

Whatever It Takes?

The Impact of Conditional Policy Promises

Valentin Haddad, Alan Moreira, and Tyler Muir*

April 7, 2022

Abstract

The announcement of an economic rescue tool often comes with implicit promises of more intense intervention if conditions worsen. We propose and implement a method to identify conditional policy promises and quantify their impact using data from options markets. When the Federal Reserve introduced corporate bond purchases during the COVID-19 crisis, markets expected five times more price support in crash scenarios relative to the median case. This implicit promise to significantly expand the size of the intervention in bad states explains half of the market response to the announcement. Furthermore, we document that the behavior of the price and tail risk of corporate bonds remains substantially distorted even after purchases have ceased. We confirm the pervasive influence of conditional promises across several policy announcements: U.S. quantitative easing, Bank of Japan asset purchases, bank equity injections in 2008, and FOMC releases.

*Haddad: UCLA & NBER, valentin.haddad@anderson.ucla.edu; Moreira: University of Rochester & NBER, alan.moreira@simon.rochester.edu; Muir: UCLA & NBER, tyler.muir@anderson.ucla.edu. We thank seminar participants at UCLA, LBS, USC, the Federal Reserve Board, UT Dallas, and Temple for comments.

Policy announcements are often perceived to come with promises, implicit or explicit, of more intense intervention if conditions worsen. The presence and role of these promises is intensely debated, particularly in the context of policies supporting financial markets and institutions as a way to stimulate the economy. For example, the Federal Reserve’s purchases of corporate bonds during the COVID-19 crisis sparked criticism due to the concern that further interventions any time markets crash could distort prices and create moral hazard issues. Others instead argue that commitment to more intervention makes policies more effective, a view exemplified by then-ECB director Mario Draghi’s “whatever it takes” approach to the 2012 sovereign debt crisis. Finally a third camp, common in much economic research, ignores promises altogether and takes policy announcements at face value. Despite the importance of these issues, quantitative evidence on promises is scarce because they are often involuntary, implicit, or simply vague.

In this paper, we propose and implement a method to measure conditional promises contained in policy announcements and their impact. Our method uses changes in the distribution implied by option prices before and after policy announcements to characterize the impact of policy across many possible states. We find pervasive evidence of the role of conditional promises, that is, larger interventions in bad states, across many financial stabilization policies including corporate bond purchases during the COVID-19 crisis, U.S. quantitative easing, asset purchases by the Bank of Japan, and bank capital injections during the 2008 crisis.

Consider the Federal Reserve’s corporate bond purchases during the COVID-19 crisis, our primary empirical setting. With the expansion of the pandemic, corporate bond markets tumbled during the first quarter of 2020. On March 23rd, the Federal Reserve made the unprecedented announcement that it would purchase up to \$300 billion of investment-grade corporate bonds. Investment-grade corporate bonds recovered by about \$500 billion in value immediately after the announcement was made, and they recovered by about \$1 trillion in total within three days. A few months later, the purchases occurred

and totaled about \$50 billion. One interpretation of these numbers is that asset purchases are an extremely potent tool to support prices, with a surprisingly large price impact of a dollar purchased between \$1.5 and \$10 depending on whether one uses the announced or realized quantities of purchases. If one uses the longer three-day window to assess the announcement effects these numbers are about doubled. However, it is also possible that markets inferred that the Federal Reserve was prepared to buy substantially more bonds than they actually did if the crisis deepened further. This promise would then make the sharp recovery less surprising given the announced quantities. Indeed, the response of bond prices to the policy announcement reflects the market's expectation of price support across all potential states. For example, suppose there is a 20% probability of a further crash and in this state the Federal Reserve purchases five times more bonds than they initially announced. In this case, the market response to the announcement is doubled relative to a situation without promises.

How can we infer the contingent policy and measure the role of promises? We need to assess how much the policy changes bond prices in each potential future state of the world. Option contracts provide a unique window into this state-contingent behavior. Consider, for example, an out-of-the-money put, a claim that only pays off in low price states. Expectations of policy in bad states entirely drive the change in the price of this contract at the announcement. Conversely, an out-of-the-money call that pays off only when prices are high reveals policy when conditions improve. This simple idea already leads to evidence supportive of the presence of promises. Using options on corporate bond ETFs, we see a more robust recovery in the price of contracts targeting poor states of the world at the announcement – implied volatilities fall much more in the left tail than the right tail of bond prices. We show how to go further than this qualitative assessment and obtain the entire state-contingent policy.

We estimate what we call a “price support function,” $g(\cdot)$, that answers: if the price moves to p absent intervention, the Federal Reserve will raise it by $g(p)$. First, we use

a result from Breeden and Litzenberger (1978). They show that observing the price of options across different strikes reveals the market perception of the distribution of the future price, that is, the risk-neutral distribution. We implement this procedure on options maturing three months from the policy announcement, roughly when the purchases were to be implemented. Comparing before and after the announcement, we can see how the perceived purchases change the entire distribution of the price at that horizon. The second step is to find a price support function that ties these two distributions together. Which prices have been moved to which new levels to obtain the post-announcement distribution? This type of mathematical question is a transport problem, and we derive its solution. The price support function is unique as long as the policy is order-preserving: the Federal Reserve does not support prices so much in poor states of the world that they exceed prices in better conditions.

The price support function we recover in the data is strongly asymmetric: while price support is relatively flat in good states of the world, it increases to much larger values as one moves towards worse states. This shape resembles a put option's payoff, lending support to the presence of a "policy put."¹ Quantitatively, we find that the market expected over five times more price support for adverse states in the corporate bond market compared to the current state. If the corporate bond market (absent any intervention) would have fallen by 30% more, the Fed was expected to boost corporate bond prices by around 40%. In contrast, if the corporate bond market had stayed in the same conditions, the Fed intervention would have raised prices only about 7%. We can infer this conditional implicit price support for many possible prices. We decompose the recovery in overall prices at announcement between a base effect and promises using these numbers. The structure of the change in option prices imply that at least 50% of the initial response in

¹See Cieslak and Vissing-Jorgensen (2021) and also Drechsler et al. (2018) for a related discussion of monetary policy and stock returns and Hattori et al. (2016) on the response of quantitative easing on stock market tail risk.

corporate bond prices is due to the conditional behavior in the left tail. That is, if the price support were constant and equal to the price support in the current state, announcement returns would have only been 50% as large. The significant effect of conditional promises supports the view expressed by Henry Paulson: “If you’ve got a bazooka, and people know you’ve got it, you may not have to take it out.”

Our results also suggest that promises are an appealing explanation for the broader finding that asset purchases by the Federal Reserve are associated with large movements in asset prices (Gagnon et al., 2018; Vissing-Jorgensen and Krishnamurthy, 2011; Haddad et al., 2021).² Further, Hesse et al. (2018) and Meaning and Zhu (2011) find a “weakening announcement effect:” initial stage announcements of asset purchases in the US and Europe have large effects on asset prices but later stage announcements do not. An interpretation consistent with our results is that early announcements convey large conditional promises to do more if the situation worsens, but in later stages of asset purchase programs these promises have already been reflected in prices.³ Thus, even a large announcement later on can appear to have zero effect.

We can go from measuring contingent price support to measuring contingent policy actions by making additional assumptions on their relation. To illustrate this process, consider the simple view that the elasticity of bond prices to bond purchases is constant across states. In this case, the asymmetry that we document implies over five times larger purchases if the price had fallen by 30%. Using the announced \$300 billion as a baseline, this corresponds to an extension of bond purchases to a whopping \$1.5 trillion if the market had further crashed.

We bring evidence from additional assets to sharpen our inference. We explore in which states the Fed was likely to make purchases. Corporate bond prices can fall either because of rising risk-free rates, increases in credit risk or because of disruptions in cor-

²See also D’Amico and King (2013), Hamilton and Wu (2012).

³See also Grinblatt and Wan (2020), which discusses anticipated effects of announcements.

porate bond markets not due to fundamentals. We infer the distribution of a synthetic corporate bond index using options on Treasuries and options on the investment-grade CDX index. Using copula methods to model their joint distribution, we find that the distribution of synthetic bonds did not show as severe of a left tail initially. Moreover, this left tail did not shrink as substantially after the announcement. These observations mean that the market expected large conditional purchases to happen in states where corporate bond markets were highly dislocated. In other words, these would be states with an even larger CDS-bond basis than at the time of announcement. Another question is whether changes in the broad pricing of risk due to the policy can drive the announcement effects instead of our inferred price support. The lack of response to the policy announcement in the closely related high-yield market suggests this is not the case.

The price support function gives a sharp measurement of the short-term implications of policy promises. But the announcement of purchases can also have long-term implications if the market believes that the Federal Reserve will now step into the bond market whenever it gets distressed. We present three pieces of evidence suggesting long-term effects of the introduction of bond purchases. The idea is to focus on properties of tail risk specifically, similarly to our main approach. First, after the new policy, tail risk in corporate bond markets becomes far less sensitive to other measures of asset price tail risk, measured using equity or CDS index options. This result suggests that the policy put is still present, dampening downside risk in this market. Second, corporate bond returns also become far less sensitive to changes in the VIX, which shows the downside dampening is relevant for the level of prices. Third, corporate bond spreads become far less responsive to changes in the “pseudo spreads” implied by equity options constructed by Culp et al. (2018). This last result demonstrates that the change in downside risk properties is restricted to, or strongest in the corporate bond market. These long-lasting effects suggests that expectations of future interventions are already impacting price dynamics today. This is consistent with the view that the possibility of future purchases in bad

states is de-linking bond prices from fundamentals what can induce moral hazard by issuers and investors. But it is also consistent with the more benign view that future interventions make the bond market less prone to dislocations during crises.

While we conduct an in-depth study of bond purchases during the COVID-19 crisis, our measurement framework is not specific to this event. We construct the conditional price support function for several other policy announcements. We focus specifically on policies that have a substantial impact on financial markets and for which we have relevant option data. One family are central bank asset purchases: by the Bank of Japan in 2013, the quantitative easing operations in the US from 2008 to 2013, and announced asset purchases by the European Central Bank in 2010 to 2012. Another type of intervention is direct support to specific institutions: the financial sector bailout in the US in 2008. Finally, we also consider regular monetary policy operations by looking at the response to FOMC announcements. We find that conditional promises play a pervasive role across all these announcements, albeit with different intensity.

Our results speak to macro finance models that assess the impact of policies that support financial market prices during crashes such as asset purchases or equity injections to the financial sector (e.g., He and Krishnamurthy (2013), Moreira and Savov (2017), Vayanos and Vila (2021)). Quantitative theoretical exercises typically treat the policy implementation as “one-off” events for simplicity and ignore promises. In contrast, our findings suggest that promises are first order to understanding the effectiveness of policy announcements.

Our results are related to broader work using information in options markets to interpret policy. Kelly et al. (2016b) focus on the price of political uncertainty associated with uncertain actions over a pre-specified date (e.g., elections). Our work instead focuses on inferring conditional policy typically inferred from an unscheduled, unexpected event. Relatedly, Kelly et al. (2016a) use options markets to evaluate government guarantees on the financial sector in the 2008 crisis. Barraclough et al. (2013) use option prices to inform

merger announcements.

1. Measuring Conditional Promises

In this section, we introduce a framework for measuring conditional policy promises. We start by a simple example illustrating how the presence of these promises affects the response of asset prices to policy announcements. The overall market response reveals the combined effect of the announced policy and conditional promises. However, the contingent nature of option contracts sheds light on the states in which promises will be fulfilled. Our method builds on this insight to quantify promises. Specifically, we show how to estimate a price support function: how much is the policy changing prices as a function of the state of the world.

1.1 The Effect of Policy Promises on Asset Prices

Consider the following stylized example with two dates, 0 and 1. At date 0, the initial price of an asset is p_0 . Under rational expectations, this price is the (risk-neutral) expected value of the date-1 price: $p_0 = E[p_1]$.

No promises. A new policy is unexpectedly announced at date 0: a quantity Q of a policy tool will be used at date 1. The per-unit effectiveness of the policy in moving prices is given by \mathcal{M} . For example, as will be our main empirical focus, the Fed unexpectedly announces in the middle of a crisis that it will purchase a quantity Q of corporate bonds (the asset) at a future date. In that interpretation, \mathcal{M} reflects the price impact per quantity of asset purchased, the inverse elasticity of demand for corporate bonds. Another example would be the announcement of a new fiscal stimulus package. There, \mathcal{M} would be the present value of the product of the fiscal multiplier with the pass-through from GDP

to corporate profits.

