants likely considered by price-setters in the FAFH industry and provides a counterfactual of the inflation rate stemming from costs pressures during the Covid19 pandemic.

The results suggest that (i) independent and multi-outlet restaurants report similar price trends; (ii) prices of soups and beverages without alcohol, potentially substituted by home-production, exhibit low price growth rates; (iii) in contrast, prices of mains and desserts have been on the rise; (iv) the heterogeneous growth rates across dish categories seem to be explained by the extensive margin; and (v) episodes associated to the escalation in Covid cases seem to increase price rigidities.

With respect to the analysis of price determinants through the lens of a Random Forest Regression, electricity fees, the real wage bill and real average wage of permanent workers, as well as wholesale prices of beans, rice and shrimps best describe the FAFH inflation rate in Mexico City at different horizons. Regarding utilities, gas LP and natural gas are not the most important inflation determinants as one might have expected. With respect to labor costs, wage indicators from permanent workers generally outperform wage statistics form temporal workers in predicting the FAFH inflation rate in Mexico City. Lastly, the frequent presence of beans and rice as sides in most meals in the Mexican cuisine could be a potential explanation on their relevance. In terms of meat, the price of shrimps in wholesale food markets in Mexico City is a robust explanatory variable across horizons. The price of beef only determines the current inflation rate, suggesting perhaps a faster cost pass-through.

Although web scraped data is increasingly used for analyzing inflation, to the author's knowledge the literature has mainly focused on goods' prices observed at supermarkets or departmental stores. In contrast, this research contributes to the literature by focusing on analyzing web scraped prices in an industry at the service sector.

The rapid adoption of online ordering and delivery platforms while on-site dinning was depressed, as well as the numerous shocks input prices suffered partially due to the change in households' consumption habits, leaves this industry as a prosperous area of research. From the study of multi-product pricing models, the analysis of price dispersion as online platforms allow greater number of alternatives and easier price comparisons, to the examination of pass-through determinants, are among some of the venues to be explored in the FAFH industry.

References

Alvarez, F. and F. Lippi (2014): "Price setting with menu cost for multiproduct firms," *Econometrica*, 82, 89–135.

Araujo, G. S. and W. P. Gaglianone (2020): "Machine learning methods for inflation forecasting in Brazil: new contenders versus classical models," Tech. rep., Mimeo.

Bils, M. and P. J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947–985.

Card, D. and A. B. Krueger (1994): "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania," *The American Economic Review*, 84, 772–793.

Cavallo, A. (2018): "Scraped data and sticky prices," *Review of Economics and Statistics*, 100, 105–119.

Cortes, P. and J. Pan (2013): "Outsourcing household production: Foreign domestic workers and native labor supply in Hong Kong," *Journal of Labor Economics*, 31, 327–371.

Coulombe, P. G., M. Leroux, D. Stevanovic, and S. Surprenant (2020): "How is machine learning useful for macroeconomic forecasting?" *arXiv preprint arXiv:2008.12477*.

Dhyne, E., J. Konieczny, F. Rumler, P. Sevestre, et al. (2009): "Price rigidity in the euro area-An assessment," Tech. rep., Directorate General Economic and Financial Affairs (DG ECFIN), European .

Diewert, W. (2020): "The Chain Drift Problem and Multilateral Indexes," Tech. rep., Discussion Paper 20-07, Vancouver School of Economics.

Diewert, W. E. and K. J. Fox (2020): "Measuring Real Consumption and CPI Bias under Lockdown Conditions," Tech. rep., National Bureau of Economic Research.

Fetzer, T. (2022): "Subsidising the spread of COVID-19: Evidence from the UKs Eat-Out-to-Help-Out Scheme," *The Economic Journal*, 132, 1200–1217.

Flower, T. (2019): "Using alternative data sources in consumer prices, UK," Tech. rep.

Fougère, D., E. Gautier, and H. Le Bihan (2010): "Restaurant prices and the minimum wage," *Journal of Money, Credit and Banking*, 42, 1199–1234.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek (2017): "Inflation dynamics during the financial crisis," *American Economic Review*, 107, 785–823.

González-Pampillón, N., G. Nunez-Chaim, and K. Ziegler (2021): *Recovering from the First Covid-19 Lockdown: Economic Impacts of the UK's Eat Out to Help Out Scheme*, Centre for Economic Performance, London School of Economics and Political Sciences.

Hobijn, B., F. Ravenna, and A. Tambalotti (2006): "Menu costs at work: restaurant prices and the introduction of the euro," *The Quarterly Journal of Economics*, 121, 1103–1131.

Hull, I., M. Löf, and M. Tibblin (2017): "Price information collected online and short-term inflation forecasts," Tech. rep., Bank for International Settlements.

Joseph, A., E. Kalamara, G. Kapetanios, G. Potjagailo, et al. (2021): "Forecasting UK inflation bottom up," Tech. rep., Bank of England.

Konny, C. G., B. K. Williams, and D. M. Friedman (2019): "Big Data in the US Consumer Price Index: Experiences and Plans," in *Big Data for 21st Century Economic Statistics*, University of Chicago Press.

Macias, P., D. Stelmasiak, et al. (2019): *Food inflation nowcasting with web scraped data*, Narodowy Bank Polski, Education & Publishing Department.

Nakamura, E. and J. Steinsson (2008): "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 1415–1464.

Peña, J. and E. Prades (2021): "Price setting in Chile: Micro evidence from consumer on-line prices during the social outbreak and Covid-19," Working Papers 2112, Banco de España.

Solórzano, D. (2021): "Stylized Facts From Prices at Multi-Channel Retailers in Mexico," *Manuscript*.

A Appendix

A.1 Data

Figure 14: Data Collection

Hourly and Daily Observations

(a) Robot's Collection (b) Daily Observations

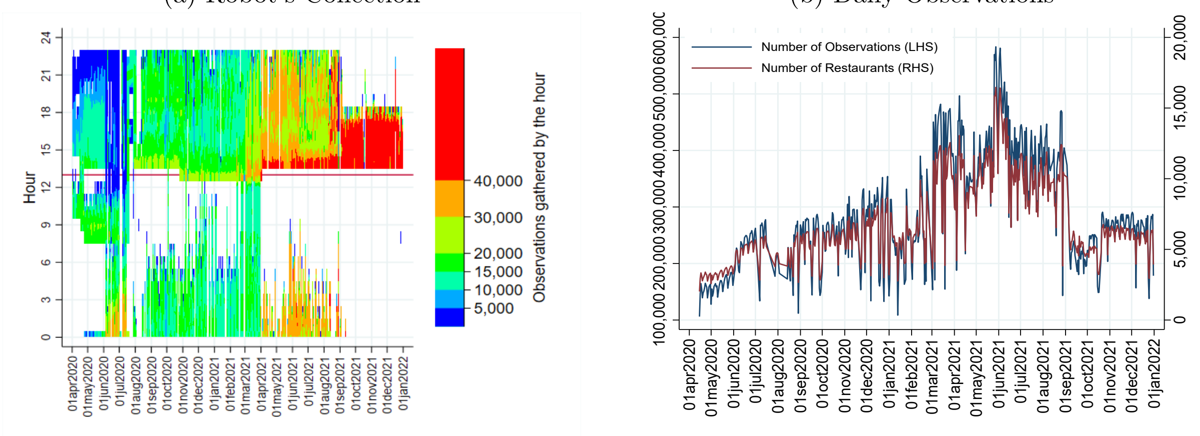

A.2 Machine Learning Dish Classification

This Appendix provides greater details on the classification of dishes (cross-section dimension of the panel) using machine learning techniques. First, it describes the construction of the training set, under which a number of algorithms are trained. Second, text cleaning procedures are outlined. Third, it sketches how dish descriptions are taken into a matrix form. Forth, classifiers are listed and, fifth, hyper-parametrised through k-fold cross-validation. Sixth, some forensic statistics of trained models are discussed. Finally, the winner and runner-up models are compared.

A.2.1 Training Set

As presented in the main text, out of the around 616,000 unique descriptions in the dataset, I manually classify more than 13,000 random dishes based on the descriptions provided by the restaurants.Thus, the manual classification considers a little more than 2% of the dishes in question. The dishes are classified into 19 categories, which are chosen on the basis of (i) well-recognized headers in many restaurants' menus, (ii) categories with direct mapping to Mexico's CPI categories and (iii) research question at hand.

The categories are: (1) Starters, (2) Salads, (3) Soups, (4) Eggs, (5) Mains, (6) Pizzas, (7) Tacos, (8) BBC, (9) Grilled and Roasted Chicken, (10) Desserts, (11) Beverages with Alcohol, (12) Beverages without Alcohol, (13) Meals with Beverages, (14) Meals without Beverages, (15) Group Combos, (16) Dessert Combos, (17) Extras, (18) Others (Non-Food) and (19) Ambiguos.[49]

A.2.2 Text Cleaning

The descriptions in training set are then parsed by text cleaning routines in order to have homogeneous notation in the dish descriptions. Removal of stopwords, special characters, hashtags; standardising numbers' and units' abbreviations, are among some of the cleaning procedures.

A.2.3 Word Tokenising

Once descriptions are clean and homogeneous, words in dish descriptions are ready to be tokenized and used as explanatory variables by the different classifiers. To that end, I convert the collection of dish descriptions into a matrix of token (words) counts. That is, the columns in the matrix represent each and every single word appearing at least once in the collection of descriptions, the rows of the matrix are the dishes in the dataset, and each matrix cell counts the number of times a word (column) appears in the description (row).[50]

I get three different sets of explanatory variables, which will be used one at the time by the classifiers in order to assess how sensitive the performances of the classifiers are to the token count specification. These specifications are: (i) the universe of words found in the descriptions i.e. complete set of single words (unigrams), (ii) subset of unigrams by cutting-off infrequent terms and (iii) subset of unigrams and bigrams cutting-off infrequent terms.[51]

The first one uses all words in the collection of descriptions. Hence, this first specification induces a matrix with over 68,000 words (columns).

The second specification is a subset of the first specification (matrix with lower columns

[49]BBC stands for Barbacoa, Birria and Carnitas, which are common taco fillings, and are considered in the Mexican CPI as a specific product category. Meals with/without Beverages consider two or three times meals. For instance, a Meal with Beverage could be a bundle of starter, salad, main and a soft drink.

[50]This type of matrix is commonly referred as a sparse matrix since each row contains a large number of columns with zeros and only a few with non-zero values.

