Whatever It Takes?
The Impact of Conditional Policy Promises

Valentin Haddad, Alan Moreira, and Tyler Muir*

August 15, 2024

Abstract

At the announcement of a new policy, agents form a view of state-contingent policy actions and impact. We develop a method to estimate this state-contingent perception and implement it for many asset-purchase interventions worldwide. Expectations of larger support in bad states—"policy puts"—explain a large fraction of the announcements' impact. For example, when the Fed introduced purchases of corporate bonds in March 2020, markets expected five times more price support had conditions worsened relative to the median scenario. Perceived promises of additional support in bad states alter asset prices, risk, and the response to future announcements.

*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 are grateful to Viral Acharya, Mike Chernov, Darrel Duffie, Andrea Eisfeldt, Robert Goldstein, Anil Kashyap, Narayana Kocherlakota, Ralph Koijen, Sydney Ludvigson, Emi Nakamura, Pierre-Olivier Weill, Stavros Panageas, Jon Steinsson, Yao Zeng, and three anonymous referees, as well as conference and seminar participants at UCLA, Chicago Booth, LBS, USC, MIT, Stockholm, Cornell, Stanford, University of Utah, University of Florida, Georgia State University, the Federal Reserve Board, New York Fed, Boston Fed, OFR, UT Dallas, Temple, NYU, Stockholm, NBER Monetary Economics, NBER Asset Pricing, AFA, Macro Finance Society, Canadian Derivatives Institute, Jackson Hole, Rochester, UT Austin, OSU, HKU, CUHK, and HKUST for helpful comments and suggestions.

Suppose we are in an economic crisis and the central bank announces asset purchases of $100 billion to support financial markets and the economy. After the announcement, asset prices strongly increase. A large response relative to the amount announced makes it tempting to conclude that asset purchases are a very effective policy tool. But an alternative view is that the strong response is not only driven by the announced quantity, but also by the perception that policymakers might do whatever it takes, in that they will go to much greater lengths to backstop markets if the situation gets worse. Distinguishing these views is important. Under the whatever-it-takes view, the expansion of the central bank's balance sheet may be much larger, the price impact per dollar of purchases much smaller, and the price response to future announced purchases much weaker, compared to the view of a one-time $100 billion asset purchase. These issues are not only important for asset purchases, an increasingly utilized tool for central banks globally, but also arise for many other policy announcements including bank bailouts, fiscal policy, or forward guidance: when a policy is announced, markets do not only learn a single headline number, but form a view of policy actions in many different states of the world.

We propose and implement a method to measure the state-contingent impact of policy actions using option prices that is valid for a wide class of mechanisms. The key idea is to use the change in the entire option-implied distribution of a given asset over the announcement to characterize the impact of policy across many possible states, rather than only the change in price. We find pervasive evidence of larger policy impact (e.g., larger policy interventions) in bad states of the world across many financial stabilization policies including corporate bond purchases during the COVID-19 crisis, U.S. quantitative easing, asset purchases by both the Bank of Japan and European Central Bank, and bank capital injections in the 2008 crisis. When central banks first step into markets, this "policy put" of additional interventions if conditions worsen explains a large share of the markets' response. These perceptions of conditional interventions appear persistent, with evidence that they affect how markets respond to the announcement of subsequent interventions.

Our primary empirical example is the Fed's announcement of corporate bond purchases in March of 2020, which immediately led to a dramatic recovery in corporate bond prices despite a small amount of ultimate purchases. We decompose this recovery in a state-by-state response. In states where corporate bond prices would have decreased 35% without intervention, the intervention raises them by about 40%.[1] However, in states where prices would have otherwise increased 20%, the intervention only increases them by an additional 2%. This striking asymmetric effect has a signature very similar to a put option. Adding up these responses across all potential states, we find that at least 50% of the large price recovery from the announcement came from additional policy in the left tail. This explains why the announcement led to an immediate increase in corporate bond prices of between \$500 billion and \$1 trillion despite ultimate purchases of less than \$15 billion: the market expected a more aggressive intervention if economic conditions had worsened.

How do we infer the state-contingent policy impact? We need to assess how much the policy changes asset prices in each potential future state of the world at a given horizon. Option contracts provide a unique window into these potential states. Intuitively, an out-of-the-money put, which only pays off in low price states, is entirely driven by expectations of policy in bad states. An out-of-the-money call reveals policy when conditions improve. Options on corporate bond ETFs show a stronger price increase for contracts targeting bad states of the world at the announcement. We show how to go further than this qualitative assessment and obtain the entire state-contingent policy.

Specifically, we estimate a "price support function," $g(.)$, that measures the state-by-state response of prices to the policy intervention: in the state that the price were to move to p absent intervention, the intervention will raise it by $g(p)$ percent so that we observe $p(1+g(p))$ in that state. First, the price of options across different strikes reveals the market perception of the distribution of the future price or risk-neutral distribution (Breeden and Litzenberger, 1978). We obtain this distribution using options maturing three months from the policy announcement, roughly when the bond purchases were 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 $g(.)$ 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. We derive its unique solution under the assumption that the policy is order-preserving: the Federal Reserve does not support prices so much in bad states of the world that they exceed prices in better states. While this measure is intuitive, its economic interpretation requires a set of assumptions, which we formalize, discuss, and assess

[1] Prices are still dropping in these states, by $1-(1-0.35)\times 1.4 = 9\%$.

robustness to. One such assumption is a separability property: policy in a state of the world does not affect outcomes in other states.

In the case of the 2020 corporate bond purchases, the price support function we recover in the data is strongly asymmetric and resembles a "policy put:" price support is low and relatively flat in good states of the world but increases to much larger values as we move toward worse states.[2] Our measurement framework is not specific to this event. We construct the conditional price support function for many other policy announcements for which we have relevant option data: equity purchases by the Bank of Japan in 2013, the quantitative easing operations in the US from 2008 to 2013, asset purchases by the European Central Bank from 2010 to 2012, and the financial sector bailout in the US in 2008. We find pervasive evidence of asymmetric price support across all these announcements, albeit with different intensity.

After documenting this asymmetric price support, we turn to the economic channels to interpret our results. First, a natural interpretation is that all price support comes from conditional policy actions: the Fed will buy more in bad states than good states. For example, consider the view that the elasticity of bond prices to bond purchases is constant across states. In this case, the asymmetry that we document implies around 20 times larger purchases if the price had fallen by 30% compared to the realized state, and around 6 times larger purchases compared to the state with no price change. Calibrating the elasticity to the literature implies purchases of about $800 billion in bad states, suggesting that additional announcements would have occurred then. This state-dependent quantity view is consistent with statements made by the chair of the Federal Reserve Jerome Powell in June 2020, for example: "markets are functioning pretty well, so our purchases will be at the bottom end of the range that we have written down."[3] We argue this view is more plausible than the view of a constant small quantity of purchases but state-dependent elasticity – under that view the price impact for each dollar of purchases would have to be around $200 in some states, which is orders of magnitude larger than any estimates elsewhere in the literature.

Our baseline calculation assumes invariance of the short-term pricing kernel to the announcement. That is, while we do not take a stand on the form of the pricing kernel, we assume that it does not change for payoffs at the horizon between the announcement of purchases and actual

[2]See 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.

[3]This quote is part of his statement to the U.S. House of Representatives on June 16, 2020.

purchases. Changes in risk-pricing after actual purchases are made are allowed by our framework and a likely justification for the price impact of purchases.[4] We extend the framework to entertain a response of the short-term pricing kernel to the purchases. We show that under the null of constant price support, investors' pricing kernel generally does not change on announcement for payoffs that occur before purchases are made. At a minimum, our estimates thus reject the null of constant price support. When purchases are state-dependent, the risk of the asset itself can change and this could affect the specialized investors' pricing kernel between announcement and purchases. We show how to adjust our price support function conservatively for this effect, and find that our main conclusions of strongly asymmetric price support still hold. In addition, other assets such as stocks or high-yield bonds do not react to the announcement, cutting against the idea of a broad change in the pricing of risk.

Mechanisms that do not fit our framework could offer alternative explanations of these findings. In particular, theories of multiple equilibria and runs do not satisfy the separability property: out-of-equilibrium beliefs about how policy plays out in a run affect the likelihood of runs on the equilibrium path. Recovering a policy plan from the price support function under such theories would require a specific structural model.

We also show how to draw inference on the states in which the Fed was likely to provide more support, by bringing evidence from options on additional assets. Corporate bond prices can fall because of rising risk-free rates, increases in credit risk or because of disruptions in corporate 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. We find that the intervention worked almost exclusively by shrinking the probability of very high dislocations — states where the gap between the synthetic corporate bond and actual corporate bond prices widened. This pattern is consistent with Powell's emphasis on "markets [...] functioning" when discussing the intervention.

Our results suggest that state-contingent policy is an appealing explanation for the broader finding that announcements of asset purchase programs are associated with large movements in asset prices (Gagnon et al., 2018; Vissing-Jorgensen and Krishnamurthy, 2011; Haddad et al.,

[4]See Appendix Section B. For example, the elasticity of asset prices to purchases could operate through a change in the amount of risk that investors bear once the central bank buys the assets as in Vayanos and Vila (2021).

2021).[5] They also help reconcile these strong initial responses with weak reactions to subsequent interventions. Hesse et al. (2018), Meaning and Zhu (2011), and Bernanke (2020) find that initial stage announcements of asset purchases in the US and Europe have large effects on asset prices but later stage announcements have little to no effect. In our framework, early announcement responses also reflect the value of state-contingent actions to do more if the situation worsens. Thus, if one does not account for the costs of these state-contingent promises, the market response to an early announcement provides an overly rosy view of the cost-benefit of these programs. Conversely, in later stages of asset purchase programs this state-contingent policy has already been reflected in prices. Thus, even a large announcement can appear to have zero effect even if it had exactly the same effectiveness as the early announcements. We expand the set of announcements in this literature and show that our results can make sense of these repeated announcement effects.

Our results speak to macro-finance models that assess the impact of policies to support financial markets 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)). Our findings suggest that state-contingent policy is first order to understanding the effectiveness of policy announcements, and we provide an entire state-contingent price response that such models could target. We highlight that communication about policy in bad states — as opposed to the baseline state only — is particularly important because it plays a sizable role in shaping how markets respond.

Our findings also relate to the literature on forward guidance and central banking communication starting with the seminal work of Gürkaynak et al. (2004). Other examples include Piazzesi (2005), Swanson (2011), Hanson and Stein (2015), and Nakamura and Steinsson (2018). Haddad et al. (2024) study the impact of asset purchase policy rules. Hubert et al. (2024) relate the communicated goal of QE operations by the ECB with their impact. Relatedly, Bianchi et al. (2022) study monetary policy rule changes and their effect on asset prices while Bauer et al. (2022) study perceptions of rule changes implied by surveys. We show that agents immediately form perceptions of a state-contingent plan of actions for policies with no prior track record, and how to measure these perceptions in real time.

Finally, our analysis also connects to broader work using information in options markets to

[5] See also D'Amico and King (2013), Hamilton and Wu (2012).

interpret policy. Kelly et al. (2016b) focus on the price of uncertainty associated with regular political events (e.g., elections). Our work instead focuses on inferring conditional policy from an unexpected event. Relatedly, Kelly et al. (2016a) use options markets to evaluate government guarantees on the financial sector in the 2008 crisis. Kitsul and Wright (2013) and Hilscher et al. (2022) use option prices to assess inflation probabilities. Barraclough et al. (2013) use option prices to inform merger announcements.[6] While we focus on asset purchase policies, we provide a blueprint for using our method in a wide variety of contexts.

1. Measuring State-Contingent Policy

In this section, we introduce a framework for measuring the state-contingent impact of policy announcements. We start by a simple example illustrating how the presence of policy 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 the state-contingent impact of policy. Specifically, we show how to estimate a price support function and provide and discuss a set of precise assumptions sufficient to interpret the price support function as measuring the state-contingent impact of a policy.

1.1 The Effect of Policy Promises on Asset Prices

Consider a stylized example with two dates, 0 and 1. At date 0, the price of an asset is p_0. Under rational expectations, this price is the (risk-neutral) expected value of the date-1 price: $p_0 = E[p_1]$. For simplicity, we assume that the risk-free rate is equal to 0 between date 0 and 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 $\mathcal{M}$, where $\mathcal{M}$ is a constant. For example, the Fed unexpectedly announces that it will purchase a quantity Q of corporate bonds (the asset) at a future date, where Q is denoted as a fraction of the total supply of corporate bonds. In that interpretation, $\mathcal{M}$ reflects the price impact per quantity of

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

asset purchased, or the inverse elasticity of demand for corporate bonds. Specifically, $\mathcal{M}$ is the percentage response in the price of the asset per percentage point of supply purchased. 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 $p_1' = p_1(1 + \mathcal{M}Q)$. Therefore the post-announcement price becomes

$$p_0' = E[p_1](1 + \mathcal{M}Q). \tag{1}$$

In other words, the return at announcement, $(p_0' - p_0)/p_0$, 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.[7]

Conditional promises. When the new policy is announced, the market might (rightfully) infer that the policymaker will take different actions depending on the state of the world. For example, it could be that the market expected the Fed to purchase even larger amounts of corporate bonds were the COVID-19 crisis to deepen. Our notion of promises is what the market expects policy makers do in different states of the world, but we are silent on whether the promises are explicitly made by policymakers or even intentional. For example, Mario Draghi's famous speech explicitly promised to do "whatever it takes" in the midst of the Euro area sovereign debt crisis in 2012, but it is unclear what the market inferred from the announcement. Conversely, announcements of U.S. quantitative easing specified limited quantities of purchases, but it is plausible that market participants made inferences about additional purchases in the future.

To illustrate the impact of conditional promises, assume that the policymaker is ready to do whatever it takes to maintain the price above a lower threshold $\underline{p}$. That is, if we are in a state at date 1 where the price would fall below $\underline{p}$ despite the baseline intervention Q, they will scale up the policy by any additional amount $\widetilde{Q}$ needed to reach $\underline{p}$. This corresponds to $\widetilde{Q} = \max\left(\frac{1}{\mathcal{M}}(\underline{p}/p_1 - 1) - Q, 0\right)$. The price at date 1 becomes $p_1' = p_1(1 + \mathcal{M}\left(Q + \widetilde{Q}\right))$,

[7] For example, Bouveret and Yu (2021) and Chaudhary et al. (2023) estimate multipliers for corporate bonds and Hamilton and Wu (2012), Greenwood and Vayanos (2014), and Busetto et al. (2022) estimate for government bonds.

and the post-announcement price is

$$p_0' = E[p_1](1 + \mathcal{M}Q) + E[\mathcal{M}\widetilde{Q}p_1]. \tag{2}$$

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 expectation of the product of the policy effectiveness, the additional policy action, and the no-policy price in states in which that policy is implemented.[8]

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:

$$\mathcal{M}_{\text{estimated}} = \mathcal{M}\left(1 + \frac{E[p_1\widetilde{Q}]}{E[p_1]Q}\right). \tag{3}$$

Because the promise provides additional price support, the effectiveness of the policy will be overestimated. Both the intensive — how often additional policy will be implemented — and the extensive margin — how much additional policy is implemented — shape the bias. First, the bias depends on how likely the promises are to be implemented, specifically how much the states where the promise is implemented contribute to the initial price. Because new policy tools are often used in difficult and uncertain conditions — think of the midst of the COVID-19 crisis — this probability is likely non-negligible. Second, the bias depends on the size of the promised policy relative to the baseline amount $\widetilde{Q}/Q$ when it is implemented. This second effect can be sizable, for example much larger than 1. Indeed, if the crash scenarios are dramatic, the policy maker might expend significantly more resources. For the main event we study in this paper, we find that 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 $0.5-1 trillion in value when the Fed announced corporate bond purchases in March of 2020 though the Fed ended up only

[8] Equation (2) as well as the next equation (3) hold for any distribution of additional policy intervention $\widetilde{Q}$.

making purchases of around $15 billion. Our estimates suggest that close to half of the response is due to a 20-fold increase in the size of the program in the lowest 20% of states.

Knowing whether the promise is present is also relevant for understanding the dynamics of policy decisions. For example, if the policymaker only announced explicitly the baseline purchases Q, market participants' perception of the potential additional purchases $\widetilde{Q}$ indicates that more interventions will be subsequently announced. Such inference depends on the content of the announcement itself. As an alternative example, it could also be that the policy maker explicitly announced both Q and the intention of maintaining the price above $\underline{p}$.[9] Then, no subsequent announcement would be needed. In practice, announcements often fall somewhere in the middle, with a cap on the quantity purchased; the question then is whether the perceived state-contingent intervention $\widetilde{Q}$ can exceed this cap. We come back to this question for the 2020 corporate bond purchases in Section 2.4.

Which data can separate promises from the baseline policy? Option contracts on the asset offer a path forward. In the simple example of this section, the price of put options with a strike below $\underline{p}$ goes to 0 after the announcement in the case with a promise because the intervention prevents the asset price from going below this threshold. More broadly, the response of option prices allow us to zoom in on different parts of the state space and measure the conditional impact of the policy for a more general set of interventions.

1.2 A Method to Estimate Conditional Policy Impact

We present a method to estimate conditional policy perceptions following the announcement of a new policy using option prices. Specifically, we estimate a conditional price support function: how much will the policy intervention change the price of the asset in each state of the world at date 1. This corresponds to identifying a function $g(.)$ such that $p_1' = p_1(1 + g(p_1))$. In the example without promises, this function is constant, $g(p_1) = \mathcal{M}Q$, while the price support is state-dependent and decreasing in the example with promises, $g(p_1) = \mathcal{M}(Q + \widetilde{Q}(p_1))$.

Our approach proceeds in two steps. 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

[9] The Bank of Japan's yield curve control policy (https://www.boj.or.jp/en/mopo/mpmdeci/mpr_2016/k160921a.pdf) is an example of such a policy.

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 risk-neutral distribution of the future price of the asset. They show that observation of option prices (calls or puts) across strikes allows 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. Formally, this observation corresponds to

$$\frac{dPut(p_1, K)}{dK} = \lim_{h \to 0} \frac{Put(p_1, K + h/2) - Put(p_1, K - h/2)}{h} = 1_{\{p_1 < K\}} \tag{4}$$

Turning back to date 0, this implies that the slope of the put prices with respect to the strike 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) $F_{p_1}^{\mathbb{Q}}(K)$. We use the superscript $\mathbb{Q}$ to indicate the risk-neutral aspect of the distribution.

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. This allows us to recover the CDFs of the no-policy price p_1 and the post-policy price p_1', which we denote $F_{p_1}^{\mathbb{Q}}$ and $F_{p_1'}^{\mathbb{Q}}$, respectively. In practice, we only observe option prices for a finite number of specific strikes, so interpolation between observed strikes is needed. Appendix A provides additional details on implementation.

The first two plots of Figure 1 illustrate these distributions for our two examples. In the case of a constant price support, Panel A, the probability density functions (PDF) shifts to the right. With the added promise of supporting the price above $\underline{p}$, the left tail of the distribution is replaced by a mass at that point.

Step 2: Solving the transport problem. Once we have the distributions with and without intervention, the next task is to find the conditional price support $g(.)$ that explains the change in distribution. This type of problem is known as a transport problem: how should we move the values of a random variable in each state to obtain a new distribution? A simple and, as we will discuss in the next section, economically appealing way to find a solution to this problem is to assume the transport is monotone. That is, we look for a price support that does not change the ranking of the asset price across states. This order-preserving property guarantees existence and uniqueness of a solution $g(.)$ up to probability-0 events. Furthermore, there is a simple method to construct this mapping. Start from a value x and compute the corresponding CDF $F^{\mathbb{Q}}_{p_1}(x)$. 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 $F^{\mathbb{Q}}_{p'_1}(y) = F^{\mathbb{Q}}_{p_1}(x)$. Once we find this price mapping, we simply have: $y = x(1 + g(x))$, 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.[10] The price support function is the change in price necessary to move from this initial value to the post-policy price. The following definition summarizes this calculation.

Definition 1. *The price support function g is the unique order-preserving transport from $F^{\mathbb{Q}}_{p_1}$ to $F^{\mathbb{Q}}_{p'_1}$. It is equal to*

$$g(p_1) = \frac{{F^{\mathbb{Q}}_{p'_1}}^{-1}(F^{\mathbb{Q}}_{p_1}(p_1)) - p_1}{p_1}. \tag{5}$$

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 $g(.)$ over the domain ${F^{\mathbb{Q}}_{p_1}}^{-1}(F^{\mathbb{Q}}_{p_1}(I) \cap F^{\mathbb{Q}}_{p'_1}(I'))$, where I and I' are the domain of strikes covered by options before and after the announcements.

The third plot of Figure 1 reports the price support functions recovered with this approach for our two examples. For Panel A, with a constant intervention, the method correctly identifies a parallel shift in the distribution (in log), and therefore a constant price support. In Panel B, the replacement of the tail by a constant value leads to a sharply decreasing function $g(p)$ below 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

[10]This is similar to constructing a Quantile-to-Quantile ("Q-Q") plot of p_1 and p'_1.