Given the new policy, the price at date 1 will be $p'_1 = p_1 + \mathcal{M}Q$. Therefore the post-announcement price becomes

$$p'_0 = E[p_1] + \mathcal{M}Q. \quad (1)$$

In other words, the price change at announcement ($p'_0 - p_0$) is exactly proportional to $\mathcal{M}Q$. A number of researchers have used this idea to back out the overall effect of purchase policies and the multiplier of prices to purchases.

Conditional promises. However, it is not that easy. When the new policy is announced, the market might (rightfully) infer that the policymaker is willing to intervene more strongly if conditions worsen. For example, it could be that the market expected the Fed to purchase even larger amounts to corporate bonds were the COVID-19 crisis to deepen. More broadly, conditional promises can be voluntary or not: on the one hand, the policymaker might want to make the new policy instrument part of their toolkit; on the other hand, they might have opened Pandora's box and lack the commitment to stop using this instrument in the future. Promises can also be implicit or explicit. On the implicit side, market participants expend large efforts trying to infer the future conduct of monetary policy following FOMC statements. On the explicit side, Mario Draghi, the then-president of the ECB, expressed clearly his willingness to do "whatever it takes" in the midst of the Euro area sovereign debt crisis in 2012.

To illustrate the impact of conditional promises, assume that the policymaker will scale up the policy by an additional amount Q^* if we are in a state at date 1 where the no-intervention price would fall below a cutoff value p^* . In this situation, the price at

date 1 becomes $p'_1 = p_1 + \mathcal{M} (Q + 1_{p_1 \leq p^*} Q^*)$. The post-announcement price is

$$p'_0 = E[p_1] + \mathcal{M}Q + \mathcal{M}P[p_1 \leq p^*]Q^*. \quad (2)$$

We see that both the baseline policy and the implicit promise shape the price response to the announcement. The promise provides an additional boost to the price equal to the product of the additional policy implemented, policy effectiveness, and the probability of additional policy.

Both effects are intertwined and, based on the price response to the announcement alone, they cannot be separated. In particular, ignoring the presence of promises leads to incorrect inference about the effectiveness of the policy. If an econometrician assumes that only the baseline policy is present and estimates the multiplier by comparing the price response to the announced purchases (or the realized purchases provided the promises are not realized), their estimate will be biased:

$$\mathcal{M}_{\text{estimated}} = \mathcal{M} \left(1 + \underbrace{P[p_1 \leq p^*]}_{\text{prob. of realized promise}} \times \underbrace{\frac{Q^*}{Q}}_{\text{rel. size of the promise}} \right). \quad (3)$$

Because the promise provides additional price support, the effectiveness of the policy will be overestimated. How large is the bias? First, the bias depends on how likely are the promises to be implemented. Of course these probabilities are always less than 1, and can be small if a massive crash is unlikely. However, because new policy tools are often used in difficult and uncertain conditions — think of the midst of the COVID-19 crisis — these probabilities are likely non-negligible. Second, the bias depends on the size of the promised policy relative to the baseline amount Q^*/Q . This second term can be sizable, for example much larger than 1. Indeed, if the crash scenarios are dramatic, the policy maker might expend significantly more resources. In the empirical evidence

we study later on in this paper, we find that indeed, both the probability of additional support and the strength of additional price support are economically significant. To be concrete, the corporate bond market increased by about \$1 trillion in value when the Fed announced corporate bond purchases in March of 2020 though the Fed ended up only making purchases of around \$50 billion. Our estimates suggest that close to half of the response is due to about a 20% probability of a 5-fold increase in the size of the program.

The information in option prices. How can one separate the promise from the baseline policy in this case? Option contracts on the asset offer a path forward. Consider for example a call option — a contract paying $\max(p_1 - K, 0)$ at date 1 — with a large strike price ($K > p^* + \mathcal{M}(Q + Q^*)$). Because this call only pays in states of the world where the promise is not realized, the change in option price at announcement is entirely driven by the baseline promise:

$$p'_{0,call} - p_{0,call} = E[\max(p_1 + \mathcal{M}Q - K, 0)] - E[\max(p_1 - K, 0)]. \quad (4)$$

Conversely, option contracts focusing on the poorest states of the world will respond to the fulfilled promise. A put option with a low strike ($K < p^* - \mathcal{M}(Q + Q^*)$) only pays off in states in which the promise is realized. Then the change in option price at announcement reflects only interventions with the promise:

$$p'_{0,put} - p_{0,put} = E[\max(K - p_1 - \mathcal{M}(Q + Q^*), 0)] - E[\max(K - p_1, 0)]. \quad (5)$$

These two cases highlight that we can learn about the conditional nature of the policy by using options to zoom in on various parts of the state space. One has to go further than this simple intuition to get to quantitative statements. For example, comparing equations (4) and (5) suggests that the presence of policy promises can be detected by “more

action” in out-of-the-money puts than out-of-the-money calls. Because these two contracts are very different to start with, more careful calculations are needed to implement such a comparisons. This is the focus of next section.

1.2 A Method to Estimate Conditional Policy Promises

We present a method to estimate conditional policy promises following the announcement of a new policy using option prices. We introduce a flexible representation of conditional policies and the assumptions underlying its estimation. Then, we explain the two steps necessary to go from option data to a conditional policy.

1.2.1 The conditional price support function.

We maintain the timing assumptions of our motivating example. At date 0, a policy is unexpectedly announced to be implemented at date 1.⁴ The policy announcement potentially contains conditional promises. That is, the policy implementation can depend on the realized state of the world at date 1. Our first assumption is that the state of the world at date 1 maps exactly to the date-1 price of the asset absent policy intervention p_1 . Given this mapping, the effect of the policy on the price can be represented by a price support function $g(\cdot)$. The price support function computes how much the price is changed by the policy in each state of the world.

Assumption 1. Price support function. *The asset price at date 1 after the policy is announced p'_1 is equal to the no-policy price p_1 plus a conditional price support $g(p_1)$,*

$$p'_1 = p_1 + g(p_1). \tag{6}$$

This assumption entertains the conditional nature of the announced policy in a flexi-

⁴We relax this assumption of a fixed timing when considering extensions of our approach.

ble way. For example a fixed policy can result in a constant g , while our example from the previous section corresponds to $g(p) = \mathcal{M}Q + \mathcal{M}Q^*1_{\{p \leq p^*\}}$. The representation of the policy by a price support function does not imply that the policy acts only on the asset price or is designed to focus on the asset price. Rather, the conditional price support is the information of the policy that we can recover using price information alone. With additional information about the policy — such as Q or \mathcal{M} in our example — one can link back the price support back to specific actions. Another aspect where our assumption has content is the assumption that the conditioning is entirely as a function of the no-policy price. Policymakers, even when they explicitly want to support prices, look at a variety of pieces of information to make decisions. The simplifying assumption reduces this to a unique dimension, capturing intuitively the difference between good and bad states of the world. Using the no-policy price as the conditioning information reflects the aspect of conditioning that is captured by option prices. As we discuss in our empirical work, for this assumption to be plausible it is important to focus on an asset that captures well the information driving the policy studied.

Second, an assumption is necessary about the pricing of the asset and options on it at date 0. We assume pricing by the same risk-neutral distribution before and after the policy announcement.

Assumption 2. Asset pricing. *The same risk-neutral distribution F_{p_1} over states of the world p_1 prices the asset and options before and after the policy announcement. That is, for all functions h , a claim paying $h(p_1)$ at date 1 has price $\int h(p_1)dF_{p_1}(p_1) = E[h(p_1)]$.*

Underlying this assumption is the simple view that policies do not change the fundamental randomness of the world. Instead, they change what happens in various states of the world. This is standard in the setup of dynamic stochastic models: start with a primitive filtration and probability measure, and derive equilibrium outcomes. We go one step further and assume a constant risk-neutral probability measure. This assump-

tion brings some flexibility: we do not impose coincidence of risk-neutral and historical probabilities. However it also has some bite: we are implicitly assuming that the stochastic discount factor between date 0 and date 1 is unaffected by the policy. When looking at the data, we explore ways to assess the plausibility of this assumption. For example, one can check whether a related asset for which the policy has no direct effect responded to the announcement. In the case of corporate bond purchases, we compare high-yield bonds, which were not initially targeted, to investment-grade bonds which were. A shift in pricing kernel perspective would imply large movements in the risk-neutral distribution of high-yield bonds while conditional policy targeted at investment-grade would not.

Finally, we need to impose some regularity on the price support function $g(\cdot)$ to be able to estimate it from the data.

Assumption 3. Order-preserving policy. *The post-policy price $p'_1 = p_1 + g(p_1)$ is increasing in the no-policy price p_1 .*

Said otherwise, we assume that the policy does not change the ranking of the asset price across states. This assumption is plausible. For example, the policy does not support the price so much in (no-policy) bad states that it becomes higher than in good states. There is also a sense in which such policies are efficient. Consider a policy-maker who targets a given distribution of the price. Multiple price support functions can lead to this distribution, but an order-preserving policy minimizes the use of large changes in prices. Another take on this assumption is that it leads to conservative estimates of the conditional nature of the policy. This is because a policy with order switching leads to more asymmetry across states; bad states have to be relatively supported even more to make them switch with good states.

1.2.2 Estimation strategy

We show how to use the behavior of option prices around the policy announcement to recover the conditional price support function $g(\cdot)$. First, we obtain the distributions of the price with and without policy, p_1 and p'_1 . Second, we solve the transport problem of inverting the price support to move from one distribution to the other.

Step 1: Recovering the future price distribution. We follow the approach of Breeden and Litzenberger (1978) to recover the distribution of the future price of the asset. They show that observation of option prices (calls or puts) across strikes allow you to infer the distribution of the price of the underlying. Let us review this result. Denote $Put(p_1, K) = \max(K - p_1, 0)$ the payoff of a put with strike K when the price is equal to p_1 . The difference between the payoff functions of two puts with close strikes approximates a step function at that point. Figure 1 illustrates this result.

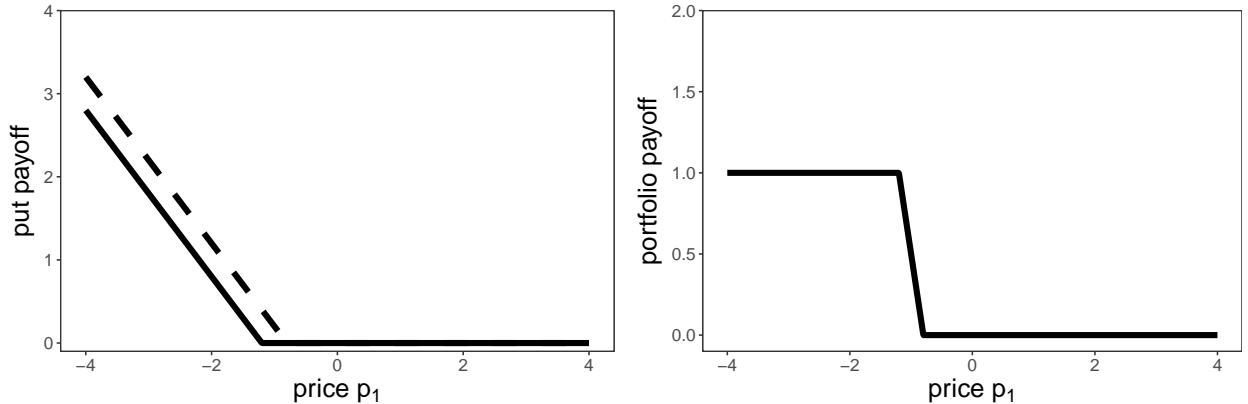


Figure 1: Using a Put Portfolio to Approximate an Indicator Function

The left panel reports the payoff of puts with strike -1.2 (solid line) and -0.8 (dashed line). The right panel reports the payoff of a portfolio long $(-0.8 - -1.2)^{-1} = 2.5$ units of the high-strike put, and short 2.5 units of the low-strike put.

The left panel plots the payoff function for two strikes close to -1 : -1.2 and 0.8 respectively. The right panel reports the difference between these payoffs scaled by the

difference in strike 0.4. Equivalently, this is the payoff of a portfolio long 2.5 units of the high-strike put and short 2.5 units of the low price put. This difference is very close to a step function equal to 1 below -1 and 0 above. Formally, this observation corresponds to

$$\frac{dPut(p_1, K)}{dK} = \lim_{h \rightarrow 0} \frac{Put(p_1, K + h/2) - Put(p_1, K - h/2)}{h} = 1_{\{p_1 < K\}} \quad (7)$$

Turning back to date 0, this implies that the slope of the put prices with respect to the price is equal to the expected value of the indicator function. This expected value is exactly the probability that p_1 is less than K , the cumulative distribution function (CDF) $F_{p_1}(K)$.

