Forward Inflation Expectations: Evidence from Inflation Caps and Floors

Karsten O. Chipeniuk*† Todd B. Walker‡

July 2021

Abstract

Using daily prices of inflation caps and floors from 2012 through 2017, we document that caps are more sensitive to statements by the Board of Governors of the Federal Reserve or its Chair, while changes in floors are typically attributable to structural economic performance (e.g., labor markets, oil prices). We adapt nonparametric estimation methods to derive forward probabilities from the empirical distribution of historical U.S. inflation and estimate a regime-switching model for dispersion in the forward distribution. We link the transitions in regime to perceived changes in monetary policy and sources of major economic uncertainty. While our analysis is not able to assign a quantitative value to the change in inflation expectations (i.e., inflation expectations have declined by X%), we can confirm the direction of change at a daily frequency. Moreover, we can assign nearly all substantial movements to specific events, documenting asymmetry in the behavior of inflation expectations. We provide a structural interpretation of our findings, emphasizing the importance of a zero-lower bound over this time period. We also connect the dispersion in inflation expectations to newly developed measures of uncertainty [Jurado, Ludvigson, and Ng (2015), Baker, Bloom, and Davis (2016)]. Innovations to financial and macroeconomic uncertainty increase the dispersion in inflation floors, while decreasing the dispersion in inflation caps.

Keywords: Time-Series Models, Inflation, Financial Markets, Asset Pricing
JEL Classification Numbers: C22, E31, E44, G12

*Economics Department, Reserve Bank of New Zealand, Karsten.Chipeniuk@rbnz.govt.nz.

†The views expressed in this article are solely those of the author and do not necessarily reflect the position of the Reserve Bank of New Zealand.

‡Department of Economics, Indiana University, walkertb@indiana.edu

1 Introduction

The market for direct hedges of the inflation rate has matured since its inception, with $1.2 trillion worth of treasury inflation protected securities (TIPS) issued in 2017[1] and over $2 billion of notional principle cleared daily in the inflation swap market.[2] Alongside the TIPS and swap markets, cap and floor contracts have been written over-the-counter since late 2009. A question of continuing interest to financial market participants and central banks is the extent to which activity in the market for inflation options provides information with respect to investor beliefs about future changes in the rate of inflation.

We analyze daily Bloomberg composite prices for zero-coupon caps and floors over the period starting in January 2012 until May 2017. We argue the substantial increase in volume over the initial years of our sample period is sufficient for identifying changes in investor sentiment (Section 2). To back out implied probability densities, we use the "canonical valuation" method introduced by Buchen and Kelly (1996) and Stutzer (1996). This approach finds the forward densities that correctly price the inflation option (no arbitrage) and are the closest to the empirical distribution of inflation, where "closest" is measured according to the Kullback-Leibler Information Criterion (KLIC). The advantage of the canonical method for our purposes here is that we can isolate the density associated with specific options; that is, we do not need the entire set of option prices to form a forward distribution that is necessary of the "derivative method" developed by Ross (1976), Breeden and Litzenberger (1978) and Aït-Sahalia and Duarte (2003). This allows us to examine dispersion measures associated with inflation caps and floors separately. More importantly, because the canonical valuation methodology is not reliant on having a full set of option prices, we can exclude specific prices that suffer from potential illiquidity concerns.

Section 3 contains our main results. We find that dispersion in the forward inflation distribution exhibits clear breaks between regimes following major economic developments. Transitions in dispersion typically take place over a period of one day up to a week, and display high inertia. The high stability and rapid transitions allow us to sidestep the critique of Fair (2002), and we find that breaks are typically accompanied by either statements from the Federal Reserve and its Chairperson, or a shift in uncertainty around a major economic events such as the European debt crisis, US debt-ceiling crises, and the Greek bailouts. In particular, dispersion in the forward measure due to movements in the prices of *inflation caps* is associated with statements by the Board of Governors of the Federal Reserve or its Chair, while changes in dispersion due to movements in the prices of *inflation floors* are typically attributable to structural economic performance (e.g., labor markets, oil prices). Sections 3.2 conducts a narrative view of changes in regime and we count only three exceptions to this stylized fact over the entire sample period. Our approach demonstrates the benefit of high-frequency methods to identify monetary policy following Gertler and Karadi (2015), Hanson and Stein (2015) and Nakamura and Steinsson (2018).

The asymmetry in inflation caps and floors is our primary finding and we offer a structural interpretation in Section 3.3. The model is a relatively standard New Keynesian model with households that are subject to discount factor shocks and a monetary policy authority constrained by the zero-lower bound (ZLB). Using the nonlinear solution method and results developed in Richter, Throckmorton, and Walker

[1] See Kowara (2017).

[2] Daily volume reported by LCH: https://www.lch.com/services/swapclear/volumes

(2014), Gavin, Keen, Richter, and Throckmorton (2015) and Plante, Richter, and Throckmorton (2016), we show that technology shocks are qualitatively different under the ZLB. Weakness in the labor market is correlated with much lower expected inflation. This is consistent with our empirical finding that statistics portending structural weakness in the economy (e.g., a weak employment report) almost always caused an increase in the dispersion of inflation floors. The model can also generate substantial inflation uncertainty, defined as time-varying second moments, which is consistent with our regime-switching specification.

Section 3.4 documents that our measures of expected inflation dispersion are significantly predicted by changes in recently developed metrics of uncertainty and compares our metric to other measures of it (e.g., VIX). Specifically, we regress dispersion on the financial and macroeconomic uncertainty of Jurado, Ludvigson, and Ng (2015) and policy uncertainty of Baker, Bloom, and Davis (2016). All measures of uncertainty are significant predictors of dispersion but the uncertainty of Jurado, Ludvigson, and Ng (2015) enters negatively for forward dispersion due to inflation cap price movements and positively for floor price movements. Thus, an increase in financial uncertainty decreases (increases) the dispersion in caps (floors). This asymmetric response is consistent with the deflationary concerns revealed in our regime-switching results, and explains why innovations to financial time series, on average, consolidated expectations in inflation cap markets while unanchoring them in floor markets.

1.1 Connections to the Literature

We examine derivative data from inflation floors and caps, as opposed to the majority of the work in the inflation derivatives literature that attempts to back out expectations from Treasury Inflation-Protected Securities (TIPS). Absent market imperfections, the yield of an inflation protected treasury will be lower than that on a vanilla treasury by an amount equal to the expected inflation rate. The literature has achieved various levels of success in being able to back-out accurate forecasts of inflation with TIPS, with the primary concern being liquidity [see, Sack and Elsasser (2004), Fleming and Krishnan (2004), Gurkaynak, Sack, and Wright (2010), Grishchenko and Huang (2013), Fleckenstein, Longstaff, and Lustig (2014), Grishchenko, Vanden, and Zhang (2016), Andreasen, Christensen, and Riddell (2018)].[3] Kitsul and Wright (2013) also examine forward distributions for inflation implicit in cap and floor prices early in the market's lifetime. Since their study, the volume for inflation caps and floors has quadrupled and thus worthy of additional study. More importantly, and unlike Kitsul and Wright (2013), we systematize the event analysis to econometrically pinpoint when and why events transmitted through to investor uncertainty via a narrative view of regime change at a daily frequency.

Our methodological approach differs from the standard literature along two dimensions. First, a majority of the literature employs a "derivative method" based on the well known result that the second derivative of the price of a call option with respect to the strike delivers the risk-neutral density [Breeden and Litzenberger (1978), Ross (1976)].[4] In contrast, we adapt the method of Buchen and Kelly (1996) and Stutzer (1996), which minimizes the Kullback-Leibler Information Criterion (KLIC) metric, to estimate the market forward distributions. We argue that this method is a relatively efficient and flexible

[3] Of course liquidity is a concern for our data as well. We discuss this issue in the following section.

[4] The appendix examines the derivative method using our data and offers a comparison with the canonical approach.

procedure. Second, our paper does not attempt to recover expectations but instead focuses on connections between dispersion in forward distributions and underlying economic uncertainty. A simulation exercise in Section 3 demonstrates that if changes in risk premia amount primarily to time variation in the location parameter, our dispersion metric will correlate strongly across the physical and forward probabilities. Our focus on dispersion is empirically motivated by substantial time variation in tails of implied distributions and the associated pricing kernels found in our data. In addition, our empirical application establishes a strong case for informational content of the difference between the 25th and 75th percentiles of forward distributions for inflation, which are constrained to correctly price individual as well as multiple options. That there is something to learn about inflation expectations from these data relies on the underlying assumption that option prices contain information about extreme events relative to macroeconomic data [Backus, Chernov, and Martin (2011)].

Our primary results complement Christensen, Lopez, and Rudebusch (2015), Grishchenko, Vanden, and Zhang (2016), Fleckenstein, Longstaff, and Lustig (2017), Gimeno and Ibanez (2018), Gertler and Karadi (2015), Hanson and Stein (2015) and Nakamura and Steinsson (2018). Christensen, Lopez, and Rudebusch (2015) employ a term structure model with stochastic volatility to back out deflation protection embedded in TIPS. They show that the model accurately reflected the deflationary concerns prior to (and throughout) the financial crisis. The option value is shown to closely follow overall market uncertainty measures (e.g. VIX). Grishchenko, Vanden, and Zhang (2016) show that the information content contained in TIPS concerning future inflation remains statistically significant even when explanatory variables include lagged inflation, gold, crude oil, the VIX, liquidity, forecasting surveys, and the yield spread between nominal Treasuries and TIPS. Fleckenstein, Longstaff, and Lustig (2017) examine the relationship between deflation risk and financial and macroeconomic tail risks found in inflation swaps and options. They find that deflation risk varies with the horizon; short-term deflation risk correlates strongly with measures of risk in the financial markets such as Libor spreads, swap spreads, stock returns, and stock market volatility; intermediate-term deflation risk correlates with structural factors such as the risk of sovereign defaults in the Eurozone; and long-term deflation risk is driven primarily by macroeconomic factors. Gimeno and Ibanez (2018) focus on how risk-neutral densities, backed out from the forward 5-on-5 year inflation rate, respond to ECB's decisions and communication since 2009. Their main finding is that these distributions have significant time-variation. Like these papers, we show the informational content embedded in the inflation-derivatives market is substantial. We also demonstrate a tight connection between our dispersion measure and financial, macroeconomic, and policy measures of uncertainty. Unlike these papers, we show [i] how inflation caps and floors respond *differently* to uncertainty measures and macroeconomic factors; [ii] we provide a narrative view of regime change in dispersion that is consistent with a compelling economic narrative; [iii.] we focus on high-frequency (daily) identification of market expectations that can provide policy makers with real-time analysis. Finally, our paper contributes to the movement toward using high-frequency identification methods in monetary policy. Gertler and Karadi (2015), Hanson and Stein (2015) and Nakamura and Steinsson (2018) all demonstrate how high-frequency methods can help improve identification. Our main results connect daily movements in option prices to monetary policy announcements and other macroeconomic news.