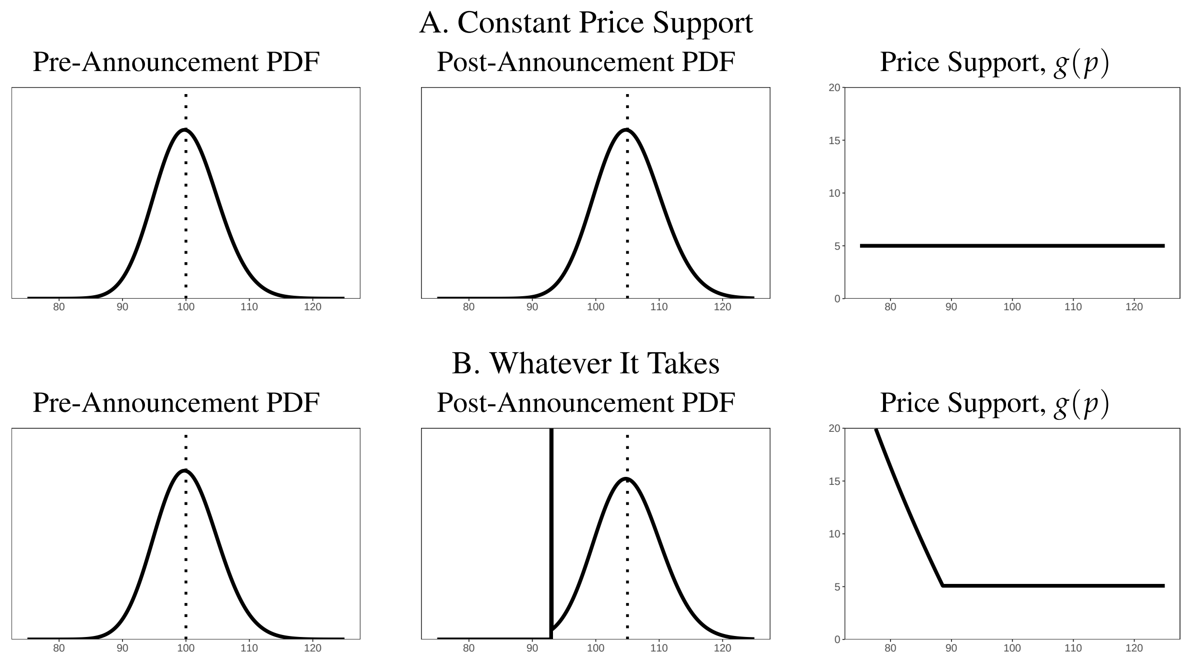

Figure 1: Examples of Distributions and Conditional Price Support Policies

The top row considers a constant price support: the price is increased by the same amount in all states. The bottom row considers a whatever-it-takes promise to maintain the price above a threshold $\underline{p}$. The left panels report the PDF of the date-1 price before the announcement. The middle panels report the PDF of the date-1 price after the announcement. The right panels report the corresponding price support functions $g(.)$.

option payoff, giving a formal way to measure the commonly used notion of a "policy put."

1.3 When Does the Price Support Function Reveal State-Contingent Policy Impact?

The price support function, as an empirical measure, can always be constructed from observations of option prices before and after an announcement or event. In the idealized examples of Section 1.1 and Figure 1, it reveals the state contingent impact of the policy. When is such an interpretation valid in real-world situations?

The next proposition provides a set of four formal conditions. We maintain the timing assumptions of our motivating example. At date 0, a policy is announced to be implemented at date 1. Market participants infer a state contingent plan from the announcement. That is, the policy implementation can depend on the realized state of the world at date 1.

Proposition 1. *The price support function defined in equation* (5) *measures the state-contingent* *impact of the policy if the following assumptions are satisfied:*

1. *(Event-study identification) Asset prices before and after the announcement reflect equilibria of economies that are identical at the announcement date except that in one case the policy plan is never implemented, and in the other the policy plan is implemented with certainty. These two economies have uncertainty at date 1 driven by a state s with identical distribution F.*

2. *(Separability) Policy acts separably across states at date 1. In the equilibrium with policy, the policy action and its impact g in a given state depends only on the value of the asset price in that state absent policy $p_1(s)$:*

$$p_1' = p_1\left(1 + g(p_1)\right). \tag{6}$$

3. *(Order-preserving) The with-policy price p_1' is weakly increasing in the no-policy price p_1.*

4. *(Invariant stochastic discount factor) The same stochastic discount factor $m(s)$, and therefore the same risk-neutral measure $E^{\mathbb{Q}}$, prices the asset and options at date 0 in the with-policy and the no-policy economy:*[11]

   i) *Without policy, for all functions h, a claim paying $h(p_1)$ at date 1 has price*

   $$\int m(s)h(p_1(s))dF(s) = E^{\mathbb{Q}}[h(p_1)]. \tag{7}$$

   ii) *With policy, for all functions h, a claim paying $h(p_1')$ at date 1 has price*

   $$\int m(s)h(p_1'(s))dF(s) = E^{\mathbb{Q}}[h(p_1')] = E^{\mathbb{Q}}[h(p_1(1+g(p_1)))]. \tag{8}$$

While these conditions are unlikely to be perfectly respected in practice, they provide guidance for what one should assess before interpreting an estimated price support function. We now discuss each of these assumptions and how to test, approximate, or relax them.

[11]To simplify notations, we are ignoring risk-free discounting between date 0 and 1. Equivalently, we are working with forward prices.

1.3.1 Event-study identification

The first condition is a form of event-study empirical design. It assumes that *observational data* about prices before and after the announcement reflect the *counterfactual experiments* of situations with or without policy. This occurs when the announcement is unanticipated and its implementation is certain. Unanticipated means that the probability of the policy plan implementation before the announcement is close to 0. Certain implementation means that after the announcement, the probability of implementation is close to 1.

In contrast, if the policy is anticipated, prices before the announcement already reflect some of its influence. Hence, by focusing on the response to the announcement, one would underestimate the policy's impact.[12] In Appendix E, we show how to adapt our method when the policy is somewhat anticipated and the implementation is uncertain. When applying this approach to our main empirical estimates, we find that accounting for mild anticipation and uncertain implementation leads to small changes and that, if anything, the baseline estimates understate not only the effectiveness of the policy plan but also its asymmetry across states. In other words, using this first assumption to approximate reality leads us to draw conclusions that are too conservative on the role of state-contingent policy.

Another natural concern with this assumption is that other news can be concurrent with the policy announcement. Understanding the specific context of each announcement is necessary to assess whether another major event is going on. For example, Section 2.1 discusses the events around corporate bond purchases in 2020. Furthermore, financial data always exhibits sizable volatility reflecting a constant flow of background news. Focusing on a tight window around the policy announcement limits the role of these other news. In this paper, we compare prices the trading day before the announcement with those at the end of the day of the announcement, the tightest our data allows. Relatedly, to draw careful statistical inference one has to assess whether the "signal" of the policy announcement dominates the "noise" of the remaining ambient news. Concretely, this means constructing appropriate confidence intervals for our estimates, something we explain precisely how to do in Section 2.3.

Finally, we also assume that the two counterfactual worlds are driven by the same source of

[12] Intuitively, if a given market response is triggered by only a 50% update in the probability of intervention, this implies that the total effect of the policy is twice as large as if the same market response resulted from a 100% update.

fundamental uncertainty, the state s. For example, s could be a state of nature (e.g. rain or shine, or how bad the COVID-19 disease is) with dynamics that do not respond to the policy.

1.3.2 Separability

The second condition is a simple way to define an economically meaningful notion of state-contingent policy impact. It departs from a monolithic view that a policy announcement conveys a single number. Instead this condition entertains different policy actions across states of the world that each translate into their own impact, represented by the price support function $g(.)$. This representation assumes that the state-contingent impact depends only on what the date-1 price of the asset would be absent policy intervention, p_1. It summarizes the total impact of the policy, and on its own does not separate the exact origin of this impact. With additional information about the policy — such as Q or $\mathcal{M}$ in our example — one can link the price support back to specific actions; we will do so in our applications. 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. Policymakers, even when they explicitly want to support prices, look at a variety of pieces of information to make decisions. 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.

While this representation is flexible, it still imposes restrictions on the family of mechanisms by which policy operates. Separability lines up well with standard models of asset purchases in the style of Gertler and Karadi (2011), Greenwood and Vayanos (2014), or the foundation we present in Appendix B. However, this interpretation prevents policy in some states to affect outcomes in other states of the world. If such cross-state effects happen, our measurement framework still correctly identifies the total effect of the policy plan in each state, but cannot link it to actions only in that state. Doing so would rely on taking a stand on the form of non-separability in the data, something that can be guided by theory.

Another family of models that does not immediately fit our framework are those in which out-of-equilibrium beliefs about policy matter. For example, in theories of runs, out-of-equilibrium beliefs about how policy plays out in a run can rule out runs on the equilibrium path. This is a concern common to all event-study approaches: the observed policy news can be associated with unobserved changes in out-of-equilibrium beliefs. Such beliefs cannot be directly recovered from observational data because by definition these beliefs are about outcomes that are out of the equilibrium. While there is no simple solution to this issue aside from comparing the price support function to a much more structural model prediction, the user of our measurement framework should keep in mind this limitation.[13]

1.3.3 Order-preserving

The order-preserving condition assumes that the policy does not change the ranking of the asset price across states. For example, the policy does not support the price so much in (no-policy) bad states that it becomes higher than in good states. Mechanically, this condition ensures that the policy acts as a monotone transport, lending itself to our estimation approach.

This assumption is economically appealing: there is a sense in which such policies are efficient. Indeed, consider a policymaker who targets a given distribution of the price after intervention. Multiple price support functions can lead to this distribution, but an order-preserving policy minimizes the use of large changes in prices, and likely minimizes the cost of the policy.

In practice, one reason that policy might not be order-preserving are costs (e.g., political) to expanding policy smoothly. Instead policymakers might prefer to increase interventions in fixed increments as conditions worsen instead of the smooth example of Section 1.1 — a practice referred to as "expanding the envelope" in policy circles. In that case, our measurement framework simply smoothes out these discrete jumps up, but does not substantively affect our conclusions.

More generally, we show that wrongfully making the order-preserving assumption leads to conservative estimates of the conditional nature of the policy. Any other price support function rationalizing the data must be more asymmetric than our estimate. Appendix D formalizes these arguments.

1.3.4 Invariant stochastic discount factor

The last assumption restricts the behavior of the stochastic discount factor between date 0 and date 1. Such an assumption is necessary because the response of prices to the announcement of the

[13] Bocola and Dovis (2019) is an example of such a structural approach in the context of debt crises.

policy (date 0) reflect a combination of changes in the actual distribution of implementation prices (date 1) and changes in the stochastic discount factor between the two dates (that is, between date 0 and date 1). Condition 4 offers the simplest possible take on this issue by assuming that the stochastic discount factors between the two dates are identical with and without the policy intervention, that is $m(s) = m'(s)$.

This assumption brings some flexibility along two dimensions. First, it does not assume coincidence of risk-neutral and historical probabilities (which would be $m(s) = m'(s) = 1$). It also does not assume that the econometrician knows the stochastic discount factor or is able to recover it from prices.

Second, this assumption does not put any restrictions on the determinants of prices at date 1 or after that. In particular, we do not take a stand on the mechanisms through which the purchases at date 1 affect the price. In the example model of the previous section, this means that we do not make assumptions about where $\mathcal{M}$ comes from, and allow it to come from changes in future stochastic discount factors. In fact, in our applications, an appealing mechanism for why asset purchases are effective is through effects of purchases on risk premia *from date* 1 *onwards*.[14]

While condition 4 is flexible on these two dimensions, it takes a strong stance on the invariance of discounting between announcement and implementation. If we go to the other extreme and put no restriction at all on the short-term stochastic discount factors, virtually any change in risk-neutral distribution can be mathematically explained by a pure change in the stochastic discount factor.[15] However, in basic models, or in unconditional data, variation in the short-term stochastic discount factor plays a modest role in asset price variation, and even more so in the shape of price movements across the state space. Appendix C.1 reviews these standard observations.

Still, for large policy announcements, the response of the SDF could have a sizable effect, and we propose two ways to quantitatively assess how this response can affect our conclusions. We rely on a much weaker form of invariance and start by the premise that changes in price in various states might pass through to the stochastic discount factor. In words, if the intervention changes

[14]Formally, if we assume that the asset pays off the random amount X_τ at a future date $\tau > 1$, we have: $p_0 = E_0[m_1(s)p_1(s)] = E_0[m_1(s)E_1(m_{1,\tau}X_\tau|s)]$. Condition 4 assumes that $m_1(s)$, the stochastic discount factor between date 0 and 1, is unaffected by the policy, but allows the policy to work through changes in $m_{1,\tau}$.

[15]Specifically, the measures implied by $F^{\mathbb{Q}}_{p_1}$ and $F^{\mathbb{Q}}_{p'_1}$ must be absolutely continous with respect to each other; that is, have the same set of probability-0 events. In this case, the change in risk-neutral distributions can be explained by $g = 0$ and $m'(s)/m(s) = dF^{\mathbb{Q}}_{p'_1}/dF^{\mathbb{Q}}_{p_1}$.

the value of prices in different future states, it might also change how investors value those states from today's perspective. For example, in a segmented market with specialized investors, changes in the price of the asset in a given state affect these investors' wealth, and therefore their marginal utility. In equilibrium this translates into a change in the stochastic discount factor for that state — a mechanism present in the model we consider in Appendix B. Within a broader view of markets the pass-through of the policy to the stochastic discount factor is mediated by the portfolio share of the marginal investor in the targeted assets, but also captures potential spillovers of the policy to other assets, or other general equilibrium effects affecting the marginal investor.

We show two results. First, irrespective of the form of this pass-through, our empirical method recovers the correct price support function g if it is constant. Second, if one takes a stand on the intensity of this pass-through, one can adjust our method to recover the correct g.

Testing the null hypothesis of a constant price support. Consider the null hypothesis that the policy provides the same proportional price support in all states of the world. We show that the approach of Definition 1 recovers the price support irrespective of how the pricing kernel responds to change in prices under this null hypothesis.

Proposition 2. *If the true price support function $g(.)$ is constant, conditions 1-3 of Proposition 1 hold, and the pricing kernel before and after the intervention can be written:*

$$m = \Theta\left(s, p_1(s)/p_0\right),$$
$$m' = \Theta\left(s, p_1'(s)/p_0'\right)$$

for the same function Θ, then equation (5) *correctly recovers the price support function.*

The empirical content of this proposition is that finding a non-constant g reveals the presence of state contingency in the price support. What is the family of pricing kernels for which this result holds? In words, they depend on two elements: the state of the world at date 1 (s), and the return of the asset between date 0 and date 1 (p_1/p_0 before the announcement, p_1'/p_0' after). This second component encodes in a flexible way the pass-through of the asset price to the pricing kernel. For example, a CRRA model with the asset representing total wealth is $\Theta(s, R) = R^{-\gamma}$, with γ the coefficient of risk aversion, and R the asset return. Many asset pricing models also feature pricing kernels determined by asset returns: other utility functions, loss aversion, etc.

Intuitively, a constant price support shifts all date-1 prices up by a certain proportion g. This does not change the nature of the risk of the asset. Indeed, the date-0 price increases by the same proportion g, the distribution of returns between dates 0 and 1 is unchanged, and so is the pricing kernel. Appendix Section C.2 derives the result more formally.

Adjusting estimates for endogenous risk premia. By taking a stand on the pass-through of asset price to the stochastic discount factor, we can go further and provide estimates of the price support function that take into account this effect. Specifically, we replace condition 4 by the following.

Assumption 5. Endogenous pricing kernel. *The pricing kernel is $m = \theta(s)\frac{p_0}{p_1}$ before the announcement, and $m' = \theta(s)\frac{p_0'}{p_1'}$ after the announcement.*

This assumption allows the announcement to affect the pricing kernel with a unitary pass-through. This endogenous part of the pricing kernel is the same as that of a log-utility investor with all her wealth invested in the asset.[16] Beyond this specific theoretical motivation, we show in our empirical application that this value implies a large risk premium, which is also very responsive to the distribution of the asset. Therefore, it leads to conservative estimates of the price support $g(.)$ accounting for risk premium effects.

The following proposition shows how to recover the price support function when replacing condition 4 by this unitary pass-through.

Proposition 3. *Under conditions 1-3 of Proposition 1 as well as Assumption 5, the price support function $g(.)$ is the unique solution to equation* (5)*, where the risk-neutral distribution is replaced by the forward-neutral distribution $F^{\mathcal{N}}$, which can also be obtained from option prices:*

$$dF^{\mathcal{N}}(p_1) = E[\theta(s)|p_1]dF(p_1). \tag{9}$$

In words, instead of focusing on matching the invariant risk-neutral distribution before and

[16] Martin (2017) and Kremens and Martin (2019) use pricing kernels taking a similar form to construct predictions for stock returns and exchange rates using option prices. Our assumptions are different: in our setting, $\theta(s)$ can take any shape as long as it is unchanged by the announcement while they impose restrictions on the shape of $\theta(s)$, but not on its time series properties.

after the announcement, we concentrate on a distribution affected only by the exogenous part of the pricing kernel.[17] Thus, this distribution is not affected by the policy. How can we recover this distribution using options? The basic idea is to use option contracts that are expressed in the numeraire of the asset return, so that they cancel out with the endogenous part of the pricing kernel. While such contracts might seem unusual, we show that they can be replicated very simply by combining the same calls and puts as in our baseline case. Appendix Section C.3 derives these results.

This leaves us with an alternative estimate to compute the price support function $g(.)$. By measuring both our baseline and this alternative, researchers can simply assess the sensitivity of their conclusions to changes in the SDF. We follow this recommendation in our empirical exercises.

2. Corporate Bond Purchases in 2020

We turn to our main empirical application, the March 2020 Federal Reserve announcement of corporate bond purchases. We start by discussing our identification setting.

2.1 The Announcement, Price Response, and Identification

On March 23, 2020 the Federal Reserve announces purchases of investment-grade corporate bonds and investment grade corporate bonds ETFs through the Secondary and Primary Market Corporate Credit Facility (SMCCF and PMCCF). The announcement immediately raises investment-grade corporate bond prices by around $500 billion in market value. The announced capacity of corporate bond purchases on March 23 was $300 billion, but the Fed left unclear the total amount they would buy. Ultimately, purchases occurred around three months later and total only around $15 billion, 0.2% of the market capitalization of investment-grade corporate bonds as of June 2020. No further purchases occur afterwards. Haddad et al. (2021) and Boyarchenko et al. (2020) offer in-depth accounts of the Fed's announcement and additional details on purchases.

[17] The forward-neutral measure prices cash-flows expressed using the asset as a numéraire (see Geman et al. (1995) for more discussion). Specifically, for a random payoff x the price can be obtained using the risk-neutral distribution $p = E[mx] = E^{\mathbb{Q}}[x]$. It can also be obtained using the forward-neutral distribution: $E[mp \times x/p] = E^{\mathcal{N}}[x/p]$, where x/p is the payoff expressed in number of units of the underlying asset. Appendix C.3 shows how to construct this distribution empirically.

The March 23 announcement constitutes the setting of our event study. As such, it is crucial to assess whether it is indeed the driver of price movements in investment-grade corporate bonds at that point, following the discussion of Section 1.3.1. We track this price using iShares' investment-grade corporate bond ETF (LQD) following Haddad et al. (2021). This large ETF captures the immediate price response for the broad universe of investment-grade corporate bonds without having to obtain the transaction level data of individual bonds which trade much less frequently.[18]

Several elements support the view that the announcement explains the change in the price of the index over this period. Figure 2 reports the cumulative return of LQD around our event. The price of the ETF immediately jumps up when markets open after the announcement and closes up 7%. This large abnormal movement is similar using the return at the market open or close on March 23, suggesting that little additional news arrives after the announcement on this day.[19] However this observation still leaves open the possibility of other economic news between the market close on Friday and the market open on Monday. If that were the case, one would expect that other economic-sensitive assets like stocks or high-yield bonds would also respond to such news. Figure 2 shows that high-yield corporate bonds and the stock market are flat or slightly down from Friday to Monday, cutting against the importance of other news. This lack of response by similar but non-targeted assets also suggests that the announcement itself did not reveal information about the broader economy. Appendix G.1 confirms this pattern formally using event-study regressions in which we control for the stock market, high-yield bonds, and Treasury bonds.

The event-study identification assumption in Proposition 1 also requires the event to be unanticipated. The context of the announcement is suggestive of this view. This policy intervention is the first time the Fed enters the corporate bond market, a policy that had previously never been explicitly considered. While it is impossible to provide sharp evidence of no anticipation, we show in Appendix E that our inference is robust to relaxing this assumption and allow a mild amount of anticipation. In fact, the intuition of standard event studies holds here: if the policy was somewhat anticipated or not perfectly credible, the price support function would be biased towards 0.

[18] Haddad et al. (2021) show a similar response in individual bond prices using TRACE data. See 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.

[19] A concern in the opposite direction is that the narrow one-day window might not leave enough time for the market to process the news. Using a longer three-day window increases the raw and abnormal excess returns to about 14% and 10%, respectively. We favor the narrower window for this event given that volatility was very high and also the fact that the CARES act was signed into law four days after this announcement.

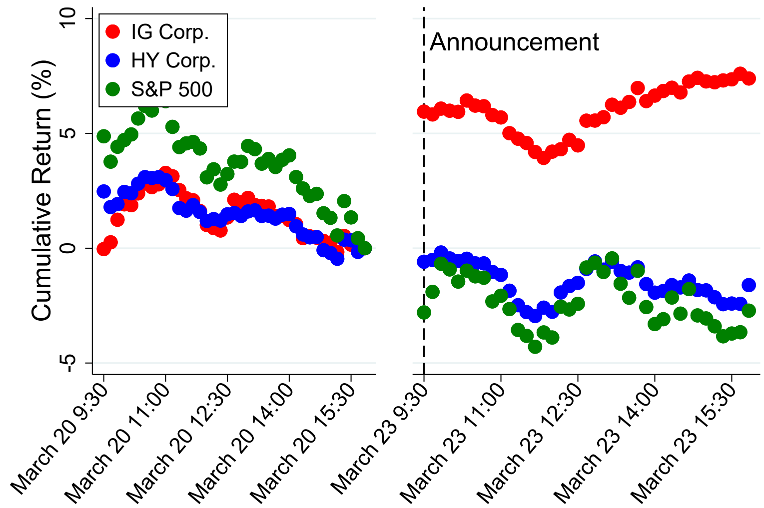

Figure 2: Event Study Price Responses.
This figure reports the cumulative return on the investment grade bond ETF, high yield bond ETF, and S&P500 at 10 minute intervals. The Fed purchase announcement occurs the morning of March 23 (vertical dashed line).