The first step of our method is to apply this idea to the option curve (the relation between strike and put price) before and after the announcement. Doing so allows to recover the cumulative distribution function of the no-policy price p_1 and the post-policy price p'_1 , which we denote F_{p_1} and $F_{p'_1}$, respectively. In practice, one cannot observe option prices for all strikes, but instead for a finite number of specific strikes. We follow the common practice of interpolating these prices between strikes, specifically as described in Malz (2014). However, we are careful to not extrapolate the curves. As a consequence we only obtain the CDF for finite intervals, which we denote I and I' . While this limitation precludes the usual application of the result of Breeden and Litzenberger (1978) (pricing arbitrary option contracts), we will see next that we are still able to recover exactly the function $g(\cdot)$, only over a finite interval.

Step 2: Solving the transport problem. Once we have the two distributions, the next task is to find the conditional price support $g(\cdot)$ that explains the change in distribution. This type of problem is known as a transport problem: how should we move all the values of a random variable to obtain a new distribution? The order-preserving property (assumption 3) imposes that this transport is monotone. This feature is enough to guarantee existence and uniqueness (up to probability 0 events) of a solution $g(\cdot)$. There is a

simple method to construct this mapping. Start from a value x and compute the corresponding CDF $F_{p_1}(x)$. Then, because the order of states of the world is unchanged, this value must map to another value y that falls at the same ranking, that is, the same CDF value. This corresponds to finding y such that $F_{p'_1}(y) = F_{p_1}(y)$. Once we find this price mapping, we simply have: $y = x + g(x)$, which reveals $g(x)$. For example, assume your initial value is the 20th percentile of the distribution of p_1 . The post-policy price corresponding to this state is the 20th percentile of the distribution of p'_1 . The price support function is the change in price necessary to move from this initial value to the post-policy price. The following proposition summarizes this calculation.

Proposition 1. *The unique order-preserving policy price support function to go from F_{p_1} to $F_{p'_1}$ is equal to*

$$g(p_1) = F_{p'_1}^{-1}(F_{p_1}(p_1)) - p_1. \quad (8)$$

Going back to the issue of implementation, we only observe the CDFs on finite intervals. Examining this formula tells us that we can only recover the function g for states for which we can measure both CDFs. That is, if we can measure the 20th percentile of both CDFs, we can obtain the mapping for this percentile. Formally, this implies that we can solve the function $g(\cdot)$ over the domain $F_{p_1}^{-1}(F_{p_1}(I) \cap F_{p'_1}(I'))$.

Figure 2 illustrates how changes in distribution map to the price support function. We consider two extreme cases of conditional promises. Panel A and B report the probability density functions (PDFs) before and after policy announcement and the price support function for a constant price support. In this case, there is no conditional promise: the price is increased by the same amount no matter what happens. The whole distribution simply experiences a parallel shift to the right. Panel C and D report the same quantities but for a price floor, that is $p'_1 = \max(p_1, \underline{p})$ for some threshold \underline{p} . This is the most extreme case of conditional promise: if the price falls too low absent intervention, the policymaker

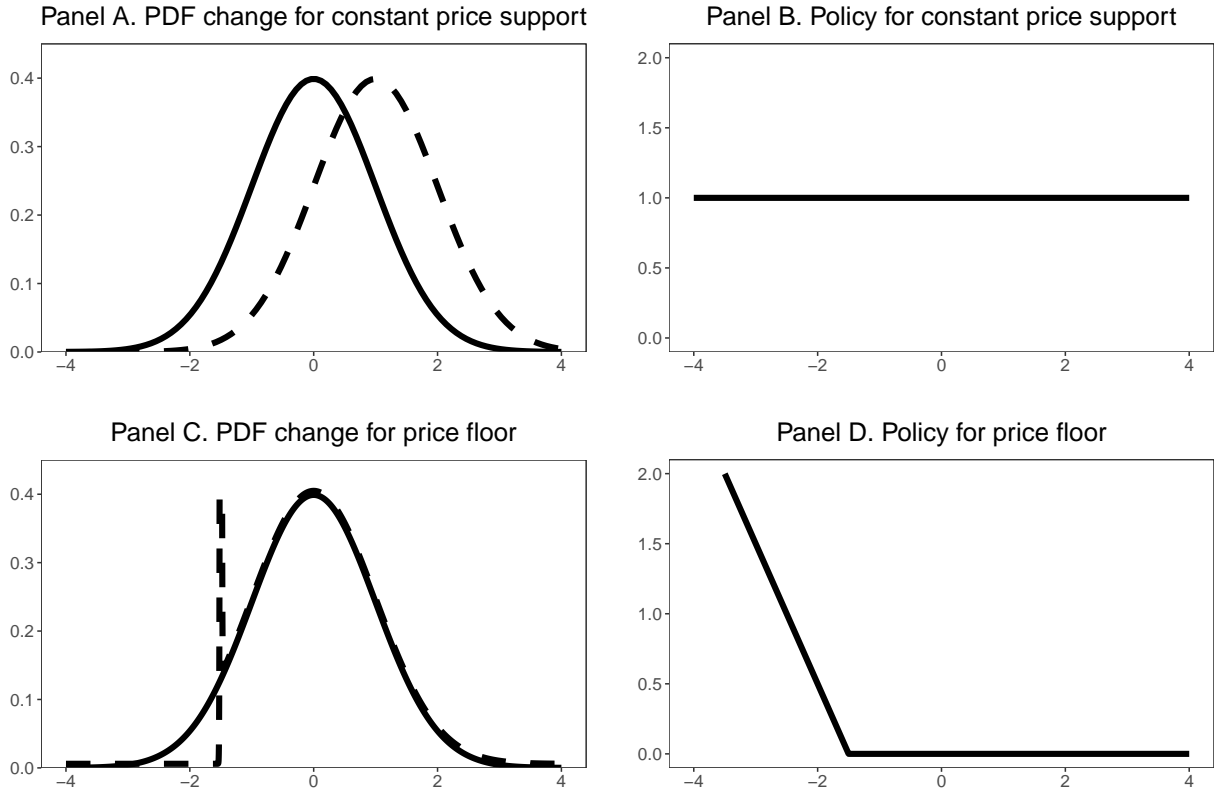


Figure 2: Examples of Distributions and Conditional Price Support Policies

The top row considers a constant price support: the price is increased by the same amount in all states. The bottom row considers a price floor: the price is forced to stay above a threshold \underline{p} . The left panels report the PDF of the date-1 price before (solid line) and after (dashed line) the policy announcement. The right panels report the corresponding price support functions $g(\cdot)$.

does whatever it takes to ensure it stays up to the threshold. In terms of distribution we see no change above the threshold but all the probability below the threshold becomes accumulated right at the threshold. This corresponds to a sharply decreasing function $g(p)$ below the threshold. For each unit that the no-policy price falls further down, the price gets supported by one more unit to stay at the threshold. This slope of -1 in this range is actually the largest permitted while maintaining the order-preserving property. Interestingly this price support function coincides exactly with a put option payoff, lend-

ing formal support for the commonly use notion of a policy put such as the “Greenspan put.”

2. Data

We use a variety of financial instruments that have traded option contracts referenced to them and were the direct target of policy announcements. These include options on the iShares investment grade corporate bond ETF (LQD), the iShares high yield corporate bond ETF (HYG), the future on the S&P500 index, the future on the ten year maturity Treasury bond, the financial sector ETF (XLF), the future on the Nikkei index, and the CDX investment grade credit basket spread. We aim to use options of maturities close to three months. Ideally one would like longer maturity options given the nature of implicit promises are often quite durable, but in practice the liquidity in the vast majority of these markets is heavily concentrated around three months.

Our approach to recover the state price density follows Malz (2014). We obtain prices and use the standard Black-Scholes formula to translate prices into implied volatilities. We then fit a cubic spline to the implied volatility curve. Armed with this function we can easily compute derivatives of option prices numerically for the range of liquid strikes. Specifically we evaluate the Black-Scholes formula for different strikes and the associated implied volatilities. We then compute first and second differences to recover the implied state-price density.

3. Corporate Bond Purchases in 2020

3.1 Background and Effect on Prices

On March 23rd, 2020 the Federal Reserve unveiled facilities that would purchase investment grade corporate bondx and corporate bonds ETFs through the Secondary and Primary Market Corporate Credit Facility (SMCCF and PMCCF). This announcement was important because it was the first time the Fed had directly targeted the corporate bond market. As Haddad et al. (2021) show, the announcement came at a time when corporate bond prices were depressed (investment grade bond indices had fallen over 20% from February to March 23rd) and dislocated (very safe investment grade debt was trading at steep discounts to Treasuries even taking into account credit risk measured from the CDS market).⁵ These facilities took time to set up so that no bonds were purchased for several months, and there was a total capacity of \$300 billion.

The announcement of the SMCCF and PMCCF had a significant and immediate impact on corporate bond prices. Table 1 shows the return response for the iShares investment grade corporate bond ETF (LQD) using a window of one to three days around the announcement. This large ETF captures the broad universe of investment-grade corporate bonds and is effectively a leading investment grade bond price index. The ETF has the advantage that it summarizes the impact of the announcement on corporate bond prices without having to obtain transaction level data of individual bonds which trade less frequently. The cumulative three-day announcement window return is 14%, meaning the Fed announcement had a large effect on corporate bond prices. In terms of abnormal excess returns (with controls for high-yield and the stock market) the response is still substantial at around 10%. The overall return of 14% translates into around \$1 trillion

⁵See also O'Hara and Zhou (2020), Boyarchenko et al. (2020), Kargar et al. (2020), Gilchrist et al. (2020), D'Amico et al. (2020) who study the effect of Fed interventions during this period on market liquidity and corporate bond prices.

Table 1: Announcement Effect

This table shows the return on an investment-grade corporate bond ETF (LQD) on the announcement on March 23rd 2020 by the Fed to purchase corporate bonds. The first two columns use a three day announcement window and the coefficient represents the cumulative daily return on the announcement. The second column uses the excess return over TLT, a long term Treasury ETF, and controls for excess returns on high yield bonds and the stock market so that the announcement effect is the cumulative abnormal return. The last two columns repeat this same exercise over a one-day window.

	(1) Three days	(2) Three days	(3) One day	(4) One day
$Announce_t$	14.17*** (3.78)	10.27** (1.25)	7.37*** (0.01)	6.63*** (0.07)
$r_t^{HighYield}$		0.54*** (0.04)		0.55*** (0.04)
r_t^{SP500}		0.03 (0.02)		0.03 (0.02)
Constant	0.02** (0.01)	-0.01 (0.01)		
N	2,988	2,988	2,988	2,988
R^2	0.11	0.87	0.09	0.87

increase in market value for investment grade corporate bonds. Using a one-day window for the announcement drops the raw return and abnormal excess return to about 7%. The shorter one-day window provides better identification at the cost that it may take the market time to process the announcement. While we use event windows here measured in days, Haddad et al. (2021) show in high frequency intraday data that prices increased right at the time of the announcement, and that other news was unlikely a factor given other assets such as high yield corporate bonds, stocks, or Treasury bonds showed little movement.

Most importantly for this paper, since the announcement of the facility was a new and unexpected intervention into corporate bonds, it is natural that the announcement shifted market participants expectations of further interventions if the situation deteriorated. While it was left unclear the total size of purchases that would be made, the

announced total capacity of purchases on March 23rd was \$300 billion. Ultimately, only around \$50 billion of purchases were made, around 0.7% of the market capitalization of investment grade corporate bonds. It is initially surprising that the corporate bond market increased by \$0.5-1 trillion on announcement given so few purchases were made. A reasonable view is that purchases would have been larger had the bond market fallen further and that part of the increase in value comes from this conditional promise. The most natural place to look for these expectations is to see how option prices on LQD changed around the announcement. A straightforward implication of the conditional purchases view is that it should reduce left tail events in corporate bond markets which show up directly in option prices.

3.2 Option Prices and Changes in the Distribution of Corporate Bonds

To explore the tail risk view of conditional purchases, Figure 3 plots implied volatility curve for three month options on the investment grade bond ETF (LQD) on the trading day before the announcement was made compared to the end of the day the announcement was made. While implied volatility dropped notably, the drop was most pronounced in the left tail (deltas below 30%). We will interpret this left tail drop as the conditional impact of Fed policy. That is, as capturing market expectations of stronger intervention if corporate bond prices fell further. This left tail drop was even more pronounced over a longer three-day window.