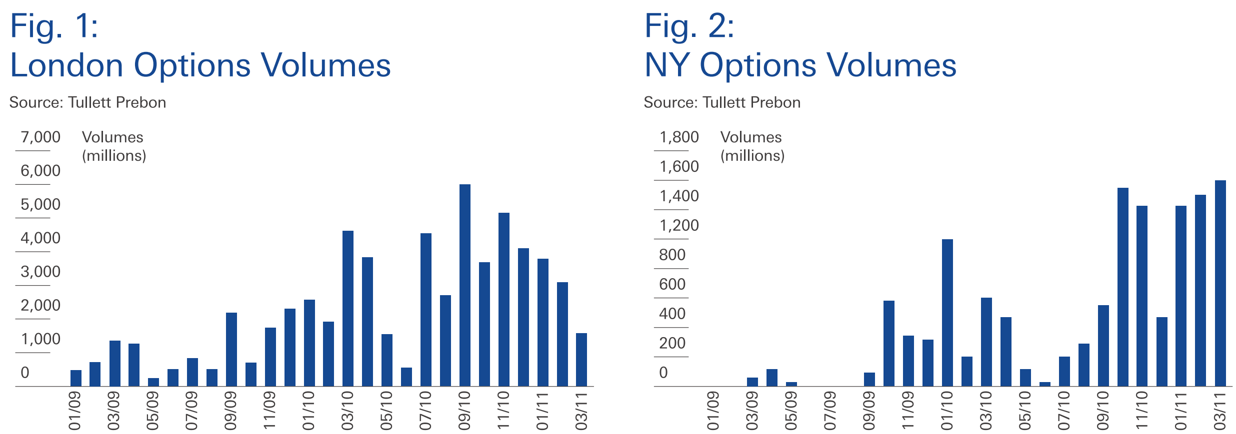

Figure 1: Volume of inflation options traded in New York and London.

2 Inflation Caps and Floors

A zero-coupon inflation cap of strike rate k and maturity h written at time t is a contract in which the seller agrees to pay the buyer the difference between actual average annualized inflation rate (headline consumer price index, non-seasonally adjusted) over the period t to $t+h$ ($\bar{\pi}_{t,t+h}$) and the strike rate in the event that this difference is positive, $\max((1+\bar{\pi}_{t,t+h})^h - (1+k)^h, 0)$. In exchange for the contract, the seller receives a payment of $V_t(k,h)$ at time t, which is a function of the strike rate and time to maturity. An inflation floor is analogous, with the payment being $\max((1+k)^h - (1+\bar{\pi}_{t,t+h})^h, 0)$.

Our analysis considers daily prices for zero-coupon caps and floors over the period starting in October 2009 until May 2017.[5] The data are Bloomberg composite prices (CMPN) which consist of averages of market quotes from various banks and brokers (e.g., Bank of America, Merrill Lynch, and BGC), with outliers removed. Figure 1 shows the substantial increase in volume over the initial years of our sample period, reaching 50bn in 2010 (up from 13bn in 2009 and 1bn in 2005). As a percentage of the overall US inflation derivative market based on interdealer volumes, inflation options grew from less than 10% of the market in 2009 to roughly 30% in 2011, exceeding the TIPS ASW market (BGC Partners). Fleming and Sporn (2013) argue that “the U.S. inflation swap market appears reasonably liquid and transparent despite the market’s over-the-counter nature and modest activity." While the option market is smaller than the swap market studied by Fleming and Sporn, many of the same participants are active in both markets. Firms that offer inflation protection typically have inflation-adjusted inflows (e.g., utilities, real estate developers, retailers). Conversely, firms and entities that buy inflation protection have inflation-linked outflows (e.g., pension funds, inflation mutual funds). Both types of firms are active traders in the options, swaps, and TIPS markets according to Kerkhof (2005). Moreover, as argued by Kitsul and Wright (2013), while the notional amount traded in these option markets is relatively small compared to, for example, the market for U.S. Treasuries, the amounts are “still big enough to presumably reflect the beliefs of traders in this market, and far bigger than those in experimental games and in prediction

[5]We do not extend the data beyond 2017 due to liquidity concerns. The period of our analysis contains several strike prices and volume was increasing, indicating a relatively liquid market. These phenomena reversed course in 2018, which explains our end-point of 2017.

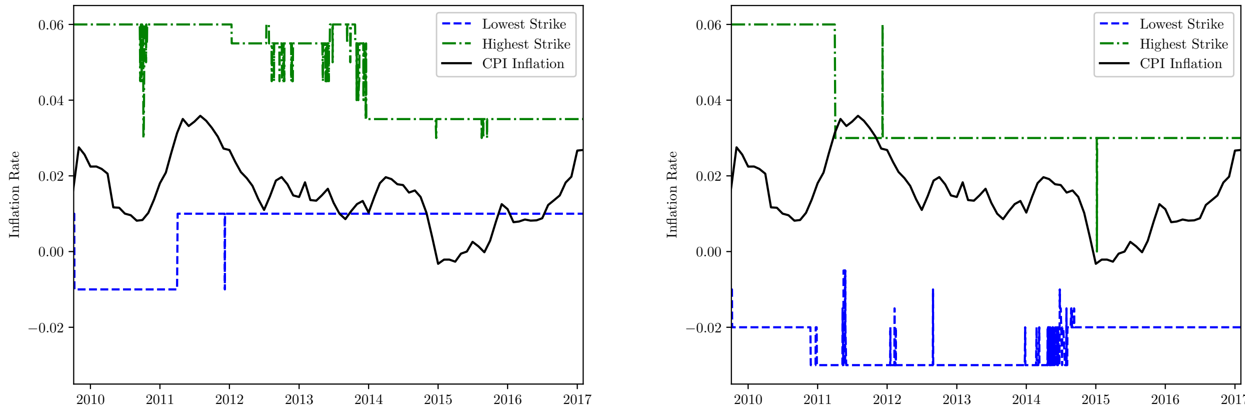

(a) Strikes for inflation caps

(b) Strikes for inflation floors

Figure 2: Maximum and minimum strike rates for traded inflation caps (a) and floors (b), along with inflation as calculated from CPI

markets" where studies have shown prices are informative.

For a given date in our sample, data will potentially contain market prices for one year caps with strikes $-1\%, -0.5\%, \ldots, 6\%$. Likewise, the possible strikes for traded one year floors could range from -3% to 5% in 0.5% increments. For nearly all dates in our sample, there are observed prices for one year caps of strikes between 1% and 3.5%, and for inflation floors between -2% and 3%. Price data for one year option strikes beyond this range are available only relatively early in the sample period, indicating an overall trend towards consolidation in the market around the middle of the initial band.

Figure 2 plots the maximum and minimum strikes which are available from late 2009 through 2016, along with the non-seasonally adjusted, year-over-year percentage change in the CPI (the base asset in the option contracts). Note that the inflation series falls below the available strike band for caps in late 2014 and stays outside this band for the following year and a half. The analogous plot for inflation floors is shown in Figure 2b. Generally the market for these options appears to be broader, with contracts being available across the range from -2% to 3% for most of the sample period. On the other hand, the availability of floor contracts has been relatively more volatile than that of caps since 2014: while not visible in the figure, only strike rates of 0% were available on multiple days near the start of the 2015 deflationary episode.

2.1 Information Content of Prices

Before turning to our more formal analysis of the data, we provide prima facie narrative evidence that the prices of inflation caps and floors are an important source of news concerning inflation expectations. The purpose of this section is twofold: first, it shows that prices of caps and floors do move in response to economic news; and second, it provides motivation for our formal analysis of Section 3.1 by demonstrating that prices with strikes far away from current inflation rates can move for reasons that have no economic interpretation. By repeating this exercise with a formal model, we are able to highlight the added identification that comes with a more formal approach. (Foreshadowing results, Table 3 contains several more entries relative to 1.)

Date	Cap	Event
Feb. 21, 2012	Down	Euro zone finance ministers agree to bail out Greece.
Apr. 4-9, 2012	Up	Draghi speech on downside risks; Fed forecasts less QE.
Sept. 13, 2012	Down	FOMC, increase to QE.
Jan. 31, 2013	Down	FOMC statement, Q4 contraction expected to be short lived.
Feb. 11, 2014	Up	Oil price crash; Yellen says negative rates are not off table.
Apr. 15-16, 2014	Up	Fed Beige Book release and Yellen's first speech.
May 26, 2015	Down	Yellen: appropriate to raise interest rates within the year.
July 5, 2015	Down	Day after Greek referendum on bailout.
Date	**Floor**	**Event**
Jan. 23, 2012	Down	Greece negotiations. Fed forward guidance.
Aug. 10, 2012	Down	Strong US labor market data.
Nov. 14, 2012	Up	Weak US retail data, FOMC minutes published.
Apr. 16, 2015	Down	Rise in oil prices, ECB rate decision.
March 2-4, 2016	Down	US labor market and manufacturing data, increased oil prices.

Table 1: Economic events that coincident with z-scores greater than one for 1% cap option maturing in 1 year (top) and 1% floor options maturing in 1 year (bottom). There are four unexplained changes in z-score that have no synchronous economic event.

Specifically, we calculate z-scores for prices of 1% caps and floors with a duration of one year, and compare significant movements (> 1.5 standard deviation) to news events from 2012 to 2017. Table 1 summarizes our findings. A search was conducted to identify any major economic news or policy developments which coincided with significant z-scores (> 1.5) in the series. Table 1 displays the z-scores that coincide with economic events. Movement in caps can be associated with a statement by the Board of Governors of the Federal Reserve or its Chair, while movements with respect to floors are typically preceded by changes in structural economic performance (e.g., labor markets, oil prices). There are only four changes in z-score that do *not* coincide with an economic announcement or event. Extending the prices beyond the 1% cap / floor to include strikes greater than or equal to 3% picks up 3/4 of the events of Table 1 but eight changes in z-score that do not appear to be consistent with economic events / announcements. If we extend out to 4% strikes, the number of changes that are unrelated to economic events is even higher at 15.[6] We view this narrative exercise as evidence in favor of our hypothesis that caps and floors respond to economic news with caps reacting more to monetary policy announcements, assuming a 1% strike. However, we need a more formal connection to the underlying inflation process, which is developed in the next section.