2.2 Option Prices and Changes in the Distribution of Corporate Bonds

We now turn to option prices on the same investment grade bond ETF around the announcement. Figure 3 Panel A plots the 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 on which the announcement was made. Three months is the maturity with most liquidity and is also the horizon at which total amount of bonds the Fed would actually buy becomes clear. Implied volatilities decrease across all strikes, but the drop is most pronounced in the left tail (deltas below 30%). This pattern implies that the risk-neutral probability of extreme low prices is particularly reduced following to the announcement. Panel B reports the risk-neutral CDFs implied by the option prices following Breeden and Litzenberger (1978), confirming a substantial drop in the probability of prices in the left tail.

2.3 Conditional Price Support

We apply our measurement framework to recover the conditional price support provided by the Fed. 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

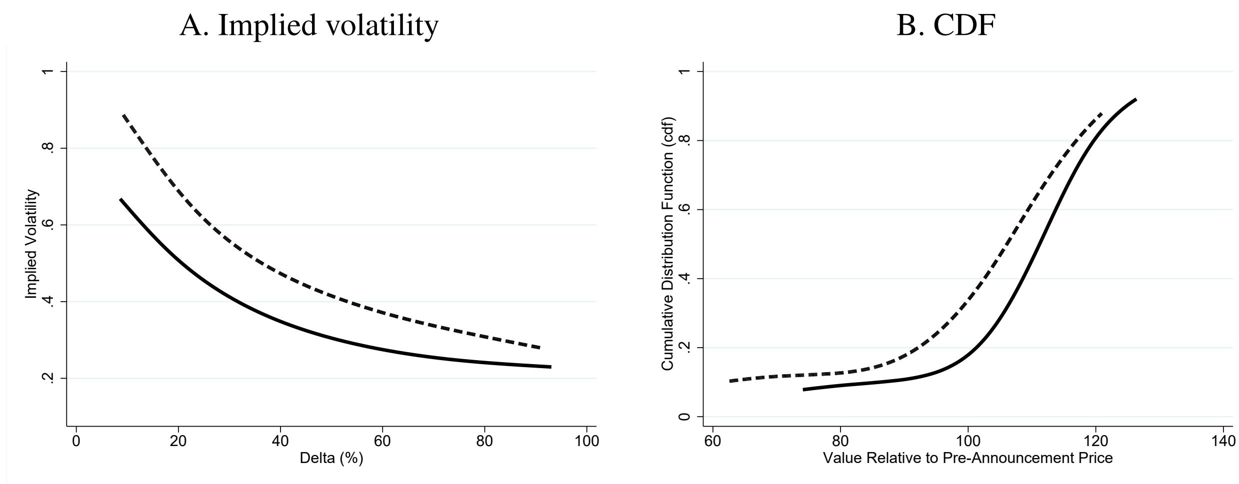

Figure 3: Option Implied Volatility and CDF Before and After the Announcement.
This figure shows the option-implied volatility (Panel A) and risk-neutral cumulative distribution function (Panel B) for the investment-grade corporate bond ETF (LQD) on March 20 (dashed line) and March 23, 2020 (solid line). The time-to-maturity of the options is 3 months.

Fed intervention and should be thought of as capturing the underlying state of the corporate bond market.

Figure 4 plots the function $g(p)$ expressed as a percentage of the no-policy price p. First, $g(p)$ is not flat like an unconditional price support would imply, but instead is strongly downward sloping, particularly for low values of the price. At values where the price drops by 20-30%, the slope of the price support function is nearly -1, which suggests a policy close to a price floor — that is, each dollar lost in this region is offset by conditional price support by the Fed. The price support function strongly resembles the payoff of a put option, lending support to the view of a "policy put." This suggests that bond investors perceived a backstop where stronger intervention would occur if corporate bond prices fell further. To gauge magnitudes it is worth picking two points on the figure. If, absent any policy intervention, prices would have increased by 20%, Fed support would push the price up by an additional 2%. If prices would have declined by 35% instead, the Fed would push the price up by around 40%. Thus, the asymmetry is economically very large.

A natural question is whether these movements in the implied volatility curve are typical and hence could have happened by chance, as opposed to being driven by the policy announcement. For example, it is well understood that tail risk movements are an important driver of returns both in equity and corporate bond markets. The 95% confidence interval in the gray shaded region in

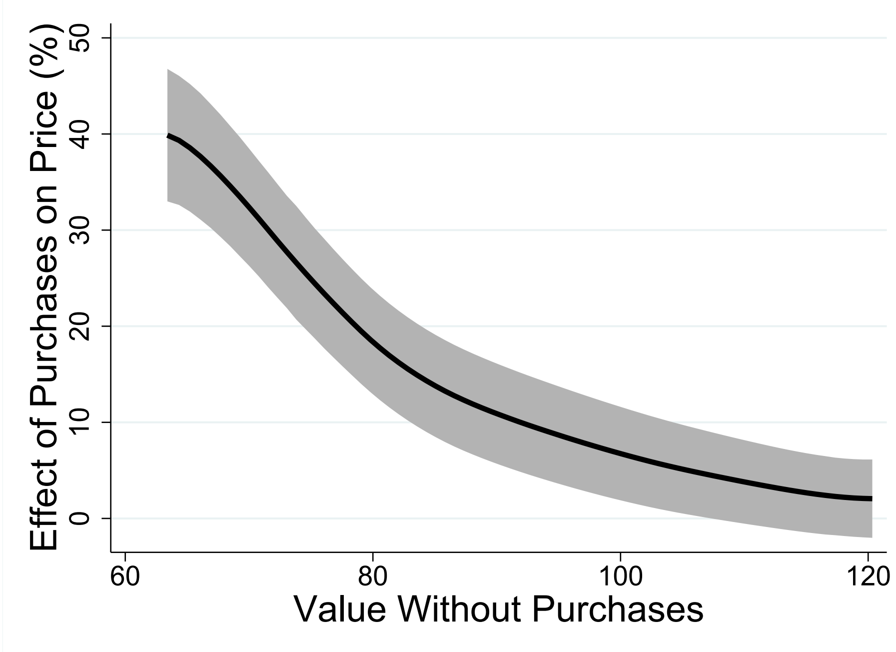

Figure 4: Conditional Price Support Function $g(p)$.
This figure shows the implied price support (expressed in percent) as a function of the pre-policy price. The pre-policy price is normalized to 100 before announcement.

Figure 4 indicates statistically significant price support for the left tail of the distribution, and only insignificant estimates of price support at the upper end. Thus, the pattern we find is extremely unlikely to emerge by chance alone, consistent with the unusually large overall announcement return of 7%. Furthermore, the null hypothesis of a constant price support (i.e., a proportional shift in distribution at this level) is also strongly rejected, with values of the price support in bad states of the world being sharply statistically higher than 7%.

To construct these confidence intervals, we start from the sample of all dates between January 2010 and February 2020. For each pair of consecutive trading days, we compute the risk-neutral CDF on each day, and find the price support function connecting them. A simple bootstrap would then compute for each point p on the x-axis of this price support function the sample distribution of $g(p)$ outcomes in the y-axis. However, we need to account for the fact that the data is highly heteroskedastic. In particular, there is much more volatility for our event than on a typical day. We follow the standard treatment of heteroskedasticity when volatility is measured: rescale the data to make it comparable across dates. In our setting, this means rescaling observations on the x and y-axis by the at-the-money implied volatility on each date to make to make them have essentially the same volatility. We then bootstrap based on these rescaled values. Appendix H provides step-by-step details of this calculation and considers alternative treatments of heteroskedasticity. Appendix H.2 shows that even large returns in the rest of the sample do not lead to as much asymmetry as on the day of the announcement, assuaging the concern that large shocks mechanically have a larger impact on the left tail of the distribution.

Appendix G presents additional tests confirming that this finding is a robust feature of the data. Appendix G.7 demonstrates that our results are robust to accounting for bid-ask spreads in option prices and discusses the liquidity of the options we use. Appendix G.5 uses a longer window for the event study to allow more time for markets to react at the cost of less tight identification. While magnitudes in this case are larger compared to the results in the one-day window, the asymmetric pattern is similar.

How much did promises contribute to the overall price movement? We next compute the fraction of the initial announcement return stemming from the left tail asymmetry we found. To do so, we construct the announcement return for a counterfactual price support function without promises. Recall that by definition $E^*[p_1 g(p_1)]/p_0$ is the one-day announcement return of 7% discussed earlier. To assess the role of downside support, we shut it down: we assume that the Fed supports the price by a constant amount in the downside, equal to the support in the case of a median change in the no-policy price, $g(p_{med})$. Formally, we define $\tilde{g}(p) = g(p)$ when $p > p_{med}$ and $\tilde{g}(p) = g(p_{med})$ when $p \leq p_{med}$. Appendix Figure OA.4 gives a graphical representation of $\tilde{g}(p)$. We compute the counterfactual announcement return $E^*[p_1 \tilde{g}(p_1)]/p_0$ using the risk-neutral probabilities of each state. Comparing this counterfactual number to the actual announcement return quantifies the effect of the policy put. We find that 50% of the overall effect on prices comes from conditional policy to support prices more heavily in adverse states. Thus, the policy put option by itself boosted the total bond market value by around 3.5% or about $250 billion.

2.4 Interpreting the Asymmetric Price Support Function

Our main methodology delivers the conditional price support, a measure of the impact of the policy in each state of the world. We now discuss how to tie this impact to specific actions by taking a stronger stance on the economic mechanism and using further sources of information. In this application, we aim to estimate the state-contingent quantity of investment-grade corporate bonds purchased by the Fed. Beyond quantifying variation in this quantity, it is particularly interesting to assess whether market participants forecasted purchases above \$300 billion, the size of the facility announced on March 23.

Baseline calculation: constant multiplier. A natural starting point is the assumption that the multiplier, or price impact, of purchases $\mathcal{M}$ is constant across states. In this case, $g(p) = \mathcal{M}Q(p)$, with $Q(p)$ the fraction of total supply purchased.

Immediately, this gives relative statements on purchase amounts across states by using the ratio of $g(p)$ for two different states. We use the state where bond prices fall 35% as a reference "bad state." The Fed would buy 6 times as many bonds in this bad state compared to the state where prices do not change ($p = 100$), and 20 times as many bonds compared to the case where prices appreciate by 20%. These are large relative magnitudes.

Estimating absolute quantities requires knowing the multiplier $\mathcal{M}$ for corporate bonds and the total supply. In 2020, the total supply of investment-grade corporate bonds is around \$7 trillion.[20] For the multiplier, we follow Chaudhary et al. (2023) who estimate $\mathcal{M} = 3.5$ for corporate bonds. This number is similar to other estimates in the literature: Darmouni et al. (2022) find a multiplier around 6, while the estimated price impacts in Bouveret and Yu (2021) correspond to values between 2.3 and 5.6.[21]

These assumptions yield that the Fed would purchase around \$800 billion of bonds in the bad state, \$135 billion in the case of no change in price, and \$40 billion in the good state with a 10% appreciation. Using values for $\mathcal{M}$ at the bottom of the range found in the literature (a more elastic bond market) would increase the implied quantities. Recall that the announced facility size on March 23 was \$300 billion, so the numbers in the bad state imply a substantial expansion of the facility size originally announced. Interestingly, they remain in line with the expansion of the facility two weeks later on April 9 to about \$900 billion. The likelihood of an expansion of the intervention appears robust: to generate the price support of 40% in the left tail with a \$300 billion purchase would require price impact of about $\mathcal{M} = 10$. Such a number is significantly higher than

[20]SIFMA reports \$8.8 trillion in US corporate bonds outstanding at the end of 2019. Given that the ratio of investment-grade to high yield bonds is approximately 6 to 1, this results in around \$7 trillion in investment-grade corporate bonds. Appendix G.4 provides several alternative data sources which produce a range of investment-grade corporate bond supply between \$6-8.5 trillion. We take \$7 trillion as the average across these sources.

[21]Estimates at the individual bond level instead of the asset class level tend to be an order of magnitude smaller. In other asset classes, Gabaix and Koijen (2021) estimate $\mathcal{M} \approx 5$ for the stock market and Greenwood and Vayanos (2014) find a multiplier around 0.4 for Treasuries using shocks to supply.

estimates in the literature.

The quantity of \$40 billion in the relatively good state is of comparable magnitude as the actual purchases of \$15 billion that the Fed implemented. This interpretation based on quantities fits the narrative provided by Jerome Powell in a testimony given in June 2020 that "markets are functioning pretty well, so our purchases will be at the bottom end of the range that we have written down."

State-contingent multiplier. The asymmetry of the price support could also be driven by a multiplier that changes across states of the world, $\mathcal{M}(p)$. As one extreme, we ask how much variation in multiplier is necessary to rationalize the price support function absent any variation in purchases. This view leads to implausible values for the multiplier. First, it requires an *average* multiplier across states of over 30, at least an order of magnitude higher than estimates in the literature.[22] Second, the multiplier in the bad state would have to be around 200 (that is, purchasing 1% of the market cap of corporate bonds raises their price by 200%), multiple orders of magnitude above other estimates and intuitively implausible. Third, the variation in multiplier across states would also be an order of magnitude more extreme than estimates in the literature.

Still, while conditional policy effectiveness alone is unlikely to explain the price asymmetry we see, it could impact the quantities we estimated above. The literature on corporate bonds does not offer estimates of variation in multipliers across states, so we rely on the work on Treasuries for a quantification. Greenwood and Vayanos (2014) estimate that the price impact of Treasury supply is 25% higher for a one-standard deviation shock to arbitrageur wealth, where these shocks are proxied for using realized bond returns. Thus, a back-of-the-envelope calculation using implied volatility of at-the-money options in our setting would imply a multiplier that is roughly 50% higher in extreme bad states and 25% lower in extreme good states than our baseline. For quantities purchased, this would imply more than \$500 billion in purchases in the extreme bad state instead of \$800 billion — still substantially over the announced cap of \$300 billion — and \$50 billion in purchases in the relatively good state instead of \$40 billion. Appendix I presents these calculations in more detail and discusses additional evidence suggesting that multipliers are relatively stable across states.

[22] For realized purchases of \$15 billion, or 0.21% of total supply, to lead to an announcement return response of 7%, the average multiplier would need to be $7/0.21 \approx 33$.

Changes in the pricing kernel. Another consideration is the role of changes in the pricing kernel in explaining the asymmetric price support. Intuitively, conditional purchases might change relative state prices.[23] The strict assumptions of Proposition 1 ruled out such a change.

Under the view that changes in the pricing kernel affects all markets, the lack of announcement response in other asset classes supports the assumption that there were no broad changes in the pricing kernel.[24] For example, a reduction in representative investors' risk-aversion would impact high-yield bonds more than investment grade due to their higher exposure to credit risk. However, pricing could also reflect factors specific to the corporate bond market as a consequence of market segmentation.

Therefore, we favor following the procedure of Section 1.3.4 to quantify the robustness of our conclusions to changes in the SDF. Recall that if the price support is invariant across states, it implies an invariant SDF for a wide class of standard models. Hence, if the price support policy was indeed constant, we would recover a constant support function. Thus, pricing kernel adjustments can affect the magnitude of our estimates, not their asymmetry.

We include a risk-premium adjustment in our calculations by following Proposition 3. Figure 5 shows the price support in the presence of an endogenous pricing kernel. The asymmetry, with a much stronger support the left tail, remains sharp in this case. The magnitude of the support is only dampened by about a third, despite an aggressive adjustment for risk premia. Specifically, our calculation implies a large risk premium for investment-grade corporate bonds of 16.5%, about 20 times the long-term average of 0.8% reported in Giesecke et al. (2011). We conclude that our baseline estimates are robust to the presence of changes in pricing kernel.

Multiple equilibria. Finally, it is important to acknowledge that all these interpretations are conditional on the broad set of assumptions outlined in Proposition 3. An alternative framework to interpret these results that is outside of these assumptions is one that features multiple equilibria. In such a setting we no longer have separability because the policy impact in a particular state depends on out-of-equilibrium actions done by the Fed. For example, if the Fed announcement can coordinate investors in a good equilibrium, this can lead investors to respond differently across different states — even if the policy action is constant across states, and even if there is no pol-

[23] For example, Pflueger and Rinaldi (2022) studies a model where policy announcements affect risk-aversion.

[24] See Appendix G.6 for the comparison with high yield and also Haddad et al. (2021) for more comparisons across asset classes.

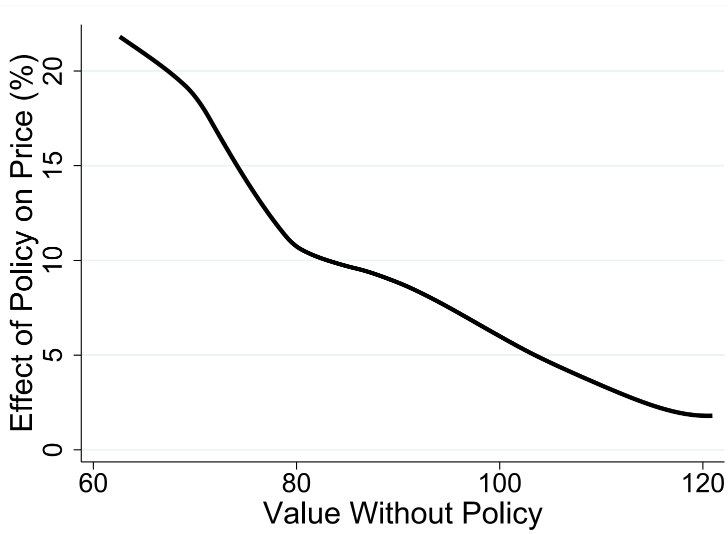

Figure 5: Risk-Premium Adjusted Price Support.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement, following Proposition 3.

icy action at all. While in this case we can no longer map state contingent price effects to state contingent policy, the function g should still be a useful moment to discipline models of this kind.[25]

2.5 In Which States Was the Fed Expected to Buy?

We have shown that the data is consistent with the market expecting the Fed to provide more price support in states where bond prices would be low. However, our analysis does not speak to whether these low price states are due to a deterioration of the credit risk of corporate bonds, high risk-free interest rates, or a high dislocation of corporate bond prices from fundamentals.

We shed light on this by using options on Treasuries (to capture shifts in risk-free interest rates) and options on a portfolio of CDS contracts (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" also referred to as 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 recover the distribution of this synthetic bond by using the option prices for Treasuries and the investment grade CDX contract (a portfolio of CDS contracts) and assuming a correlation between the

[25] Ma et al. (2020) and Eisenbach and Phelan (2022) on bond market fragility are two recent examples of this view.

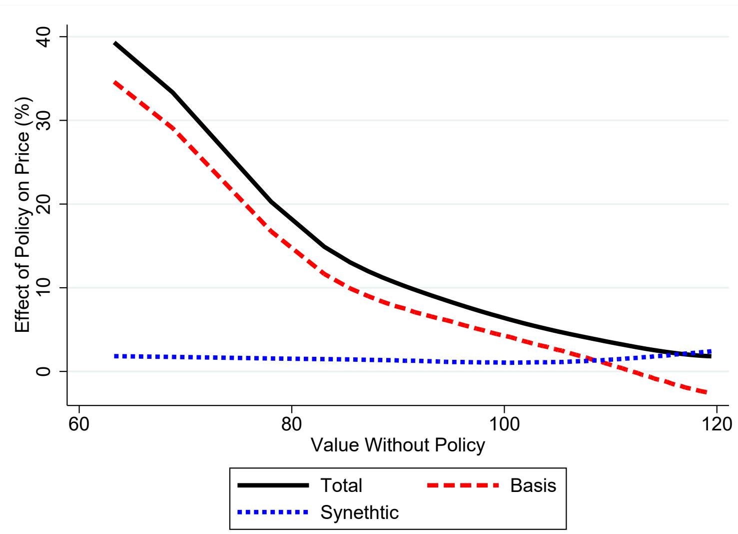

Figure 6: Decomposition of Announcement Effects: Basis vs Synthetic Bond.
This figure plots the decomposition of price support coming from the synthetic corporate bond vs the basis between corporate bond prices and the synthetic corporate bond (constructed using Treasuries and CDS).

CDX and Treasuries equal to the historical average and using copula functions following Haugh (2016).[26] This enables us to decompose our earlier finding into effects driven by movements in the synthetic bond and the dislocation component. Specifically, Figure 6 shows the conditional expectation of the share of the corporate bond return due to movements in the basis and the synthetic bond for different states. We find that the asymmetric effect of the announcement on prices is entirely driven by a sharp reduction in the basis in these bad states of the world. This is consistent with stronger Fed intervention in states where the corporate bond market is highly dislocated. This pattern of buying more in states of high dislocations is in line with with the statements from Powell emphasizing "markets [...] functioning."

3. Conditional Promises Everywhere

We expand our analysis to a large set of events worldwide in which policymakers announce interventions in financial markets, and study perceptions of state-contingent policy. We study equity injections in the US financial sector during 2008, an announcement of large asset purchases by the Bank of Japan in 2013, the various dates associated with the implementation and unwinding of quantitative easing (QE) in the United States from 2008-2013, and a set of asset purchase an-

[26]See Appendix J for details on implementation and robustness to correlation assumptions.

Table 1: Downside Support Across Policy Announcements.

This table lists the policy announcements we consider and applies our methodology to decompose the announcement response. We compute the fraction of each announcement return explained by state-dependence in the left tail (additional price support below the median). The specific events, methodology, and financial instruments are provided in the text.