[51]A bigram is defined as the pair of consecutive words. For instance, the unigram representation of "Today is Monday" is ["Today", "is", "Monday"], while the bigram representation is ["Today is", "is Monday"].

dimension). As words are the set of explanatory variables to be used by the algorithms, which might lead to the curse of dimensionality and intensive computational work, this second matrix comprehends words appearing in the collection of descriptions at least 3 times.[52] Thus, around 23,000 words are considered after implementing this cut-off approach.

The third specification adds bigrams to the matrix of unigrams. This is due to concerns arising from, for example, $description_1$ = "chicken with salad" and $description_2$ = "salad with chicken".[53] The gain of using bigrams is obvious in the previous simple example but it is less clear if, on the aggregate and in the presence of more unigrams, the potential improvement in accuracy outweights the greater number of explanatory variables i.e. matrix dimension. In order to keep dimensions attainable, I also keep only bigrams showing up at least 3 times in the corpus. The matrix of unigrams and bigrams has over 32,000 columns.

In sum, there are three different specifications of matrices of token counts: (i) complete set of words (or unigrams), (ii) subset of unigrams and (iii) subset of unigrams and bigrams. These matrices are used, one by one, in the classifiers, which are then compared in terms of their accuracy on the training set.

A.2.4 Classifiers

The classifiers used for this analysis are (i) decision tree, (ii) random forest, (iii) multinomial naive Bayes and (iv) logistic regression.

A.2.5 Hyper-parameters Tuning

All classifiers require some form of hyper-parameter selection prior to estimation. To that end, I use k-fold cross validation procedures. That is, 80% of the training set is divided

[52] That is, I drop terms that have a frequency lower than 0.0005% of the 616,000 descriptions. The aim of this approach is to neglect restaurant-specific terms that might not be relevant for the classification task. For instance, suppose a fictional restaurant named XXYY offers a restaurant-specific dish with the description (including stop words) "Burger XXYY with bacon and avocado"; the word "XXYY" might not be representative for the broad classification task and adds an extra column to the matrix of explanatory variables. The threshold of 3 in order to be considered in the analysis is chosen with the goal of neglecting very rare words and interfering the less possible with the vocabulary.

[53] In this simple example, the unigram representation leads to the same vector representation. That is, without loss of generality and assuming no stop words, if column 1 counts the word "chicken" and column 2 counts the word "salad", the unigram representation would be $description_1 = [1, 1]$ and $description_2 = [1, 1]$. By using bigrams, without loss of generality, column 1 counts "chicken", column 2 counts "salad", column 3 counts the bigram "chicken salad" and column 4 counts "salad chicken", resulting in a vector representation of $description_1 = [1, 1, 1, 0]$ and $description_2 = [1, 1, 0, 1]$.

intro k folds, after which the model is fitted using observations in k-1 folds under a specific set of parameters and compute the accuracy in the k-th fold. This process is repeated k times in order to compute the accuracy in every fold. The average accuracy is used to select the hyper-parameter configuration maximizing the performance of each classifier.

The grid of parameters used for this search is detailed in Table 5. Note there are *balanced* versions of the decision tree and random forest classifiers. These versions take into account that category sizes are highly unbalanced in the training set. This could be a problem since the parameters and costs functions developed in these algorithms could end up focusing (over specialising) on large categories only. Hence, I impose greater penalties on errors made in smaller categories.[54] The penalties come in the form of weighting observations, where the weights are inversely proportional to category frequencies in the data.

Table 5: Grid of Parameters

By Classifier

Classifier	Parameter	Values
Decision Tree	Max Depth	[500, 550, 600, 625, 650, 675, 700, 750, 800, 900, 1000]
	Min Samples Split	[4,5,6,7,8,10]
	Criterion	[Gini, Entropy]
Balanced Decision Tree	*Same as Above*	*(Various)*
	Class Weight	[Balanced]
Random Forest	Max Depth	[600, 750, 900, 1050, 1200, 1350]
	Min Samples Split	[3,4,5,6,7,10]
	Criterion	[Gini, Entropy]
	N Estimators	[75, 100, 125, 150, 175, 200, 300, 500]
Balanced Random Forest	*Same as Above*	*(Various)*
	Class Weight	[Balanced]
Multinomial Naive Bayes	α	[0.00001,0.0001,0.001,0.01,0.1,0,1,2,3,4,5,10]
	Fit Prior	[True,False]
Logistic Regression	C	[0.00001,0.0001,0.001,0.01,0.1,1,2,3,4,5,10,20]

Note: The grid-search implements an exhaustive search over specified parameter values for each classifier. As mentioned above, k-fold cross validation is used for hyper-parameter selection. That is, 80% of the training set is divided intro k folds (stratified by category sizes), after which the model is fitted using observations in k-1 folds under a specific set of parameters and compute the accuracy in the k-th fold. This process is repeated k times in order to compute the accuracy in every fold. The average accuracy is used to select the hyper-parameter configuration maximizing the performance of each classifier.

[54]Previous versions of this paper included an exercise upsampling small categories by bootstrapping with replacement. However, there seems to be no gain in accuracy at the expense of computation time (as the dataset grows due to the bootstrap with replacement). Results not reported but available upon request.

A.2.6 Classifiers Forensics

As a result of the k-fold cross-validation, Table 6 reports the hyper-parameters that maximize the accuracy score in classifying 80% of the training set. The accuracy scores of these models are depicted in Figure 15a.

Figure 15a shows that, for this specific task for classifying dishes, the use of different specifications on the matrix of token counts generate little gains in terms of accuracy. However, as seen in the various panels of Table 6, the hyper-parameter configuration does change depending the matrix specification as expected.

Moreover Figure 15a highlights that the *logistic regression* and *balanced decision tree* (both with unigrams and bigrams) achieve the greatest and lowest accuracy scores, respectively. Nonetheless, there are only minor differences across models' performance over the sample under which they were trained.

Table 6: Parameters Maximizing Accuracy in Testing Set
By Matrix of Token Counts and Classifiers

Classifier	Parameters
	A. Unigrams
Decision Tree	'max˙depth': 625, 'min˙samples˙split': 5
Balanced Decision Tree	'max˙depth': 750, 'min˙samples˙split': 4
Random Forest	'max˙depth': 600, 'min˙samples˙split': 3, 'n˙estimators': 200
Balanced Random Forest	'max˙depth': 900, 'min˙samples˙split': 3, 'n˙estimators': 500
Naive Bayes	'alpha': 0.1, 'fit˙prior': True
Logistic Regression	'C': 1
	B. Unigrams (Cut-off)
Decision Tree	'max˙depth': 1000, 'min˙samples˙split': 6
Balanced Decision Tree	'max˙depth': 800, 'min˙samples˙split': 4
Random Forest	'max˙depth': 1200, 'min˙samples˙split': 3, 'n˙estimators': 200
Balanced Random Forest	'max˙depth': 1200, 'min˙samples˙split': 3, 'n˙estimators': 175
Naive Bayes	'alpha': 0.1, 'fit˙prior': True
Logistic Regression	'C': 2
	C. Unigrams and Bigrams (Cut-off)
Decision Tree	'max˙depth': 700, 'min˙samples˙split': 7
Balanced Decision Tree	'max˙depth': 500, 'min˙samples˙split': 6
Random Forest	'max˙depth': 900, 'min˙samples˙split': 4, 'n˙estimators': 125
Balanced Random Forest	'max˙depth': 900, 'min˙samples˙split': 4, 'n˙estimators': 150
Naive Bayes	'alpha': 0.1, 'fit˙prior': True
Logistic Regression	'C': 2

Trained algorithms are then deployed over the remaining (unseen) 20% of the manually constructed training set. The models' accuracy scores classifying unseen data are highlighted in Figure 15b.

Similarly as in the sample over which models were trained, there are no stark differences

Figure 15: Accuracy Score

By Classifier and Explanatory Variables

(a) Training Set: 80% of manually classified dishes (b) Test Set: 20% of manually classified dishes

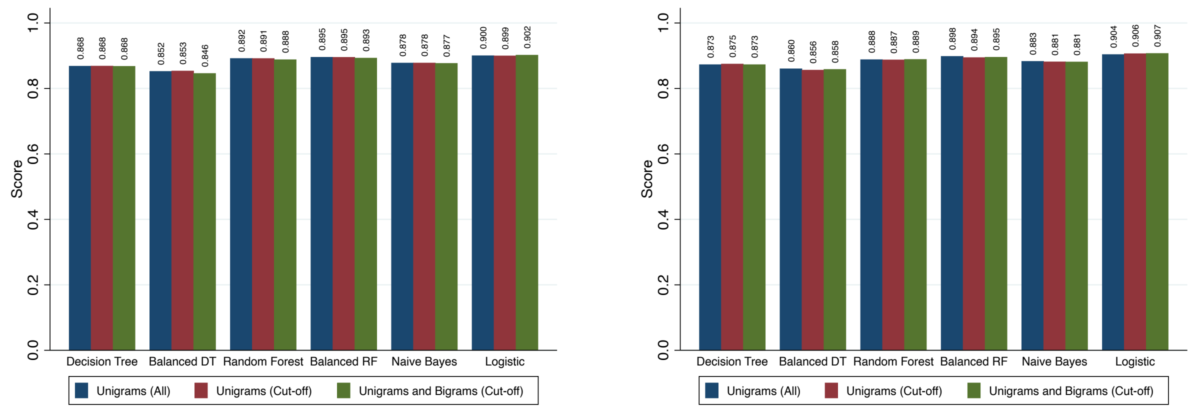

in terms of accuracy over the unseen sample. This is the case neither across classifiers nor specification of tokens.

A.2.7 Model Selection

Since it is the one with greatest accuracy (average point estimate), as well with the lowest computational time, the *logistic regression* using unigrams and bigrams is picked as the winner across models. The runner up is the *balanced random forest*, also using unigrams and bigrams, but with significantly more computational time.

Figure 4 in the main text depicts the confusion matrix on the prediction of dish labels using the logistic regression fitted under the complete training set. It provides a graphical representation on whether the prediction matches with the true value. Each cell reports the share of each instance such that every row (true labels) adds up to one. Correct predictions lay in the diagonal, values outside the diagonal highlight prediction errors. As shown in Figure 4, most cells on the diagonal report values close to one.

As a bypass, Figure 16 shows the confusion matrix on the prediction of dish labels using the balanced random forest trained under the complete training set.