[6]Of course, it is impossible to rule out all events that could cause prices to fluctuate. In order to help determine "non-events" we looked at the correlation with spikes in the VIX.

2.2 Canonical Valuation We now describe a nonparametric method for extracting implied probability distributions for inflation from option strike-price curves based on the work of Buchen and Kelly (1996) and Stutzer (1996).[7] Consider the empirical distribution $(\boldsymbol{\pi}, \boldsymbol{p}) = \{\pi_i, p_i\}_{i=1}^N$ of h year annualized inflation $\pi_{t,t+h}$. That is, $\boldsymbol{\pi}$ contains observations of h year annualized inflation over the sample period while $p_i = 1/N$ for every i. Buchen and Kelly (1996) and Stutzer (1996) provide a simple method for estimating the risk-neutral distribution $(\boldsymbol{\pi}, \boldsymbol{p}^*)$ from $(\boldsymbol{\pi}, \boldsymbol{p})$. Specifically, this method minimizes the Kullback-Leibler Information Criterion (KLIC) metric[8]

$$I(\boldsymbol{p}, \boldsymbol{p}^*) = \sum_{i=1}^N p_i^* \log(p_i^*/p_i) \tag{1}$$

subject to the constraints

$$\begin{aligned} p_i^* &\geq 0, \quad i = 1, \ldots, N \\ 1 &= p_1^* + \ldots + p_N^* \\ V_t &= B_{t,t+h} \sum_{i=1}^N p_i^* V_{i,t+h} \end{aligned}$$

where V_t denotes the time t value of an option expiring in h years, $V_{i,t+h}$ denotes its expiration value if annualized inflation over the option's lifetime is π_i, and $B_{t,t+h}$ denotes the price of an h year bond at time t. Hence the third constraint stipulates that the probabilities $\boldsymbol{p}^*$ should correctly price a particular option on a particular date. This constraint can be tightened to include multiple options.

When applied with a uniform distribution across N observed outcomes, as occurs for an empirical distribution, the KLIC minimization is equivalent to maximizing the Shannon entropy of the estimated forward distribution, given by

$$-\sum_{i=1}^N p_i^* \log(p_i^*) \tag{2}$$

The solution of the Shannon entropy constrained maximization problem is known to give a multivariate canonical distribution

$$p_i^* = \frac{\exp\left(\sum_{m=1}^M \lambda_m^* V_{i,t+h}^{(m)} B_{t,t+h}\right)}{\sum_{j=1}^N \exp\left(\sum_{m=1}^M \lambda_m^* V_{j,t+h}^{(m)} B_{t,t+h}\right)} \tag{3}$$

in the case of M option pricing constraints where the corresponding options have time $t+h$ payoff $V_{i,t+h}^{(m)}$ for each prevailing inflation state $i = 1, \ldots, N$. Here the Lagrange multipliers $\boldsymbol{\lambda}^* = (\lambda_1^*, \ldots, \lambda_M^*)$ are given as

[7] The appendix contains an analysis of our data using the derivative method of Aït-Sahalia and Duarte (2003) and Kitsul and Wright (2013).

[8] Information theoretical approaches to statistics and econometrics have a long history (see Kullback (1968) and Judge and Mittlehammer (2012)). Generalizations of the method of Stutzer (1996) and their interpretation appear in Haley and Walker (2010) and Haley, Mcgee, and Walker (2013).

solutions to the unconstrained minimization problem

$$\boldsymbol{\lambda}^* = \arg\min_{\boldsymbol{\lambda}} \sum_{i=1}^{N} \exp\left(\sum_{m=1}^{M} \lambda_m \left[V_{i,t+h}^{(m)} B_{t,t+h} - V_t^{(m)}\right]\right) \tag{4}$$

Given h year treasury yields and prices for inflation derivatives with maturities of h years on a given date, we can therefore numerically solve (4). Once the Lagrange multipliers are found, the estimated forward probabilities are obtained by substitution into (3). The minimization in (4) is tractable for a low number M of options, but will face the curse of dimensionality as the number of constraints is increased.

3 Uncertainty in Inflation Options Markets

Applying the method described in the previous section to floor and cap options, we turn now to examining the dynamics and drivers of uncertainty in the forward distributions for inflation, as measured by the difference between the 25th and 75th percentiles. Our interest in this particular metric for uncertainty follows from the observation that model implied pricing kernels estimated using data from floors and caps display substantial time variation in tail behavior. In particular, using the approach of Rosenberg and Engle (2002), Kitsul and Wright (2013) specify the pricing kernel as a nonlinear function of average annual inflation over the next h years, $\pi(h)$

$$M_t(\pi(h)) = \theta_{0t} T_0(\pi(h)) + \theta_{1t} T_1(\pi(h)) + \theta_{2t} T_2(\pi(h)) + \theta_{3t} T_3(\pi(h))$$

where $T_j(\cdot)$ are Chebyshev polynomials defined over the range of annual inflation (-2% to 6%) and the vector of parameters θ_t are estimated by minimizing the distance between the actual price of the inflation floor or cap and the model-implied price. Figure 3 is taken from Kitsul and Wright (2013) and plots the empirical pricing kernels for various maturities and years. The distributions are all centered around 1.5% to 2% but the behavior in the tails of the distributions [< 0%; > 2.5%] show substantial variation with time.

While our results apply to the subjective forward distributions rather than to the actual physical likelihood of (dis)inflationary events, simulation evidence suggests that under some circumstances uncertainty may correlate well across both. For example consider pricing a European call option with expiration date T and strike price X. The price C discounted at the risk-free rate of interest r is given by

$$C = \mathbb{E}_t^Q \left\{ \frac{\max[S_T - X, 0]}{(1+r)^T} \right\}, \tag{5}$$

where S_T is the price of the underlying asset at date T, and $\mathbb{E}_t^Q$ implies that the expectation is taken with respect to the risk-neutral (equivalent-martingale) measure. Suppose further that the stock price follows a normal mixture, $f(S) = \sum_{i=1}^{k} [\lambda_i f_i(s)]$, where f_i denotes the normal distribution with moments μ_i and σ_i, and $\sum_i \lambda_i = 1$. It is well known [Cox and Ross (1976)] that the one-period-ahead option price follows a mixture of Black-Scholes formulas $C(S_t, t) = \sum_i \lambda_i C_i(S_{i,t}, t)$. If the option expires several periods ahead, the weights are drawn according to the multinomial distribution. We then conduct the following steps:

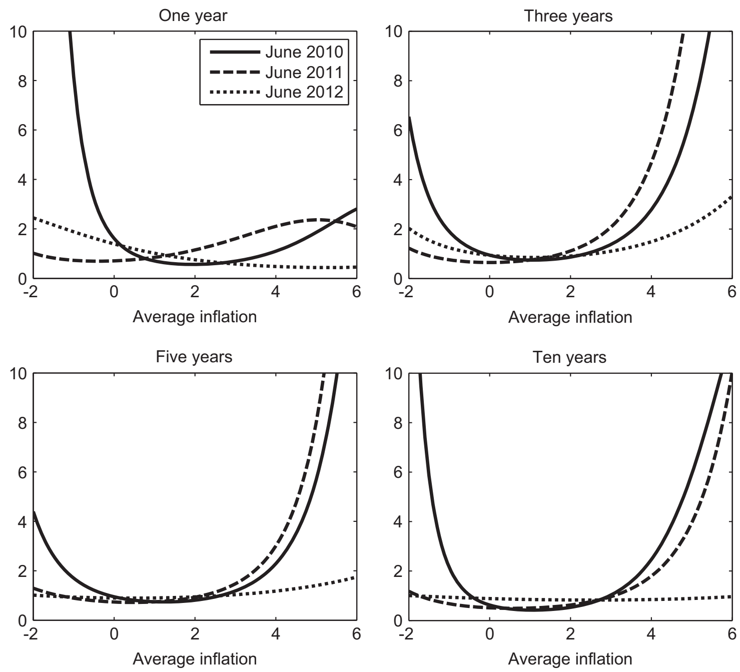

Figure 3: Empirical pricing kernels of Kitsul and Wright (2013)

1. Simulate data from the model over a period equal to roughly one year, using a high (0.25) and low (0.1) variance calibration with the means also equal to a high (0.3) and low (0.025) setting. The calibration roughly matches the CPI process over our sample period.

2. Calculate the empirical distribution of returns, $R_{t,h} = (S_{t+h}/S_t)$, at horizon h.

3. Using the Canonical Valuation methodology of Section 2.2, compute the forward probabilities, p_t^*, that correctly price the European call option, (5).

4. Compute the 25th-75th dispersion in both the physical and forward densities, and calculate the statistical discrepancy between the two measures.

Note that our calibration assumes the discrepancy in our first moments exceeds that of the second moment. Therefore time-variation in the mean exceeds that of the variance in our simulations. If we assume 250 observations, the correlation between the mean of the physical and forward distribution is 0.67 while the correlation in the 25th-75th dispersion between the two distributions is 0.91. If we increase the number of observations to 1000, the correlations become 0.63 and 0.94, respectively. A caveat to our results is that the calibration, which follows from the underlying inflation data, assumes much of the variation is due to changes in the mean. If we reverse this assumption and impose more variation in the volatility, the correlations in the dispersion metric fall below 0.9.