	Fraction Explained by Left Tail	
Event	Risk-Neutral Measure	Forward Measure
Mar 23 2020	50%	30%
Apr 9 2020	8%	3%
Oct 13 2008 (Paulson Plan)	32%	33%
Apr 3 2013 BoJ Purchase Speech	10%	11%
US Quantitative Easing Events:		
Nov 25 2008	3%	2%
Dec 16 2008	28%	28%
Mar 18 2009	21%	21%
Jun 19 2013	7%	7%
ECB Events:		
May 10 2010	18%	15%
Jul 26 2012	8%	7%
Aug 2 2012	38%	37%
Sep 6 2012	16%	16%
Average	20%	17%

nouncements by the ECB in 2010-2012 associated with the European debt crisis. While this list is non-exhaustive, it seeks to illustrate both the role of state-dependence 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. The validity of the event-study assumption for each of these events is discussed in the papers we refer to hereafter. We investigate quantitatively the sensitivity of our conclusions to changes in risk pricing.

Table 1 lists all the announcement dates, and reports the fraction of the announcement response explained by downside support. We compute this number using our baseline approach (using the risk-neutral measure) and adjusting for a response of risk pricing (forward-neutral measure). Figure 7 represents the price support functions. Promises are pervasive: the vast majority of announcements have an asymmetric price support function with stronger support in the downside, and this price support typically explains a sizable fraction of the announcement return. This downside support is substantial even after we adjust for risk-premia following the approach of Proposition 3 (Appendix G.3 describes the implementation).

2008 US financial sector bailout. We first study the October 10, 2008 announcement of large

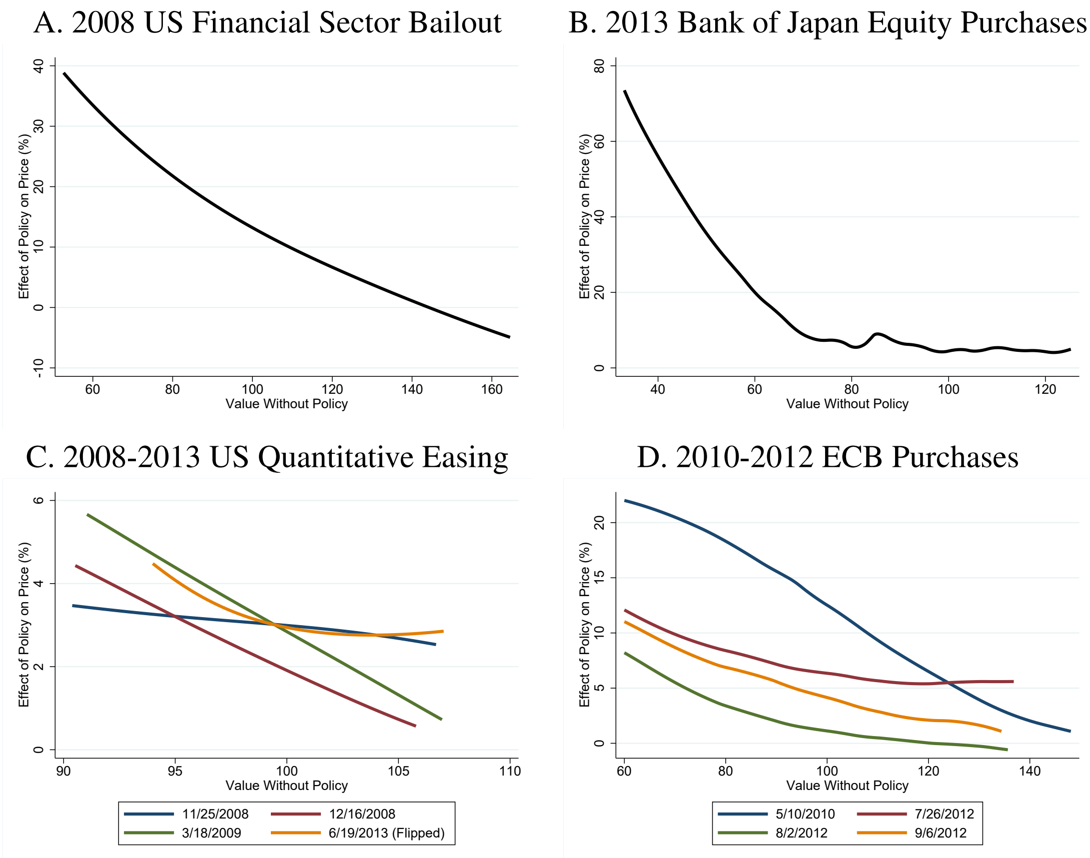

Figure 7: Price Support Function Across Policy Announcements
This figure reports price support as a function of the pre-policy price, normalized to 100 before the announcement across policy announcements. The events are described in the text.

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." This announcement was widely perceived as communicating a promise to backstop the financial sector. Consistent with this interpretation, Veronesi and Zingales (2010) find large positive responses of both bank equity and bank debt.[27] Our approach offers a way to quantify the state-contingent impact of this promise.

We use option prices on a financial sector index, the Financial Select Sector SPDR Fund. The price support function — Panel A of Figure 7 — reveals that the policy acted as a put option to the financial sector. For example, we see price support of 40% if the equity of the financial sector were to fall by 50% and no price support if the equity of the financial sector increased by 50%.

This asymmetry suggests that the policy was largely effective because of larger interventions in poor states of the world. Still, we cannot rule out that this stabilization is driven by out-of-

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

equilibrium actions along the lines of statements by then-Treasury Secretary Paulson that "if you've got a bazooka, and people know you've got it, you may not have to take it out."

2013 Bank of Japan equity purchases. The top right panel of Figure 7 studies Japan on April 4, 2013 after a speech given by Bank of Japan governor Haruhiko Kuroda in which he outlined a plan to use "every means available" to drive up inflation through large purchases of government bonds and equities. 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 announcement to estimate the price support function. Here again we find a strongly asymmetric response, with large price support for adverse states and a flat price support of around 5% for good states above the current price.

2008-2013 US quantitative easing. Panel C of Figure 7 looks at several US quantitative easing (QE) announcements 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 likely depend on state variables besides the price of the 10-year Treasury, which our measurement does not capture.

We first focus on three announcements introducing QE following Vissing-Jorgensen and Krishnamurthy (2011), each of which contained significant news of increased asset purchases. These are the initial announcement of large scale asset purchases (LSAPs) on November 25, 2008 and the FOMC statements of December 16, 2008 and March 18, 2009. All three policy announcements see an increase in Treasury prices (fall in yields) and a price support function that is downward sloping. The magnitudes are fairly similar with about 3-5% price support in cases where prices fall, roughly 3% at current prices, and 1-3% when prices increase. The importance of downside support is also consistent with the observation of Greenwood and Vayanos (2014) that the response to the initial QE1 program is larger than implied by the multiplier they estimate using long-term variations in supply.

Interestingly, the initial policy announcement on November 25 is less asymmetric than the others. One interpretation is that, because the policy was new, it took multiple announcements for agents to solidify views that the policy was not a one-off and would be increased should conditions worsen.

We also consider the "Taper Tantrum" on June 19, 2013, where the Fed announced that purchases would decline and markets had a strong negative reaction. The sign is flipped in the plot so that the magnitudes across events are more easily comparable. 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 also an associated decline in conditional price support from future purchases. This highlights that our approach can not only identify the presence of a policy put, but also identifies cases where this policy put is removed or weakened.

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 findings suggest that one channel that leads to powerful announcement effects of QE is the expectation of stronger interventions in bad states.

2010-2012 ECB asset purchases. We now look at announcements of asset purchases by the European Central Bank from 2010 to 2012. We use 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 high risk countries in the Eurozone, which is most directly where the asset purchase announcements were aimed. These options are not available for the 2010-2012 period. However, Krishnamurthy et al. (2018) show a very broad asset price response to the announcements. In fact, they estimate that for the vast majority of the events stock markets respond strongly and in line with debt markets.[28] For the set of events we consider, where the stock market tracks developments in debt markets, we use options on the Euro Stoxx index to assess conditional purchases. The underlying assumption is that any conditional purchases will be correlated with the stock index value and that the stock index will respond to purchases in those states.

Panel D of Figure 7 shows strong effects of conditional policy from these announcements. Across all four announcements, the extra support in the left tail explains about 20% of the overall reaction on average. This is expected as part of the goal of the announcements was to promise to do

[28]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. These broader effects on asset prices allow us to use stock options in place of options on sovereign bonds. Relative to Krishnamurthy et al. (2018), we omit the August 7 2011 announcement where we see a weaker response in sovereign bonds and no response in Euro Stoxx.

"whatever it takes." While the exact quantities implied by this promise were vague, the intention to commit to promises in this case was explicit.

Interestingly, the July 26, 2012 speech by Mario Draghi has relatively weaker asymmetry in the price support function compared to the follow on announcements on August 2 and September 6. The July 26 speech was strong in language but did not give a specific policy plan. Once this speech was followed with concrete announcements by the ECB a week later (the August 2 announcement), the market priced in stronger conditional price support. The August 2 and September 6 announcements included the Outright Monetary Transactions (OMT) which were aimed at purchasing Eurozone sovereign bonds under certain conditions. These findings emphasize that our method captures the markets reaction to announcements, rather than simply what the announcement states.

April 9, 2020 high-yield announcement. The announcement of corporate bond purchases by the Federal Reserve on March 23 focused on investment grade bonds. However, the Fed made an additional announcement on April 9, 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. This announcement is useful because we should expect the opposite patterns for high-yield versus investment-grade bonds compared to March 23. This is exactly what we see. Appendix Figure OA.7 plots the price support from this announcement following our same methodology applied to options on a 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 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.

4. Implications of Promises for Market Dynamics: Are the Effects of Asset Purchases Getting Weaker?

We have shown that many policy announcements appear to convey a put option of more intervention in bad states. If the perception of these policy puts is persistent, this will have long-term implications for market dynamics. For example, the market might believe that the Federal Reserve is now more likely to step into the corporate bond market whenever it gets distressed, which would affect the response of bond prices to bad economic news, and lower ex-ante tail risk in these markets.[29] Furthermore, promises in early policy announcement can lead to a weaker response of asset prices to subsequent announcements, despite no change in policy multiplier. This occurs because latter announcements are at least partially anticipated and because they do not contain as much new information about future interventions.

The broad pattern of response to QE announcements is consistent with this prediction. Figure 8 collects estimates of this response from the literature for the US, UK, and Europe: black bars report the yield change at announcements and gray bars the multiplier when the announcement comes with specific quantities. Both yield response and multiplier are strongest at early announcement and quickly decay, an observation also made by Hesse et al. (2018), Meaning and Zhu (2011), and Bernanke (2020).[30]

A challenge to interpreting this evidence with the promises view is that expectations of interventions for later announcements might not be driven by other factors than the realization of the initial promises. In the UK setting, one can take advantage of survey data to measure expectations and focus on the unexpected component of the announcements. Specifically, Busetto et al. (2022) construct measures of surprise intervention and relate them to yield responses for Gilts. We reproduce their data below, fitting a regression line through all announcements except QE1.[31] We see that while the linear regression fits well the data, the response to QE1 appears about twice as powerful as predicted by this relation. This observation supports our second mechanism, whereby the presence of promises in early announcements leads to seemingly inflated multipliers.

While this interpretation is consistent with our evidence from option prices, we cannot completely rule out that multipliers have changed over time. For example, it could be that multipliers are larger when the economy is in particularly bad shape, and this corresponds with when initial QE announcements occured. Some of the evidence around the COVID-19 shock of 2020, another period of significant market distress, cuts against this view as fully explaining the patterns in the

[29] Appendix Section K.2 provides suggestive evidence that the covariance of corporate bond returns with risk and economic news has changed after the Fed announced their first ever direct intervention in this market.

[30] Relatedly, Fabo et al. (2021) finds that estimates of the effect of QE vary significantly in the literature, in part due to which announcements are used.

[31] We thank these authors for sharing their data with us.

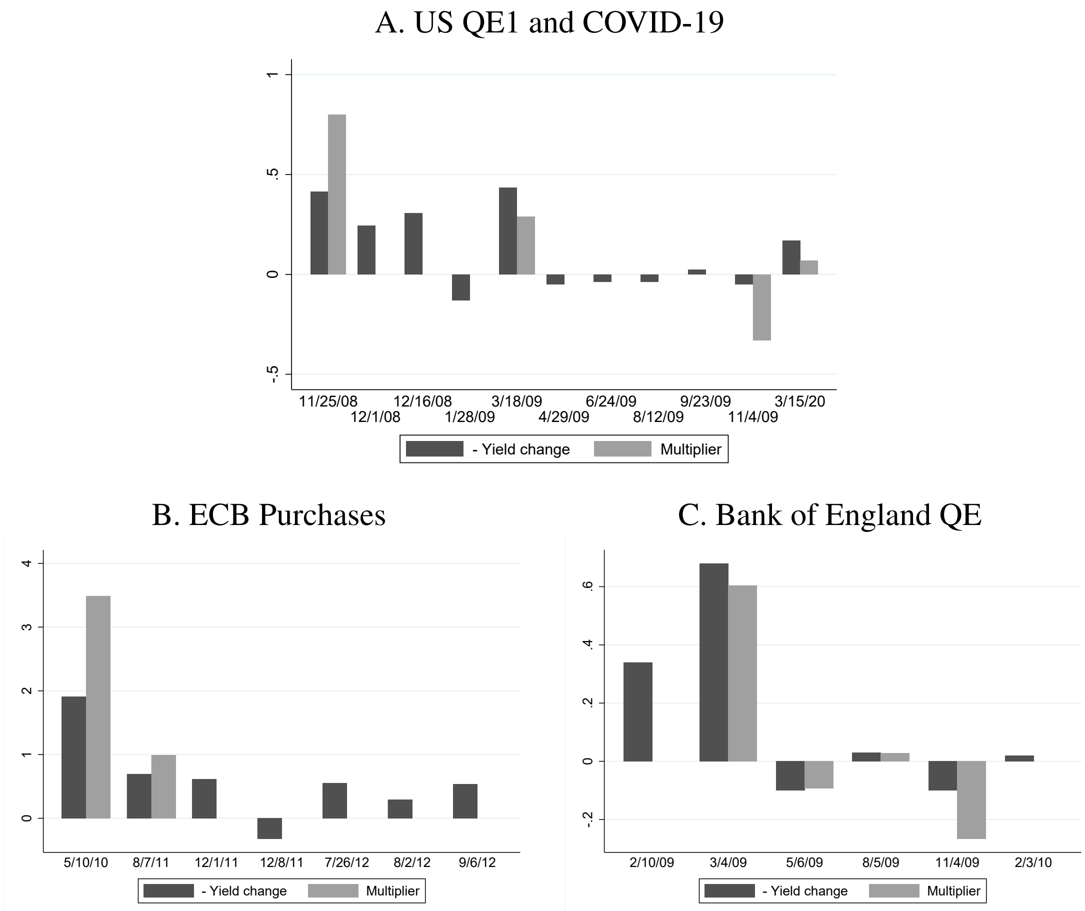

Figure 8: Weakening Announcement Effects of Asset Purchases

This figure plots announcement effects to asset purchase announcements made by the US Federal Reserve, Bank of England, and the ECB. See Appendix Section K.1 for details; numbers come from various studies including Joyce and Tong (2012), Meaning and Zhu (2011), Gagnon et al. (2018), Vissing-Jorgensen (2021), and Krishnamurthy et al. (2018).

data. While within three weeks of the initial March 15 announcement, the Fed had purchased over $1 trillion in Treasuries, the price response remained modest. For example, Vissing-Jorgensen (2021) reports that "an increase of 0.1 (buying 10% of supply) leads to a 5.35 bps larger decline in yields", at least an order of magnitude lower than multiplier estimates focusing on other events.[32] Furthermore, these results for Treasuries during COVID contrast sharply with the large response for corporate bonds during the same period. A natural interpretation is that the key difference between these two interventions is that the Fed had never before purchased corporate bonds and thus the announcement was a surprise.

[32] These estimates accumulate not only the response to announcements but also to realized purchases. A caveat is that these interventions might have been of market-functioning nature, and worked through a different mechanism.

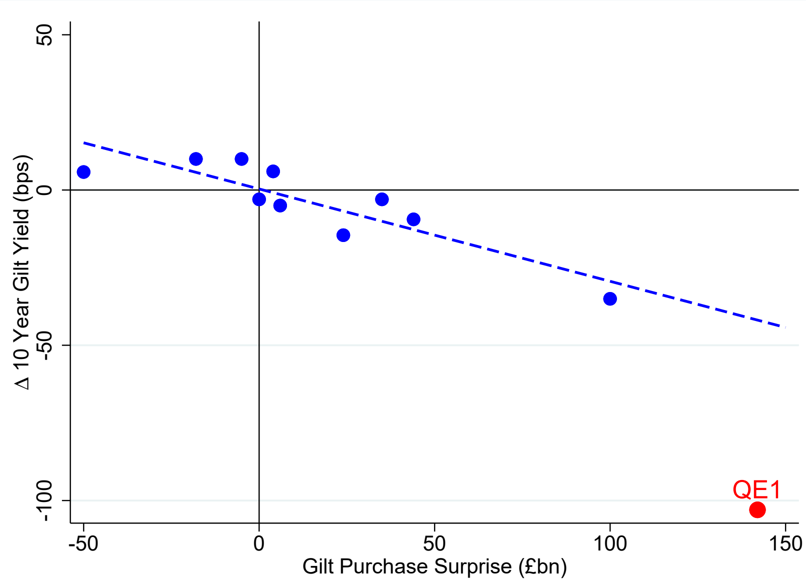

Figure 9: Surprise Purchases and Response of Gilt Yields. This figure uses data from Chart 5 of Busetto et al. (2022). Each point represents the the change in 10-year Gilt yield and the purchase surprise formed using survey data for a Bank of England QE announcement (see Busetto et al. (2022) for details). The dashed blue line corresponds to a regression of the yield surprise on the purchase surprise, excluding the first announcement (red point labeled QE1).

5. Conclusion

We provide a framework and methodology to evaluate the state-contingent impact of policy, which we apply to several policy announcements. We find a large role for state-contingent policy that indicates more intervention if conditions worsen. In our main empirical setting, the announcement of corporate bond purchases during the COVID-19 crisis, we find evidence that markets expected five times more price support in crash scenarios relative to the median state and significantly more relative to good states. This policy put of significant expansion in the size of the intervention in bad states explains a large share of the market response to the announcement. We extend our analysis to many other policy announcements in the US, UK, Euro area, and Japan, and find that perceptions of downside support are pervasive. The state-contingent policy view also helps rationalize jointly the apparently large multipliers from initial policy announcements and weakening or disappearing announcement effects for later announcements of the same policy.

References

Acharya, Viral V, Ryan Banerjee, Matteo Crosignani, Tim Eisert, and Renée Spigt, 2022, Exorbitant privilege? quantitative easing and the bond market subsidy of prospective fallen angels, Working paper, National Bureau of Economic Research.

Arora, Rohan, Sermin Gungor, Joe Nesrallah, Guillaume Ouellet Leblanc, and Jonathan Witmer, 2021, The impact of the bank of canada's government bond purchase program, Working paper, Bank of Canada.

Asquith, Paul, Andrea S Au, Thomas Covert, and Parag A Pathak, 2013, The market for borrowing corporate bonds, *Journal of Financial Economics* 107, 155–182.

Balthrop, Justin, and Jonathan Bitting, 2022, Implied future policy promises and firm leverage .

Barraclough, 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.

Bauer, Michael D, Carolin Pflueger, and Adi Sunderam, 2022, Perceptions about monetary policy, Working paper, National Bureau of Economic Research.

Becker, Bo, and Efraim Benmelech, 2021, The resilience of the us corporate bond market during financial crises, Working paper, National Bureau of Economic Research.

Bernanke, Ben S, 2020, The new tools of monetary policy, *American Economic Review* 110, 943–983.

Bianchi, Francesco, Martin Lettau, and Sydney C Ludvigson, 2022, Monetary policy and asset valuation, *The Journal of Finance* 77, 967–1017.

Bocola, Luigi, and Alessandro Dovis, 2019, Self-fulfilling debt crises: A quantitative analysis, *American Economic Review* 109, 4343–77.

Bouveret, Antoine, and Jie Yu, 2021, *Risks and Vulnerabilities in the US Bond Mutual Fund Industry* (International Monetary Fund).

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.

Bretscher, Lorenzo, Lukas Schmid, Ishita Sen, and Varun Sharma, 2022, Institutional corporate bond pricing, *Swiss Finance Institute Research Paper* .

Busetto, Filippo, Matthieu Chavaz, Maren Froemel, Michael Joyce, Iryna Kaminska, and Jack Worlidge, 2022, Qe at the bank of england: a perspective on its functioning and effectiveness, *Bank of England Quarterly Bulletin Q* 1.

Campbell, John Y, and Robert J Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of financial studies* 1, 195–228.

Carr, Peter, and Liuren Wu, 2009, Variance risk premiums, *The Review of Financial Studies* 22, 1311–1341.

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.

Chaudhary, Manav, Zhiyu Fu, and Jian Li, 2023, Corporate bond multipliers: Substitutes matter, Working paper, Working paper, Columbia Business School.

Cieslak, Anna, and Annette Vissing-Jorgensen, 2021, The economics of the fed put, *The Review of Financial Studies* 34, 4045–4089.

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* .

Darmouni, Olivier, Kerry Siani, and Kairong Xiao, 2022, Nonbank fragility in credit markets: Evidence from a two-layer asset demand system, *Available at SSRN* .

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.

Eisenbach, Thomas M, and Gregory Phelan, 2022, Fragility of safe asset markets, *FRB of New York Staff Report* .

Fabo, Brian, Martina Jančoková, Elisabeth Kempf, and L'uboš Pástor, 2021, Fifty shades of qe: Comparing findings of central bankers and academics, *Journal of Monetary Economics* 120, 1–20.

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* .