Finally, Table 1 included in the main text adds on the impact of the machine learning techniques used in this research. The first bloc of columns reports the composition of the manually classified dataset. The second bloc of columns summarizes the outcome labels generated through the logistic regression.

Table 7: Dishes' Labels

By Classification Approach

	Course Type	Manual Count	Manual Share (%)	Machine Learning Count	Machine Learning Share (%)
1	Starters	640	4.76	16,625	3.06
2	Salads	996	7.41	15,049	2.77
3	Soups	64	0.48	5,604	1.03
4	Eggs	447	3.32	8,819	1.62
5	Mains	3,948	29.36	251,969	46.32
6	Desserts	1,063	7.91	55,010	10.11
7	Beverages wo/Alcohol	2,639	19.63	89,490	16.45
8	Beverages w/Alcohol	208	1.55	9,539	1.75
9	Tacos	429	3.19	23,256	4.28
10	Pizzas	1,444	10.74	20,819	3.83
11	Grilled/Roasted Chicken	15	0.11	345	0.06
12	BBQ, Birria, Carnitas	23	0.17	788	0.14
13	Combo wo/Beverage	55	0.41	355	0.07
14	Combo w/Beverage	100	0.74	7,647	1.41
15	Group Combo	124	0.92	3,171	0.58
16	Dessert Combo	25	0.19	175	0.03
17	Extras	1,041	7.74	33,846	6.22
18	Others	158	1.17	1,220	0.22
19	Ambiguous	28	0.20	193	0.05
	Total	13,447	100.00	543,920	100.00

Note: Extras, Others and Ambiguous are also considered in the classification exercise but not reported.

Figure 16: Balanced Random Forest Confusion Matrix

Predictions Over the Entire Training Set

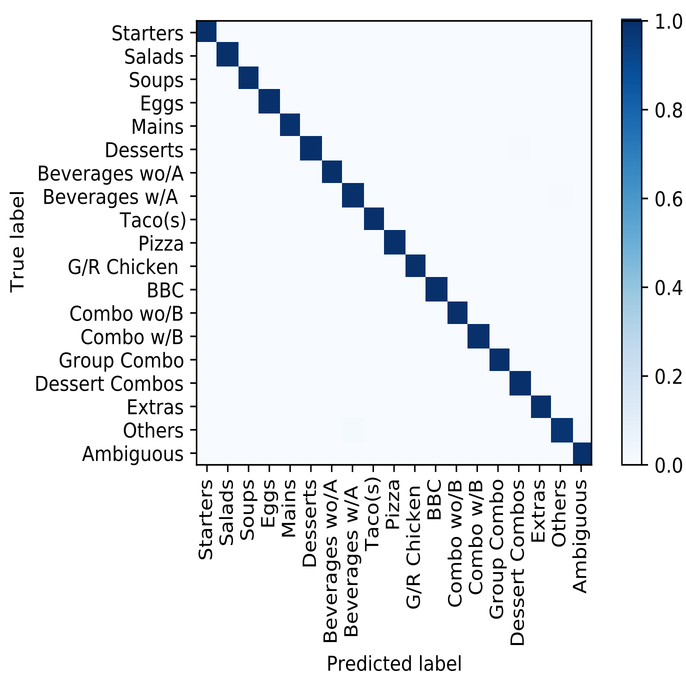

A.2.8 Restaurant Calssification

Figure 17: Dataset Composition by Type of Restaurant

Breaking Up Multi-Outlet Restaurants

(a) Share of Restaurants

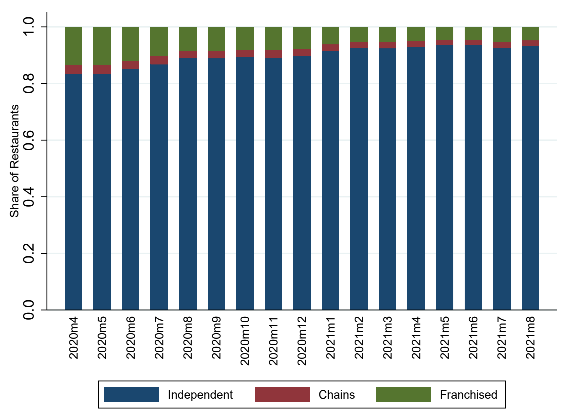

(b) Number of Meals

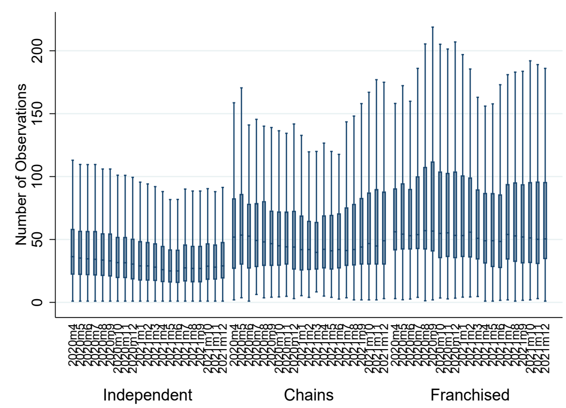

A.3 Experimental Price Indices

Figure 18: Experimental Price Indices

(a) API and AVI Comparison (Representtive Items) (b) Representative Items vs All Items

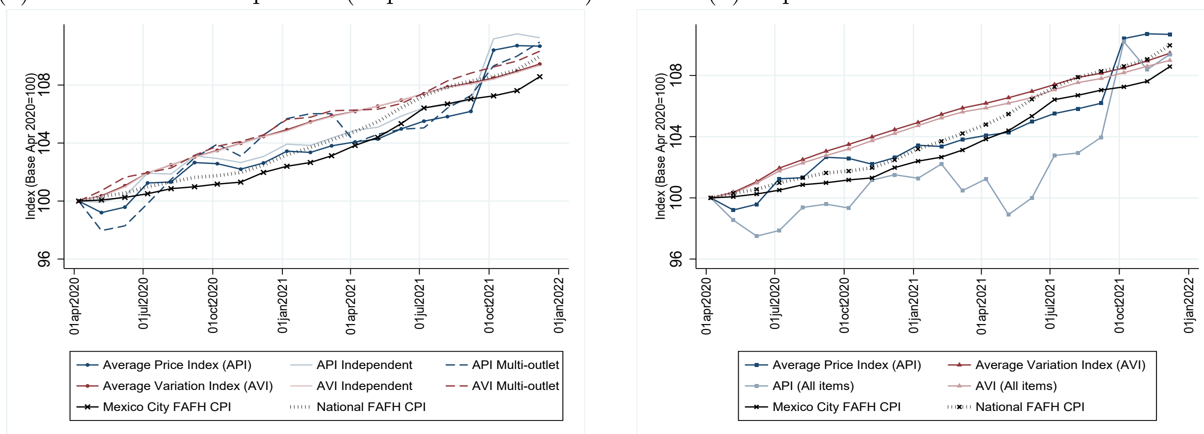

Table 8: Cross-Correlation of Monthly Inflation Up to September 2021
From Experimental Price Indices and CPI

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Average Price Index (API)	1.000							
(2) API Independent	0.972***	1.000						
(3) API Multi-outlet	0.515***	0.330*	1.000					
(4) Average Variation Index (AVI)	0.481***	0.443***	0.329*	1.000				
(5) AVI Independent	0.413**	0.382**	0.252	0.970***	1.000			
(6) AVI Multi-outlet	0.307*	0.266	0.366**	0.222	-0.020	1.000		
(7) Mexico City FAFH CPI	-0.029	-0.034	-0.145	-0.153	-0.120	-0.128	1.000	
(8) National FAFH CPI	0.057	0.068	-0.033	-0.224	-0.214	-0.024	0.684***	1.000

*** p<0.01, ** p<0.05, * p<0.1

Table 9: Cross-Correlation of Monthly Inflation in 2020
From Experimental Price Indices and CPI

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Average Price Index (API)	1.000							
(2) API Independent	0.977***	1.000						
(3) API Multi-outlet	0.652***	0.498**	1.000					
(4) Average Variation Index (AVI)	0.590**	0.545**	0.416*	1.000				
(5) AVI Independent	0.501**	0.457*	0.370	0.977***	1.000			
(6) AVI Multi-outlet	0.403*	0.398	0.210	0.093	-0.121	1.000		
(7) Mexico City FAFH CPI	-0.117	-0.154	0.006	0.168	0.179	-0.043	1.000	
(8) National FAFH CPI	0.126	0.181	-0.041	0.271	0.310	-0.178	0.654***	1.000

*** p<0.01, ** p<0.05, * p<0.1

Figure 19: Experimental Price Indices

By Dish Type

(a) Average Price Index

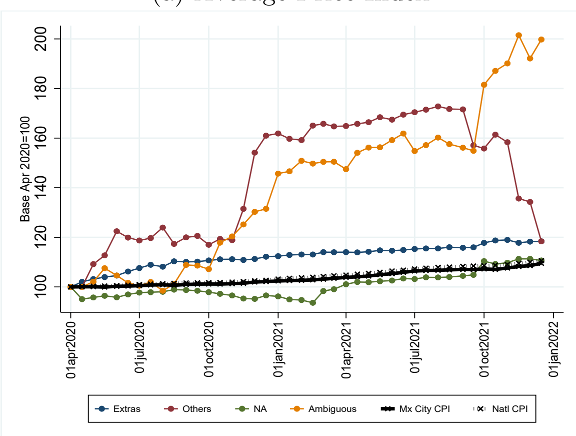

(b) Average Variation Index

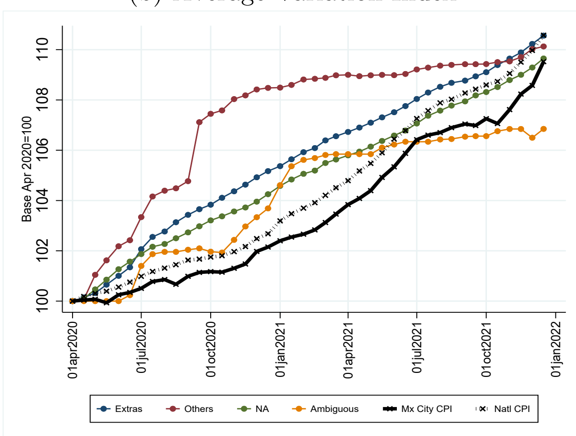

A.4 Stylized Facts

This subsection summarizes (i) using seasonal fixed effects as instead of time controls in the benchmark regressions and (ii) a comparison of results using "representative dishes" with respect to estimates while using the complete dataset. The regressions take form of:

$$P\left(y_{i,j,t}=1|x\right)=\beta_1 x_{dishtype}+\beta_2 x_{dow}+\beta_3 x_{day}+\beta_4 x_{month}+\beta_5 x_{year}+\beta_6 x_j+\varepsilon_{i,j,t}$$

where $y_{i,j,t}=1$ is a dummy variable if the price of item i at restaurant j at time t changed with respect to day $t-1$, or zero otherwise. x_{dow}, x_{day}, x_{month}, x_{year} represent day of the week, calendar day, month and year fixed effects, respectively. x_j represents a restaurant fixed effects.[55] Likewise, the second equation studies the (absolute value) of size of price adjustments, given a price change:

$$\mid \Delta y_{i,j,t} \mid = \beta_1 x_{dishtype}+\beta_2 x_{dow}+\beta_3 x_{day}+\beta_4 x_{month}+\beta_5 x_{year}+\beta_6 x_j+\varepsilon_{i,j,t}$$