Below we will examine dispersion in the estimated forward distributions for individual one year cap and floor markets separately. That is we set $M = 1$ in the notation of the previous section, and compute the resulting dispersion metric when the option is a one year floor with strike rate near the middle of the historically available range, repeating the exercise with a one year cap. To understand the implications of this selection for the resulting dispersion dynamics, consider the case in which the empirical distribution of inflation outcomes contains only two observations ($N = 2$), π_1 and π_2. More specifically, suppose that the option pricing constraint is that for a one period floor of strike rate k, $\pi_1 < k < \pi_2$, so that the option is out of the money in outcome 2, $V_{2,t+1} = 0$. Under these assumptions, the minimization in (4) admits an analytic solution

$$\lambda^* = \frac{1}{V_{1,t+1}B_{t,t+1}} \log\left(\frac{V_t}{V_{1,t+1}B_{t,t+1} - V_t}\right)$$

with corresponding forward probabilities

$$p_1^* = \frac{\frac{V_t}{V_{1,t+1}B_{t,t+1}-V_t}}{1 + \frac{V_t}{V_{1,t+1}B_{t,t+1}-V_t}} = \frac{V_t}{V_{1,t+1}B_{t,t+1}}, \qquad p_2^* = 1 - p_1^*.$$

While this simple discrete distribution does not allow a calculation of dispersion as the difference between fixed percentiles, a model of the cummulative distribution function which is continuous and piecewise linear while maintaining these values does. In particular, if we let $\pi_{\mathrm{lo}} = \pi_1 - p_1^* \left(\frac{\pi_2-\pi_1}{p_2^*}\right)$ and estimate the forward CDF as

$$F(\pi) = \begin{cases} 0 & \pi < \pi_{\mathrm{lo}} \\ p_1^* + (\pi - \pi_1)\left(\frac{p_2^*}{\pi_2-\pi_1}\right) & \pi_{\mathrm{lo}} \le \pi < \pi_2 \\ 1 & \pi \ge \pi_2 \end{cases}$$

then dispersion in the forward distribution is given by

$$d_t^* = \frac{1}{2}\left(\frac{\pi_2 - \pi_1}{1 - \frac{V_t}{V_{1,t+1}B_{t,t+1}}}\right) \tag{6}$$

Hence dispersion under these assumptions is a function of $V_t/B_{t,t+1}$, the prevailing market price of the floor empirically measured at time t compounded one period forward. As the price of the floor increases, so does dispersion, and vice versa. In other words, the value of insurance increases if and only if forward uncertainty increases in this simplified framework. While weakening the assumption on N prevents the derivation of a closed form expression, it is clear from the structure of the canonical method optimization that dispersion will remain a function of the compounded option price. Moreover, weakening the assumption on M results in a dispersion metric which fluctuates with each of the included option prices, once again compounded forward.[9]

As a consequence of these considerations, we observe a number advantages to considering the forward distributions associated with individual option prices. In the event that different information drives the various markets, as could be the case if markets are segmented[10] or participants condition

[9] Closed form expressions are once again available for multiple options with $N = 2$.

[10] One potential explanation for segmented markets for inflation floors or caps is the possibility that differing portfolio types need to be hedged against differing risks. For example, a fixed income portfolio manager is likely to be more concerned about

on market-dependent information sets, including multiple options could obscure information about what is drives the individual markets. Based on the evidence in Section 2, to be strengthened below, we believe it is likely that inflation cap and floor markets react to different sources of uncertainty. Moreover, illiquidity in markets for very high or low strike rates potentially introduces a further source of noise, motivating the omission of these prices.

3.1 Regime-Switching Model

We consider a regime-switching model that disciplines the hypothesis of asymmetry between the canonical floor and cap measures. In particular, we specify a model for dispersion of the form,

$$\mathrm{Disp}_t = \mu_{s_t} + \epsilon_{s_t}$$

in which the mean μ of the model stochastic process depends on the regime $s_t \in \{0,1\}$. Likewise, the innovation variances are regime specific, so that $\epsilon_{s_t} \sim N(0, \sigma_{s_t})$. If in state 0, the process remains in state 0 with probability p_{00}, so that the expected duration of the state is $1/(1-p_{00})$ and similarly for other states. With probability p_{01}, the state transitions from 0 to 1.

We estimate four iterations of the model using the Canonical measure of dispersion: 1. one-year inflation cap with two states; 2. one-year inflation cap with three states; 3. one-year inflation floor with two states; 4. one-year inflation floor with three states. We focus on floors and caps separately in order to determine if dispersion behaves differently for changes in expected inflation vis-a-vis expected disinflation. The model is estimated using using quasi-maximum likelihood estimation [White (1982), Cho and White (2007)]. Algorithms used to determine the predictive, filtered, and smoothed probabilities permit a “quasi" likelihood estimation that maximizes the log-likelihood of the weighted average of Gaussian distributions. Given an initial value for the state, the quasi-likelihood function is given by

$$\mathscr{L}_T(\Theta) = \frac{1}{T}\sum_{t=1}^{T} \ln(f(x_t|x^{t-1};\Theta))$$
$$f(x_t|x^{t-1};\Theta) = pr(s_t = 0|x^{t-1};\Theta)f(x_t|s_t = 0, x^{t-1};\Theta) + pr(s_t = 1|x^{t-1};\Theta)f(x_t|s_t = 1, x^{t-1};\Theta)$$

where $\Theta = \{\mu_s, \sigma_s, p_{00}, p_{11}\}$ and $pr(s_t = j|x^{t-1};\Theta)$ is the predictive probability of being in state j conditional on information (dispersion) available through $t-1$, x^{t-1}. The estimated filtered, smoothed and predictive probabilities are obtained in the usual recursive fashion [see, Hamilton (1994)].

The estimated parameters are given in Table 2. With two states, the estimated means are 1.24% and 1.44% respectively for caps, and 1.64% and 3.01% for floors. Note that mean dispersion is twice as high when using floor data. Moving to a three-state model does not change the estimated means substantially for the caps with a range of 1.16% to 1.33%. However, the estimated means for the floor are quite different in the three-state regime model. The mean in the high state is nearly three times as large as the low state mean (4.44% vs. 1.63%). Volatility is also substantially different between floors and caps. The low

high inflation eroding returns, whereas a risky equity portfolio manager may be more concerned with macroeconomic developments correlated with disinflation.

	Cap-2		Floor-2		Cap-3			Floor-3		
	State 0	State 1	State 0	State 1	State 0	State 1	State 2	State 0	State 1	State 2
μ	0.0124	0.0144	0.0164	0.0301	0.0116	0.0133	0.0144	0.0163	0.0265	0.0441
σ	1.91(-3)	2.32(-4)	9.02(-4)	8.56(-3)	2.45(-3)	5.34(-4)	1.99(-4)	8.58(-4)	5.86(-3)	1.19(-3)
p_{0j}	0.978	0.022	0.994	5.96(-3)	0.985	6.95(-3)	7.75(-3)	0.994	5.56(-3)	0.00
p_{1j}	0.028	0.972	1.16(-2)	0.988	4.31(-3)	0.963	3.26(-2)	1.27(-3)	0.976	1.08(-2)
p_{2j}	NA	NA	NA	NA	4.43(-3)	2.32(-2)	0.972	0.00	5.43(-2)	0.946

Table 2: Parameter estimates for a two and three state regime model for dispersion in 1 year inflation expectations for floors and caps. The numbers in parenthesis represent a shorthand for scientific notation; that is, $(-3) \equiv 10^{-3}$.

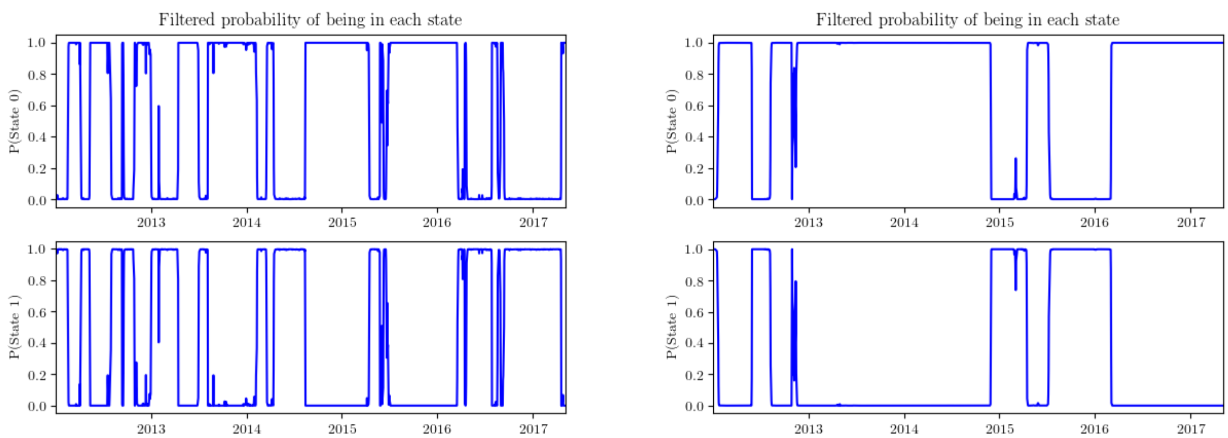

(a) Caps-2

(b) Floors-2

Figure 4: Filtered probabilities in two state regime switching models for dispersion in one year ahead inflation expectations, as implied by the canonical distributions which correctly price a cap or a floor with strike 1% at each date.

dispersion state is an order of magnitude more volatile than the high dispersion state(s) for caps, while the opposite is true for floors. For the two-state regime model, the high dispersion state is an order of magnitude more volatile for floors, and five times more volatile in the three-state model.

Regimes are persistent. Duration in each regime is slightly higher for floors relative to caps, with the low-mean state being the most persistent. The 3-state floor estimates show that regime progresses sequentially from low (State 0) to medium (State 1), and from high (State 2) to medium (State 1), but never from high to low or vice versa. For caps, transitions primarily occur between the high and medium dispersion states, while the low dispersion state decays to either of the other states with approximately equal probability. Figure 4 plots the two-state filtered probabilities over the sample period for caps and floors. The additional persistence of the floor regime is evident. For all of 2013 and most of 2014, the dispersion measure due to floors stayed in the low-mean regime, while the dispersion of caps changed regime several times. Visually, the transitions occur very quickly, while the states persist for a relatively long time with a few exceptions, consistent with the estimated mean durations.