Geman, Hélyette, Nicole El Karoui, and Jean-Charles Rochet, 1995, Changes of numeraire, changes of probability measure and option pricing, *Journal of Applied probability* 32, 443–458.

Gertler, Mark, and Peter Karadi, 2011, A model of unconventional monetary policy, *Journal of monetary Economics* 58, 17–34.

Giesecke, Kay, Francis A. Longstaff, Stephen Schaefer, and Ilya Strebulaev, 2011, Corporate bond default risk: A 150-year perspective, *Journal of Financial Economics* 102, 233 – 250.

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.

Greenwood, Robin, and Dimitri Vayanos, 2014, Bond supply and excess bond returns, *The Review of Financial Studies* 27, 663–713.

Grinblatt, Mark, and Kam-Ming Wan, 2020, State pricing, effectively complete markets, and corporate finance, *Journal of Corporate Finance* 60, 101542.

Gürkaynak, Refet S, Brian P Sack, and Eric T Swanson, 2004, Do actions speak louder than words? the response of asset prices to monetary policy actions and statements, *The Response of Asset Prices to Monetary Policy Actions and Statements (November 2004)* .

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.

Haddad, Valentin, Alan Moreira, and Tyler Muir, 2024, Asset purchase rules: How qe transformed the bond marke, Working paper, Working paper.

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.

Hanson, Samuel G, and Jeremy C Stein, 2015, Monetary policy and long-term real rates, *Journal of Financial Economics* 115, 429–448.

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.

Hilscher, Jens, Alon Raviv, and Ricardo Reis, 2022, How likely is an inflation disaster? .

Hubert, Paul, Christophe Blot, Caroline Bozou, and Jérôme Creel, 2024, Same actions, different effects: The conditionality of monetary policy instruments, *Journal of Monetary Economics* 103596.

Joyce, Michael AS, and Matthew Tong, 2012, Qe and the gilt market: a disaggregated analysis, *The Economic Journal* 122, F348–F384.

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.

Kitsul, Yuriy, and Jonathan H Wright, 2013, The economics of options-implied inflation probability density functions, *Journal of Financial Economics* 110, 696–711.

Kremens, Lukas, and Ian Martin, 2019, The quanto theory of exchange rates, *American Economic Review* 109, 810–43.

Krishnamurthy, Arvind, Stefan Nagel, and Annette Vissing-Jorgensen, 2018, Ecb policies involving government bond purchases: Impact and channels, *Review of Finance* 22, 1–44.

Ma, Yiming, Kairong Xiao, and Yao Zeng, 2020, Mutual fund liquidity transformation and reverse flight to liquidity, Working paper.

Martin, Ian, 2017, What is the expected return on the market?, *The Quarterly Journal of Economics* 132, 367–433.

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* .

O'Hara, Maureen, and Xing Alex Zhou, 2021, Anatomy of a liquidity crisis: Corporate bonds in the covid-19 crisis, *Journal of Financial Economics* 142, 46–68.

Pflueger, Carolin, and Gianluca Rinaldi, 2022, Why does the fed move markets so much? a model of monetary policy and time-varying risk aversion, *Journal of Financial Economics* 146, 71–89.

Piazzesi, Monika, 2005, Bond yields and the federal reserve, *Journal of Political Economy* 113, 311–344.

Rachev, Svetlozar T, and Ludger Rüschendorf, 1998, *Mass Transportation Problems: Volume I: Theory*, volume 1 (Springer Science & Business Media).

Swanson, Eric T, 2011, Let's twist again: a high-frequency event-study analysis of operation twist and its implications for qe2, *Brookings Papers on Economic Activity* 2011, 151–188.

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, 2021, The treasury market in spring 2020 and the response of the federal reserve, *Journal of Monetary Economics* 124, 19–47.

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.

Online Appendix

A. From Option Prices to Risk-Neutral Densities: Implementation.

Our numerical approach to recover the state price density is standard. 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. However, we are careful to not extrapolate the curves. We thus only recover the option-implied risk neutral density for the range where we have option prices. 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 use a spline smoother on the implied volatility curve and choose a step size and smoothing parameters that produce non-negative risk-neutral densities. The required step and smoothing parameter varies according to the quality of the underlying option data. We then compute first and second differences to recover the implied risk-neutral cumulative distribution function or probability density function. As a consequence of our choice to not extrapolate the implied volatility curves we only obtain the CDF for finite intervals. While this limitation precludes the usual application of the result of Breeden and Litzenberger (1978) (pricing arbitrary option contracts), we we are still able to recover exactly the function $g(.)$, only over a finite interval.

B. Economic Model

We introduce a simple model in the style of Vayanos and Vila (2021) and Greenwood and Vayanos (2014) to understand the economic effects of purchase announcements and to further clarify the assumptions made in our main empirical section. We adapt the model from Vayanos and Vila (2021) because it is the leading framework the literature has used to think about the direct asset pricing effects of asset purchases (e.g., Bernanke (2020)).

There are three dates, 0, 1, and 2. There is a risky asset in unit supply paying off X at date 2 where we assume X is lognormal with $ln(X) \sim N(\mu, \sigma^2)$. There are three agents: a specialized arbitrageur, inelastic investors, and a policy maker (e.g., a central bank). At date 0, the policy maker announces asset purchases to be implemented at date 1.

The specialized arbitrageur has log utility over final wealth and chooses their portfolio allocation in periods 0 and 1 between the risky asset and a risk-free asset. We take the risk-free rate as exogenous and label its gross return R_f; denote $r_f = ln(R_f)$. We keep the risk-free rate constant for simplicity but this isn't necessary for our conclusions. The arbitrageur is endowed with shares of the risky asset worth W_0 at date 0.

Inelastic investors have W_I dollars of the risky asset at date 0 and are price inelastic. They can be thought of as insurance companies, pension funds, or other institutions that hold a large fraction of the bond market but do not trade frequently or are inattentive. In contrast, the arbitrageur should be thought of as a dealer bank, hedge fund, or other active trader. Inelastic investors have a stochastic demand shock at date 1 that leads them to sell $\tilde{B}$ dollars of the asset. It is convenient to define $\tilde{b} = \tilde{B}/W_1$ as the dollar sales made by the inelastic investors as a fraction of the arbitrageurs' date 1 wealth.[33] This fire sale shock is the only source of date 1 uncertainty. The fire sale shock depresses prices but is independent of fundamentals of the asset payoff. While the COVID-19 episode fits well with this fire-sale interpretation, one could alternatively consider fundamental cash flow shock at date 1, that is a shock to date 1 cash-flow expectations. Such an approach might be more in line with other episodes with asset purchase announcements such as quantitative easing.

We solve for date 1 prices and quantities, then use these to arrive at date 0 prices. The arbitrageur's first order condition at date 1 can be approximated by

$$\alpha_1 = \frac{E_1[ln(X/P_1)] - log(R_f)}{Var_1(ln(X/P_1))} = \frac{\mu - p_1 - r_f}{\sigma^2} \tag{OA.1}$$

where α_1 is the arbitrageur's portfolio share in the risky asset, X/P_1 denotes the gross return on the asset from date 1 to date 2, $p_1 = ln(P_1)$, and $E_1[.]$ denotes the conditional expectation taken at time 1.

The central bank purchases q of the asset at date 1, where we denote q as a fraction of the arbitrageur's date 1 wealth. We allow this amount q to be stochastic from the perspective of time 0, and correlated with the fire sale $\tilde{b}$. For example, the central bank could purchase more in states where the fire sale shock is larger to dampen price dislocations.

Because the arbitrageur absorbs the net supply imbalance, market clearing for the asset at date 1 implies that $\alpha_1 - \tilde{b} + q = 1$ so that $\alpha_1 = 1 + \tilde{b} - q$. Combining this with the arbitrageur's first order condition, and solving for p_1, gives

$$p_1 = \sigma^2(1 - \tilde{b} + q) + \mu - r_f \tag{OA.2}$$

This equation gives a multiplier σ^2 for the effect of asset purchases q on the (log) price p_1. Higher purchases q remove the asset from the arbitrageur's balance sheet and raise prices, and vice versa for fire sales $\tilde{b}$. Since q is normalized by the arbitrageur's wealth, $\frac{P_1}{W_1}\sigma^2$ gives the multiplier in the more standard units of a fraction of total market capitalization. If the central bank purchased 1 percent of the total market capitalization of the asset, the price would increase by $\frac{P_1}{W_1}\sigma^2$ percent. If arbitrageur capital is a small portion of the wealth invested in the risky asset, the multiplier will be large because purchases or sales are absorbed by a relatively small amount of active capital. We also note that the multiplier is constant and does not depend on the realization of the state b at date 1.

The date 1 pricing equation shows that this framework can naturally explain the "weakening" effect of follow on purchase announcements. Consider the difference between p_1, the price of date 1 after the actual purchases are implemented, and $E[p_1|b]$, the price of the asset right after the selling shock b is realized but just before the date 1 purchases q are implemented,

$$p_1 - E[p_1|b] = \sigma^2(q - E[q|b]) \tag{OA.3}$$

Only purchases that deviate from what was expected given the announcement in date 0 have any effect, and when the policy maker simply fulfills their promises the effect is exactly zero. This

[33] This shock can also be interpreted as affecting total supply instead of sales by the inelastic investors. For example, it could be the consequence of firms' large debt issuance needs during COVID-19.

zero effect does not mean that the date 1 intervention is ineffective, but simply that it was already reflected in the date-0 price response.[34]

Purchases have no effect on the exogenous asset fundamentals X in the model, and thus they move prices only through their affect on the asset risk premium from date 1 to 2. This also implies that date 2 pricing kernel will change with asset purchases q. Because the agent has log utility, the pricing kernel is given by W_1/W_2 or the inverse return on the arbitrageur's wealth from date 1 to date 2. Labeling the pricing kernel as m_2 we have

$$m_2 \;=\; \left(\alpha_1 R_2 + (1-\alpha_1) r_f\right)^{-1} = \left((1+\tilde{b}-q)R_2 + (q-\tilde{b})r_f\right)^{-1} \qquad \text{(OA.4)}$$

Intuitively, the pricing kernel changes when purchases q are made because this is when risk is actually removed from the arbitrageur's balance sheet. This pricing kernel effect will be reflected in date 0 prices, even in the case where the pricing kernel from 0 to 1 remains unchanged. To see why, note that the time 0 price is the discounted time 1 price but the discounting between 0 and 1 remains unchanged. Since the time 1 price has risen, the time 0 price will also rise. This shows how asset purchases can impact prices by moving future risk-premium even if they do not impact the pricing kernel between the announcement date and the asset purchase date as we assume in our baseline analysis.

At date 0, the arbitrageur's first order conditions for the risky and risk-free asset, respectively, give $E_0\left[\frac{W_0}{W_1}\frac{P_1}{P_0}\right] = 1$ and $E_0\left[\frac{W_0}{W_1}R_f\right] = 1$. Market clearing at date 0 implies $\alpha_0 = 1$ so that the arbitrageur invests fully in the risky asset. Since the inelastic agents do not buy or sell at date 0, the arbitrageur must hold on to the shares they are endowed with in equilibrium. This implies $W_1/W_0 = R_1 \equiv P_1/P_0$ where R_1 is the risky asset return. This trivially means that $E_0\left[\frac{W_0}{W_1}\frac{P_1}{P_0}\right] = 1$. Using $E_0\left[\frac{W_0}{W_1}R_f\right] = 1$, we have

$$P_0 = \frac{1}{R_f}\frac{1}{E_0\left[\frac{1}{P_1}\right]} \qquad \text{(OA.5)}$$

It follows from the expression above that a date 0 announcement by the central bank to purchase a constant share of the asset at date 1 does not change the pricing kernel between dates 0 and 1. Because the intervention pushes up prices proportionally both at date 0 and date 1 it does not change asset risk between these dates. Thus, our framework recovers the correct price support function.[35]

Deterministic purchases do not affect the risk premium at date 0 for two reasons. The first is that purchases, whether deterministic or state-dependent, do not remove risk from the arbitrageur

[34]While we consider here one announcement date paired with one purchase date, it is straightforward to extend the model to have multiple announcement dates paired and multiple implementation dates dates. The key intuition is that the first announcement reveals the policy rule that the policy maker will adopt in the follow on announcements. In this way the response to the first announcement is driven not only by the immediate follow-on purchases but also by the follow-on announcements which themselves will appear ineffective.

[35]Note that $g(p_1) = (F^{-1}_{p_1'}(F_{p_1}(p_1)) - 1)$, where F and F' are the risk-neutral distributions of prices in date 1 before and after the announcement. Given the kernel implied by the specialist model we have $F_{p_1}(y) = F^P(y)\frac{1}{R_f y E_0[\frac{1}{y}]}$ and $F_{p_1'}(y) = F^P_{p_1'}(y)\frac{1}{R_f y E[\frac{1}{y}]}$ where subscript P stands for the natural probability distribution. Plugging an intervention that buys a constant share of the asset market capitalization $p_1' = (g_a+1)\times p_1$ to the equation above recovers $g(p_1) = g_a$.

balance sheet until date 1. Removal of risk only impacts the pricing kernel from date 1 forward. The second is that deterministic purchases do not change the risk of the asset because it moves prices uniformly up. Stochastic purchases can impact the pricing kernel between date 0 and date 1 through the effect they have on the risk of the asset between dates 0 and 1. Plugging in the date-1 price of the asset gives the date-0 price as

$$p_0 = \mu + \sigma^2 - ln\left(E_0\left[exp\left(\sigma^2(\tilde{b} - q)\right)\right]\right) \tag{OA.6}$$

It is immediate from this expression that the date 0 price reflects the announcement of purchases made at date 1.

In summary, we have provided a model in the style of Vayanos and Vila (2021) where: (1) prices may be initially "dislocated" or depressed because of fears of future fire sales rather than cash flows (though the source of depressed prices is effectively irrelevant), (2) purchases affect asset prices through their affect on future risk premiums, (3) announcements of purchases affect prices even if purchases happen later, (4) constant purchases of assets require no additional risk adjustment between announcement and purchases, and (5) state-dependent purchases (state-dependent q) can alter the pricing of risk between announcement and purchases through their effect on the risk of the asset. In this last case, one needs to adjust our methodology to account for changes in the risk of the asset following Proposition 3.

C. Relaxing the Stochastic Discount Factor Invariance

We first review some standard model results and facts about how changes in short-term stochastic discount factor affect asset prices. We then derive the results of Section 1.3.4 which generalize our approach to recover the conditional price support.

C.1 The Typical Effect of Changes in Short-Term Stochastic Discount Factor

We review a couple of standard results on how changes in short-term SDF affect asset prices, in some simple models, then in terms of unconditional variation in the data.

A simple benchmark. A natural intuition is that because SDFs tend to overweigh bad outcomes, changes in SDF that affect disproportionately bad outcomes will automatically create price support functions g that are strongly asymmetric. This intuition turns out to be incorrect. We illustrate why through an analytical example with standard assumptions.

Assume the true distribution of the price at date 1 is log-normal, $\log(p_1) \sim \mathcal{N}(\mu, \sigma^2)$, and that the risk-free rate is equal to 0. To consider the effect of a pure change in SDF, we assume that this distribution is unchanged by the policy. We assume that there are two values $\gamma > \gamma'$ such that the SDF before and after the policy are $ap_1^{-\gamma}$ and $a'p_1^{-\gamma'}$. These functional forms correspond both to what happens in an equilibrium model with CRRA preferences and to the pricing kernel in the Black-Scholes model. The two constants a and a' must be such that the expectation of the pricing kernels are equal to 1, to coincide with a constant risk-free rate: $a = E(p_1^{-\gamma})^{-1} = \left[e^{-\gamma\mu + \frac{1}{2}\gamma^2\sigma^2}\right]^{-1} = e^{\gamma\mu - \frac{1}{2}\gamma^2\sigma^2}$.

The pricing kernels are asymmetric, with much larger values for lower values of p_1, and lowering the "risk aversion" γ makes the pricing kernel less asymmetric. However, we show that this does not result in an asymmetric price support.

Indeed, the risk neutral distribution of p_1 without policy is:

$$f^{\mathbb{Q}}(p_1) = \frac{1}{p_1}\frac{1}{\sqrt{2\pi\sigma^2}}e^{-\gamma(\log(p_1)-\mu)-\frac{1}{2}\gamma^2\sigma^2}e^{-(\log(p_1)-\mu)^2/2\sigma^2} \tag{OA.7}$$

$$= \frac{1}{p_1}\frac{1}{\sqrt{2\pi\sigma^2}}\exp\left[-\frac{1}{2\sigma^2}\left(\log(p_1)-\left(\mu-\gamma\sigma^2\right)\right)^2\right] \tag{OA.8}$$

Under the risk-neutral distribution, the price p_1 follows a log-normal distribution with parameters $\mu - \gamma\sigma^2$ and σ^2. A similar result obtains after policy with parameters $\mu - \gamma'\sigma^2$ and σ^2. The two distributions are log-normal with different means, so a transport from one to the next is simply multiplying the values by the exponential of difference in means. This gives immediately:

$$g(p_1) = e^{(\gamma-\gamma')\sigma^2} - 1 \tag{OA.9}$$

We see that despite the change in discount factor disproportionately affecting the left tail of the distribution, the price support function is constant. Not surprisingly, in the real world, returns are not exactly log-normal, and the SDF need not follow this functional form, so we now discuss the unconditional data.[36]

Typical empirical variation in asset prices. Empirically, changes in *short-term* discount rates only explain a small fraction of the variation in asset prices. For the price level, one can decompose: $\log(P_t) = \log(E_t(P_{t+1})) + \log(E_t(R_{t+1})$. Most variation comes from the first term relative to the second term, even if we scale the price by a quantity like dividends to make it stationary). This statement is equivalent to the observation that, at short horizons, prices behave very closely to random walks. This might seem at odds with the modern view (see, e.g. Campbell and Shiller (1988)) that "most price variation is driven by discount rates" but it is not. In this latter type of analysis, it is the cumulative *long-term* effect of discount rates that drives price variation.

If one is interested in the shape of the distribution, a simple moment to look at is is the risk-neutral variance, for example backed out from variance swap rates. Variation in this quantity is driven by changes in the actual variance — often referred to as realized variance — and in the variance risk premium. Carr and Wu (2009) compute these quantities across multiple securities and find systematically that the standard deviation of realized variance is close to that of risk-neutral variance (see their Table 2).

Of course, our specific settings reflect somewhat unusual market conditions, which leads us to construct additional measures to assess the impact of variation in short-term risk premia. See Section 1.3.4 for this discussion.

[36]If we stay within the same family of SDFs, the change in risk neutral probability at a price p_1 relative to at a reference level p_ref is: $\frac{f^{\mathbb{Q}'}(p_1)}{f^{\mathbb{Q}}(p_1)} / \frac{f^{\mathbb{Q}'}(p_{ref})}{f^{\mathbb{Q}}(p_{ref})} = \left(\frac{p_1}{p_{ref}}\right)^{\gamma-\gamma'}$. This still requires large changes in γ to obtain strong changes in probability. For example to have probabilities drop by a factor of 10 at 80% of the current price level relative to at the current price level, we need $\gamma' - \gamma \approx 10$, a very large change. Therefore, in such models, changes in SDF only generate asymmetric changes in the distribution if probabilities fall much slower than a normal distribution in the left tail.

C.2 Testing a Constant Price Support

We prove Proposition 2. First, notice that if, in equilibrium, the pricing kernel does not change, we are back to the setting of Proposition 1. We can then correctly recover the price support function. We show this is the case with a constant support for the family of pricing kernel introduced in Proposition 2.

To do so, we take a guess-and-verify approach. Denoting $g(p_1) = \bar{g}$, then $p_1'(s) = p_1(s)(1+g)$. If the pricing kernel is unchanged, then the value of the asset at date 0 increases by the same amount: $p_0' = E^{\mathbb{P}}[mp_1 \times (1+g)] = p_0(1+g)$. Therefore, in each state, we have:

$$m'(s) = \Theta\left(s, \frac{p_1'(s)}{p_0'}\right) \tag{OA.10}$$

$$= \Theta\left(s, \frac{p_1(s)(1+g)}{p_0(1+g)}\right) \tag{OA.11}$$

$$= \Theta\left(s, \frac{p_1(s)}{p_0}\right) \tag{OA.12}$$

$$= m(s). \tag{OA.13}$$

This confirms our guess and concludes the proof.

C.3 Adjusting the Estimates for a Response of the Stochastic Discount Factor

We prove Proposition 3.

The first part of the result is to notice that, as long as we find a distribution for $p_1(s)$ which is unaffected by the announcement, we can use the same idea as in our baseline of matching the quantiles of $p_1(s)$ and $p_1'(s)$. In our baseline setting, the risk-neutral distribution of $p_1(s)$ was invariant. This is not the case anymore under Assumption 5 because the pricing kernel is affected by the intervention. In contrast, the physical distribution remains unchanged, but we cannot recover it from option contracts. Instead, we focus on an intermediate distribution, only affected by the exogenous part of the pricing kernel $\theta(s)$, the forward-neutral distribution. We define this distribution in equation (9), which we repeat here:

$$dF^{\mathcal{N}}(p_1) = E[\theta(s)|p_1]dF(p_1). \tag{OA.14}$$

If we can measure this distribution, we can apply the same reasoning as in our baseline method. Let us show how to measure this distribution using options.