We follow Malz (2014) to convert this change in the implied volatility curve directly to changes in the distribution of corporate bond prices. This approach follows the insight in Breeden and Litzenberger (1978) that option prices can be used to recover risk-neutral distributions (see also Figlewski (2010)).⁶ Figure 4 shows the implied cumulative dis-

⁶Specifically, the method of Malz (2014) fits a cubic spline through the observed points on the implied volatility curve and then differentiates this to arrive at the CDF. The spline is clamped as explained in Malz (2014) to avoid potential arbitrages.

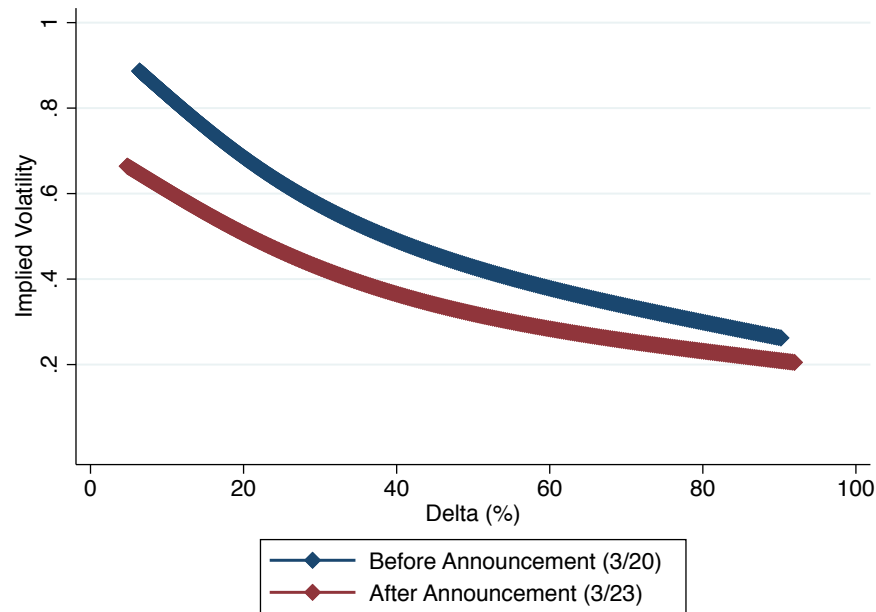


Figure 3: Implied volatility before and after announcement.

This figure provides implied volatility from options on an Investment-Grade corporate bond ETF (LQD) on March 20th and 23rd, 2020 as a function of the option delta. Time to maturity is 3 months.

tribution function (CDF) for future values of the LQD ETF on these two days based on option prices. More specifically, we capture the (risk-neutral) distribution of the potential price for LQD in three months using options with a three month maturity but with varying moneyness. We compare this distribution before and after the announcement was made. The figure reveals a clear rightward shift in this overall distribution, but again most notably there is significantly more action in left tail events. Before the event there was about a 15% chance that the value of LQD would drop by 30% or more. This state of the world is vastly reduced after the policy is announced.

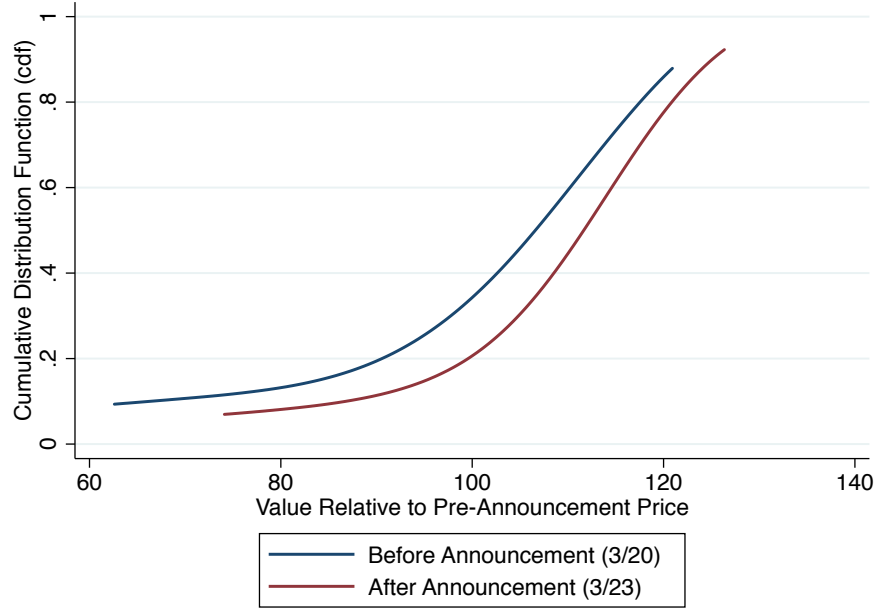


Figure 4: CDF based on option prices.

This figure shows the implied CDF of future returns on corporate bonds extracted from option prices.

3.3 Conditional Price Support

We now use this information to directly capture the conditional price effects of the Fed intervention. Let $g(p)$ denote the conditional price support of the Fed policy as a function of the non-intervention price p . That is, p denotes the price of investment grade corporate bonds absent any Fed intervention and should be thought of as capturing the underlying fundamental state of the corporate bond market. We interpret $g(p)$ as the total price effects of any Fed intervention. For example, if the Fed purchases more bonds in distress periods when the price p is low, then $g(p)$ will be declining in p . Notice in this case $g(p)$ will shrink the left tail of the distribution of bond prices. In contrast, an unconditional promise to purchase bonds with a constant price impact of purchases would mean $g(p)$ were a constant, which would simply shift the entire distribution in parallel to the right. The new CDF that includes the effects from conditional purchases is based on the post in-

intervention price which we call p' . We then have $p' = g(p) + p$. Thus, by using the behavior of the CDF pre and post announcement we can find $g(p)$ and hence assess the impact of the Fed as a function of the non-intervention price p . We have $F_{pre}(p) = F_{post}(p + g(p))$, where F is the CDF. This simply says the impact of the conditional purchases $g(p)$ shifts us to the new distribution. We can recover $g(p)$ as $g(p) = F_{post}^{-1}(F_{pre}(p)) - p$.

Figure 5 shows the implied effect of the policies on prices by plotting the function $g(p)$ expressed as a percentage of the no-policy price p (e.g., we plot $g(p)/p$ which gives more natural units). We can use this plot to gauge the magnitudes of conditional Fed intervention. First, $g(p)$ is not flat as an unconditional promise would imply, but is downward sloping particularly for low values of the price. The price support resembles a put option. This suggests that the bond market believed stronger intervention would occur if the situation in the corporate bond market deteriorated. To gauge magnitudes it is worth picking a couple points on this graph. If, absent any policy intervention, prices would have increased by 20%, the Fed would push the price up by an additional 3%. If, prices would have declined by 30% instead, the Fed would push the price up by an additional 40%.

The gray shaded region in Figure 5 gives the 95% confidence interval. We construct this by bootstrapping daily pairs of implied volatility curves for options on the invest-grade bond ETF. From these we can construct our price support function $g(p)$ at each day. We compute the price support in units of standard deviation based on at the money implied volatility on the initial day. This deals with the fact that implied volatility is much higher around the Fed announcement than on other days. We then scale these up based on at the money implied volatility on the trading day before the announcement was made. The confidence interval indicates statistically significant price support for the left tail of the distribution, and only insignificant estimates of price support at the upper end.

Revisiting Announcement Effects

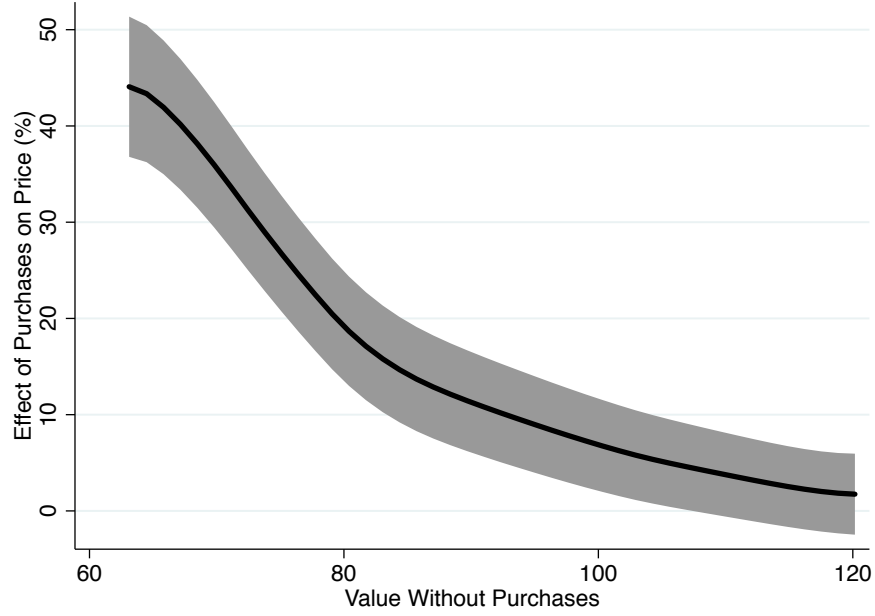


Figure 5: Effect on prices.

This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

We next compute the fraction of the initial announcement return stemming from conditional promises. Recall that $E[g(p)]$ is equal to the one-day announcement return of 7.4% shown earlier. To gauge the effects of this return coming from promises, suppose the Fed policy was not conditional. Specifically, assume that in the left tail the Fed simply supported prices by a constant amount equal to $g(100)$, the price support expected for no change in the no-policy price. Let $\tilde{g}(p) = g(p)$ when $p > 100$ and $\tilde{g}(p) = g(100)$ when $p \leq 100$. We compute $E[\tilde{g}(p)]$ using the implied probabilities of each state which gives the change in price with no promises, and thus the effect due to promises is $E[g(p)] - E[\tilde{g}(p)]$. We find that about 53% of the overall effect on prices comes from conditional promises to buy more heavily in adverse states. Thus, conditional promises *alone*, with no commitment to purchase anything unless the situation deteriorated, boosted the total bond

market value by around 3.9% or about \$275 billion.

From Price Support to Quantities

Our main methodology delivers conditional price support, but does not specify how this price support is achieved in terms of quantities of bonds purchased in different states. We propose a simple framework to make some progress on this dimension and interpret our numbers in terms of quantities purchased. The price impact of purchases in the bond market is \mathcal{M} so that the total change in bond market value of purchases are $g(p) = \mathcal{M} \times q(p)$ where $q(p)$ represents the total quantity of bonds or ETFs purchased by the Fed as a function of the no-policy price p . We express the quantity purchased $q(p)$ as a fraction of the total market value of bonds outstanding (because we express price moves in percent rather than total changes in bond market value, the purchases are then also in percent of total value of bonds outstanding). First, we focus on relative statements that don't involve taking a stand on \mathcal{M} . As a baseline, $g(p)$ is around 7% when the price without purchases remains at its current level of 100 (e.g. no change in underlying state). But the effect on prices $g(p)$ is 40% for losses of corporate bonds in the range of 30%. With a constant multiplier \mathcal{M} , this implies over 5 times as many bonds being purchased in these bad states of the world compared to the baseline case of no change. Hence, there is a very strong expectation for significantly more intervention if the bond market were to deteriorate.

Gauging overall quantitative magnitudes on the dollar amount of bonds purchased by the Fed in these states requires taking a stronger stand on the multiplier \mathcal{M} . A low choice for this number (elastic bond market) implies very large purchases to achieve such dramatic price effects. We proceed by (1) using estimates for \mathcal{M} elsewhere in the literature, (2) using the actual amount of bonds purchased and the realized state three months later to infer \mathcal{M} . For (2) we note that in June, corresponding to three months after the announcements were made, LQD was at a value of 120 relative to the baseline of 100, while Fed purchases of bonds and bond ETFs were only about \$50 billion. Since $g(p)$

in this state is about 3%, this suggests a purchase of \$50 billion had a 3% effect on bond prices. The total size of the corporate bond market is about \$ 7 trillion, thus implying a multiplier $\mathcal{M} = g(p)/\text{purchase}(p) = 3\%/(50 \text{ billion}/7 \text{ trillion}) \approx 4.2$. This implies a large multiplier, but note that using this number then implies smaller required purchases to achieve large price effects in the left tail. Applying this number to the state where the no-policy price fell by 30% would imply purchases of close to \$650 billion in these states.