Date	Cap	Event
Feb. 21, 2012	Down	Euro zone finance ministers agree to bail out Greece.
Apr. 4-9, 2012	Up	Draghi speech on downside risks; Fed forecasts less QE.
May 14, 2012	Down	+ Fed stress test results; Greek cabinet approves bailout.
July 26 - Aug. 3, 2012	Up	Fed indicates slowing economy, potential future policy action.
Sept. 13, 2012 Spike	Down	FOMC, increase to QE.
Oct. 29, 2012	Down	Hurricane Sandy (7th costliest disaster worldwide).
Jan. 2, 2013	Up	US lawmakers avert fiscal cliff; Debt-ceiling crisis starts.
Jan. 31, 2013 Spike	Down	FOMC statement, Q4 contraction expected to be short lived.
Apr 16, 2013	Down	No clear event.
Feb. 11, 2014	Up	Oil price crash; Yellen says negative rates are not off table.
Mar. 19-21, 2014	Down	Yellen says stimulus over by Fall, rate hike in early 2015.
Apr. 15-16, 2014	Up	Fed Beige Book release and Yellen's first speech.
Aug. 14-15, 2014	Down	No clear event.
Apr. 15, 2015	Up	Fed Beige Book.
May 26, 2015	Down	Yellen: appropriate to raise interest rates within the year.
July 5, 2015	Down	Day after Greek referendum on bailout.
Mar. 16-21, 2016	Up	Fed monetary policy statement.
Apr. 19, 2017	Down	Fed Beige book suggests rate increases.
Date	**Floor**	**Event**
Jan. 23, 2012	Down	Greece negotiations. Fed forward guidance.
May 31, 2012	Up	Weak US labor market and manufacturing data.
Aug. 10, 2012	Down	Strong US labor market data.
Oct. 31, 2012	Up	Hurricane Sandy (7th costliest disaster worldwide).
Nov. 14, 2012 Spike	Up	Weak US retail data, FOMC minutes published.
Dec. 1, 2014	Up	Fall in oil prices.
Apr. 16, 2015	Down	Rise in oil prices, ECB rate decision.
July 7-14, 2015	Up	Weak US labor market data, oil price volatility.
March 2-4, 2016	Down	US labor market and manufacturing data, increased oil prices.

Table 3: Economic events that coincident with breaks in the regime switching models for dispersion based on pricing of a 1% cap option maturing in 1 year (top) and 1% floor options maturing in 1 year (bottom). Notes: April 16, 2013 was the Tuesday after the Boston bomber. Equities dropped on Monday but recovered on Tuesday. There is an unexplained break to high dispersion in early July 2013 and returning to low dispersion in early August. In June of 2015, there were a few short lived switches, possibly following developments in negotiations between Greece and the EU. There are several additional breaks in July-Sept. 2016, which all coincide with Fed publication release dates.

3.2 A Narrative View of Regime Change

We next conduct a narrative study of events occurring at the same time as transitions in the two state model for both caps and floors. We focus on the two state model both for simplicity and because it is selected over the three period model by the Akaike information criterion. Specifically, a break in the series is identified as a period over which the filtered probabilities of the high dispersion state transition from values above 0.99 to below 0.01 or vice versa. A typical break in the model occurs over a period of one to 6 business days, and with a few exceptions states persist for a month or more. The clear breaks and long lived states gives us a fair degree of confidence that we are identifying events which do indeed lead to transitions, in response to the critique of Fair (2002).

A search was conducted to identify any major economic news or policy developments which coincided with a break in the series. Events of interest primarily occurred the day of or prior to either the start of the break, as well as the day of or prior to any particularly large jump in the filtered odds. Table 3 displays the results of the study. The consistent timing of identified events relative to probability changes gives us further reassurance in our interpretations.

As in Section 2.1, nearly every transition in the model relating to caps can be associated with a statement by the Board of Governors of the Federal Reserve or its Chair, while transitions with respect to floors are typically preceded by changes in structural economic performance (e.g., labor markets, oil prices). There are only a few exceptions to this observation, in April 2013, August 2014, and October 2012, as well as a high dispersion event which lasted through July of 2013. The former break happened shortly following the bombing of the Boston marathon on Sunday, April 14, 2013. Markets reacted by turning sharply downward the day following, and then recovering on the day of the break. The August 2014 event, on the other hand, appears to possibly be misidentified. First, in Figure 4 the 2014 break corresponds to a series drop from ≈ 0.014 to ≈ 0.013, which was followed by a further, more drastic fall in mid September. The three state model identifies the first drop as a transition from High to Middle dispersion, and finds a break from Middle to Low dispersion from September 17 to 19. The largest jump in the latter transition occurred on September 18, accompanied by a decision by the FOMC to not raise interest rates. Finally, the July 2013 event does not have a clear explanation, although there is some indication that investor fears about the rollback of quantitative easing emerged at this time.

Looking more closely at the events associated with monetary policy, dispersion in the forward distributions generally decreases when policy is relatively tight and increases when it is loose, as is sensible for a period which saw persistently low inflation and stimulative policy. The series transitions several times in the wake of statements by former Federal Reserve Chairwoman Janet Yellen, particularly in the weeks following her swearing in on February 3, 2014. Late in the sample period, switches are more associated with official Fed documentation, in particular the release of beige book summaries of economic conditions. Not listed in the table are breaks on July 28, August 18, August 30, and September 8-13 of 2016; each of these occurred in line with Fed releases and statements.

Early in the series history, breaks appear to be associated with actions by European leaders during the debt crisis. This feature reemerges in mid-2015, the period during which Greece considered leaving the European Union. Additional events associated with a transition include the oil price crash of 2014 and the day on which Hurricane Sandy made landfall in New York. The latter is notable as the second costliest disaster in the United States at the time, following Hurricane Katrina,[11] and the seventh costliest worldwide.

[11] It has since been surpassed by Hurricane Harvey in 2017.

3.3 A Structural Interpretation We now offer a structural interpretation of our results following the model, calibration and solution strategy of Gavin, Keen, Richter, and Throckmorton (2015) and Plante, Richter, and Throckmorton (2016). The model is (now) a well-known New Keynesian framework used to study zero-lower bound constraints.

Households. Households choose consumption, labor and bond holdings to maximize expected lifetime utility, $E_0 \sum_{t=0}^{\infty} \tilde{\beta}_t [\log c_t - \xi n_t^{1+\eta}/(1+\eta)]$ where $1/\eta$ is the Frisch elasticity of labor supply, c_t is consumption, n_t is labor hours, b_t is the real value of a 1-period nominal bond, E_0 is the expectation operator and $\tilde{\beta}_t = \prod_{j=1}^{t} \beta_j$ for $t > 0$. β is a time-varying subjective discount factor that satisfies, $\beta_t = \tilde{\beta}(\beta_{t-1}/\beta)^{\rho_\beta} \exp(\varepsilon_t)$. The household's budget constraint is $c_t + b_t = w_t n_t + r_{t-1} b_{t-1}/\pi_t + d_t$, where w_t is the real wage, $\pi_t = p_t/p_{t-1}$ is the inflation rate, d_t are profits of the intermediate firms, and r_t is the nominal interest rate set by the central bank.

Firms. The monopolistically competitive intermediate goods firms produce a continuum of differentiated inputs and a representative firm produces a final good. Each intermediate firm $f \in [0,1]$ produces a differentiated good $y_t(f)$ with identical technologies, $y_t(f) = z_t n_t(f)$, where $n_t(f)$ is the level of employment used by firm f and z_t represents the level of technology that is common across all firms and follows, $z_t = z(z_{t-1}/z)^{\rho}_{z} \exp(v_t)$. Each intermediate firm chooses labor supply to minimize operating costs, $w_t n_t(f)$ subject to its production technology. The final good firm purchases $y_t(f)$ units from each intermediate firm to produce the final good according to a Dixit-Stiglitz aggregator, $y_t = [\int_0^1 y_t(f)^{(\theta-1)/\theta} df]^{\theta/(\theta-1)}$, where $\theta > 1$ is the elasticity of substitution between intermediate goods. The firm's optimality condition yields the demand function for intermediate inputs given by $y_t(f) = (p_t(f)/p_t)^{-\theta} y_t$. Following Rotemberg, each firm faces a cost to adjusting its price, $a_t(f)$, using the functional form, $a_t(f) = \psi[p_t(f)/(\pi p_{t-1}(f) - 1]^2 y_t/2$. Firm f chooses its price, $p_t(f)$, to maximize the expected discounted present value of real profits.

Monetary Policy. Each period, the central bank sets the gross nominal interest rate according to $r_t = \max\{1, r^\star(\pi_t/\pi)^{\theta_\pi}(y_t/y)^{\theta_y}\}$, where π is the inflation target and θ_π and θ_y represent the extent of the policy response to inflation and the output gap. Thus, the central bank is subject to a zero-lower bound constraint.

Solution and Calibration. The model is solved nonlinearly following a policy function iteration approach of Richter, Throckmorton, and Walker (2014). See Gavin, Keen, Richter, and Throckmorton (2015) and Plante, Richter, and Throckmorton (2016) for a complete description of the model, solution method and quarterly calibration.

Figure 5 reports the generalized impulse response function to a technology shock when the ZLB binds (dashed) and when it does not (solid).[12] The figure shows that there are qualitative differences between the two scenarios with respect to the real interest rate and real wage rate. A technology shock lowers the rate of inflation and—away from the ZLB—this prompts a decline in the nominal interest rate through the policy rule. However at the ZLB, the nominal rate cannot go lower and the Fisher equation adjusts through lower expected inflation, the real interest rate must increase to keep the nominal rate

[12]The ZLB binding case is achieved by initializing the economy at the average state vector conditional on the ZLB binding in a 500,000 quarter simulation. The figure replicates Figure 8 of Gavin, Keen, Richter, and Throckmorton (2015).