Figure 20: Stylized Facts of Price Adjustments

Representative Dishes Using Different Set of Time FE
Price Changes Regardless Sign of Adjustment

(a) Frequency of Changes (b) Size of Adjustments

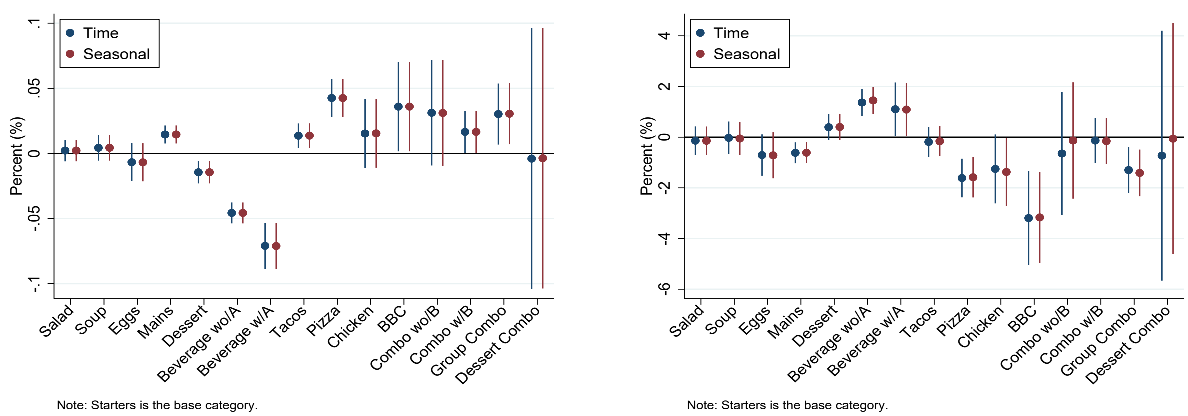

[55] Day of the week and calendar day coefficients are not reported on basis of the informant confidentiality.

Table 10: Stylized Facts of Pride Changes
Linear Probability Model
Representative Dishes

	(1)	(2)	(3)	(4)	(5)	(6)
	$\mathbb{1}\_\Delta p\_it \neq 0$	$\mathbb{1}\_\Delta p\_it \neq 0$	$\mathbb{1}\_\Delta p\_it > 0$	$\mathbb{1}\_\Delta p\_it > 0$	$\mathbb{1}\_\Delta p\_it < 0$	$\mathbb{1}\_\Delta p\_it < 0$
Starter	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)
Salad	0.00	0.00	0.00	0.00	0.00	0.00
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Soup	0.00	0.00	0.00	0.00	0.00	0.00
	(0.005)	(0.005)	(0.005)	(0.005)	(0.002)	(0.002)
Eggs	-0.01	-0.01	-0.01	-0.01	-0.00	-0.00
	(0.007)	(0.007)	(0.007)	(0.007)	(0.002)	(0.002)
Mains	0.01***	0.01***	0.01***	0.01***	0.00*	0.00*
	(0.004)	(0.004)	(0.003)	(0.003)	(0.001)	(0.001)
Dessert	-0.01**	-0.01**	-0.02***	-0.02***	0.00	0.00
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Beverage wo/A	-0.05***	-0.05***	-0.04***	-0.04***	-0.00*	-0.00*
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Beverage w/A	-0.07***	-0.07***	-0.06***	-0.06***	-0.01**	-0.01**
	(0.009)	(0.009)	(0.007)	(0.007)	(0.004)	(0.004)
Tacos	0.01**	0.01**	0.01*	0.01*	0.00*	0.00*
	(0.005)	(0.005)	(0.004)	(0.004)	(0.002)	(0.002)
Pizza	0.04***	0.04***	0.03***	0.03***	0.01***	0.01***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.003)	(0.003)
Chicken	0.02	0.02	0.02	0.02	-0.00	-0.00
	(0.013)	(0.013)	(0.013)	(0.013)	(0.004)	(0.004)
BBC	0.04*	0.04*	0.03	0.03	0.01	0.01
	(0.018)	(0.018)	(0.017)	(0.017)	(0.005)	(0.005)
Combo wo/B	0.03	0.03	0.03	0.03	0.00	0.00
	(0.021)	(0.021)	(0.019)	(0.019)	(0.008)	(0.008)
Combo w/B	0.02*	0.02*	0.01	0.01	0.01*	0.01*
	(0.008)	(0.008)	(0.007)	(0.007)	(0.003)	(0.003)
Group Combo	0.03*	0.03*	0.02*	0.02*	0.01	0.01
	(0.012)	(0.012)	(0.010)	(0.010)	(0.005)	(0.005)
Dessert Combo	-0.00	-0.00	-0.01	-0.01	0.01	0.01
	(0.051)	(0.051)	(0.052)	(0.052)	(0.016)	(0.016)
Extras	-0.01***	-0.01***	-0.01***	-0.01***	-0.00	-0.00
	(0.004)	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)
Others	-0.03	-0.03	-0.02	-0.02	-0.00	-0.00
	(0.023)	(0.023)	(0.019)	(0.019)	(0.007)	(0.007)
NA	-0.04***	-0.04***	-0.03***	-0.03***	-0.00	-0.00
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Ambiguous	-0.04	-0.04	-0.05***	-0.05***	0.02	0.02
	(0.021)	(0.021)	(0.014)	(0.014)	(0.014)	(0.014)
Observations	82615078	82615078	82615078	82615078	82615078	82615078
Adjusted R^2	0.005	0.004	0.004	0.004	0.003	0.003
Restaurant FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	.	Yes	.	Yes	.
DOW FE	.	Yes	.	Yes	.	Yes
DOC FE	.	Yes	.	Yes	.	Yes
Month FE	.	Yes	.	Yes	.	Yes
Year FE	.	Yes	.	Yes	.	Yes

DOW and DOC stand for day of the week and calendar respectively.
Estimates multiplied by 100 for ilustration purposes.
Standard errors clustered at restaurant level. *** p<0.01, ** p<0.05, * p<0.1

Table 11: Stylized Facts of Pride Changes
Size of Price Adjustments
Representative Dishes

	(1) Abs Price Adj	(2) Abs Price Adj	(3) Abs Price Adj	(4) Abs Price Adj	(5) Abs Price Adj	(6) Abs Price Adj
Starter	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)
Salad	-0.14	-0.14	0.05	0.07	-2.14*	-2.27*
	(0.287)	(0.288)	(0.275)	(0.270)	(0.894)	(0.893)
Soup	-0.03	-0.05	0.22	0.19	-1.80	-1.32
	(0.329)	(0.330)	(0.339)	(0.336)	(1.057)	(1.092)
Eggs	-0.71	-0.72	-0.64	-0.70	-0.94	-0.92
	(0.416)	(0.464)	(0.390)	(0.416)	(1.044)	(1.253)
Mains	-0.62**	-0.61**	-0.45*	-0.47*	-1.64*	-1.59*
	(0.211)	(0.213)	(0.204)	(0.205)	(0.707)	(0.707)
Dessert	0.39	0.40	0.59*	0.56*	-0.83	-0.83
	(0.262)	(0.267)	(0.255)	(0.259)	(0.828)	(0.852)
Beverage wo/A	1.37***	1.45***	1.52***	1.57***	0.17	-0.21
	(0.268)	(0.273)	(0.255)	(0.261)	(0.880)	(0.923)
Beverage w/A	1.10*	1.09*	1.39*	1.20*	-1.95	-1.83
	(0.538)	(0.534)	(0.554)	(0.543)	(2.170)	(1.952)
Tacos	-0.19	-0.16	-0.19	-0.16	-1.10	-1.05
	(0.299)	(0.302)	(0.281)	(0.283)	(1.038)	(1.068)
Pizza	-1.61***	-1.58***	-1.31***	-1.33***	-1.34	-1.07
	(0.389)	(0.407)	(0.357)	(0.381)	(1.361)	(1.369)
Chicken	-1.25	-1.37*	-1.25	-1.42*	-0.88	-0.58
	(0.694)	(0.680)	(0.655)	(0.652)	(3.023)	(2.922)
BBC	-3.19***	-3.17***	-2.59**	-2.62**	-10.12	-11.44
	(0.943)	(0.914)	(0.863)	(0.813)	(7.307)	(7.753)
Combo wo/B	-0.64	-0.13	-0.75	-0.02	1.04	2.06
	(1.239)	(1.171)	(1.131)	(1.063)	(3.360)	(2.998)
Combo w/B	-0.13	-0.16	-0.22	-0.26	0.21	0.03
	(0.455)	(0.463)	(0.396)	(0.401)	(1.715)	(1.698)
Group Combo	-1.30**	-1.41**	-1.03*	-1.28**	0.62	1.48
	(0.461)	(0.470)	(0.478)	(0.463)	(1.482)	(1.794)
Dessert Combo	-0.73	-0.06	3.27	3.44	-1.02	-1.02
	(2.516)	(2.326)	(2.171)	(1.910)	(1.320)	(1.284)
Extras	0.85***	0.86***	1.11***	1.10***	-1.19	-1.18
	(0.255)	(0.258)	(0.252)	(0.252)	(0.790)	(0.805)
Others	3.08*	3.11*	3.29**	3.55**	3.53	1.27
	(1.210)	(1.466)	(1.023)	(1.198)	(4.642)	(5.026)
NA	0.58**	0.64**	0.74***	0.78***	-0.83	-0.87
	(0.224)	(0.229)	(0.214)	(0.220)	(0.720)	(0.723)
Ambiguous	-0.00	-0.13	-0.46	-0.45	-1.61	-2.01
	(1.818)	(1.714)	(2.292)	(2.246)	(2.075)	(1.831)
Observations	158365	158366	137139	137140	20415	20436
Adjusted R^2	0.462	0.432	0.500	0.463	0.621	0.563
Restaurant FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	.	Yes	.	Yes	.
DOW FE	.	Yes	.	Yes	.	Yes
DOC FE	.	Yes	.	Yes	.	Yes
Month FE	.	Yes	.	Yes	.	Yes
Year FE	.	Yes	.	Yes	.	Yes