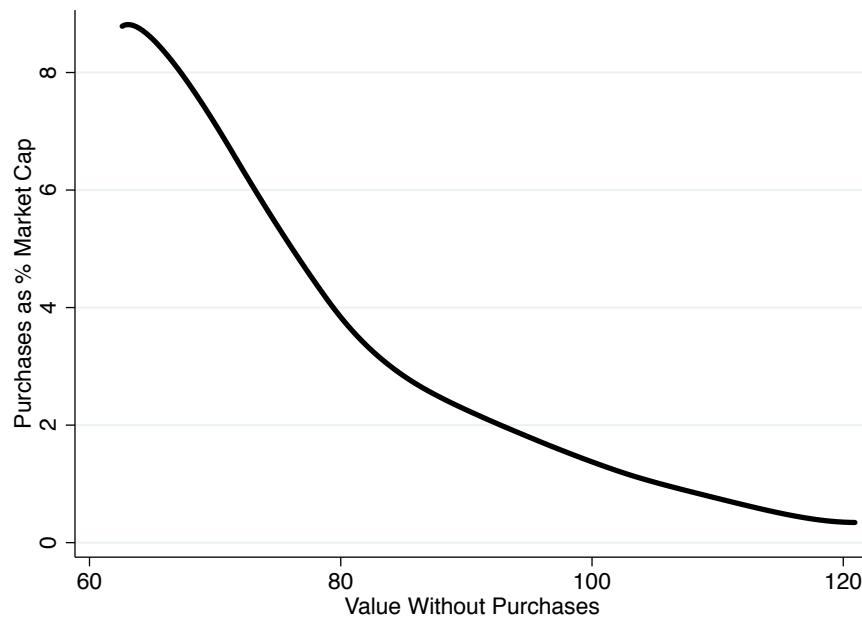


Figure 6: Implied Purchases Based on Multiplier.

This figure plots implied purchases (in billions of USD) as a function of the future distribution of prices before the intervention (normalized to be 100 before the announcement). The multiplier is chosen to match the observed price increase in the state where the Fed actually made purchases of \$50 billion.

We note that this value of the multiplier \mathcal{M} matches fairly well to estimates taken elsewhere in the literature. Gabaix and Koijen (2021) estimate $\mathcal{M} \approx 5$ in the stock market. If we use this value instead, it requires roughly similar, but slightly lower, implied purchases in dollar terms. That is, in the state where bond prices fell by 30% purchases would need to be around \$600 billion which would be close to 7% of the total corporate bond

market and a dramatic increase in the size of the Fed programs. Figure 6 plots implied purchases in billions of dollars using this multiplier.

Calibrating to smaller values for \mathcal{M} (a more elastic bond market) would imply even larger numbers. For example, $\mathcal{M} = 1$ would imply purchases close to 40% of the total corporate bond market in bad states of the world which seems implausible. We acknowledge that a specific value for \mathcal{M} is potentially controversial. But by providing the numbers of the total response in various states, readers can easily infer how much conditional purchases they would need to believe for a given multiplier. Given all the evidence we have provided, the most natural takeaway is for a highly inelastic bond market combined with a commitment to dramatically increase the size of the program conditional on bad realizations for corporate bonds.

3.4 In Which States was the Fed Expected to Buy?

We have outlined conditional purchases by the Fed as a function of the underlying price of the corporate bond market. However, so far, we have taken no stand on whether purchases in low price states are due to deterioration in credit risk, high risk-free interest rates, or high dislocations of corporate bond prices from fundamentals.

We are able to shed light on this by using options on Treasuries (to capture shifts in risk-free interest rates) and options on CDS indices (to capture shifts in underlying credit risk). It is straightforward to map the price of a corporate bond into the price of an equivalent duration Treasury bond, a credit risk component (captured by prices of CDS), and a component which we call “dislocations” affected by moves in the CDS bond basis. We denote the synthetic corporate bond as the bond price implied by the Treasury yield curve and corresponding CDS prices. We can recover the distribution of this synthetic bond by using the option prices for Treasuries and the investment grade CDX index by assuming a correlation between the CDX and Treasuries equal to the historical average

and using copula functions following Haugh (2016). Our main finding is that the intervention massively reduces the basis risk – the risk that the cash-synthetic arbitrage would widen substantially.

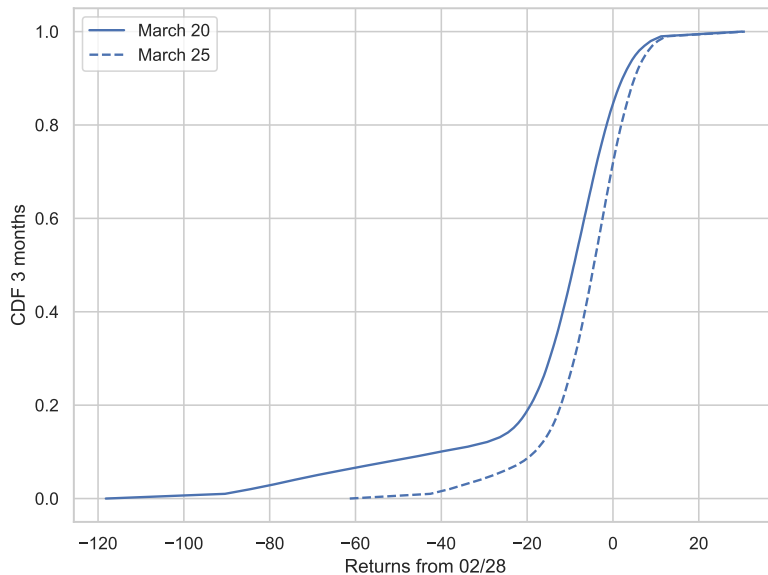


Figure 7: Distribution of the CDS bond basis.

This figure plots the CDF of corporate bond prices minus the price of a synthetic corporate bond constructed from CDS and Treasury options.

Figure 7 plots the distribution of the basis, defined as the dislocation between the cash market and the synthetic. Specifically, let $basis = p_{cash} - p_{Treas} - p_{CDS}$ where p_{cash} is the price of the corporate bond ETF, p_{Treas} is the price of a duration matched Treasury, and p_{CDS} is the price of the credit risk component of investment grade bonds inferred from CDS prices. The main finding highlighted in the Figure is that the intervention dramatically shifts the tail of the basis distribution, while also shifting the whole distribution to the right (shrinking the overall basis). This is consistent with stronger Fed intervention in states where the corporate bond market is highly dislocated so that a large basis opens

up. The 10th percentile of the distribution shifts from a -60% to -20% drop in corporate bond prices relative to the synthetic price of the corporate bond.

3.5 Robustness

Importantly, we contrast the effects on investment grade bonds with those for high yield, using options on the HYG ETF (the exact counterpart to the LQD ETF that captures investment grade). The upper right panel shows that, over a one-day window, the overall returns for high yield are actually slightly negative. If anything, the pattern is also upward sloping, meaning less price support at low prices and vice versa. These results are useful for two reasons. First, they suggest that the announcement didn't coincide with other macroeconomic news affecting corporate bond markets, since the effects are strongly concentrated in investment-grade bonds which were the target of the purchases. Second, and more importantly, they speak to the issue of changes in the pricing kernel partially driving our results.

These results cut strongly against a broader change in pricing kernel or price of risk view to understand our results. Specifically, since we work with risk-neutral distributions, a concern is whether our results reflect implicit promises or a change in the pricing kernel (e.g., of a representative agent) that dramatically lowers the broad price of risk for bad outcomes. A lowering of the price of risk would show up in high-yield bonds as well, which we do not see over this short window.

We next show robustness to using a longer window in our event study. Our main results use one day which tightens identification. However, it could also be reasonable to allow more time for markets to react at the cost of tighter identification since a longer period means that other shocks could be affecting markets. The lower left panel of Figure 8 shows our results are similar if we restrict announcement effects to a one-day window. While magnitudes are slightly larger compared to the results in the one-day window (re-

produced in the upper left panel), the asymmetric effect is similar. In our earlier analysis, we attributed about 53% of the announcement effect to additional promises in the left tail. Using a three-day window this number falls to about 40% because the right tail remains elevated. However, the dollar value from additional left tail promises increases from about \$250-300 billion using the one-day window to about \$400 billion when we expand to the three-day window. This comes from the overall return on corporate bonds being larger over three days compared to one day. This shows our choice of event-window size doesn't have a large effect on these results.

Over a three-day window, there is some evidence of asymmetry in high-yield (lower right panel). However, comparing the investment grade, the magnitudes are about half as large. This is the opposite of what we would expect from a price of risk view, based on the fact that high yield has a much higher beta compared to investment grade. That is, a lowering of the price of risk will boost the value of the riskiest claims (high-yield) compared to safer claims (investment-grade). These results also help control for information effects that might be revealed from the Fed announcement about the macroeconomy (Nakamura and Steinsson, 2018).

3.6 Evidence from April 9th High-Yield Announcement

The announcement of corporate bond purchases on March 23rd focused on investment grade bonds. However, the Fed made an additional announcement on April 9th, 2020, that expanded the facilities to include high-yield bonds. If this announcement contains implicit promises, we would expect them to show up particularly in high-yield bonds. Figure 9 plots the price support from this announcement following our same methodology applied to options on the high-yield ETF (HYG). We see the same effects of asymmetry: price support is very high for low prices, peaking at over 10%, but is much lower at around 5% for higher levels of prices. This provides strong support for implicit promises

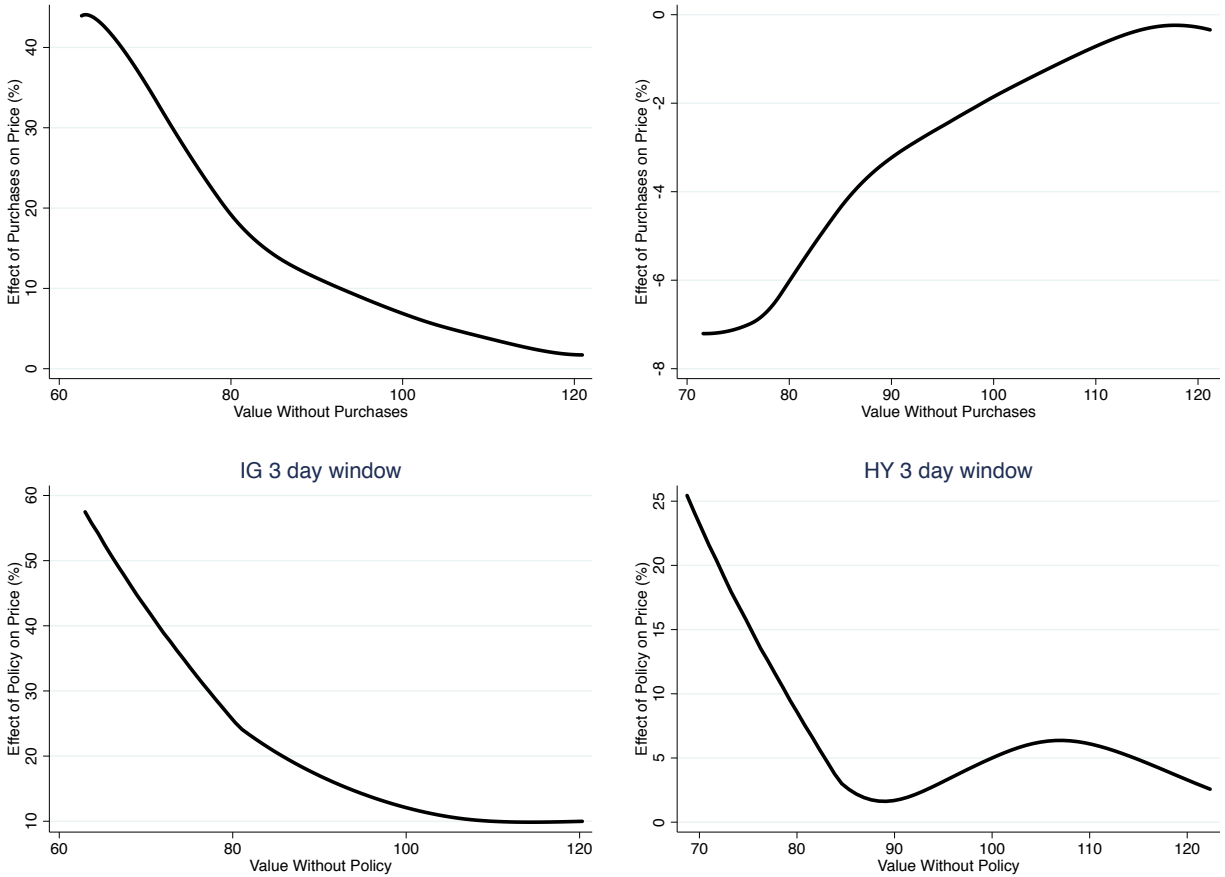


Figure 8: Investment Grade vs High Yield.

This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

boosting the value of high-yield bonds. In contrast, investment-grade is now much flatter, consistent with this announcement not reflecting any additional promises to investment-grade.