Consider a contract that pays off $C_K(s)$ defined by

$$C_K(s) = \begin{cases} p_1(s)/p_0 & \text{if } p_1(s) \leq K \\ 0 & \text{if } p_1(s) \leq K. \end{cases} \tag{OA.15}$$

We define $C_K'(s)$ similarly after the announcement. The price of this contract coincides with the forward-neutral CDF:

$$E[m(s)C_K(s)] = E[E[m(s)C_K(s)|p_1]] \tag{OA.16}$$
$$= \int_{-\infty}^{K} E[\theta(s)|p_1]\frac{p_0}{p_1}\frac{p_1}{p_0}dF(p_1) \tag{OA.17}$$
$$= F^{\mathcal{N}}(K) \tag{OA.18}$$

Similarly, if we note $K' = K(1+g(K))$, we obtain:

$$E[m'(s)C'_{K'}(s)] = E[E[m'(s)C'_{K'}(s)|p_1]] \tag{OA.19}$$
$$= \int_{-\infty}^{K} E^{\mathbb{P}}[\theta(s)|p_1]\frac{p'_0}{p'_1(p_1)}\frac{p'_1(p_1)}{p'_0}dF^{\mathbb{P}}(p_1) \tag{OA.20}$$
$$= F^{\mathcal{N}}(K) \tag{OA.21}$$

Replicating the contract C_K**.** Finally, the remaining task is how to replicate the contract $C_K(s)$ using calls and puts. We show it is not much different from the case of the risk-neutral measure. We can rewrite:

$$C_K(s) = \frac{K}{p_0}1_{\{p_1 \leq K\}} - \frac{1}{p_0}\max\left(K - p_1, 0\right). \tag{OA.22}$$

To see this, note that both terms on the right-hand-side are equal to 0 when $p_1 \geq K$. When $p_1 < K$, the right-hand side becomes: $K/p_0 - K/p_0 + p_1/p_0 = p_1$.

The formula gives us a simple way to replicate the contract: purchase K/p_0 of a digital option with strike K and short $1/p_0$ of a put with strike K. Recall that the price of the digital is simply the derivative of the put price with respect to the strike — this is the contract we used to replicate the risk-neutral measure. So, we have:

$$F^{\mathcal{N}}(K) = E\left[m(s)C_K(s)\right] \tag{OA.23}$$
$$= \frac{K}{p_0}\frac{\partial Put}{\partial K}(K) - \frac{1}{p_0}Put(K), \tag{OA.24}$$

where $Put(K)$ is the price of put options as a function of the strike K. Figure OA.1 illustrates the comparison of this replicating portfolio with that of the risk-neutral measure.

D. The Order-Preserving Condition

We discuss the order-preserving condition of Proposition 1. We show a simple family of economic problems under which order-preserving interventions are optimal. We demonstrate that focusing on order-preserving estimates leads to conservative estimates of the asymmetry of a policy. Finally, we conduct a numerical example to illustrate the effect of using our method in a plausible setting in which the policy is not order-preserving.

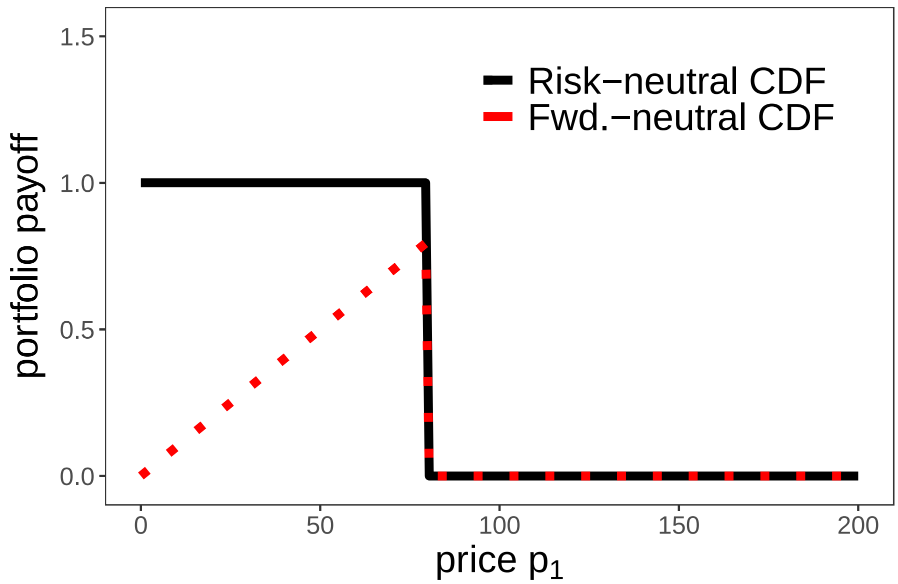

Figure OA.1: Estimating the risk-neutral and forward-neutral measures.
This figure reports the payoffs for contracts replicating the risk-neutral measure (solid black line) and forward-neutral measure (dotted red line) for the value $K = 80$.

A foundation for order-preserving policies. Assume that the policymaker has an objective function that depends on the post-policy distribution of the price $F_{p_1'}$, $\mathcal{U}(F_{p_1'})$, and the cost of changing the value in each state is $h(p_1' - p_1)$, with h an increasing convex function. This corresponds to a total cost $\int h(p_1'(s) - p_1(s))dF(s)$.

Among all policies reaching a given distribution $F_{p_1'}$, the policymaker will always choose one minimizing this total cost, that is the one minimizing $\int h(p_1'(p_1) - p_1)dF_{p_1}(p_1)$. Irrespective of the shape of h, the optimal transport is given by equation 5, and is order-preserving (see, for example, Rachev and Rüschendorf (1998)). More generally, monotonicity is a property that obtains generically for L^2-Wasserstein optimal transport problems; our problem is a special case with one dimension.

Why the order-preserving transport minimizes asymmetry. Consider a specific cutoff value K for the price p_1. One simple way to measure the asymmetry of the policy is to compute $E[p_1'(p_1)|p_1 > K] - E[p_1'(p_1)|p_1 < K]$. This quantity measures how high the with-policy price in "bad" states of the world is relative to in "good" states of the world. Across all functions $p_1'(p_1)$ generating the same distribution $F_{p_1'}$, this difference is maximized for the order-preserving transport. Indeed, the order-preserving transport puts all the highest values of p_1' to the right of the threshold, and the lowest values to the left. Higher values of this criterion reflect a lack of policy asymmetry, as it maintains bad states as far away from good states as possible.

A numerical example. We consider a variation of the whatever-it-takes example, where the true price support function has discrete jumps up at a series of cut-offs. Figure OA.2 reports the result of this exercise. The black line represents our assumption about the true price support function. With an assumption about the distribution of the underlying state (we assume a simple log-normal), we can compute an estimated price support function following our approach. This corresponds to the red line on the figure. The estimated price support function smoothes out the discrete increments, but otherwise exhibits a close quantitative behavior.

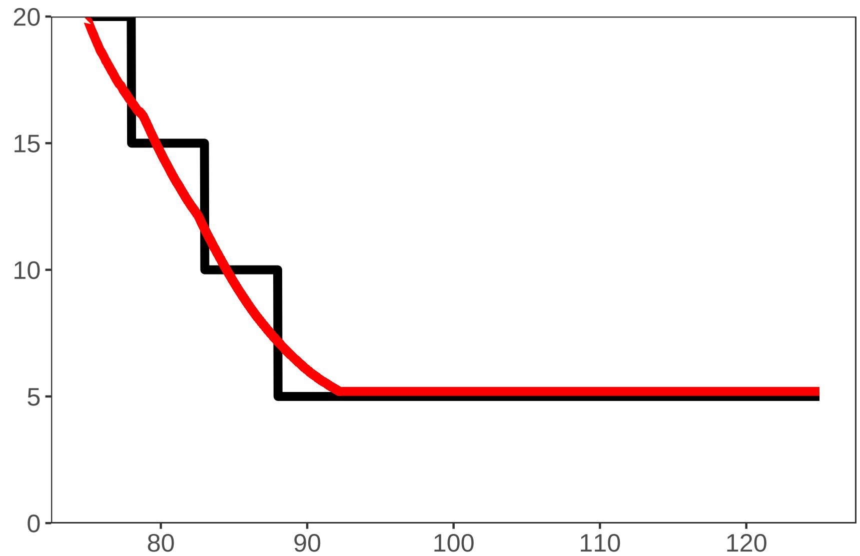

Figure OA.2: Price support function estimation with discrete jumps in policy.
The black line represents the actual price support function. We assume the underlying state is log-normal with mean level 100 and volatility of 7.5% and follow our method to estimate the price support function. The red line reports the estimate price support function.

E. The Effect of Anticipation

Our baseline interpretation of the results relies on a strong form of event-study assumption: just before the announcement, investors never thought about this policy, and right after they are sure of its implementation. This is what is colloquially referred to as an "MIT shock." Here we discuss how to evaluate whether our results are sensitive to the assumption by generalizing to a case where the policy is somewhat expected.

Adjusting for anticipated and uncertain policy. Let the probability of intervention be $\theta^- \leq 0$ before the announcement, and θ^+ after the announcement, with $\theta^- \leq \theta^+ \leq 1$. Naturally, we assume that the announcement increases the probability of intervention. Denote F_{p_1} the distribution of the date-1 price if no policy happens, and $F_{p_1'}$ the distribution of that price if the policy happens for sure. In this appendix, we drop the superscript $\mathbb{Q}$ for simplicity. Those two are the relevant counterfactuals that we need to recover the price support function g.

Then we have:

$$F^- = \left(1-\theta^-\right) F_{p_1} + \theta^- F_{p_1'} \tag{OA.25}$$

$$F^+ = \left(1-\theta^+\right) F_{p_1} + \theta^+ F_{p_1'}. \tag{OA.26}$$

Rearranging these equations allow to inverse the two distributions:

$$F_{p_1} = \frac{\theta^+ F^- - \theta^- F^+}{\theta^+ - \theta^-} \tag{OA.27}$$

$$F_{p_1'} = \frac{1}{\theta^+ - \theta^-}\left[(1-\theta^-)F^+ - (1-\theta^+)F^-\right]. \tag{OA.28}$$

The data allow us to measure F^- and F^+ which are both compounded lotteries of the risk-neutral distribution with policy, $F_{p_1'}$, and without, F_{p_1}. If we know the probabilities θ^- and θ^+, we can recover F_{p_1} and $F_{p_1'}$. We can then find the price support function g that transforms F_{p_1} into $F_{p_1'}$.

Special cases. In the case of an unanticipated and certain announcement $\theta^- = 0$ and $\theta^+ = 1$, we get back to our baseline approach:

$$F_{p_1} = F^-$$
$$F_{p_1'} = F^+.$$

If the announcement is anticipated but not uncertain, then $\theta^- = \theta_{ant}$, and $\theta^+ = 1$, which gives:

$$F_{p_1} = \frac{F^- - \theta_{ant} F^+}{1-\theta_{ant}}$$
$$F_{p_1'} = F^+.$$

Application to the 2020 corporate bond purchases. Figure OA.3 reports the price support function implied by a few alternative combinations of θ^- and θ^+ for our main empirical example (the March 23 announcement). The results are as one would expect. Anticipation makes the effect of the policy larger. Residual uncertainty about the implementation has a similar effect. However, both features do not strongly change the estimated asymmetry in price support across states.

F. 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 Treasury bond, the financial sector ETF (XLF), the future on the Nikkei index, the future on the European stock index (STOXX), and the CDX investment grade credit basket spread. We aim to use options of maturities close to three months which

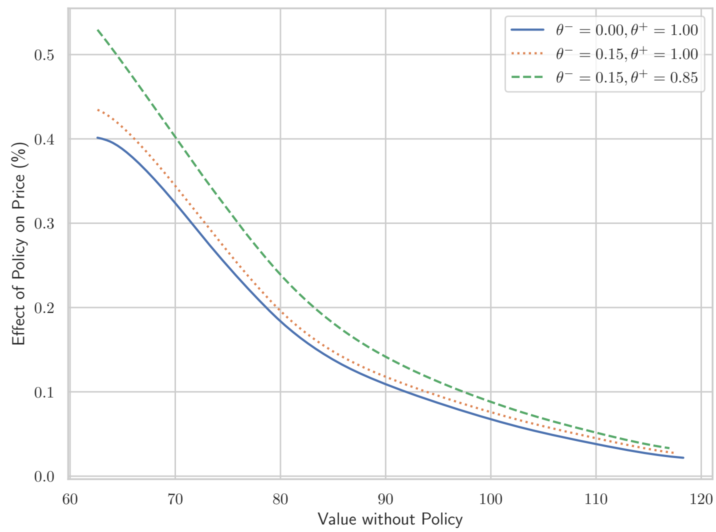

Figure OA.3: Effects of anticipation on price support. This figure shows the price support function under alternative assumptions about the probability of policy intervention before (θ^-) and after the announcement (θ^+).

are frequently the most liquid. Data from exchange traded instruments come from OptionMetrics. Data on CDX and their options come form Bloomberg and Markit.[37]

G. Details of Corporate Bond Purchase analysis

G.1 Event Study for Corporate Bond Purchases

The announcement of the SMCCF and PMCCF had a significant and immediate impact on corporate bond prices. Table OA.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 summarizes the announcement effect 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%, and the abnormal excess return is 10% (with controls for high-yield bonds and the stock market). The 14% return translates into around a $1 trillion 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%. A shorter one-day window provides better identification at the cost that it may take

[37] Ideally, one would like to use longer maturity options as well to study whether implicit promises are longer-term in nature. However, in practice, liquidity in the vast majority of these markets is heavily concentrated around or below three months.

Table OA.1: Bond Price Response to the March 23 Announcementt

This table shows the return on an investment-grade corporate bond ETF (LQD) on the announcement on March 23 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***	10.27**	7.37***	6.63***
	(3.78)	(1.25)	(0.01)	(0.07)
$r_t^{HighYield}$		0.54***		0.55***
		(0.04)		(0.04)
r_t^{SP500}		0.03		0.03
		(0.02)		(0.02)
N	2,988	2,988	2,988	2,988
R^2	0.11	0.87	0.09	0.87

the market time to process the announcement.[38] Haddad et al. (2021) show in higher 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.

G.2 How much did promises contribute to the overall price movement?

Figure OA.4 plots the implied price support function in our main exercise along with a dashed line for flat price support below the median price absent policy. We compute the expectations of each of these two price support functions to assess how much the extra support in the left tail increased the price at the time of announcement.

G.3 Computing the contribution of the left tail support to the announcement response

We show how to construct the estimates of contribution of support in the left tail to the announcement return that we report in Table 1.

We start by recovering the risk-neutral CDF $F^{\mathbb{Q}}$, the forward-neutral CDF $F^{\mathcal{N}}$ (as described in Appendix C.3), and the respective price support functions g and $g^{\mathcal{N}}$.

Baseline case. Consider first the case where the SDF is assumed to be invariant, as also computed in Section 2.3. I this case, we have:

[38] For this event, it is particularly desirable to have a narrow window given that volatility was very high and, in addition, the fact that the CARES acts was signed into law four days after the announcement.

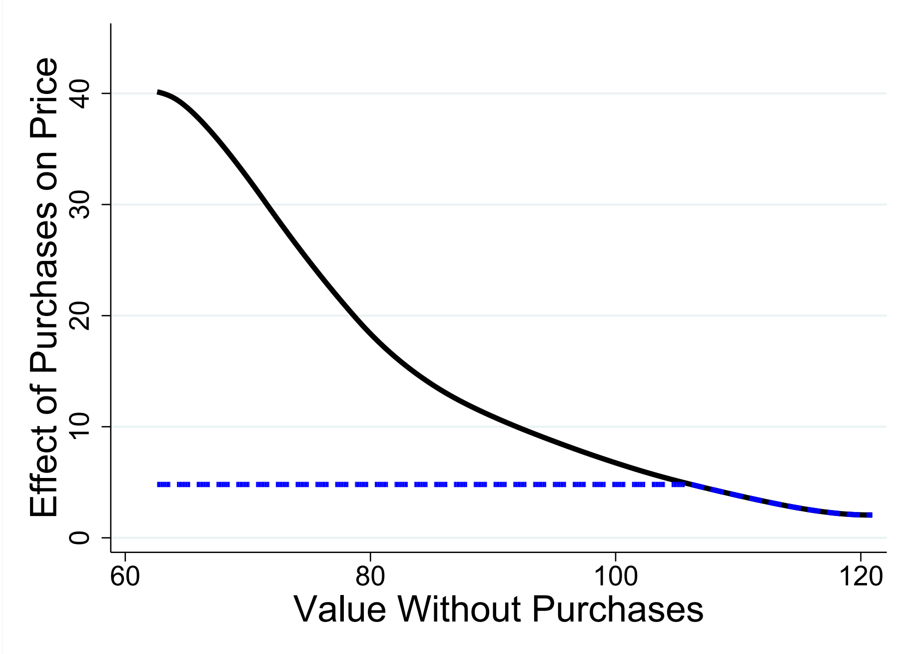

Figure OA.4: Counterfactual Price Support Function Without Downside Support.
This figure shows the implied price support (expressed in percent) as a function of the pre-policy price. The pre-policy price is normalized to 100 before announcement. The dashed line indicates a flat price support function below the median price absent policy.

$$p_0 = \int p_1(s) dF^{\mathbb{Q}}(s), \tag{OA.29}$$

$$p_0' = \int p_1(s)(1 + g(s)) dF^{\mathbb{Q}}(s), \tag{OA.30}$$

$$\tag{OA.31}$$

where p_0 and p_0' are the date-0 prices of the asset before and after the policy is announced. The after-announcement price if a counterfactual policy $\tilde{g}(s)$ was announced follows immediately:

$$\tilde{p}_0 = \int p_1(s)(1 + \tilde{g}(s)) dF^{\mathbb{Q}}(s) \tag{OA.32}$$

We compute this price for a price support function truncated at is median as in Figure OA.4. Table 1 reports

$$1 - \frac{\tilde{p}_0 - p_0}{p_0' - p_0}, \tag{OA.33}$$

where the fraction is the share of the price response we would have observed without the abnormal price support in the left of the price at announcement p_0, i.e. with $\tilde{g}(p) = g(p)$ when $p > p_{med}$ and $\tilde{g}(p) = g(p_{med})$ when $p \leq p_{med}$.

Endogenous SDF case. We now have that the SDF m and forward measure $dF^{\mathcal{N}}$ are given by

$$m(s) = \theta(s)\frac{p_0}{p_1(s)} \tag{OA.34}$$

$$dF^{\mathcal{N}}(s) = \theta(s)dF(s), \tag{OA.35}$$

where dF is the natural probability measure. It then follows that

$$dF^{\mathcal{N}}(s) = \frac{p_1(s)}{p_0}dF^{\mathbb{Q}}(s), \tag{OA.36}$$

or equivalently $dF^{\mathbb{Q}}(s) = \frac{p_0}{p_1(s)}dF^{\mathcal{N}}(s)$. Thus

$$1 = \int dF^{\mathbb{Q}}(s) \tag{OA.37}$$

$$1 = \int \frac{p_0}{p_1(s)}\theta(s)dF(s) \tag{OA.38}$$

$$\frac{1}{p_0} = \int \frac{1}{p_1(s)}\theta(s)dF(s) \tag{OA.39}$$

$$\frac{1}{p_0} = \int \frac{1}{p_1(s)}dF^{\mathcal{N}}(s). \tag{OA.40}$$

Analogously we obtain the post-announcement and counterfactual policy prices:

$$\frac{1}{p_0'} = \int \frac{1}{p_1(s)(1+g(s))}dF^{\mathcal{N}}(s), \tag{OA.41}$$

$$\frac{1}{\tilde{p}_0^{\mathcal{N}}} = \int \frac{1}{p_1(s)(1+\tilde{g}(s))}dF^{\mathcal{N}}(s). \tag{OA.42}$$

We can then simply calculate the contribution of left tail support:

$$1 - \frac{\tilde{p}_0^{\mathcal{N}} - p_0}{p_0' - p_0}. \tag{OA.43}$$

As before, the numerator is the price change without abnormal price support at the tail. So when we subtract this from 1, we obtain the component of the price movement that is exclusively due to abnormal support at the tail of the distribution. Further note that p_0' and p_0 are date-0 observables so they are identical under both measures.

G.4 The Size of the Investment Grade Corporate Bond Market

In our calculations, we use $7 trillion as the size of the investment graded corporate bond market. The $7 trillion number for the supply of investment-grade bonds is supported by multiple sources. First, SIFMA reports total US corporate bonds outstanding as of 2019 at $8.8 trillion (https://www.sifma.org/resources/research/us-corporate-bonds-statistics/). However, this number includes both investment-grade and high-yield bonds but does not separate between the two. In order to split between the two, we follow one of two approaches, both of which yield about the same number of a 6-to-1 ratio of investment-grade to high-yield. These numbers are also reported in O'Hara and Zhou (2021), who study the effects of the COVID-19 movements in corporate bond markets. They state: "The US corporate bond market totals almost \$8.8 trillion, with investment-grade bonds approximately six times larger than high-yield bonds." Before discussing this split between investment-grade and high-yield, note that this ratio implies the total size of investment-grade as \$8.8 $\times$ 6/7= \$7.5 trillion.

To split between investment-grade and high-yield, we take the market capitalization of the Bloomberg US corporate bond indices for investment-grade and high-yield respectively (Bloomberg ticker "LUACTRUU INDEX"). These indices put the market capitalization of investment grade at \$6 trillion and high-yield at \$1.2 trillion as of February of 2020. One might want to use the Bloomberg US investment-grade market capitalization directly, which would put the number at \$6 trillion, but this is likely slightly low given not all bonds are included in the index.