DOW and DOC stand for day of the week and calendar respectively.
Estimates multiplied by 100 for ilustration purposes.
Standard errors clustered at restaurant level. *** p<0.01, ** p<0.05, * p<0.1

Table 12: Stylized Facts of Pride Changes
Linear Probability Model
All Dishes (Representative or Not)

	(1)	(2)	(3)	(4)	(5)	(6)
	$\mathbb{1}\_\Delta p\_it \neq 0$	$\mathbb{1}\_\Delta p\_it \neq 0$	$\mathbb{1}\_\Delta p\_it > 0$	$\mathbb{1}\_\Delta p\_it > 0$	$\mathbb{1}\_\Delta p\_it < 0$	$\mathbb{1}\_\Delta p\_it < 0$
Starter	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)
Salad	0.00	0.00	0.00	0.00	0.00	0.00
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Soup	0.00	0.00	0.01	0.01	-0.00	-0.00
	(0.005)	(0.005)	(0.004)	(0.004)	(0.002)	(0.002)
Eggs	-0.01	-0.01	-0.00	-0.00	-0.00	-0.00
	(0.007)	(0.007)	(0.006)	(0.006)	(0.003)	(0.003)
Mains	0.01***	0.01***	0.01***	0.01***	0.00	0.00
	(0.003)	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)
Dessert	-0.01***	-0.01***	-0.01**	-0.01**	-0.00	-0.00
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Beverage wo/A	-0.05***	-0.05***	-0.04***	-0.04***	-0.01***	-0.01***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Beverage w/A	-0.07***	-0.07***	-0.06***	-0.06***	-0.01***	-0.01***
	(0.008)	(0.008)	(0.006)	(0.006)	(0.004)	(0.004)
Tacos	0.02**	0.02**	0.01**	0.01**	0.00*	0.00*
	(0.005)	(0.005)	(0.004)	(0.004)	(0.002)	(0.002)
Pizza	0.05***	0.05***	0.04***	0.04***	0.01**	0.01**
	(0.007)	(0.007)	(0.006)	(0.006)	(0.003)	(0.003)
Chicken	0.01	0.01	0.02	0.02	-0.00	-0.00
	(0.014)	(0.014)	(0.012)	(0.012)	(0.005)	(0.005)
BBC	0.03	0.03	0.03	0.03*	0.00	0.00
	(0.016)	(0.016)	(0.015)	(0.015)	(0.005)	(0.005)
Combo wo/B	0.01	0.01	0.02	0.02	-0.01	-0.01
	(0.020)	(0.020)	(0.017)	(0.017)	(0.008)	(0.008)
Combo w/B	0.01	0.01	0.01	0.01	0.01*	0.01*
	(0.008)	(0.008)	(0.007)	(0.007)	(0.003)	(0.003)
Group Combo	0.02	0.02	0.01	0.01	0.01	0.01
	(0.011)	(0.011)	(0.009)	(0.009)	(0.005)	(0.005)
Dessert Combo	-0.02	-0.02	-0.04	-0.04	0.02	0.02
	(0.047)	(0.047)	(0.042)	(0.042)	(0.022)	(0.022)
Extras	-0.02***	-0.02***	-0.01***	-0.01***	-0.00	-0.00
	(0.004)	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)
Others	-0.04*	-0.04*	-0.03	-0.03	-0.01*	-0.01*
	(0.019)	(0.019)	(0.017)	(0.017)	(0.005)	(0.005)
NA	-0.04***	-0.04***	-0.04***	-0.04***	-0.01**	-0.01***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Ambiguous	-0.03	-0.03	-0.05**	-0.05**	0.03	0.03
	(0.026)	(0.025)	(0.020)	(0.020)	(0.018)	(0.018)
Observations	116756785	116756785	116756785	116756785	116756785	116756785
Adjusted R^2	0.008	0.008	0.008	0.007	0.005	0.005
Restaurant FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	.	Yes	.	Yes	.
DOW FE	.	Yes	.	Yes	.	Yes
DOC FE	.	Yes	.	Yes	.	Yes
Month FE	.	Yes	.	Yes	.	Yes
Year FE	.	Yes	.	Yes	.	Yes

DOW and DOC stand for day of the week and calendar respectively.
Estimates multiplied by 100 for ilustration purposes.
Standard errors clustered at restaurant level. *** p<0.01, ** p<0.05, * p<0.1

Table 13: Stylized Facts of Pride Changes

Size of Price Adjustments
All Dishes (Representative or Not)

	(1) Abs Price Adj	(2) Abs Price Adj	(3) Abs Price Adj	(4) Abs Price Adj	(5) Abs Price Adj	(6) Abs Price Adj
Starter	0.00	0.00	0.00	0.00	0.00	0.00
	(.)	(.)	(.)	(.)	(.)	(.)
Salad	-0.12	-0.04	0.04	0.13	-1.60	-1.11
	(0.284)	(0.282)	(0.270)	(0.265)	(0.875)	(0.888)
Soup	0.06	0.04	0.32	0.26	-1.91*	-1.15
	(0.322)	(0.320)	(0.332)	(0.324)	(0.963)	(0.976)
Eggs	-0.82*	-0.84*	-0.72	-0.77*	-1.26	-0.84
	(0.392)	(0.425)	(0.372)	(0.388)	(0.978)	(1.240)
Mains	-0.49*	-0.45*	-0.37	-0.36	-1.40*	-1.14
	(0.206)	(0.207)	(0.198)	(0.198)	(0.620)	(0.639)
Dessert	0.50	0.57*	0.61*	0.65**	-0.27	-0.05
	(0.258)	(0.261)	(0.249)	(0.250)	(0.748)	(0.776)
Beverage wo/A	1.35***	1.40***	1.49***	1.50***	0.12	0.15
	(0.256)	(0.259)	(0.245)	(0.248)	(0.780)	(0.810)
Beverage w/A	0.79	0.82	0.80	0.73	-1.08	-1.59
	(0.513)	(0.513)	(0.510)	(0.502)	(1.863)	(1.827)
Tacos	0.18	0.25	0.13	0.20	-0.68	-0.60
	(0.306)	(0.307)	(0.282)	(0.284)	(0.928)	(0.949)
Pizza	-1.86***	-1.92***	-1.53***	-1.69***	-0.85	-0.47
	(0.433)	(0.492)	(0.414)	(0.506)	(1.146)	(1.125)
Chicken	-1.11	-1.13	-1.29	-1.32*	-0.48	-0.48
	(0.711)	(0.687)	(0.662)	(0.655)	(2.716)	(2.728)
BBC	-2.26*	-2.05*	-2.19**	-2.17**	-4.82	-6.47
	(0.994)	(0.977)	(0.825)	(0.801)	(6.306)	(6.176)
Combo wo/B	-0.82	-0.23	-0.77	0.01	0.66	2.22
	(1.164)	(1.148)	(1.096)	(1.036)	(3.692)	(2.862)
Combo w/B	0.06	0.16	-0.32	-0.15	1.98	1.53
	(0.436)	(0.445)	(0.395)	(0.412)	(1.399)	(1.420)
Group Combo	-0.74	-0.61	-0.33	-0.45	-0.15	0.38
	(0.534)	(0.565)	(0.512)	(0.523)	(1.538)	(1.728)
Dessert Combo	2.41	3.27	4.35*	4.17*	2.39	7.02**
	(1.949)	(2.034)	(2.078)	(1.964)	(1.511)	(2.550)
Extras	1.00***	0.97***	1.21***	1.21***	-0.54	-0.61
	(0.249)	(0.249)	(0.245)	(0.243)	(0.705)	(0.735)
Others	3.36**	3.16*	3.44***	3.58***	-0.02	0.40
	(1.202)	(1.367)	(0.958)	(1.085)	(4.260)	(4.676)
NA	0.78***	0.92***	0.86***	0.98***	-0.45	-0.21
	(0.219)	(0.226)	(0.210)	(0.217)	(0.638)	(0.666)
Ambiguous	-1.39	-1.20	-2.31	-2.41	-2.21	-2.47
	(1.980)	(1.795)	(2.291)	(2.453)	(1.952)	(1.752)
Observations	228072	228072	190681	190681	35859	35873
Adjusted R^2	0.476	0.451	0.509	0.479	0.619	0.567
Restaurant FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	.	Yes	.	Yes	.
DOW FE	.	Yes	.	Yes	.	Yes
DOC FE	.	Yes	.	Yes	.	Yes
Month FE	.	Yes	.	Yes	.	Yes
Year FE	.	Yes	.	Yes	.	Yes

DOW and DOC stand for day of the week and calendar respectively.
Estimates multiplied by 100 for ilustration purposes.
Standard errors clustered at restaurant level. *** p<0.01, ** p<0.05, * p<0.1