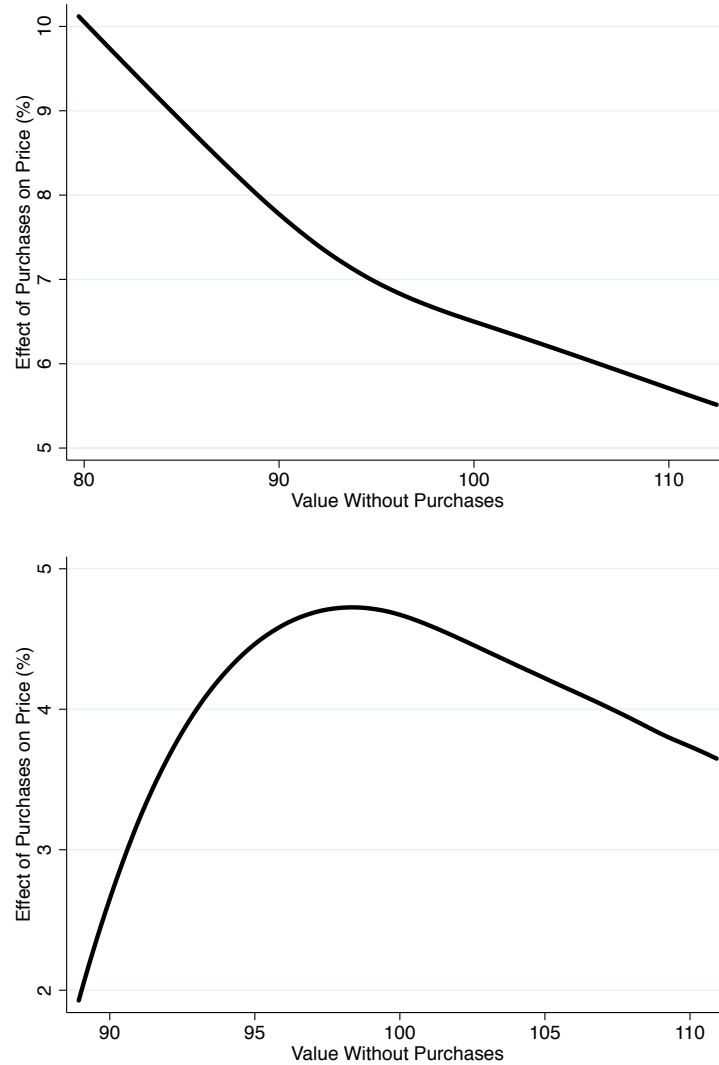


Figure 9: High-Yield Announcement, April 9th, 2020. Effect on prices.

This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

4. Conditional Purchases and the Long-run Effects in Bond Markets

We now turn to the long-run effects of the Fed intervention by studying the behavior of corporate bond tail risk after the programs are implemented. While we have already shown this tail risk falls when the Fed initially announces it will intervene, we now assess whether tail risk is less sensitive to economic conditions going forward. These longer-term effects are not easily captured in our earlier framework which focuses on conditional promises over shorter maturities at which we have option data. We take a simple approach to address this. First, we compute a tail risk index for corporate bonds using the slope of the implied volatility curve. We take implied volatility for delta of 90 minus the implied volatility for delta of 10 as a measure of tail risk. We then take the same tail risk measure but using S&P500 index options and options on the investment grade CDX index.

Table 2 shows that tail risk changed after the announcements and that corporate bond markets appeared less sensitive to tail risk in equity or CDS markets. Specifically, we regress the corporate bond tail risk on tail risk in equities and CDS markets using daily data from 2010 onward (for CDS, we only have data from 2015 onward). We then include a “post” dummy interaction term for the period after April 9th, 2020 when the Fed had already announced the expansion of its corporate bond facilities. Notably, in the period prior to this, corporate bond tail risk and equity market tail risk co-move strongly so that tail risk in corporate bonds was highly sensitive to tail risk in equity markets. After the interventions occurred, this sensitivity dropped dramatically. The sum of the two coefficients represents the total sensitivity in the post period and actually goes slightly negative. But the main effect we focus on is that the post interaction term is strongly negative and statistically significant so that corporate bond tail risk becomes much less

Table 2: Long Term Effects on Corporate Bond Tail Risk

This table measures the sensitivity of tail risk in corporate bond markets to tail risk in the stock market (using S&P500 index options) and CDS market (using options on the investment grade CDX index) in daily data from 2010-2021. The dummy “post” equal 1 after April 9th, 2020, the dummy “covid” equals 1 from February 1st, 2020 to April 9th, 2020. Interaction effects capture whether this sensitivity is lower after Fed interventions. Robust standard errors given in parentheses.

	(1) $Tail_t^{CorpBond}$	(2) $Tail_t^{CorpBond}$	(3) $Tail_t^{CorpBond}$	(4) $Tail_t^{CorpBond}$
$Tail_t^{SP500}$	0.59*** (0.05)	0.43*** (0.02)		
$Tail_t^{SP500} \times post$	-0.78*** (0.07)	-0.63*** (0.05)		
$Tail_t^{SP500} \times covid$		0.68*** (0.15)		
$Tail_t^{CDS}$			0.27*** (0.04)	0.14*** (0.02)
$Tail_t^{CDS} \times post$			-0.37*** (0.04)	-0.24*** (0.02)
$Tail_t^{CDS} \times covid$				0.90*** (0.16)
$post$	0.16*** (0.01)	0.14*** (0.01)	-0.06*** (0.01)	-0.02* (0.01)
$covid$		-0.12*** (0.03)		0.35*** (0.06)
Constant	-0.04*** (0.01)	-0.02*** (0.00)	0.11*** (0.01)	0.06*** (0.01)
Observations	2,769	2,769	1,510	1,510
R-squared	0.25	0.29	0.26	0.44

sensitive to broader tail risk in the economy. This is consistent with the view that there would be future further interventions to support corporate bond prices if there were a crash in financial markets. We find similar results using the CDS index in place of the stock market as a gauge of overall economic tail risk. Finally, the results in the pre-period are not driven by extreme behavior during the acute phase of COVID on financial markets where all tail risk measures spike. To show this, we add a COVID dummy, equal to 1 for the period of February 1st, 2020 to April 9th, 2020. Including an interaction with this dummy doesn’t change our conclusions, and in this case the non-interacted coefficient measures the sensitivity of corporate bond tail risk to other tail risk excluding the COVID episode (e.g., from 2010-early 2020).

Panel A of Table 3 does a similar exercise but uses corporate bond excess returns in place of tail risk, and gauges the sensitivity of these returns to changes in the VIX. Excess returns are taken as the daily return on LQD over the daily return on TLT (an ETF that tracks Treasuries). Before interventions, corporate bond returns were very sensitive to changes in the VIX: higher VIX was associated with lower returns. After the intervention returns became about half as sensitive to changes in the VIX. Again, this is suggestive of longer term effects of interventions on expectations. If the Fed is expected to intervene in bad states then returns become less sensitive to factors that adversely affect corporate bond prices. These results also related to Collin-Dufresne et al. (2021) who use a structural model to assess credit index options compared to equity options. They find a divergence between the two during COVID, consistent with our reduced form evidence here.

Panel B of Table 3 instead uses monthly data on option-based pseudo credit spreads from Culp et al. (2018). These are credit spreads directly implied by equity options and Treasuries. We use the two year maturity investment grade pseudo credit spread from Culp et al. (2018) (see [The Credit Risk Lab](#)). We then compare this to the Bank of America investment grade option adjusted spread index for maturities between one and three years taken from Fred. We take changes in actual credit spreads and regress them on changes in pseudo spreads. In the period after interventions, credit spread changes are far less sensitive to changes in option-based credit spreads, implying credit risk reflected in equity options is not reflected in actual investment grade credit spreads after interventions. A drawback of this data is that it is only available monthly, though an advantage is that the measure corresponds exactly to how credit spreads should behave based on the risk reflected in equity markets.

We plot the difference in actual credit spreads vs pseudo spreads in Figure 10. From 2010-2020 the two spreads track each other quite well. During early 2020 when the COVID-19 crisis hit, actual spreads for investment grade spiked well beyond those implied by equity market options. This is consistent with investment grade bond prices

Table 3: Long Term Effects on Corporate Bond Prices

Panel A measures the sensitivity of daily corporate bond excess returns to daily changes in the VIX. Panel B measures the sensitivity of monthly changes in corporate bond spreads to changes in pseudo bond spreads implied by equity options from Culp et al. (2018). The dummy “post” equal 1 after April 9th, 2020, the dummy “covid” equals 1 from February 1st, 2020 to April 9th, 2020. Interaction effects capture whether this sensitivity is lower after Fed interventions. Robust standard errors given in parentheses.

Panel A: Corp Bond Returns		
	(1) $r_t^{CorpBond,e}$	(2) $r_t^{CorpBond,e}$
ΔVIX_t	-0.21*** (0.02)	-0.20*** (0.02)
$\Delta VIX_t \times post$	0.10*** (0.03)	0.08*** (0.03)
$\Delta VIX_t \times covid$		-0.05 (0.04)
<i>post</i>	0.04 (0.03)	0.03 (0.03)
<i>covid</i>		-0.12 (0.34)
Constant	-0.01 (0.01)	-0.01 (0.01)
Observations	2,987	2,987
R-squared	0.26	0.26
Panel B: Credit Spreads and Option-Based Pseudo Spreads		
	(1) $\Delta spread_t$	(2) $\Delta spread_t$
$\Delta pseudo_t$	0.41** (0.19)	0.16*** (0.05)
$\Delta pseudo_t \times post$	-0.62*** (0.23)	-0.37*** (0.13)
$\Delta pseudo_t \times covid$		1.74*** (0.38)
<i>post</i>	-0.11** (0.05)	-0.10** (0.05)
<i>covid</i>		0.08 (0.23)
Constant	0.00 (0.02)	-0.01 (0.01)
Observations	135	135
R-squared	0.22	0.69

becoming abnormally depressed in this episode. However, following the Fed’s intervention and the recovery, investment-grade spreads became quite low, and in fact reached their lowest point at any time over to 2010-2020 window. In contrast, equity markets still featured substantial volatility, implying higher than usual default risk. This keeps pseudo

spreads elevated even after the Fed intervention. This large gap is consistent with a belief about future interventions into the investment-grade bond market in the case of a crash.

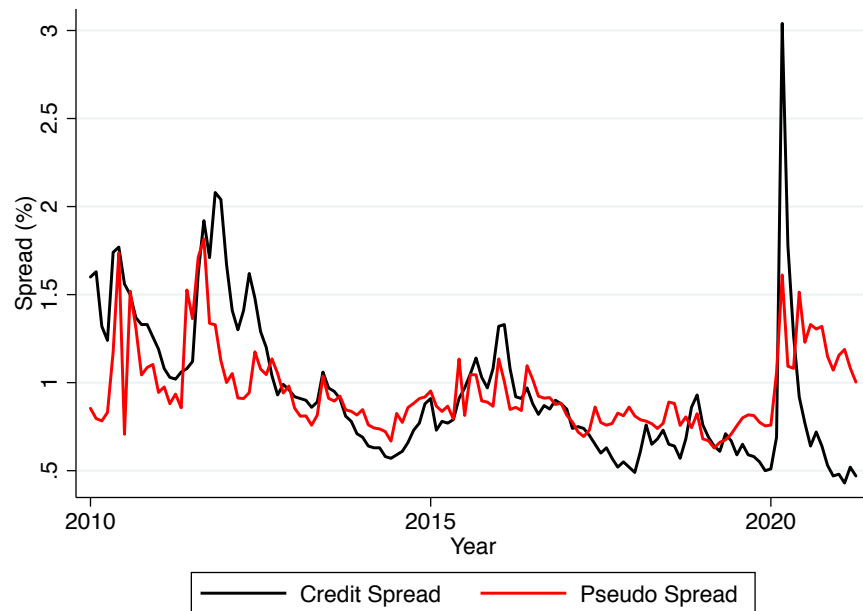


Figure 10: Spreads vs pseudo spreads.

This figure plots actual credit spreads vs pseudo spreads from Culp et al. (2018).

5. Additional Evidence from Other Announcements

Here we provide select additional evidence of key policy announcements to study conditional promises. Our list is non-exhaustive but seeks to illustrate both the role of promises and the uses of our methodology to study announcements. We are also limited to events where we have option data on relevant asset prices for the policy in question. We study equity injections in the US financial sector during 2008, an announcement of large asset purchases by the Bank of Japan in 2013, and various dates associated with the implementation and unwinding of quantitative easing (QE) in the United States from 2008-2013.