These numbers roughly agree with numbers in the Z1 release that are also on FRED, though they do not map exactly to the variables provided there. The closest series to an overall corporate bond supply is CBLBSNNCB (https://fred.stlouisfed.org/series/CBLBSNNCB), "Non-financial Corporate Business; Corporate Bonds; Liability, Level." For 2020 Q1, it reports \$6.014 trillion which is close to the number we provide. This series has two key differences with our goal: it omits financials, which understates the total a bit, but includes both investment-grade and high-yield, which overstates it. The series ASCFBL (https://fred.stlouisfed.org/series/ASCFBL), which is "All Sectors; Corporate and Foreign Bonds; Liability, Level" reports \$13.339 trillion for the same time period. However, it is conceptually further away from what we are trying to measure for a few different reason. One is foreign bonds, the other is "all sectors" (see the Flow of Funds table at https://fred.stlouisfed.org/release/tables?rid=52&eid=809162&od=2020-01-01#). Foreign bonds are playing a big role here, as the "Rest of the World" category accounts for \$3.008 trillion. Our understanding is that these are foreign bonds held by US investors, which we do not include in the total supply numbers. Removing foreign bonds would take the total down to about \$10 trillion. Then, we have to split off high yield which, using the ratio from earlier, would put us around \$8.5 trillion for investment-grade which is within the ballpark of our reported number. This number may still be slightly high because it includes other types of securities such as convertible bonds.

Bretscher et al. (2022) take yet another approach and argue for a market size of \$6.5 trillion in 2019: "To construct the U.S. corporate bond universe, we follow an approach similar to Asquith et al. (2013) and identify corporate bonds in FISD that are denominated in U.S. dollars, are issued by firms domiciled in the U.S., and are publicly traded. We exclude convertible bonds and bonds that had no outstanding amount in a given quarter. This definition of the U.S. corporate bond universe, which we refer to as the publicly traded bond universe, yields a total outstanding of 6.5 trillion U.S. dollars in 2019 (by par value)."

Thus, the range of numbers for the total investment-grade corporate bond supply around early 2020 is about \$6.5-8.5 trillion across a variety of sources, and our preferred estimate is within this range at \$7 trillion. Note that implied purchases would not change dramatically using these other figures. Moving through this range of estimates leads to a factor between 0.93 and 1.2 in terms of implied purchases.

G.5 Longer Event Window

We demonstrate 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 less tight identification since a longer period means that other shocks could be affecting markets. Figure OA.5 shows that the results are similar if we expand the event window to three days. While magnitudes are slightly larger compared to the results in the one-day window, the asymmetric effect remains pronounced. Our main analysis attributes 50% of the announcement effect to additional promises in the left tail. Using a three-day window this number falls to about 30% because the right tail remains elevated. However, the dollar value from additional left tail promises increases from about $250 billion using the one-day window to about $300 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 that the choice of event-window length does not have a substantial effect on these results.

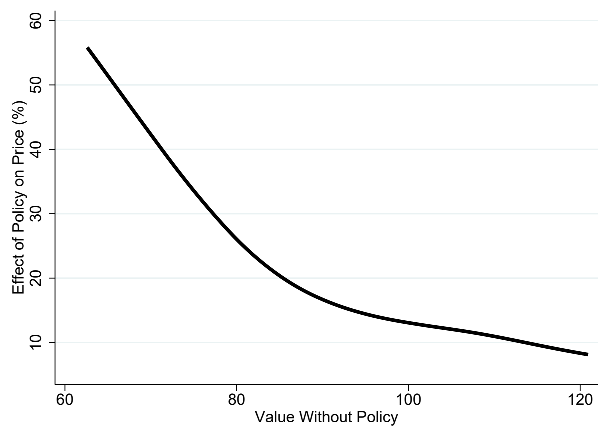

Figure OA.5: Price Support Function for the March 23 Announcement: 3-Day Window.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

G.6 Comparison to High Yield Bonds

Figure OA.6 contrasts the effects on investment grade bonds with those for high yield, using options on the largest and most liquid high yield bond ETF (HYG). Noticing the vastly different y-axes, the overall returns for high yield are basically flat. This result suggests that the announcement did not 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 against the possibility that changes in the pricing kernel are driving our results as a lowering of the price of credit risk should show up disproportionately in high-yield bonds, which is not the case.

Figure OA.7 plots the price support from the April 9, 2020 announcement that expanded the facilities to include high-yield bonds. We follow the same methodology applied to options on

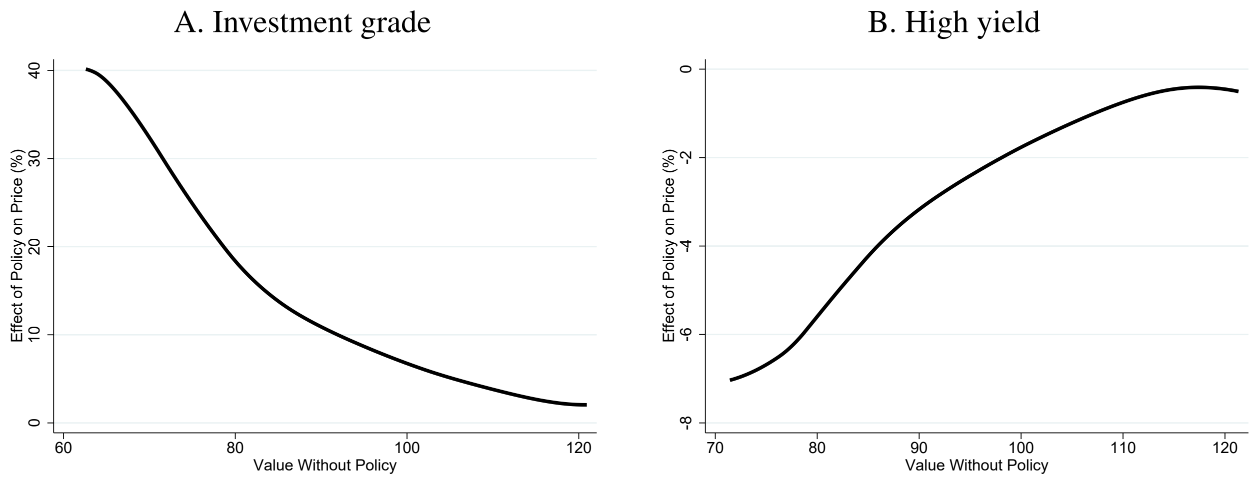

Figure OA.6: Comparison of Imvestment Grade and High Yield for March 23.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

the high-yield ETF (HYG). Not surprisingly, for this announcement, the asymmetric pattern is concentrated in high-yield bonds.

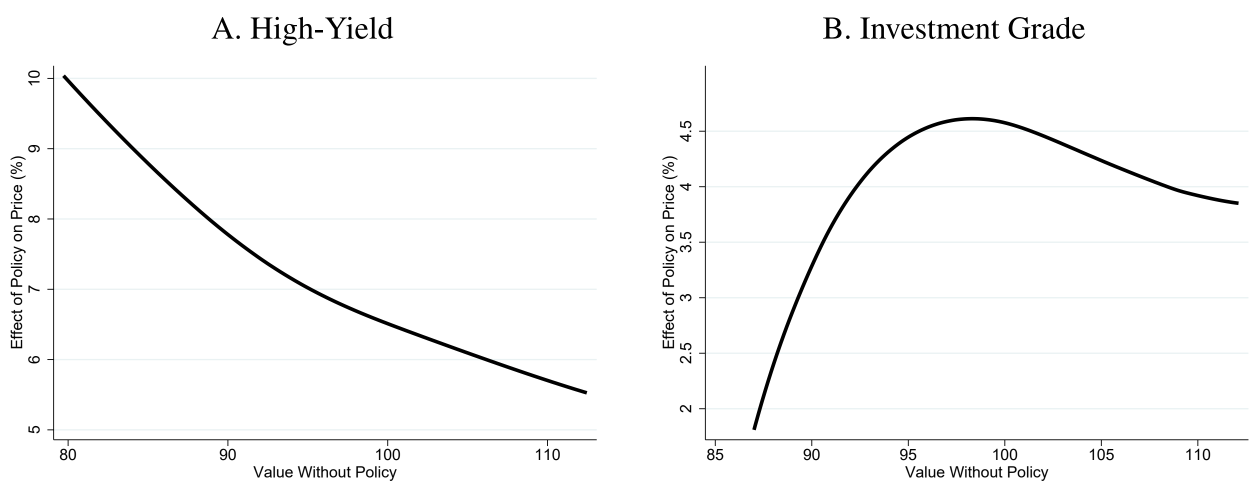

Figure OA.7: High-Yield Announcement on April 9, 2020: Price Support Functions.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

G.7 Liquidity

It is well known that liquidity in option markets—and derivatives markets more broadly—is heavily skewed. Trading volumes in the instruments we consider are not extremely high. For example, the options on the investment grade corporate bond fund we investigate in section 2.3 have an overall trading volume in 2020 (in terms of contract notionals) around \$100 billion. Bid-ask spreads are about 3% but grow considerably for low strike prices, reaching values as high as 30%. Thus, it is natural to investigate the robustness of the patterns we document to liquidity costs. In Section 2.3 we show that the negatively slopped support function we recover is very unlikely to have happened by chance. It therefore cannot be driven the overall level of liquidity in this market, since if this pattern was liquidity-driven we should expect it to show up recurrently in the data. Still, one could be concerned that liquidity disappeared exactly around the announcement since those were unprecedented times. To evaluate this possibility we replicate our recovery procedure but using bid and ask quotes. Such quotes are firm offers to buy and sell and therefore less subject to liquidity concerns. In Figure OA.8 below we report the recovered price support function with all four possible combinations. Of particular interest is the line that depicts the price support function implied by the bid-ask pair since it reflects the prices at which investors could have bought (in small quantities) options before the announcement and sold after the announcement. Thus, the implied price support function corresponds to actual returns for an investor even if we account for the illiquidity implied by a wide bid-ask spread. While accounting for these price differences has a visible impact on the price support function, the broad magnitude and shape are mostly unaltered.

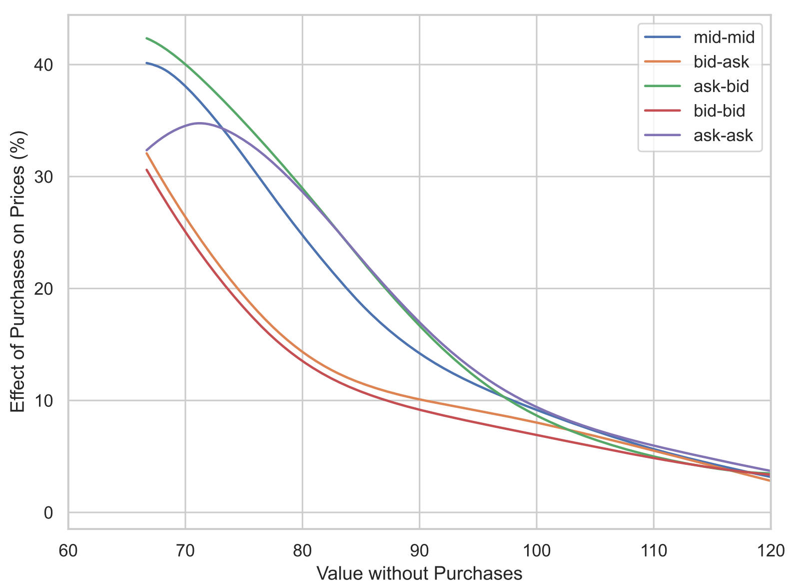

Figure OA.8: Price Support Function Using Bid and Ask Quotes.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement. Here we construct the price support function using both the bid (quote at which investors were willing to buy the options), the ask (quote at which investors were willing to sell), and the mid (the mid point between these quotes which we use in our baseline analysis)

H. Confidence Intervals

H.1 Bootstrap procedure

To produce the confidence intervals shown in Figure 4, we estimate a price support function for any pair of consecutive trading days, like we do for the policy announcement. That is, we compute the risk-neutral CDF on each of the trading days, and find the price support function connecting them. We then assess at each point on the x-axis of this price support function if the value on the y-axis for the policy date is outside of the typical range of values for other days (i.e. a 95% confidence interval). Our baseline follows the standard treatment of heteroskedasticity when volatility is measured: rescale the data to make it comparable across dates. Our procedure is:

1. Start with daily data from January 2010 to February 2020 (just before the COVID-19 shock). This sample contains roughly 2500 trading days.

2. Construct the price support function for each sequential pair of dates (t and $t+1$) in the sample. The price support function is expressed in return space, that is, g maps a realized return from date t to the option maturity, which is the strike price, ($x_t = \frac{strikeprice_t}{forwardprice_t} - 1$) into a return of the asset in that state from date t to date $t+1$ ($g_t(x) = \frac{F_{t+1}^{-1}(F_t(strikeprice_t))}{strikeprice_t} - 1$).

3. Adjust for heteroskedasticity. Because the range of both of these axis naturally stretches out with volatility, we normalize both by the at-the-money implied volatility in date t. Specifically we compute $x_{z,t} = \frac{x_t}{\sigma_t^{atm}}$ and $g_{z,t}(x_{z,t}) = \frac{g_t(x_{z,t})}{\sigma_t^{atm}}$. This puts all days in the same standard-deviation scale. A nice finance intuition to think about this procedure is that we are leveraging up each observation to have the same implied volatility as our event date. Thus it creates a counterfactual sample where the options and the underlying have similar risk characteristics as during our event.

4. Sample with replacement out of these days to construct a sample of 100k days. We could alternatively just use the sample; because we are using ten years of data, results are virtually indistinguishable in this case, but bootstrapping producer smoother confidence intervals when taking a more non-parametric approach. We call this new sample the bootstrapped sample.

5. Discretize the x-axis of the price support function. We want to compute standard errors at various points of the function, so we have to discretize the x-axis. We create 100 bins of the variable x_z where the bins are chosen to be consistent with the unconditional distribution of x_z in the bootstrapped sample. We then take the average of g_z within each date-bin pair.

6. Compute the range from the 2.5th percentile to the 97.5th percentile of the value of these price support functions for each bin. This is the confidence interval for the normalized price support function.

7. Compare the normalized price support function for the policy announcement to these normalized confidence intervals. To provide a more intuitive representation, we rescale the axis

of this comparison by the at-the-money volatility on the day of the event — this way, the point estimate of the price support function is the unnormalized one. We also follow the usual tradition of plotting the confidence interval around the estimates instead of around 0. The strict interpretation is that if this interval does not contain 0, our estimate is statistically different from 0.

H.2 A different statistical hypothesis: Is the asymmetric shape of g unusual conditional on a large realized return?

Our baseline statistical test focuses on measuring "how unusual the change in the price support function is", without the conditioning aspect. This test helps evaluating the causal statement: can we attribute the changes we see in the CDFs to the announcement of the policy? Indeed, we select the date we study based on the occurrence of the policy announcement, not the large change in bond price. The observation that the change in bond price is unusually large suggests that the announcement caused the return. The observation that the price support functino is unusually large suggests also suggest that the announcement caused this change.

In this section, we add the clause "conditional on the large change in the bond price" to the test. What is the goal of this approach? In a standard event study, once the econometrician has overcome the burden of identification that the shock actually caused the market response, then it becomes irrelevant whether the pattern of responses across assets or across the different strikes is the same as usual or not. That said, a plausible concern is whether there is a deep mechanical or structural force other than the policy that ties down various parts of the distribution. So it is independently useful to test whether large shocks are always accompanied by as much action in the left tail as we found. As we explain below, we find the answer to be no, which comforts us in our interpretation.

We ask whether when the return on a given day is an outlier, an asymmetric g obtains systematically. We repeat our baseline procedure, but look at outcomes only in the sample of return outliers. Specifically, starting from the whole sample, we only keep dates for which the normalized return ($r_{z,t} = \frac{1}{\sigma_t^{atm}} \left(\frac{forwardprice_{t+1}}{forwardprice_t} - 1 \right)$) is either below the 1st or above the 99th percentile. To make the negative returns comparable to our sample, we flip the sign of g for these, that is we replace the observations by $\tilde{g}_{z,t} = sign(r_{z,t}) \times g_{z,t}$. Then, in the same way as our baseline, we compare the range of outcomes for the normalized price support in the bootstrapped sample to the normalized price support on the event day.

Figure OA.9 reports the results. We now plot the 95% confidence bands in their initial position. These bands are not centered around 0: by construction we select large value of returns, so the average value of g is also large. The average level of g obtained by selecting these outliers (the dotted red line) is very close to the realized return for our event (the horizontal black line). This indicates that we are looking at broadly the same intensity of outliers. The blue band reveals that a strongly negative slope of the price support function is not a feature of the simulated data. Said otherwise, we reject the hypothesis that the negative slope of g we see for the announcement is explained by the fact that there was a large return.

One limitation of this approach is that it focuses on outliers relative to implied volatility. It could be that large returns irrespective of the volatility on that day trigger a different behavior of the price support function. Assessing such a behavior is challenging in the daily data because the

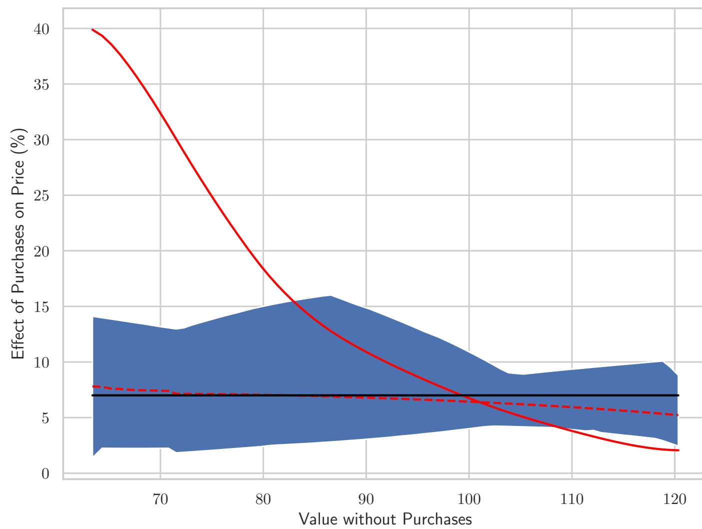

Figure OA.9: Standard Errors: Only Sampling Large Returns. The figure shows out baseline result with 95% confidence bands that now adjusts the baseline procedure to sample only from extreme (both positive and negative) 1% (implied-volatility) normalized return days. See text above for detailed description.

magnitude of the event day return is so large. To get at larger magnitudes of returns, we extend the time window to construct the "usual" behavior.[39]

Looking at longer windows allows us to get closer to matching the returns of our event without relying on leverage. However, it does not give us a way to directly construct a counterfactual for our price support function as we did in the analysis before. The basic issue is that the support of g for each day is (partially) determined by the range of traded strikes that day, which is influenced by the amount of expected volatility until the option maturity. So while we see more volatility in realized returns at longer horizon, the range of g on the x-axis does not increase.

One way around this challenge is to look at the relationship between the slope of the price support function and the realized returns on the underlying. In this way we can see whether the asymmetry we detect is unusual given the high absolute return. Our procedure is as follows: we expand the window up to 60 days. We then compare the risk neutral density in date t with the risk-neutral density in date $t+h$ where $h \in 1, 5, 10, 21, 30, 45, 60$ to recover the function $g_{t,h}$. We then compute for each date-horizon pair the slope of g. Specifically we calculate

$$Slope_{t,h} = \frac{g_{t,h}(x_t)|_{F(x_t)=\underline{F}} - g_{t,h}(x_t)|_{F(x_t)=0.5}}{x_t|_{F(x_t)=\underline{F}} - x_t|_{F(x_t)=0.5}}, \tag{OA.44}$$

where $\underline{F}$ is the percentile of the most extreme point in our event day support function ($\underline{F}$ =11%). So, we are computing the slope of g using two specific percentiles (11% and 50%) across all dates.

[39] We thank one of the anonymous referees for suggesting this approach.

We choose the left tail on our event date and the median to compute this slope, but other percentiles in the left tail or the right tail produce similar results.

Figure OA.10 reports a scatter plot of all slope and realized return pairs, i.e. all date-horizon observations. Our event is the blue star in the bottom right corner with a slope close to -1 associated with a return of about 7%. The different colors show the different horizons. The range of return realizations grows as the horizon increases, getting progressively closer to the 7% return.

Two key observations come out of this picture. First, while there is on average a negative relation between slope of g and realized return, the value of the slope for the policy event is not explained by that relation. The colored lines estimate this relation for each horizon.A linear regression predicts an observation with a 7% return has a slope of g around -0.2 associated with it, with this prediction fairly stable across horizons h. This value is far from the value of -1 for the event. Observing the realized values of the plot this difference is clearly statistically significant. Going back in terms of g, the usual relation would predict a wedge of $-0.2 \times (x_t|_{F(x_t)=\underline{F}} - x_t|_{F(x_t)=0.5}) \approx 8\%$ between values of g the median and the tail of the return distribution. The estimated g for the event instead has a wedge of about 40%, almost 5 times higher.

The second key observation is that there is quite a wide range of variation in the relation between the slope of g and realized returns. This feature strongly alleviates the concern that there could be a "mechanical" relationship that strongly ties these two objects overall.

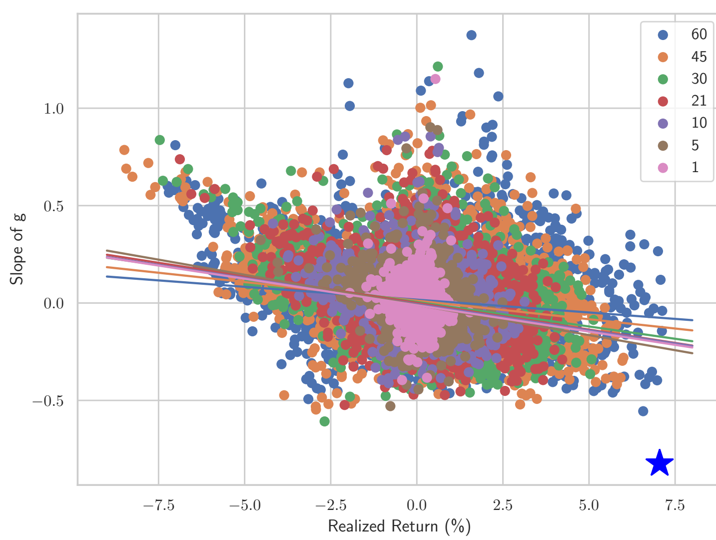

Figure OA.10: Relation Between Realized Return and Slope of the Price Support Function. Each dot represents a date-horizon pair. Different colors are different horizons ranging from 1 up to 60 trading days. The different lines are ordinary least-square predictions for each horizon. The March 23 event is shown with a star marker in the bottom right of the plot. See text for details.