Figure 21: Stylized Facts of Price Adjustments

Comparison By Dish Sample
Price Changes Regardless Sign of Adjustment

(a) Frequency of Changes

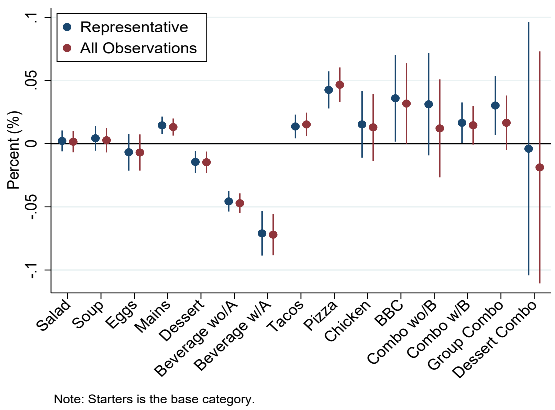

(b) Size of Adjustments

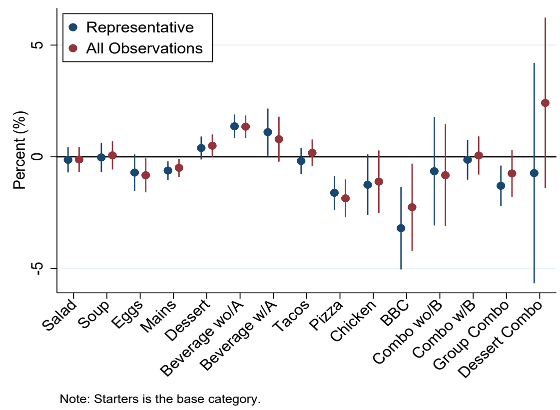

Figure 22: Stylized Facts of Price Adjustments

All Observations (Not Only Representative) Dishes
By Sign of Price Adjustment

(a) Frequency of Changes

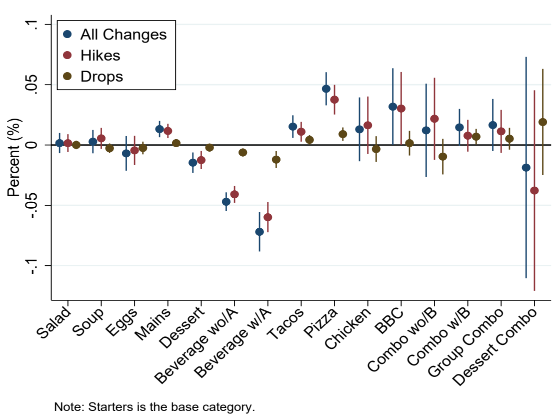

(b) Size of Adjustments

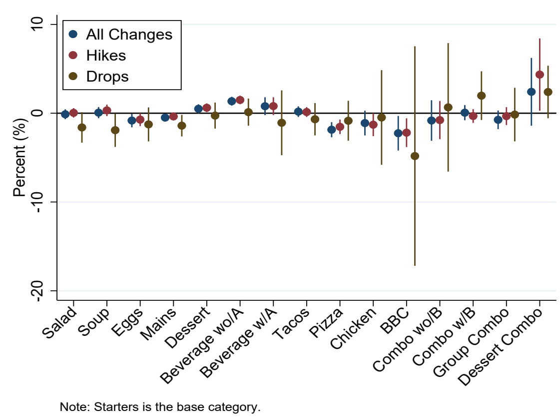

Figure 23: Stylized Facts of Price Hikes at Different Stages of the Pandemic

Representative Dishes

(a) Frequency of Changes (b) Size of Adjustments

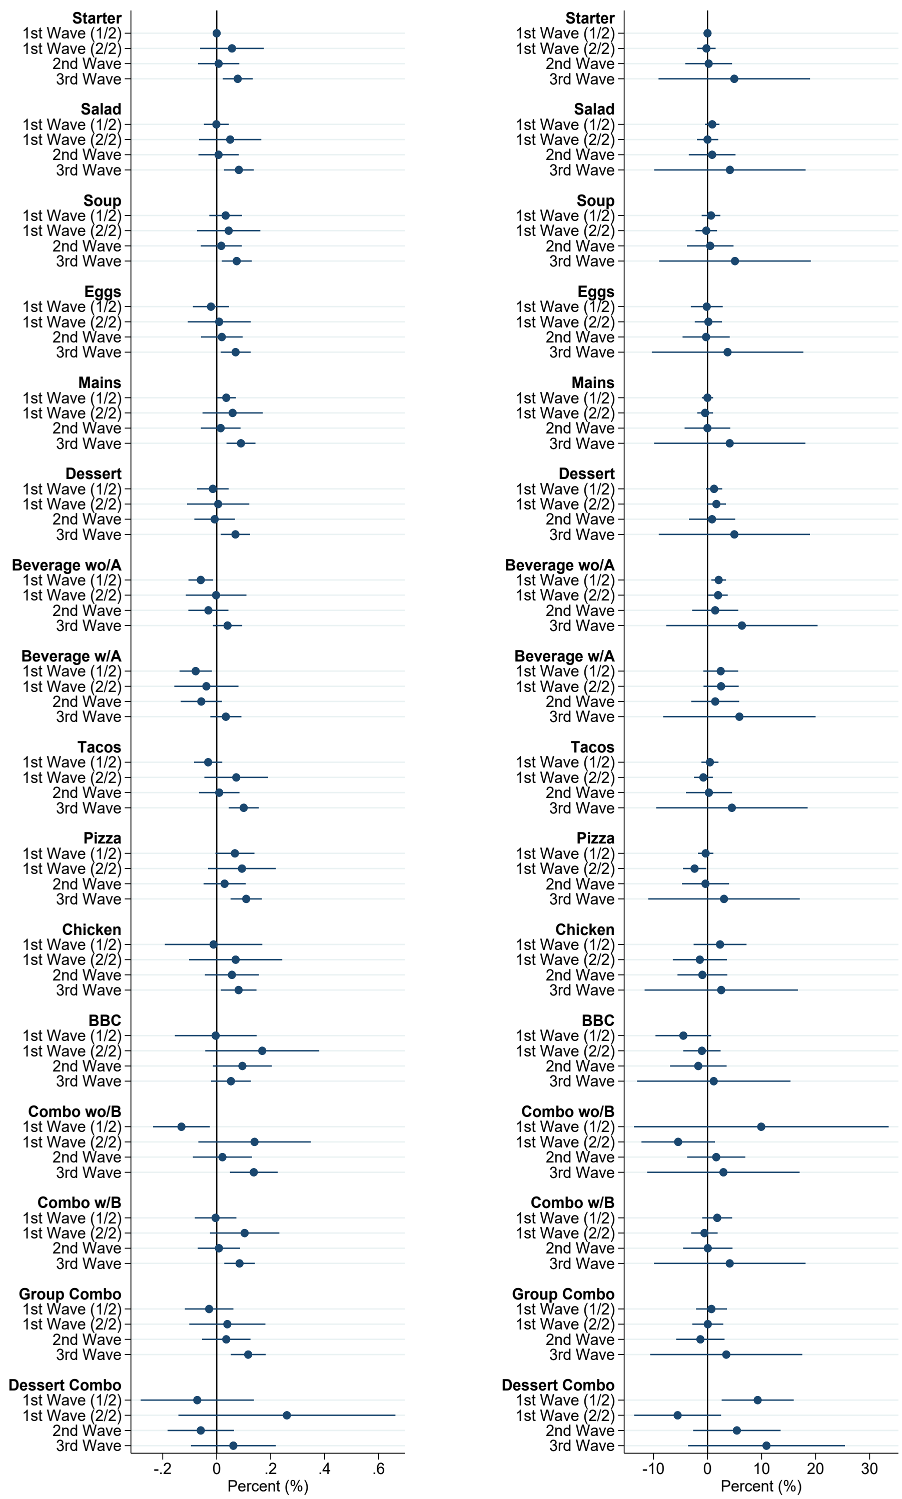

Figure 24: Stylized Facts of Price Drops at Different Stages of the Pandemic

Representative Dishes

(a) Frequency of Changes (b) Size of Adjustments

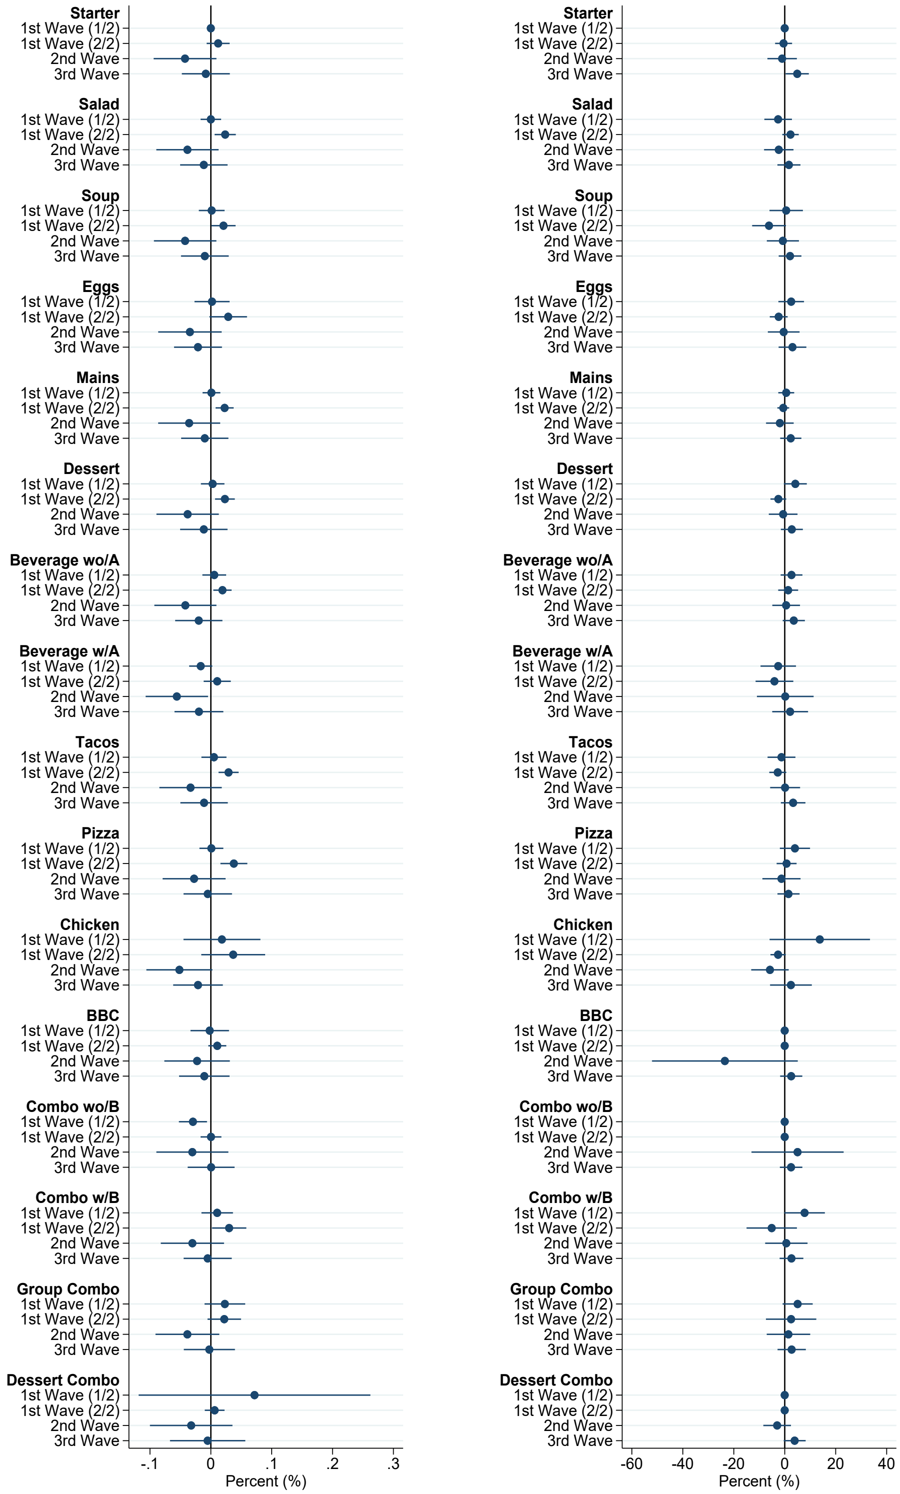

A.5 In-Sample FAFH Inflation

Table 14: Descriptive Statistics
By Explanatory Variable from 2006 to 2017

	Mean	SD	Min	Max
Electricity	3.49	4.87	-6.58	14.47
LP Gas	5.22	6.32	-7.85	37.76
Natural Gas	5.28	8.58	-18.17	30.58
Perm N Workers	3.22	3.27	-5.03	8.58
Temp N Workers	5.45	22.93	-62.58	68.26
Perm Wage Bill	7.31	3.73	-2.80	13.70
Perm Mean Wage	4.09	1.40	0.61	7.27
Temp Wage Bill	7.78	27.15	-65.50	86.48
Temp Mean Wage	2.33	8.25	-15.97	21.14
Perm Real Wage Bill	3.31	4.22	-8.20	10.08
Perm Mean Real Wage	0.09	1.53	-3.46	3.39
Temp Real Wage Bill	7.78	27.15	-65.50	86.48
Temp Mean Real Wage	2.33	8.25	-15.97	21.14
(City) Avocado	8.83	28.65	-68.68	96.52
(City) Rice	6.14	15.93	-28.95	59.85
(City) Sugar	7.60	28.27	-57.66	90.70
(City) Zucchini	4.00	45.72	-158.60	166.56
(City) Shrimp	2.97	9.69	-21.47	24.93
(City) Onion	6.74	49.44	-143.35	151.00
(City) Squash	2.97	42.98	-150.63	163.90
(City) Poblano Chili	3.86	40.77	-108.43	113.07
(City) Dried Chili	2.95	18.93	-52.51	50.23
(City) Serrano Chili	4.99	56.98	-152.48	151.37
(City) Peach	3.83	33.66	-76.04	90.16
(City) Green Beans	3.09	45.86	-135.25	119.89
(City) Beans	7.33	27.21	-43.14	72.99
(City) Other Fruits	3.21	20.57	-50.86	58.63
(City) Guava	4.99	25.08	-73.91	90.63
(City) Eggs	7.25	22.23	-57.06	72.37
(City) Tomato	5.11	42.28	-123.60	129.39
(City) Lettuce	5.57	15.44	-45.03	69.90
(City) Legumes	7.25	18.40	-27.19	57.28
(City) Lemon	10.52	46.29	-136.37	148.76
(City) Corn	10.83	19.65	-15.96	61.10
(City) Apple	8.78	21.25	-46.30	57.38
(City) Fish Others	1.42	4.59	-8.83	13.23
(City) Melon	4.87	23.67	-76.62	69.24
(City) Orange	8.72	31.86	-98.89	74.40
(City) Nopales	2.97	48.92	-163.32	135.55
(City) Potatoe	2.35	30.70	-79.95	67.24
(City) Papaya	4.75	35.78	-84.28	89.72
(City) Cucumber	4.16	38.67	-125.12	116.32
(City) Pear	7.40	25.01	-72.53	106.91
(City) Fish	3.76	4.93	-10.12	22.19
(City) Pineapple	5.49	22.79	-98.02	62.25
(City) Bananas	4.25	22.96	-65.23	74.72
(City) Chicken	6.20	17.81	-40.18	56.25
(City) Watermelon	5.34	26.51	-58.96	91.30
(City) Green Tomato	0.69	66.85	-169.96	177.54
(City) Grape	8.89	22.89	-55.94	87.55
(City) Carrot	4.95	39.47	-123.70	137.96
(Met) Avocado	8.59	24.36	-69.09	86.52
(Met) Rice	6.93	15.45	-21.88	62.32
(Met) Sugar	7.50	28.39	-58.43	90.59
(Met) Zucchini	5.26	40.43	-143.49	136.48
(Met) Onion	5.00	57.02	-191.35	165.32