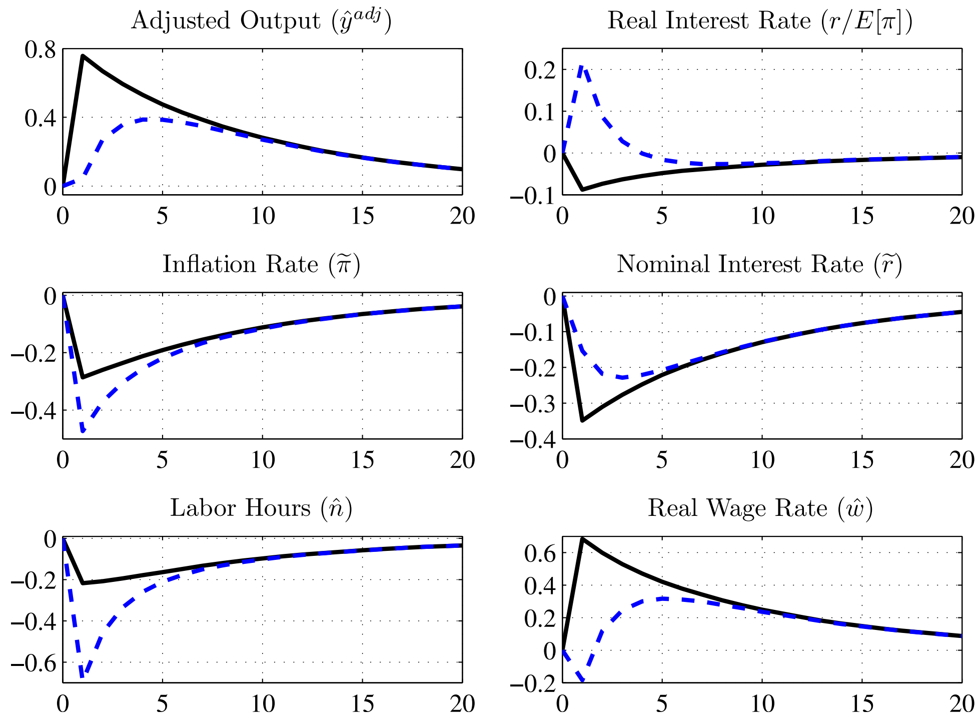

Figure 5: Generalized Impulse Response Functions to Technology Shock: ZLB vs Non-ZLB

constant. This increase in the real rate has a negative impact on labor hours and the real wage. Therefore, weakness in the labor market is correlated with a decline in the inflation rate. Consistent with this intuition, Table 1 and Table 3 show that the price and dispersion of an inflation floor would increase (decrease) with weakness (strength) in the labor market. Oil price movements also triggered substantial changes in prices of floors over this time period. While we do not model the oil market, one could think of it as an input good similar to labor. As the real wage increases under the ZLB, the inflation rate also increases. Thus a positive oil price shock would cause an expected increase in inflation and the price of floors would decline, consistent with Tables 1 and 3.

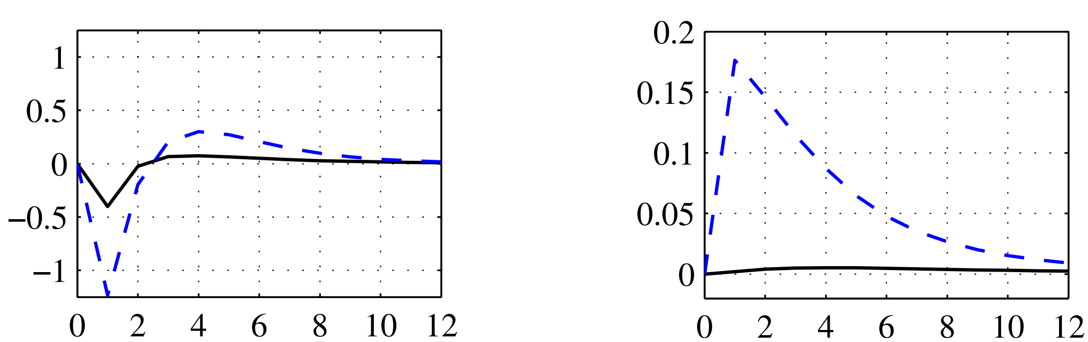

(a) Output

(b) Uncertainty

Figure 6: Generalized Impulse response function to Discount Factor Shock: ZLB(dashed) vs Non-ZLB

Figure 6 reports the impulse response to a discount factor shock at the ZLB (dashed) and away from the ZLB (solid). Panel (a) shows that output falls substantially more as the monetary authority is constrained. Panel (b) shows that uncertainty, defined as time-variation in second moments ($\sqrt{E}_t(\sigma^2_{t+1})$). Plante, Richter, and Throckmorton (2016) refer to this as endogenous uncertainty and show that a constrained monetary authority is unable to effectively respond to future shocks, which leads to an increase in household uncertainty, shown in Panel (b). Plante, Richter, and Throckmorton (2016) also show how this uncertainty extends to inflation. To summarize, the model shows how the ZLB leads to an increase in uncertainty in inflation and how the downside risks of inflation are more likely. These two forces combine to make the price of floors more volatile relative to caps at any given strike price.

3.4 Connection to Measures of Uncertainty

Recently, several measures of uncertainty have been shown to be important drivers of business and financial cycles [Jurado, Ludvigson, and Ng (2015), Baker, Bloom, and Davis (2016)]. In this section, we examine the extent to which three such measures—financial, macro and policy—are able to explain the variation in dispersion of inflation expectations.

Before turning to our regression analysis, we first document correlation between our dispersion measure and other measures of uncertainty. We focus on three such measures—VIX, dispersion in the Survey of Professional Forecasters inflation forecast, and the investor sentiment of Baker and Wurgler (2006). Comparison with the VIX is useful because it represents a measure of uncertainty available at the same frequency as our measure—daily. Our dispersion measure positively correlates with the VIX over our sample, with a correlation coefficient of 0.586, suggesting that our metric is picking up many of the same events that move the stock market. While a useful measure of uncertainty, the VIX cannot distinguish between shocks that increase/decrease expected inflation. In order to compare our dispersion measure to the SPF and the investor sentiment of Baker and Wurgler (2006), we take a monthly average. At this frequency, dispersion is positively correlated with SPF dispersion and weakly correlated with investor sentiment at 0.49 and 0.19, respectively. Despite the lack of a strong correlation, we believe our measure has value added relative to these other metrics due to its higher frequency and ability to identify directional changes in inflation expectations.

As described in Ludvigson and Ng (2019), the financial uncertainty measure is constructed from 148 monthly financial series that consists of a number of indicators "measuring the behavior of a broad cross-section of asset returns, as well as some aggregate financial indicators. These data include valuation ratios such as the dividend-price ratio and earnings-price ratio, growth rates of aggregate dividends and prices, default and term spreads, yields on corporate bonds of different ratings grades, yields on Treasuries and yield spreads, and a broad cross-section of industry equity returns." The "Macro" uncertainty measure is compiled from a database of 135 monthly U.S. indicators, taken from FRED-MD and described in McCracken and Ng (2014). These data included industrial production, weekly hours, personal inventories, monetary aggregates, interest rates and interest-rate spreads, stock prices, and consumer expectations. In both cases, the uncertainty measure taken from Jurado, Ludvigson, and Ng (2015) is the h-period ahead uncertainty in the variable y_{jt} defined as the conditional volatility of the purely unforecastable component of the future value of the series,

$$\mathcal{U}_{jt}(h) \equiv \sqrt{E[(y_{jt+h} - E[y_{jt+h}|I_t])^2|I_t]}$$

where I_t is the information available to the economic agents at t and is formulated according to a dynamic factor analysis. The large set of predictors described above are used to span the information set of the agent, while the volatility in the forecast error is estimated using a parametric stochastic volatility model. For further details see Jurado, Ludvigson, and Ng (2015).

We also examine the three-component policy uncertainty index of Baker, Bloom, and Davis (2016). The first component consists of monthly search results from 10 large newspapers with keywords 'uncertainty' or 'uncertain', the terms 'economic' or 'economy' and one or more of the following terms: 'congress', 'legislation', 'white house', 'regulation', 'federal reserve', or 'deficit'. The second component of the index compiles a list of temporary federal tax code provisions as reported by the Congressional Budget Office, the idea being that "temporary tax measures are a source of uncertainty for businesses and households because Congress often extends them at the last minute, undermining stability in and certainty about the tax code." The final component is a measure of dispersion in the individual level forecasts of variables directly influenced by government policy (e.g., purchases of goods and services by the federal government) contained in the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters.

Variable	**Mean**	**Std. Dev.**	**Min**	**Max**					
Canon. (Cap)	0.013	0.002	0.010	0.016			**Correlation**		
Canon. (Floor)	0.019	0.006	0.015	0.045	h=1	h=3	h=12	Macro	Policy
Fin. (h=1)	0.796	0.052	0.680	0.920	1.00	0.98	0.97	0.32	-0.01
Fin. (h=3)	0.856	0.043	0.760	0.958		1.00	0.98	0.34	0.00
Fin. (h=12)	0.950	0.016	0.912	0.988			1.00	0.40	0.04
Macro.	0.750	0.031	0.700	0.813				1.00	0.03
Policy	0.119	0.033	0.072	0.195					1.00

Table 4: Summary Statistics for Uncertainty and Dispersion Measures

Table 4 reports the descriptive statistics for the dispersion measures used in the regression analysis, along with the uncertainty measures. We scaled the policy uncertainty measure by 1/100 to place it roughly on the same scale as the dispersion statistics. The data are monthly observations from January 2012 through May 2017 ($N = 65$). For the financial uncertainty measure, we include three forecast horizons ($h = 1, 3, 12$), even though these measures are highly correlated. Our dispersion statistics are derived from one-year ahead strikes $h = 12$ and we want to test if the timing is relevant. Also noteworthy is the lack of correlation between the financial / macro uncertainty measures and the policy uncertainty. Meanwhile, the macro and financial uncertainty measures are weakly positively correlated.

Table 5 presents the results of our regression analysis. We regress the uncertainty measures on floor and cap dispersion separately to investigate if floors and caps can be explained by different measures of uncertainty as suggested in Section 3.2. First note that financial uncertainty measures are statistically significant predictors of dispersion in both caps and floors, increasing in significance and magnitude as the horizon increases. That the most significant and sizable response comes from the one-year horizon ($h = 12$) should not be a surprise given that the options in Table 5 are all evaluated at strikes of one-

Dependent Variable: Floor-Dispersion						
Fin. (h=1)	0.047***					
	(0.013)					
Fin. (h=3)		0.058***				
		(0.016)				
Fin. (h=12)			0.161***			0.102*
			(0.041)			(0.041)
Macro.				0.101***		0.079***
				(0.021)		(0.022)
Policy					0.016	
					(0.010)	
Constant	-0.018	-0.030*	-0.133***	-0.06***	0.017**	-0.136***
	(0.011)	(0.014)	(0.041)	(0.016)	(0.020)	(0.035)
R^2	0.165	0.175	0.200	0.272	0.01	0.339
Dependent Variable: Cap-Dispersion						
Fin. (h=1)	-0.011**					
	(0.003)					
Fin. (h=3)		-0.014**				
		(0.004)				
Fin. (h=12)			-0.034**			-0.024*
			(0.011)			(0.011)
Macro.				-0.018**		-0.013*
				(0.006)		(0.006)
Policy					0.007	
					(0.006)	
Constant	0.022***	0.025***	0.046***	0.027***	0.012***	0.046***
	(0.003)	(0.004)	(0.011)	(0.004)	(0.000)	(0.010)
R^2	0.154	0.150	0.135	0.137	0.02	0.194

Table 5: Regression Analysis

year. Note also that the policy uncertainty metric is not significant for either caps or floors, while the financial uncertainty remains significant and macro uncertainty is also significant. Moreover, financial uncertainty enters *negatively* for cap dispersion and positively for floor dispersion. Thus, an increase in financial uncertainty decreases the dispersion in caps. Recall that over this time period, disinflation was more of a concern than inflation which is why innovations to financial time series, on average, consolidated expectations in inflation caps while increasing dispersion in inflation floors. Finally, the R-square associated with macro uncertainty is significantly higher for the floor dispersion relative to the cap (0.272 vs 0.137), which is consistent with our narrative.