Table 4: Summary of Additional Events

This table applies our methodology to many other announcement events. We compute the fraction of each announcement return explained by promises (additional price support below the median). The specific events, methodology, and financial instruments are provided in the subsections below.

Event	Fraction Explained by Promises
High-Yield April 9th 2020	9%
Oct 13th 2008 (Paulson Plan)	37%
BoJ Purchase Speech	11%
<i>US Quantitative Easing Events:</i>	
Nov 25th 2008	2%
Dec 16th	14%
March 19th	14%
June 19th, 2013 (Tantrum)	9%
<i>ECB Announcements:</i>	
May 10, 2010	24%
Aug 7, 2011	26%
July 26, 2012	9%
Aug 2, 2012	39%
Sep 6, 2012	17%
Average	18%

We summarize our results for these announcements in Table 4. The subsections below discuss each set of announcement dates.

5.1 2008 Financial Sector Bailout

Figure 11 studies the “Paulson Gift” (Veronesi and Zingales (2010)) on October 10th 2008 which announced large equity injections to the banking sector, as well as guarantees on various forms of bank debt, in an effort to “restore confidence in the financial system.” Veronesi and Zingales (2010) find large effects of this intervention on bank equity and bank debt.⁷ To analyze this case, we use option prices of the Financial Select Sector SPDR Fund which focuses exposure on the financial sector. Our method reaches similar conclusions but maps out the price support of this policy as a function of the underlying equity

⁷See also Kelly et al. (2016a) who use options markets to evaluate government guarantees on the financial sector.

value. For example, we see price support of 40% if the equity of the financial sector were to fall by 50% and price support less than 10% if the equity of the financial sector increased by 50%. This provides strong evidence of large conditional promises to the financial sector. This effect was likely a goal of the policy itself – by providing strong conditional promises that the US government would do whatever it takes to keep the financial sector solvent, both debt and equity prices rose substantially.

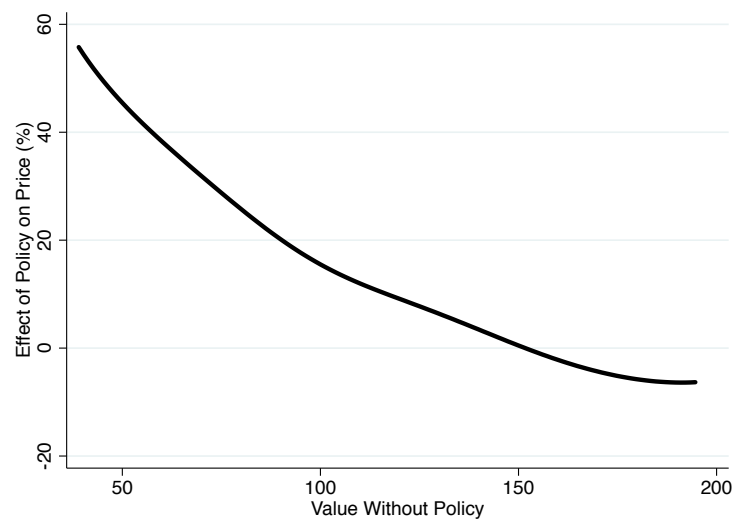


Figure 11: “Paulson’s Gift:” Banking Sector Bailout in October 2008. Effect on prices. This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

5.2 Bank of Japan Asset Purchases, 2013

Figure 12 studies Japan on April 4th 2013 after a speech given by Haruhiko Kuroda, the head of the central bank, in which he outlined large purchases of government bonds and equities to drive up asset prices and inflation. Charoenwong et al. (2021) systematically study equity purchases by the Bank of Japan and find that they increase equity prices. We use three-month options on the Nikkei index in a three day window around the an-

nouncement. This episode points to conditional promises that provide price support of around 10% for adverse states vs about 6% for good states. The conditional promises were very much intended from the announcement by Kuroda as the speech outlined a “whatever it takes” type approach to support prices and achieve growth.

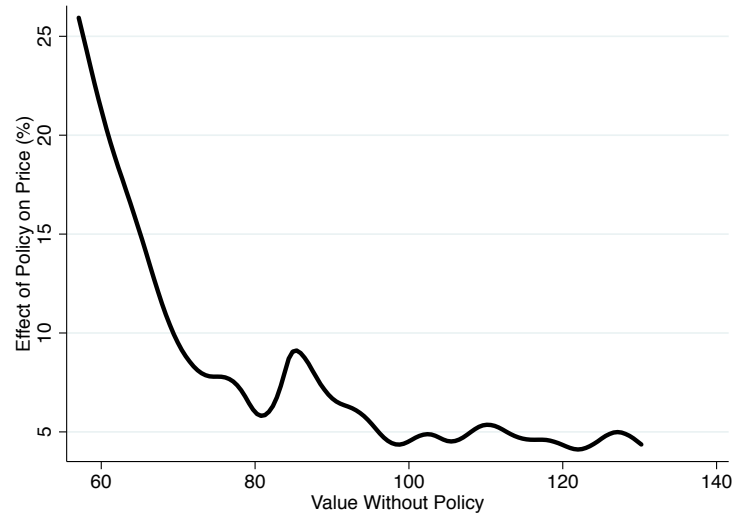


Figure 12: Announcement of Purchases by Bank of Japan in April, 2013. Effect on prices.

This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

5.3 Quantitative Easing, 2008-2013

Figure 13 looks at US quantitative easing (QE) announcements on four separate dates using three-month options on the 10 year Treasury Note futures contract. Treasuries are likely an imperfect asset to study the conditional impact of QE, because future conditional purchases may depend on many other state variables besides the price of the 10 year Treasury, which our measurement will not capture. We focus our analysis on four QE announcements. The first is the initial announcement of large scale asset purchases (LSAPs)

on November 25th, 2008. The second is the FOMC statement on December 16th, 2008. The third is the FOMC statement on March 18th, 2009. The last is the “Taper Tantrum” on June 19th, 2013. The first three policy announcements see an increase in Treasury prices (fall in yields) and the price support function is strongly downward sloping. All of these announcements are associated with increased purchases. The magnitudes are fairly similar with about 4-6% price support in cases where prices fall. The downward slope indicates a market belief for more purchases or price support should prices fall by 15% or more, roughly 3% price support at current prices, and roughly 2-3% for a 10% increase in prices. The last announcement (lower right panel) is the “Taper Tantrum” where it was announced purchases would decline. We see an overall decline in prices (sharp increase in yields) with an upward, rather than downward slope. This announcement is thus associated with not only a tapering of purchases but an associated decline in promises of future purchases. Overall, the events of QE are associated with conditional promises, though the effects are more mild than other events. These results speak to a broader literature that estimates the channels through which QE operates by using event studies (e.g., Vissing-Jorgensen and Krishnamurthy (2011)). Our results suggest that part of the large price response on announcement comes from promises.

5.4 ECB, 2010-2012

Figure 14 looks at announcements of asset purchases by the European Central Bank. We use five announcement dates found by Krishnamurthy et al. (2018) to have substantial asset price effects. Ideally, for these announcements, we would like to have options on sovereign bonds for countries in the Eurozone (or options on a sovereign bond index). This is most directly where the asset purchase announcements were aimed. However, Krishnamurthy et al. (2018) show a very broad asset price response to the announcements. In fact, they estimate most of the increase in total value from these announcements came

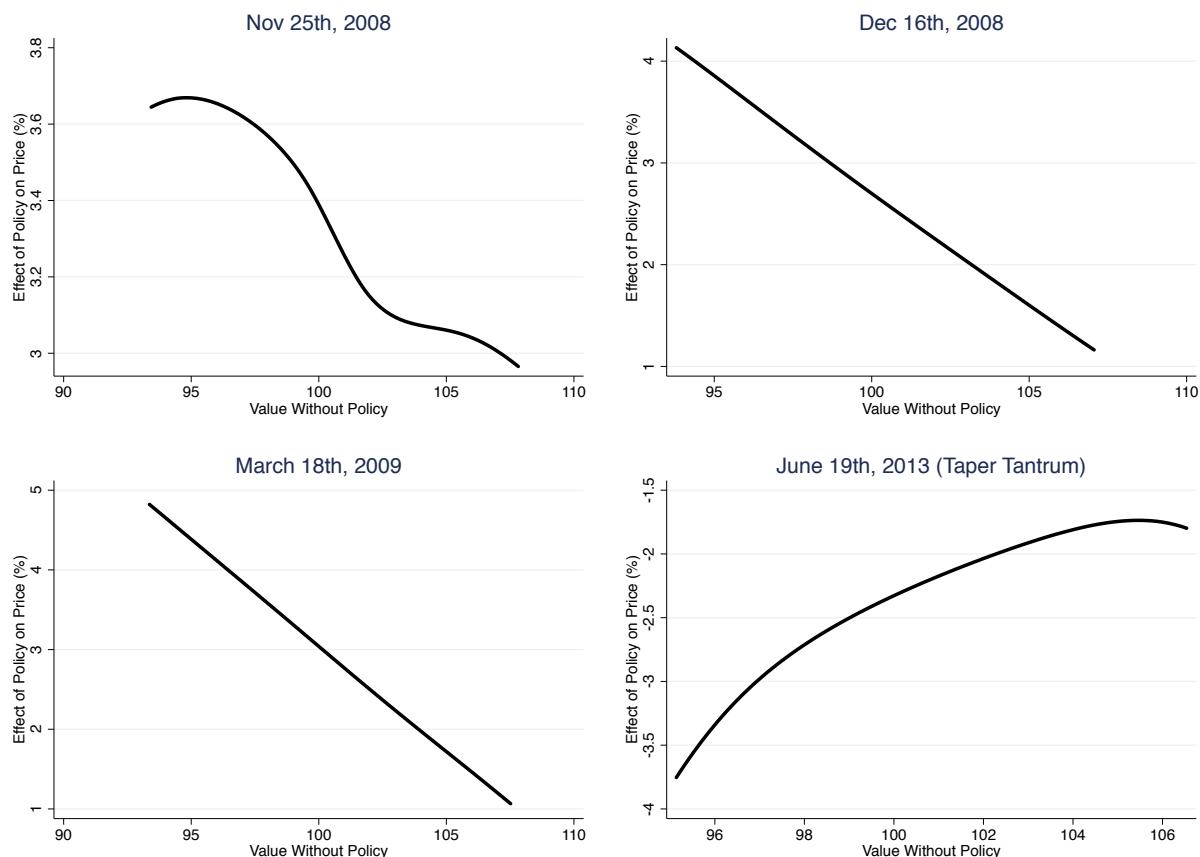


Figure 13: Quantitative Easing. Effect on prices.

This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

from the reaction of stock markets. This is likely due to different channels than what we saw for investment-grade bonds during COVID. One interpretation is for Europe there was a feared “doom loop” that sovereigns would be unable to pay creditors or roll over debt and this would lead to substantial declines in economic activity. This could be coming from higher taxation, strong fiscal adjustments, or losses born by holders of sovereign bonds which included a large portion of the banking sector in Europe. These losses could lead to a substantial credit crunch that would result in a decline in economic activity. The fact that the asset price effects were broad also means we can use options on the Euro

Stoxx index to assess promises. This assumes that any conditional purchases will be correlated with the stock index value and that the stock index will respond to purchases in those states.