(Met) Pork	5.52	10.43	-16.01	34.41
(Met) Squash	4.92	33.92	-91.82	153.38
(Met) Poblano Chili	3.98	32.32	-76.22	94.82
(Met) Dried Chili	3.54	15.50	-31.85	39.83
(Met) Serrano Chili	5.25	50.10	-134.32	137.65
(Met) Peach	2.85	23.70	-50.79	67.76
(Met) Green Beans	4.66	38.37	-117.15	98.96
(Met) Beans	7.74	26.42	-44.07	71.08
(Met) Other Fruits	4.03	13.29	-40.73	44.06
(Met) Guava	2.37	18.37	-41.76	45.59
(Met) Eggs	7.28	22.53	-55.81	75.87
(Met) Tomato	4.60	48.00	-109.21	123.33
(Met) Lettuce	5.38	13.76	-48.31	47.40
(Met) Lemon	10.23	47.40	-145.52	153.60
(Met) Corn	11.03	26.84	-46.60	94.14
(Met) Apple	8.46	15.68	-31.33	55.54
(Met) Melon	4.99	20.18	-53.20	61.22
(Met) Orange	6.96	28.49	-92.74	70.26
(Met) Nopales	5.50	42.10	-143.35	120.71
(Met) Potatoe	3.55	30.43	-58.78	80.90
(Met) Papaya	3.44	30.91	-71.46	78.87
(Met) Cucumber	6.57	33.03	-82.20	99.73
(Met) Pear	5.89	17.38	-32.42	57.96
(Met) Pineapple	9.01	30.78	-109.78	91.43
(Met) Bananas	5.33	18.34	-47.69	63.39
(Met) Beef	6.21	7.50	-7.20	24.84
(Met) Watermelon	4.80	22.62	-56.68	85.37
(Met) Green Tomato	0.86	63.06	-160.67	169.66
(Met) Grape	6.15	22.05	-59.88	74.52
(Met) Carrot	5.42	32.34	-114.08	116.99

Table 15: Grid of Hyper-parameters for Cross-Validation

Model	Parameters
Elastic Net	
	$\alpha \in \{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 0.75, 0.9, 1, 1.1, 1.25, 1.5, 2, 5, 10, 20, 30, 40, 50\}$
	L1 ratio $\in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$
Lasso	
	$\alpha \in \{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 0.75, 0.9, 1, 1.1, 1.25, 1.5, 2, 5, 10, 20, 30, 40, 50\}$
Ridge	
	$\alpha \in \{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 0.75, 0.9, 1, 1.1, 1.25, 1.5, 2, 5, 10, 20, 30, 40, 50\}$
Random Forest Regression	
	N Estimators $\in \{15, 25, 50, 100, 150, 200, 250, 300\}$
	Max Depth $\in \{3, 5, 10, 15, 20, 30, 40\}$
	Min Sample Split $\in \{5, 10, 20, 30, 40, 50\}$
Support Vector Machine Regression	
	Kernel $\in \{rbf, linear, sigmoid, poly\}$
	Gamma $\in \{scale, auto\}$

Note: Cross-validation through expanding window approach using fortnightly observations from 2006 to 2017. Benchmark results are computed using 24 folds.

Table 16: Hyper-parameter Selection

By Model and Horizon

		Horizon (Months)				
Model	Parameter	t	$t+3$	$t+6$	$t+9$	$t+12$
Elastic Net						
	α	0.0005	0.001	0.005	0.001	0.05
	L1 ratio	0.9	0.9	0.1	0.9	0.6
Lasso						
	α	0.0005	0.001	0.0005	0.001	0.05
Ridge						
	α	0.1	20	5	5	50
Random Forest Regression						
	N Estimators	200	300	50	300	100
	Max Depth	5	10	3	10	20
	Min Sample Split	5	5	10	5	5
Support Vector Machine Regression						
	Kernel	*rbf*	*linear*	*linear*	*linear*	*rbf*
	Gamma	*auto*	*scale*	*scale*	*scale*	*auto*

Table 17: Models' Root Mean Squared Error

By Horizon of Prediction. Training period using determinants from 2006 to 2017.

	Horizon (Months)				
Model	t	$t+3$	$t+6$	$t+9$	$t+12$
EE	0.53940	0.57309	0.72316	0.30331	0.83982
Lasso	0.44220	0.47105	0.72179	0.41822	0.77364
Ridge	0.34456	0.37075	0.33558	0.29300	0.39168
RFR	0.14271	0.13770	0.29559	0.37066	0.21291
SVMR	0.38481	0.37851	0.34564	0.34869	0.37814

Figure 25: Models' Fit Over Training Period

By Horizon of Prediction. Only first fortnights of every month are plotted for illustration purposes.