4 Conclusion

The frequency and richness of financial markets data makes it an attractive resource for market participants and policy making institutions, however linking asset price fluctuations to investor perceptions remains a difficult topic. While forward distributions for inflation implicit in cap and floor prices need not provide a direct measure about investor perceptions regarding the actual likelihood of events, we have argued that dispersion in these distributions may nonetheless provide a sense of broad uncertainty.

Having adapted the method of Buchen and Kelly (1996) and Stutzer (1996) to back out such distributions from prices, we have shown that changes in the level and volatility of dispersion can be captured by a simple regime switching model. Exploiting the daily frequency of the data and the cleanliness of the breaks in the model, we have been able to closely link these breaks to Federal Reserve policy and major global and domestic sources of uncertainty. At a monthly frequency, we have shown that our dispersion metric is also linked to nascent measures of economic and financial uncertainty.

References

AÏT-SAHALIA, Y., AND J. DUARTE (2003): "Nonparametric option pricing under shape restrictions," *Journal of Econometrics*, 116, 9–47.

ANDREASEN, M. M., J. H. E. CHRISTENSEN, AND S. RIDDELL (2018): "The TIPS Liquidity Premium," Working Paper Series 2017-11, Federal Reserve Bank of San Francisco.

BACKUS, D., M. CHERNOV, AND I. MARTIN (2011): "Disasters Implied by Equity Index Options," *The Journal of Finance*, 66(6), 1969–2012.

BAKER, M., AND J. WURGLER (2006): "Investor Sentiment and the Cross-Section of Stock Returns," *The Journal of Finance*, 61(4), 1645–1680.

BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 131(November), 1593—-1636.

BREEDEN, D. T., AND R. H. LITZENBERGER (1978): "Prices of State-contingent Claims Implicit in Option Prices," *The Journal of Business*, 51(4), 621–651.

BUCHEN, P. W., AND M. KELLY (1996): "The Maximum Entropy Distribution of an Asset Inferred from Option Prices," *Journal of Financial and Quantitative Analysis*, 31(1), 143–159.

CHO, J. S., AND H. WHITE (2007): "Testing for Regime Switching," *Econometrica*, 75(November), 1671–1720.

CHRISTENSEN, J. H. E., J. A. LOPEZ, AND G. D. RUDEBUSCH (2015): "Pricing Deflation Risk with US Treasury Yields," *Review of Finance*, 20(3), 1107–1152.

COX, J. C., AND S. ROSS (1976): "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3(1-2), 145–166.

DYKSTRA, R. L. (1983): "An Algorithm for Restricted Least Squares Regression," *Journal of the American Statistical Association*, 78(384), 837–842.

FAIR, R. C. (2002): "Events That Shook the Market," *The Journal of Business*, 75(4), 713–731.

FAN, J., AND I. GIJBELS (1996): *Local Polynomial Modelling and Its Applications*. Chapman and Hall, London.

FLECKENSTEIN, M., F. A. LONGSTAFF, AND H. LUSTIG (2014): "The TIPS-Treasury Bond Puzzle," *Journal of Finance*, 69(5), 2151–2197.

——— (2017): "Deflation Risk," *The Review of Financial Studies*, 30(8), 2719–2760.

FLEMING, M. J., AND N. KRISHNAN (2004): "The Microstructure of the TIPS Market," *Federal Reserve Bank of New York Economic Policy Review*, 18(1), 27–45.

FLEMING, M. J., AND J. SPORN (2013): "Trading Activity and Price Transparency in the Inflation Swap Market," *Economic Policy Review*, 19(1), 1–13.

GAVIN, W., B. KEEN, A. W. RICHTER, AND N. A. THROCKMORTON (2015): "The Zero Lower Bound, the Dual Mandate, and Unconventional Dynamics," *Journal of Economic Dynamics and Control*, 55, 14–38.

GERTLER, M., AND P. KARADI (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

GIMENO, R., AND A. IBANEZ (2018): "The Eurozone (Expected) Inflation: An Option's Eyes View," *Journal of International Money and Finance*, 86, 70–92.

GRISHCHENKO, O. V., AND J.-Z. HUANG (2013): "The Inflation Risk Premium: Evidence from the TIPS Market," *The Journal of Fixed Income*, 22(4), 5–30.

GRISHCHENKO, O. V., J. M. VANDEN, AND J. ZHANG (2016): "The Informational Content of the Embedded Deflation Option in TIPS," *Journal of Banking and Finance*, 65, 1–26.

GURKAYNAK, R., B. SACK, AND J. H. WRIGHT (2007): "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, 54(8), 2291 – 2304.

GURKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2010): "The TIPS Yield Curve and Inflation Compensation," *American Economic Journal: Macroeconomics*, 2(1), 70–92.

HALEY, M. R., M. K. MCGEE, AND T. B. WALKER (2013): "Disparity, Shortfall, and Twice-Endogenous HARA Utility," *Econometric Reviews*, 32(4), 524–541.

HALEY, M. R., AND T. B. WALKER (2010): "Alternative Tilts to Nonparametric Option Pricing," *Forthcoming, Journal of Futures Markets.*

HAMILTON, J. D. (1994): *Time Series Analysis.* Princeton university press, Princeton, N.J.

HANSON, S. G., AND J. C. STEIN (2015): "Monetary Policy and Long-Term Real Rates," *Journal of Financial Economics*, 115, 429–448.

JUDGE, G. C., AND R. C. MITTLEHAMMER (2012): *An Information Theoretic Approach to Econometrics.* Cambridge University Press.

JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105(March), 1177–1216.

KERKHOF, J. (2005): "Inflation Derivatives Explained: Markets, Products, and Pricing," Discussion paper, Lehman Brothers.

KITSUL, Y., AND J. H. WRIGHT (2013): "The Economics of Options-Implied Inflation Probability Density Functions," *Journal of Financial Economics*, 110(3), 696–711.

KOWARA, M. (2017): "20 Years In, Have TIPS Delivered?," https://www.morningstar.com/articles/801922/20-years-in-have-tips-delivered.html, Morningstar.

KULLBACK, S. (1968): *Information Theory and Statistics*. Dover Publishing, Mineola, NY.

LUDVIGSON, S. C., AND S. NG (2019): "Appendix for Updates of Uncertainty Data," Discussion paper, New York University.

MCCRACKEN, M. W., AND S. NG (2014): "FRED-MD: A Monthly Database for Macroeconomic Research," Discussion paper, Columbia University.

NAKAMURA, E., AND J. STEINSSON (2018): "The Power of Forward Guidance Revisited," *Quarterly Journal of Economics*, 133(3), 1283–1330.

PLANTE, M. D., A. W. RICHTER, AND N. A. THROCKMORTON (2016): "The Zero Lower Bound and Endogenous Uncertainty," *The Economic Journal*, 128, 1730–1757.

RICHTER, A., N. THROCKMORTON, AND T. B. WALKER (2014): "Accuracy, Speed and Robustness of Policy Function Iteration," *Computational Economics*, 44, 445–476.

ROSENBERG, J. V., AND R. F. ENGLE (2002): "Empirical Pricing Kernels," *Journal of Financial Economics*, 64, 341 – 372.

ROSS, S. A. (1976): "Options and Efficiency," *Quarterly Journal of Economics*, 90, 75–89.

SACK, B., AND R. ELSASSER (2004): "Treasury Inflation-Indexed Debt: A Review of the U.S. Experience," *Federal Reserve Bank of New York Economic Policy Review*, 10(1), 47–63.

STUTZER, M. (1996): "A Simple Nonparametric Approach to Derivative Security Valuation," *Journal of Finance*, 51, 1633–1652.

WHITE, H. (1982): "Maximum Likelihood Estimation of Misspecified Models," *Econometrica*, 50(January), 1–25.

5 Appendix A: Derivative Method

This method is based on the well known result that the second derivative of the price of a call option with respect to the strike delivers the risk-neutral density [Breeden and Litzenberger (1978), Ross (1976)]. To that end, let V_t denote the time t price of an option linked to the annual inflation rate $\pi_{t,t+h}$ from time t to time $t+h$ measured in years. Let p^* denote the time t forward probability density for $\pi_{t,t+h}$. Then V_t satisfies the pricing formula

$$V_t = B_{t,t+h} \int_{-\infty}^{\infty} p^*(\pi) V_{t+h} \, d\pi \tag{7}$$

where we follow Kitsul and Wright (2013) in discounting by an h year zero coupon bond $B_{t,t+h}$ taken from Gurkaynak, Sack, and Wright (2007). In particular, for an inflation cap of strike rate k and a one-year horizon, the value at time $t+1$ is the expiry payoff, giving

$$\begin{aligned} V_t &= B_{t,t+1} \int p(\pi) \max\left[(1+\pi)-(1+k), 0\right] d\pi \\ &= B_{t,t+1} \int_k^{\infty} p(\pi)\left[(1+\pi)-(1+k)\right] d\pi \end{aligned} \tag{8}$$

Differentiating twice with respect to k, we obtain

$$\begin{aligned} \frac{\partial V_t}{\partial k} &= -B_{t,t+h} \int_k^{\infty} p(\pi) \, d\pi \\ \frac{\partial^2 V_t}{\partial k^2} &= B_{t,t+h} p(k) \end{aligned}$$

Hence, given prices for inflation caps with several different strikes on a given date, we can back out an implied forward cumulative distribution P and the corresponding density p at values of annualized inflation in the range covered by these strikes by using the prices to estimate the above derivatives and setting

$$P(k) = 1 + \frac{1}{B_{t,t+h}} \frac{\partial V_t}{\partial k} \tag{9}$$

$$p(k) = \frac{1}{B_{t,t+h}} \frac{\partial^2 V_t}{\partial k^2} \tag{10}$$

An analogous argument allows us to derive these objects from a collection of inflation floor prices.[13]

[13] With horizons longer than one year, the expressions are more complicated, namely

$$\frac{\partial V_t}{\partial k} = -h(1+k)^{h-1} B_{t,t+h} \int_k^{\infty} p(\pi) \, d\pi$$

and

$$\frac{\partial^2 V_t}{\partial k^2} = -h(h-1)(1+k)^{h-2} B_{t,t+h} \int_k^{\infty} p(\pi) \, d\pi + h(1+k)^{h-1} B_{t,t+h} p(k)$$

As observed in Aït-Sahalia and Duarte (2003), the requirements

$$0 \leq P(k) \leq 1, \quad p(k) \geq 0 \tag{11}$$

combine with equations (9) and (10) to place sign and magnitude restrictions on the derivatives of the strike price curves. Financial market imperfections may result in these conditions being violated, with the consequence that the derived probability functions will be mathematically insensible. To address this issue, Aït-Sahalia and Duarte (2003) employ the algorithm of Dykstra (1983) to provide a method for minimizing the squared distance between the observed curves and the set of those which obey the constraints.