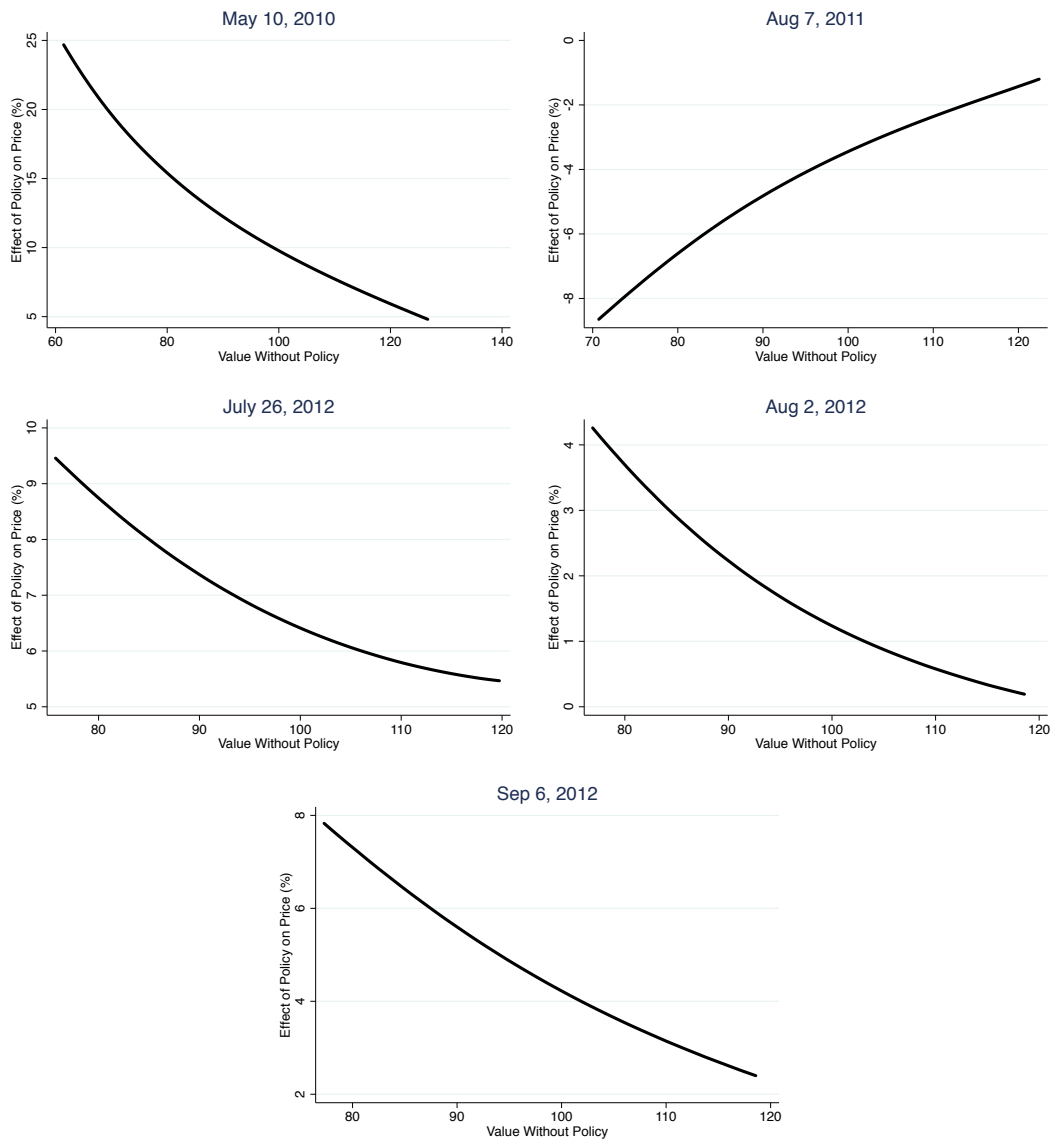


Figure 14: ECB Announcements. Effect on prices.

This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

Figure 14 shows strong effects of conditional promises from these announcements. Across all the five announcements, promises explain about 20% of the overall reaction on average. This is likely not surprising as part of the goal of the announcements was to promise to do “whatever it takes.” While the exact quantities implied by the promises were vague, the intention to commit to promises in this case was explicit.

5.5 FOMC meetings

We next explore whether news about promises is present around FOMC meetings by studying the Fed put. Cieslak and Vissing-Jorgensen (2021) show that low stock returns predict accommodating policy by the Federal Reserve, supportive of the notion of the Fed put. Here, we ask whether news about the Fed put (how much the Fed is willing to use policy to boost declines in stock prices) is revealed around FOMC meetings and press releases. The announcements we studied earlier had clear directional implications: an announcement to support prices that includes promises, such as an asset purchase program, should both boost returns and feature asymmetric price support in the left tail if promises are present. Assessing whether news about promises are revealed through Fed communication more broadly is more challenging because a given statement could increase or decrease expectations of promises. Our goal is to overcome this challenge by using the overall return response to gauge if news was good or bad, and then to Fed communication is more associated with movements in the tail of distribution compared to normal days.

Specifically, for the S&P500 we construct $E[g(p)|p < p_{median}]$ and $E[g(p)|p \geq p_{median}]$ which we refer to as “bottom” and “top,” respectively. Recall that by construction $E[g(p)]$ is equal to the return of the S&P500 each day. Our bottom and top measures assess the daily return coming from the left and right tails of the distributions, respectively. An increase in the Fed put (e.g., news that the Fed will put more weight on the stock market

Table 5: FOMC Meetings

This table studies promises around FOMC meetings. The dependent variable is the bottom minus top return on the S&P500 index. Standard errors computed using Newey-West with 5 lags.

	(1) (<i>Bottom – Top</i>) _{<i>t</i>}
<i>return_t</i>	44.12*** (1.39)
<i>FOMC_t</i>	0.01 (0.03)
<i>FOMC_t × return_t</i>	-1.84 (3.69)
<i>press_t</i>	0.10** (0.05)
<i>press_t × return_t</i>	9.83*** (4.62)
Constant	-0.03 (0.01)
<i>N</i>	4,652
Adj. <i>R</i> ²	0.81

in setting policy) should lead to two things. First, the return on stocks should be high as this represents good news for the stock market. Second, the bulk of this return should be driven by the bottom of the distribution – where the Fed is providing support through the put – relative to the top. We view press announcements – regularly scheduled FOMC meetings that are followed by a press conference – as being most informative for news about Fed policy and stock prices. Cieslak and Vissing-Jorgensen (2021) conduct a textual analysis of Fed transcripts and find strong evidence of discussion about the stock market, consistent with this view.

Table 5 regresses the “bottom minus top” return on the daily stock return and includes interactions for both regularly scheduled FOMC meetings as well as scheduled meetings that include a press announcement. Stock returns are strongly positively related to bottom minus top returns. The coefficients imply that a 1% daily rise in stock returns is associated with a 0.44% increase in the bottom minus top return. This is not surprising as news about crash risk more generally plays a role in driving returns. The interaction term of returns

and FOMC press meetings shows there is a substantially stronger relation on FOMC press release days. The magnitude here means a 1% increase in returns on FOMC press release days is associated with a 10 bps increase in the bottom minus top return so the relation is 22% stronger than on normal days. We interpret this as suggestive evidence that news about left tail risk is especially informative about returns around FOMC press release days, consistent with these days revealing news that shifts market expectations of the Fed put. We also only find significant results on scheduled FOMC meetings that include a press announcement (roughly 20% of meetings), consistent with these days providing communication about the Fed put.

6. Conclusion

We provide a framework and methodology to evaluate conditional policy promises, which we apply to several policy announcements. We find a large role for conditional promises that indicate more intervention if conditions worsen. In our main empirical setting, the announcement of corporate bond purchases during the COVID-19 crisis, we find evidence markets expected five times more price support in crash scenarios relative to the median case. This implicit promise to significantly expand the size of the intervention in bad states explains a large share of the market response to the announcement. Our methodology focuses on promises primarily at shorter horizons as this is the maturity that relevant options expire. We also find longer-run effects of conditional promises in corporate bond markets after announcements of corporate bond purchases, consistent with expectations of future intervention if a future crash should occur. We extend our analysis to several other policy announcements as well and find support for conditional promises. We view future research that measures the role of promises as being especially useful to shed light on the merits of promises as a policy tool.

References

- Barracclough, Kathryn, David T. Robinson, Tom Smith, and Robert E. Whaley, 2013, Using Option Prices to Infer Overpayments and Synergies in M&A Transactions, *The Review of Financial Studies* 26, 695–722.
- Boyarchenko, Nina, Anna Kovner, and Or Shachar, 2020, It's what you say and what you buy: A holistic evaluation of the corporate credit facilities, Working paper, Federal Reserve Bank of New York.
- Breeden, Douglas T, and Robert H Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of business* 621–651.
- Charoenwong, Ben, Randall Morck, and Yupana Wiwattanakantang, 2021, Bank of japan equity purchases: The (non-) effects of extreme quantitative easing, *Review of Finance* 25, 713–743.
- Cieslak, Anna, and Annette Vissing-Jorgensen, 2021, The economics of the fed put, *The Review of Financial Studies* 34, 4045–4089.
- Collin-Dufresne, Pierre, Benjamin Junge, and Anders B Trolle, 2021, How integrated are credit and equity markets? evidence from index options, *Evidence From Index Options (February 16, 2021). Swiss Finance Institute Research Paper* .
- Culp, Christopher L, Yoshio Nozawa, and Pietro Veronesi, 2018, Option-based credit spreads, *American Economic Review* 108, 454–88.
- D'Amico, Stefania, Vamsidhar Kurakula, and Stephen Lee, 2020, Impacts of the fed corporate credit facilities through the lenses of etfs and cdx, *Available at SSRN 3604744* .
- Drechsler, Itamar, Alexi Savov, and Philipp Schnabl, 2018, A model of monetary policy and risk premia, *The Journal of Finance* 73, 317–373.
- D'Amico, Stefania, and Thomas B King, 2013, Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply, *Journal of Financial Economics* 108, 425–448.
- Figlewski, Stephen, 2010, Estimating the implied risk-neutral density for the us market portfolio, *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle, Edited by Tim Bollerslev, Jeffrey Russell, and Mark Watson* .
- Gabaix, Xavier, and Ralph SJ Koijen, 2021, In search of the origins of financial fluctuations: The inelastic markets hypothesis, Working paper, National Bureau of Economic Research.

- Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack, 2018, The financial market effects of the federal reserve's large-scale asset purchases, *24th Issue (Mar 2011) of the International Journal of Central Banking* .
- Gilchrist, Simon, Bin Wei, Vivian Z Yue, and Egon Zakrajšek, 2020, The Fed takes on corporate credit risk: An analysis of the efficacy of the SMCCF, Working paper, NBER.
- Grinblatt, Mark, and Kam-Ming Wan, 2020, State pricing, effectively complete markets, and corporate finance, *Journal of Corporate Finance* 60, 101542.
- Haddad, Valentin, Alan Moreira, and Tyler Muir, 2021, When selling becomes viral: Disruptions in debt markets in the covid-19 crisis and the fed's response, *The Review of Financial Studies* 34, 5309–5351.
- Hamilton, James D, and Jing Cynthia Wu, 2012, The effectiveness of alternative monetary policy tools in a zero lower bound environment, *Journal of Money, Credit and Banking* 44, 3–46.
- Hattori, Masazumi, Andreas Schrimpf, and Vladyslav Sushko, 2016, The response of tail risk perceptions to unconventional monetary policy, *American Economic Journal: Macroeconomics* 8, 111–36.
- Haugh, Martin, 2016, An introduction to copulas, *IEOR E4602: quantitative risk management. Lecture notes. Columbia University* .
- He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *The American Economic Review* 103, 732–770.
- Hesse, Henning, Boris Hofmann, and James Michael Weber, 2018, The macroeconomic effects of asset purchases revisited, *Journal of Macroeconomics* 58, 115–138.
- Kargar, Mahyar, Benjamin Lester, David Lindsay, Shuo Liu, Pierre-Olivier Weill, and Diego Zúñiga, 2020, Corporate bond liquidity during the covid-19 crisis, Working paper, working paper.
- Kelly, Bryan, Hanno Lustig, and Stijn Van Nieuwerburgh, 2016a, Too-systemic-to-fail: What option markets imply about sector-wide government guarantees, *American Economic Review* 106, 1278–1319.
- Kelly, Bryan, L'uboš Pástor, and Pietro Veronesi, 2016b, The price of political uncertainty: Theory and evidence from the option market, *The Journal of Finance* 71, 2417–2480.
- Krishnamurthy, Arvind, Stefan Nagel, and Annette Vissing-Jorgensen, 2018, Ecb policies involving government bond purchases: Impact and channels, *Review of Finance* 22, 1–44.

- Malz, Allan M, 2014, A simple and reliable way to compute option-based risk-neutral distributions, *FRB of New York Staff Report* .
- Meaning, Jack, and Feng Zhu, 2011, The impact of recent central bank asset purchase programmes, *BIS Quarterly Review*, December .
- Moreira, Alan, and Alexi Savov, 2017, The macroeconomics of shadow banking, *Journal of Finance* .
- Nakamura, Emi, and Jón Steinsson, 2018, High-frequency identification of monetary non-neutrality: the information effect, *The Quarterly Journal of Economics* 133, 1283–1330.
- O'Hara, Maureen, and Xing Alex Zhou, 2020, Anatomy of a liquidity crisis: Corporate bonds in the covid-19 crisis, *Available at SSRN 3615155* .
- Vayanos, Dimitri, and Jean-Luc Vila, 2021, A preferred-habitat model of the term structure of interest rates, *Econometrica* 89, 77–112.
- Veronesi, Pietro, and Luigi Zingales, 2010, Paulson's gift, *Journal of Financial Economics* 97, 339–368.
- Vissing-Jorgensen, Annette, and Arvind Krishnamurthy, 2011, The effects of quantitative easing on interest rates: Channels and implications for policy, *Brookings Papers on Economic Activity* 43, 215–287.