(a) EE

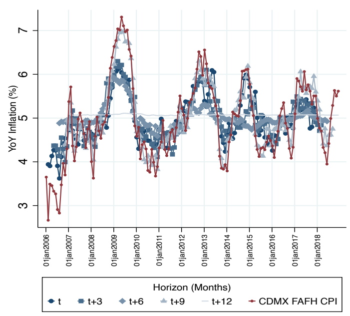

(b) Lasso

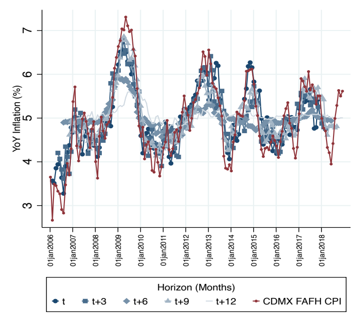

(c) Ridge

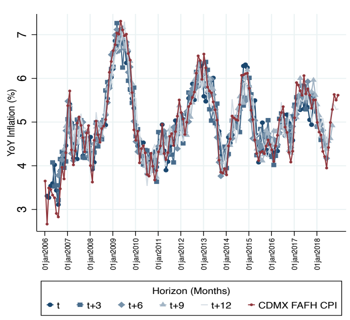

(d) RFR

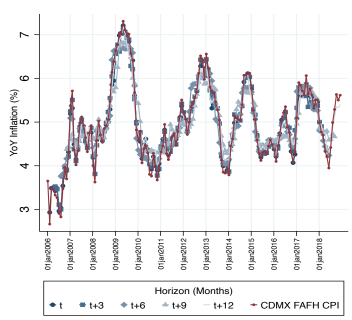

(e) SVMR

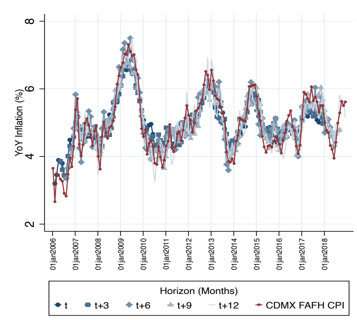

Figure 26: Complete Set of Features

Features on the vertical axis and horizons in the horizontal axis.[a]
Selected features

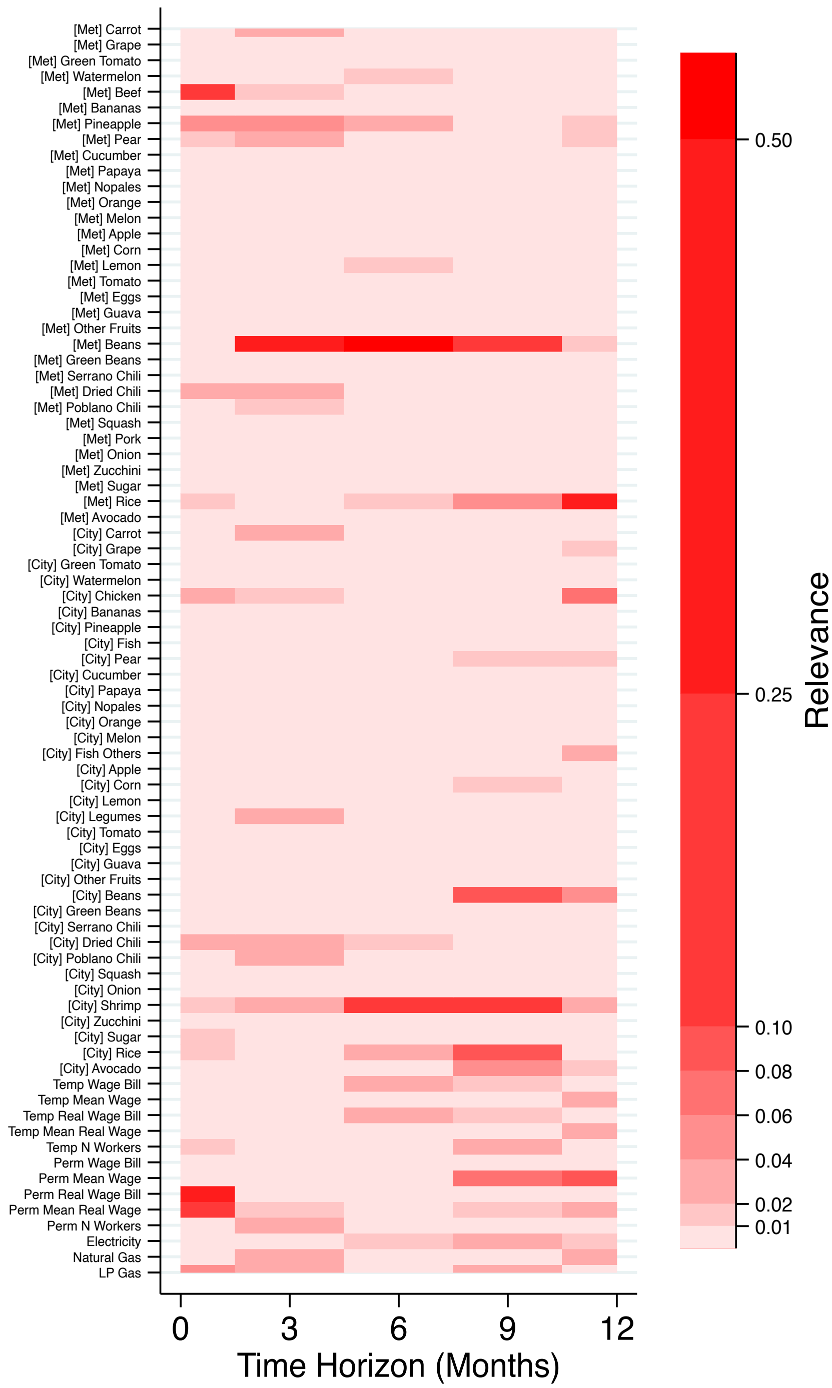

[a]Features' relevance by horizon adds up to 1.

A.6 Out-of-Sample FAFH Inflation Goodness-of-Fit

Figure 27: Models' Root Mean Squared Error from 2018 to 2019

By Horizon of Prediction.

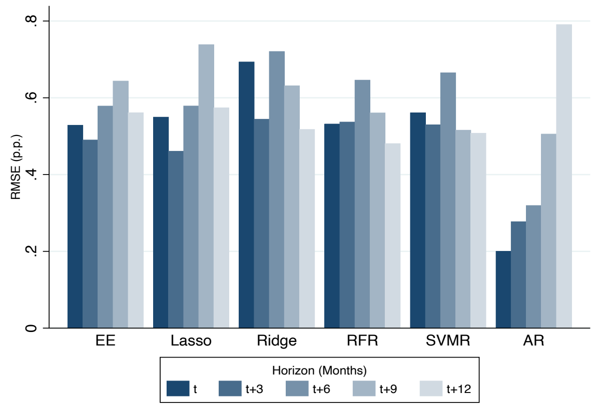

Figure 28: Models' Predictions

By Horizon of Prediction. Only first fortnights of every month are plotted for illustration purposes.

(a) EE

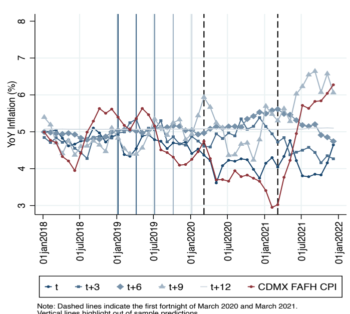

(b) Lasso

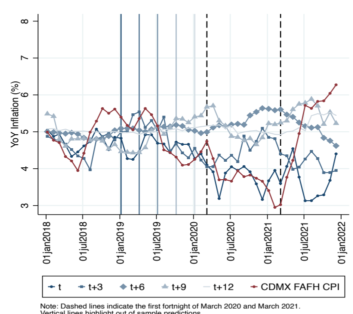

(c) Ridge

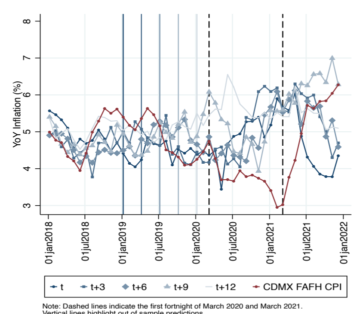

(d) RFR

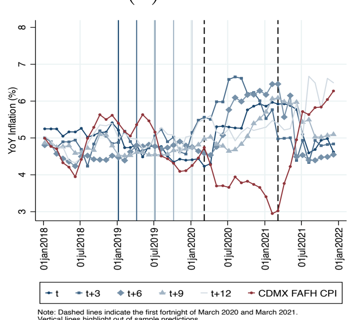

(e) SVMR

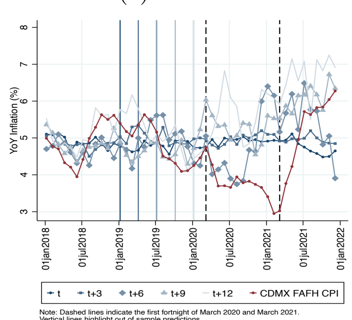

(f) AR

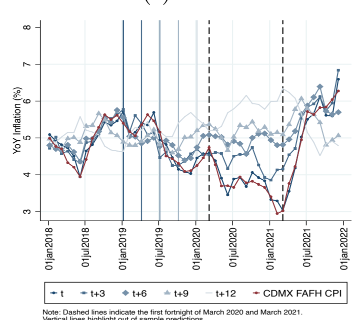

Table 18: Diebold-Mariano's Forecasting Accuracy 2018-2019[a]

Horizon t + 0

	EE	Lasso	Ridge	RFR	SVMR	AR
EE	- o -					
Lasso		- o -				
Ridge			- o -			
RFR		RFR*		- o -		
SVMR					- o -	
AR	AR***	AR***	AR*	AR***		- o -

Horizon t + 3 months

	EE	Lasso	Ridge	RFR	SVMR	AR
EE	- o -					
Lasso	Lasso***	- o -				
Ridge		Lasso**	- o -			
RFR				- o -		
SVMR					- o -	
AR	AR*	AR***	AR**			- o -

Horizon t + 6 months

	EE	Lasso	Ridge	RFR	SVMR	AR
EE	- o -					
Lasso		- o -				
Ridge			- o -			
RFR			RFR*	- o -		
SVMR					- o -	
AR			AR***	AR***	AR***	- o -

Horizon t + 9 months

	EE	Lasso	Ridge	RFR	SVMR	AR
EE	- o -					
Lasso		- o -				
Ridge			- o -			
RFR			RFR*	- o -		
SVMR			SVMR*		- o -	
AR						- o -

Horizon t + 12 months

	EE	Lasso	Ridge	RFR	SVMR	AR
EE	- o -					
Lasso	EE**	- o -				
Ridge	EE***		- o -			
RFR				- o -		
SVMR					- o -	
AR						- o -

[a]Note: Each cell reports the model with better accuracy between the column-model and row-model in question if statistically significant. Empty spaces in the lower diagonal imply no statistically significant difference in the accuracy between the column-model and the row-model. *, **, *** represent $p < 0.10$, $p < 0.05$ and $p < 0.01$, respectively, in the Diebold-Mariano test.