Since we only have prices $\{V_t(k_i)\}$ for a finite selection of strikes $k_i,\ i = 1, ..., n$ at each date, obtaining the derivatives in (9) and (10) is accomplished through local polynomial smoothing. This method, proposed in Aït-Sahalia and Duarte (2003) and applied to inflation options in Kitsul and Wright (2013), constructs Taylor coefficients for the price-strike curve of the options at each possible strike value by weighted least squares. In particular, for each strike k the objective of this problem is to choose coefficients $\beta_j(k)$ to minimize

$$\sum_{i=1}^{n} \left[V_t(k_i) - \sum_{j=0}^{p} \beta_j(k)(k_i - k)^j \right]^2 K_w(k_i - k) \tag{12}$$

where K_w is the weighting function. The derivatives of interest can subsequently be determined at each k from the estimated $\beta_j(k)$.

The main advantage of the above procedure is that it derives empirical distributions directly from the underlying economic theory using all available information. There are two main disadvantages. First, the method requires selection of a bandwidth parameter b for the weighting function K_w. While the literature proposes using an estimate of the asymptotic optimum (see Aït-Sahalia and Duarte (2003) and Fan and Gijbels (1996)), in practice one must take care to not under-smooth or over-smooth. Second, this method will also provide poor estimates of the forward distribution outside of the strike rates which were traded on a given day.

The former difficulty becomes an issue of practical importance in the latter half of our option price sample. Specifically, the number of strike rates being traded declines, so that estimated asymptotically optimal bandwidth does not sufficiently smooth between neighboring observations, resulting in severely multimodal estimated distributions or a break down of the method altogether. To adjust the bandwidth away from the theoretical optimum in a parsimonious way in our application below, we take the minimum bandwidth for which the distribution is unimodal.

To address the latter difficulty, we incorporate data from both cap and floor options on each date in order to construct a distribution over the widest range of inflation rates possible. Specifically, on a typical date there is price data available for caps of strikes $k_1^C, ..., k_m^C$ and floors of strikes $k_1^F, ..., k_n^F$, where $k_1^F < k_1^C$ and $k_n^F < k_m^C$, but $k_n^F > k_1^C$. In particular, the last inequality implies that there are strike rates for which prices of both a cap and a floor are available, say $k_1^{CF}, ..., k_\ell^{CF}$. Since the implied forward distributions need not agree in the overlapping region, we compute the distribution function $F(k)$ at such a point by linearly interpolating as follows. Denoting the floor-implied distribution function as $G(k)$ and the cap-implied distribution as $H(k)$, we set $\delta = (k - k_1^{CF})/(k_\ell^{CF} - k_1^{CF})$ and

$$F(k) = (1-\delta)G(k) + \delta H(k) \tag{13}$$

The strike-price curve estimated by applying constrained least squares and smoothing to five year inflation floors with strikes between -3% and 3% is shown in Figure 7. In contrast to the theory described above, the observed strike price curve on this date displays a non-convexity at strike rate -0.5%. This feature is smoothed over in the constrained least squares nearest neighbor, while the smoothed curve sits noticeably above the observed prices between strikes of 0% to 2.5%.

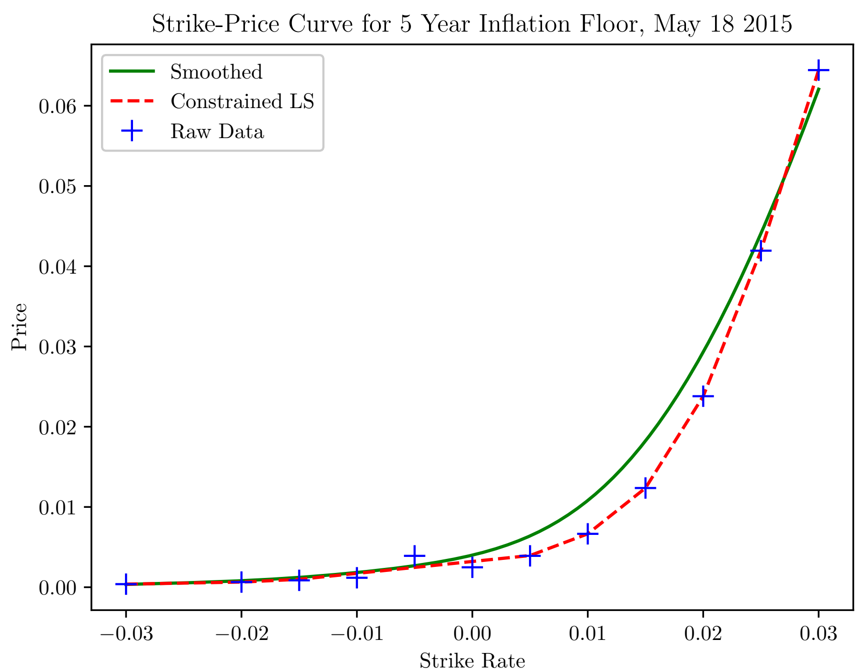

Figure 7: Observed, constrained, and smoothed strike-price curve for floors on five year inflation on May 18, 2015.

Figure 8 compares the forward cumulative distribution functions obtained from the Canonical and Derivative methodologies. Figure 8a and 8b plot the maximum and mean difference in the forward CDFs. We used the price of a one year cap with strikes of 1% and 2%, and a one year floor with a strike of 1% as constraints in the Canonical valuation. There are no clear trends in the figures, nor are there a significant number of outliers.

Figures 8c and 8d plot the one-year-ahead inflation on January 15, 2014 and January 13, 2015, obtained using each of the two methods described above for inflation caps and floors. The empirical distribution for post-1985 inflation, which is used in the application of Stutzer (1996)'s method, is plotted as

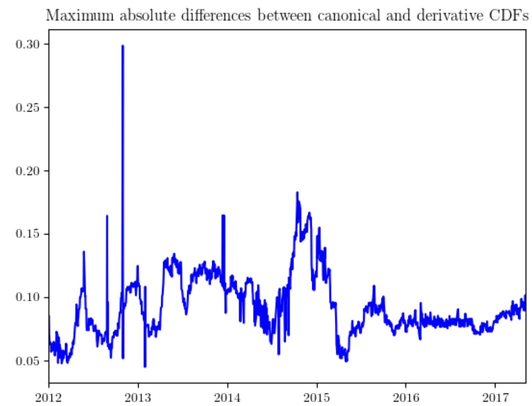

(a) Maximum Difference

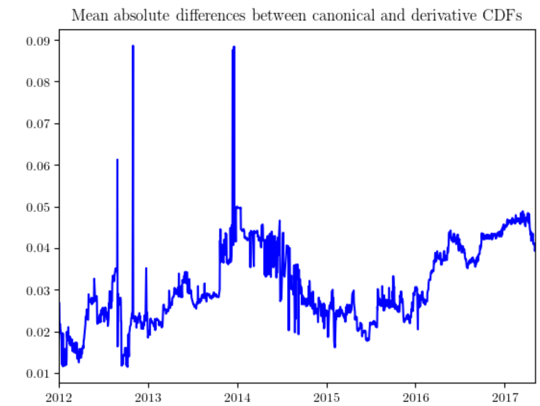

(b) Mean Difference

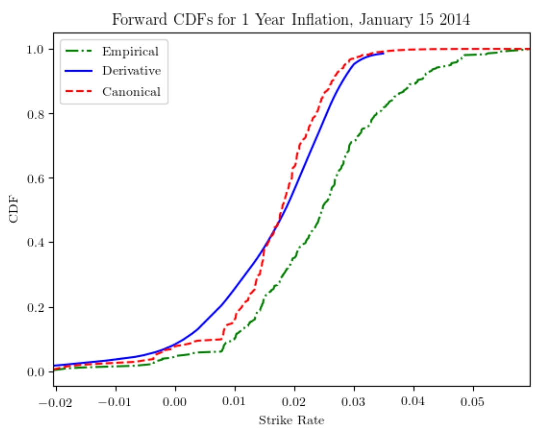

(c) Forward CDF Jan. 15, 2014

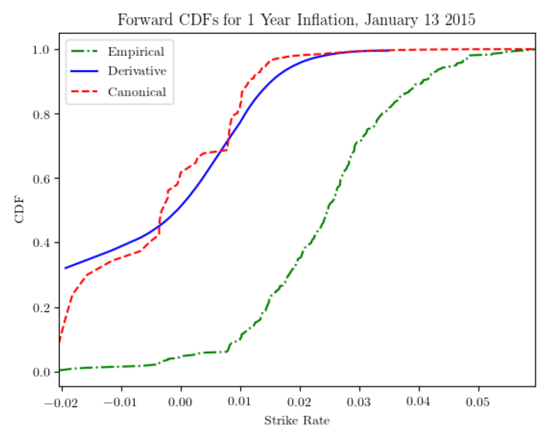

(d) Forward CDF Jan. 13, 2015

Figure 8: Forward cumulative distribution functions obtained using the Canonical and Derivative methodologies.

well. These dates are of particular interest because they represent a period in which the one-year-ahead forward rate is roughly consistent with current levels of inflation (January 15, 2014) and a date in which market beliefs suggest substantial deviation from current inflation (January 13, 2015). The figure shows that the approaches have substantial overlap during both periods (with the exception of the far left tail for January 13, 2015). We take this as prima facie evidence that our methodologies are broadly consistent, with additional evidence provided below.