In short, all of our estimates conditioning on large realized returns point to the same conclusion. While large returns are somewhat associated with asymmetry in the price support function g, this asymmetry is small when compared to what we measure for the policy announcement. The amount of asymmetry for the event day is high even when compared to the most extreme days in the sample.

I. State-dependent Multiplier Calculations

We discuss additional results on multipliers that are state-dependent. Greenwood and Vayanos (2014) interact Treasury supply with a measure of arbitrageur wealth (the amount of arbitrage capital in their model determines the price impact of changes in supply, since excess supply must be absorbed by these agents). They find that a 1 standard deviation change in arbitrageur wealth leads to a 25% increase in their coefficient of bond supply on bond yields and returns. This means a more extreme 2 standard deviation shock to arbitrageur wealth would only change the multipliers by 50%. They proxy for arbitrageur wealth using realized bond returns, so we can easily map their results to our setting. Specifically, they state: "According to Hypothesis 4, the interaction terms should have a negative coefficient: supply and slope predict returns positively, and more so when arbitrageur wealth decreases. The results confirm this prediction in the case of one-year returns. The coefficients of the interaction terms are economically significant. Consider, for example, the interaction term between supply and our first measure of arbitrageur wealth, (22). This term has a coefficient of –4.436 for one-year returns. From Table 1, the standard deviation of the wealth measure is 0.0015. Therefore, a one-standard-deviation movement in the wealth measure changes the coefficient of maturity-weighted debt to GDP by 4.436 × 0.0015 = 0.0067. This is approximately a one-quarter percentage change, since the coefficient is 0.026 (Tables 2 and 7)." These results suggest that the multiplier does change across states, but again this variation is too small to explain our findings.

To address this directly, we calibrate the Greenwood and Vayanos (2014) results to our setting based on the standard deviation of returns on investment-grade bonds implied by option prices. Specifically, we define $\mathcal{M}(p) = \mathcal{M} \times \left(1 - \frac{1}{4}R(p)/\sigma\right)$. That is, we take our baseline constant value of $\mathcal{M}$ and multiply by $\left(1 - \frac{1}{4}R(p)/\sigma\right)$ where $R(p)$ is the (net) return on corporate bonds in a given state p and σ is the standard deviation of returns using option implied volatility. This gives us quantitative variation in multipliers consistent with what Greenwood and Vayanos (2014) find empirically.[40] We then plot the associated quantities $Q(p)$ that we back out from using $g(p) = \mathcal{M}(p) \times Q(p)$. Figure OA.11 below plots the quantities when we use a multiplier that depends on the realized state.

As shown in the figure, this adjustment still produces a highly asymmetric price support function. One could also ask the question in reverse: what value of the multiplier would we need in each state to make the price support function flat? Since our price support is 20 times higher in the worst states relative to the best states, we would need a multiplier that is 20 times larger in these extreme bad states to match the asymmetry in the data. The evidence above suggests much milder variation, that the multiplier in bad states would be up to 2 times larger which is an order of magnitude from the variation needed to produce a flat price support function.

The QE literature comes to a relatively similar conclusion on this issue. For example, Bernanke (2020) states: "The evidence described so far suggests that, once we control for the fact that mar-

[40] E.g., a zero return, or no change in prices, gives the baseline multiplier, while a return of $-\sigma$ gives a multiplier that is 25% larger than the baseline. Positive returns result in a lower multiplier relative to the baseline.

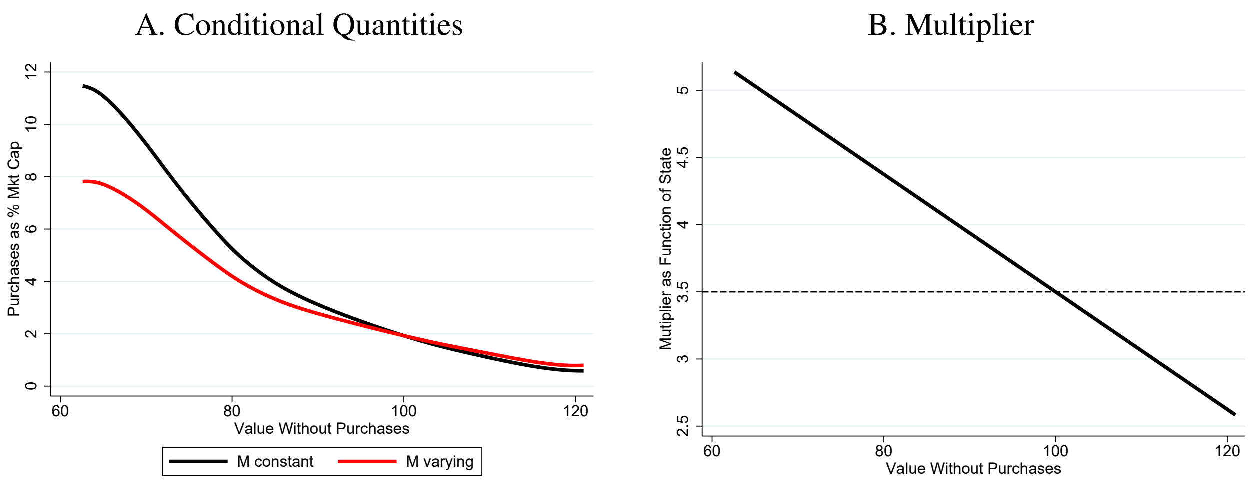

Figure OA.11: Price Support Function with a State-Dependent Multiplier

We model a state-dependent multiplier $\mathcal{M}(p) = 3.5(1 - 0.25 \times R(p)/\sigma)$. The right panel plots $\mathcal{M}(p)$. Panel A shows conditional quantities $q(p) = g(p)/\mathcal{M}(p)$ compared to our baseline of a constant multiplier. Quantities are expressed in percentage of the total supply of investment grade corporate bonds. Panel B reports the multiplier.

ket participants substantially anticipated later rounds of QE, the impact of asset purchases did not significantly diminish over time, as financial conditions calmed, or as the stock of assets held by the central bank grew." With regard to the ECB announced purchases in 2015, Bernanke (2020) states: "This reduction is economically significant and, when adjusted for the size of the program, comparable to estimates from event studies of early QE programs in the United States and the United Kingdom, even though in early 2015 European financial markets were functioning normally." Thus, the available evidence suggest that the impact of purchases per dollar are reasonably similar in calm or stressed financial market conditions.

The evidence from Busetto et al. (2022) that we discuss in Section 4 complements Bernanke's decision. Figure 9 reports the relation between surprise quantity purchased and the response of 10-year gilts across different QE announcements. Of course, the point of our paper is that these announcements might contain messaging about state-contingent intervention rather than a single number. But still, while the relation between price and quantity does not appear perfectly linear when excluding the QE1 announcement, it is not too far. The R-squared from the fitted linear relationship is high at about 0.8. Clearly, whatever nonlinearity is present here is nowhere near enough to explain the state-contingent variation we find using options.

J. In Which States was the Fed Expected to Buy? Details on the Copula method

We discuss in more detail how we pinpoint what exactly drove the price response. We start by using options on the ten year Treasury futures (formally, the security name is Ultra 10-year T-Note futures) and options on the CDX North America investment grade index. This CDX index tracks the CDS spreads of 125 of the most liquid investment grade corporates. These options allow to recover the risk-neutral density for the distributions of credit spreads and interest rates. The options on the corporate bond ETF give the risk-neutral density on the cash instrument. To separately recover the movements in the distribution of the synthetic component of moments in the price of corporate bonds (interest rate plus credit spreads) from the movements in the distribution of the basis we need to recover how the joint distribution of these three prices (and how this distribution changed). The options only give us the marginal distribution of each one. To recover the joint distribution we rely on a copula model.

It works as follows. We have three variables x, y, z each with known marginal distribution F_x, F_y, F_z. We also know the correlations between these variables which are assumed to be constant. Let this correlation matrix be given by C. We then sample $\tilde{x}_i, \tilde{y}_i, \tilde{z}_i \sim \mathcal{N}(0, C)$, the multivariate normal distribution with zero mean and covariance matrix C — and therefore unitary individual standard deviations. We label each draw using the index i. We then compute, for each realization, the corresponding quantile in the standard normal distribution. For example $fx_i = F(\tilde{x}_i)$ where F is the cdf of a standard normal distribution. We then have $\{fx_i, fy_i, fz_i\}_{i=1}^{N}$ where N is the number of draws. Finally, we use the original marginal densities to invert back the realization, i.e., $\{x_i, y_i, z_i\} = \{F_x^{-1}(fx_i), F_z^{-1}(fz_i), F_z^{-1}(fz_i)\}$. This procedure lets us to simulate from the joint distribution in a way that is consistent with the marginal distributions recovered form options prices. Therefore, it allows us to also recover the distribution of any function of these variables.

More specifically, we apply the method in two steps because it is more intuitive to think about the correlation between the synthetic and the cash instrument then to think about the correlation of the cash instrument and the different pieces of the synthetic. We first apply the copula method to the CDX and treasury options. We use a correlation of -0.25 which is consistent with the historical data. The results are quantitatively similar if we change the correlations to -0.75 or 0.25 (see Figure OA.12 Panel A). We set the correlation between the synthetic and the cash instrument to 0.8 which is consistent with the historical average. This correlation tends to go down during crisis (for example a 30-day moving average estimator has the lowest realization of 0.25 in our 9-year sample) so we show results with 0.4 and 0 as well (see see Figure OA.12 Panel B). The key result is unchanged.

Why is the result so robust to correlation assumptions? The marginal CDFs are informative about the range of each variable and only for very extreme correlations would the range of the synthetic distribution be sufficiently large to be able to account for the wide range of the cash instrument distribution.

K. Implications of Promises for Market Dynamics

K.1 Are the Effects of Asset Purchases Getting Weaker: Empirical Details

This subsection provides more detail on the data sources used in Section 4.

K.1.1 United Kingdom

Table OA.2 contains the response of 10 year Gilt yields to six announcements of purchases between February 2009 and February 2010. The Bank of England is unique in that most of these announcements contain a fairly narrow and specific quantity range. First, note that only the first

A. Varying the correlation between interest rates and credit spreads

B. Varying the correlation between synthetic and cash bond

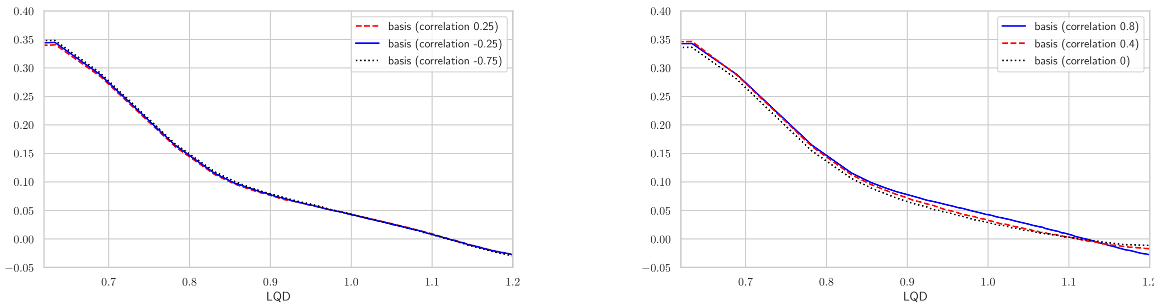

Figure OA.12: Decomposition of Announcement Effects: Robustness

In this figure we look at how the decomposition in Figure 6 depend on the correlation between interest rate, credit risk and financial dislocations. The Figures shows in the x-axis the value of the asset in different states of the world absent policy. The y-axis shows the effect due to movements in the basis. The different lines show different correlations.

Table OA.2: UK Announcement Effects

This table shows data for UK. The yield numbers are for the 10 year Gilt. Sources: Joyce and Tong (2012), Meaning and Zhu (2011), and author's calculations. Quantities are given in billions (£). The column "multiplier" indicates the percentage change in market value of the securities divided by the percentage of market capitalization purchased.

Date	Gilt Yield	Announcement	Quantity Low	Quantity High	Multiplier Range
2/10/09	-34	QE "likely"			
3/4/09	-68	75 billion	75	75	0.60
5/6/09	10	50-125 billion	50	125	[-0.13, -0.05]
8/5/09	-3	50-125 billion	50	125	[0.02, 0.04]
11/4/09	10	25 billion	25	25	-0.27
2/3/10	-2	Maintain 200	0	0	
Total	-87		200	350	[0.17, 0.29]

two announcements had any effect at all on yields, together resulting in about a 100bps decline in the 10 year Gilt yield. The first announcement, on February 10, 2009, did not contain concrete information but suggested that purchases were likely. On March 4, purchases of £75 billion lead to a decline of 70bps in the 10 year Gilt yield. In contrast, the next three announcements featured no changes in yields at all despite similar magnitudes of purchases. We can convert the yield changes and quantities into a price elasticity that gives the price impact, which we provide in the column "multiplier." The multiplier for the first announcement of 0.6 says that by purchasing 10% of the supply of Gilts the price of Gilts would fall by 6% (for a security with a duration of 10, this means a decline in yields of 60bps). When the announcement comes with a quantity range, we provide the range for the multiplier as well. The main finding is that the multiplier is much higher in the early announcement and is then quickly goes to zero.

These patterns fit well with the promises view of state-contingent policy. A natural interpretation is that upon hearing the early announcements, investors form expectations that the Bank of England would buy more Gilts if the economy remained weak. Thus the "promises" view explains both the high initial multipliers and the zero in the follow on interventions. The Bank of England implemented a second period of purchase announcements in October 2011, but Meaning and Zhu (2011) find these to have a negligible effect on yields. Unlike for other countries, we don't have reliable option data for Gilts over this period to test whether the early announcements effects were driven by the promises component.

An alternative explanation for the declining multiplier effect above is that the multiplier depends on economic conditions, and the initial announcements occurred in periods when the economy was in worse shape (after all that is when they decided to pursue this policy for the first time). We will return to this argument in each of the subsections. For the UK data, we note that the multiplier goes from 0.6 in March, 2009 to -0.1 in May, 2009. Thus economic conditions would have to change quite rapidly for the multiplier to go from high and positive to zero in only two months.

K.1.2 United States

Table OA.3 provides announcement effects for Quantitative Easing (QE) in the US, specifically QE1 which was implemented November, 2008 to November, 2009. Announcements of later QE programs, QE2 and QE3, have been shown to have had essentially no effect on yields. The Fed purchased Treasuries, Agency debt, and Mortgage-backed-securities, and we use numbers from Gagnon et al. (2018) on the yield responses. The last column "multiplier" converts average yield changes to price movements and then divides by the total amount of assets purchased as a fraction of the supply of these securities outstanding.[41] The initial announcement, which stated the Fed would purchase "up to" $600 billion across these categories, led to an average decline in yields of about 40 bps. This equates to a multiplier of about 0.8 (e.g., for a purchase sized at 1% of market cap, prices would increase by 0.8%). The next significant announcement in QE1 came in March 2009, where the Fed expanded quantities. While yields moved by about the same amount, the quantities were larger. This leads to a lower multiplier. Later announcements, for example dropping the "up to" language and effectively confirming the Fed would purchase the maximum stated amount, had no effect on yields. These patterns also fit the support function results we present in Section 3 for the initial announcements, which indicate the presence of implied policy puts in these initial announcements.

These results contrast to QE2 and QE3, where no announcement effects are found (see Meaning and Zhu (2011)). A potential concern with comparing these impacts across time periods is that perhaps the multiplier is much higher in times of more severe economic stress and economic uncertainty such as the period where QE1 was unleashed. The price response to the interventions in the treasury market at the outset of the covid shock are particularly informative to distinguish the "promises" view from the economic uncertainty view as explanations for the time-variation in multiplier.

These announcements are studied extensively in Vissing-Jorgensen (2021), who find that the announcements had no effect on Treasury yields using high frequency data from Treasury futures markets. The first announcement on March 15 stated purchases of "at least" $500 billion of Treasuries and $700 billion in total long duration assets. This is sizable not only on its own but also because the "at least" language indicated potentially much larger purchases. This was confirmed on March 23 when the purchase amounts shifted to "unlimited" and the Fed continued to purchase

[41] We find similar results using a weighted average of the yield responses where weights are given by the relative supply of each.

Table OA.3: US Announcement Effects

This table shows data for US. Sources: Gagnon et al. (2018), Vissing-Jorgensen (2021), and author's calculations. Quantities are given in billions (USD). The column "multiplier" indicates the percentage change in market value of the securities divided by the percentage of market capitalization purchased.

	Yield Responses					**Quantities**				
Date	Treas	Agy	MBS	Avg		Treas	Agy	MBS	Total	**Multiplier**
11/25/08	-22	-58	-44	-41.33	Up to	0	100	500	600	0.80
12/1/08	-19	-39	-15	-24.33	May expand					
12/16/08	-26	-29	-37	-30.67	Expanding					
1/28/09	14	14	11	13.00	Expanding					
3/18/09	-47	-52	-31	-43.33	Up to	300	200	1250	1750	0.29
4/29/09	10	-1	6	5.00						
6/24/09	6	3	2	3.67						
8/12/09	5	4	2	3.67	Drop "up to"					
9/23/09	-3	-3	-1	-2.33						
11/4/09	6	8	1	5.00			175		175	-0.33
Total QE1	-76	-153	-106	-111.67		300	475	1750	2525	0.51
3/15/20	-17			-17.00	At *least*	500		200	700	0.07
3/23/20	0			0.00	Unlimited	500			500	0.00

large quantities. These announcements quickly translated into actual purchases – within three weeks of the initial March 15 announcement the Fed had purchased over $1 trillion in Treasuries. Still, the announcements had no effect on yields as shown in Vissing-Jorgensen (2021).

Vissing-Jorgensen (2021) argues that the purchases *themselves*, rather than the announcements, had an impact in March 2020, possibly because of large frictions and selling pressure in Treasury markets at the time. However, even this effect is modest. Vissing-Jorgensen (2021) states "that an increase of 0.1 (buying 10% of supply) leads to a 5.35 bps larger decline in yields." Using a duration of ten years would then imply a 50 bps price increase, or a multiplier of about 0.05. Thus, regardless of whether one uses announcements or actual purchases, the COVID-19 period features a very low multiplier relative to QE1. The natural interpretation is that the bond market expected large purchases of Treasuries given the prior experience of QE. Under this view, it is not that purchases were not effective, just that the market already expected them to occur so the announcement is not informative about effectiveness.

The results for Treasuries during COVID-19 also contrast sharply with what we document for corporate bonds. The key difference is that the Fed had never before purchased corporate bonds and thus the announcement was a surprise. Further, once the corporate bond announcement was made, the market understood the implications for future state-contingent purchases more immediately compared to quantitative easing in 2008 where learning appeared to occur over a few announcements.

This experience also contrasts with the Bank of Canada (Arora et al., 2021) during the same time period. The Bank of Canada announced purchases of government bonds on March 27, 2020. Government bond yields declined immediately on the announcement as shown in Arora et al. (2021). Importantly, this was the first time the Bank of Canada implemented a large-scale asset purchase program involving government securities, contrasting with the US experience where such purchases were made in the global financial crisis.

Table OA.4: ECB Announcement Effects

This table shows data for ECB. Sources: Krishnamurthy et al. (2018) and author's calculations. Quantities are given in billions (Euros). We use average yield responses across maturities for each sovereign in Krishnamurthy et al. (2018) and the 10 year yield if the average is not available. The column "multiplier" indicates the percentage change in market value of the securities divided by the percentage of market capitalization purchased.

Type	Date	Italy	Spain	Portugal	Ireland	Greece	Avg	Quantity	Multiplier
SMP1	5/10/10	-47	-62	-219	-127	-500	-191	75	3.49
SMP2	8/7/11	-84	-92	-120	-49	-3	-69.6	145	0.99
OMT1	7/26/12	-72	-89	-12		-78	-62.75	unspecified	
OMT2	8/2/12	-23	-41	-8		-67	-34.75	unspecified	
OMT3	9/6/12	-31	-54	-98		-36	-54.75	unspecified	
LTRO	12/1/11	-46	-61	-27		-147	-70.25	lend to banks	
LTRO	12/8/11	35	30	9		90	41	lend to banks	

In summary, the evidence from asset purchases in the United States is quite clear: earlier announcements of a particular policy appear to have the largest impact on prices. This is apparent even in the early stages of quantitative easing ("QE1"). Beyond QE1, announcement effects have effectively disappeared for Treasuries, Agency debt, and MBS. This does not seem to be due to variation in the economic conditions around the announcements.

K.1.3 Eurozone

Table OA.4 gives results for the European Central Bank announcements in 2010-2011 during the European sovereign debt crisis. We use yield data from Krishnamurthy et al. (2018) (see their Table 3). It is difficult to immediately compare yield changes and tie them to quantities as specific quantities are only given for the first two announcements. The first announcement in May of 2010 had the largest effect on sovereign yields, with an average decline in yields of 190 bps. Given the quantity announced of €75 billion, this large decline in yields suggests a multiplier of around 3.5, where we construct this number using the total debt of the five countries considered and the average duration of the bonds purchased from Krishnamurthy et al. (2018). The next announcement in August saw a much smaller, though still substantial, decline in yields of about 70 bps. This translates to a significantly smaller multiplier.

Next, we note that there were three separate programs for the ECB sovereign crisis. The Securities Markets Programme (SMP), the Outright Monetary Transactions (OMT), and the Long-Term Refinancing Operations (LTROs). Each program was different. The SMP was the only one that involved direct purchases. As discussed, the first SMP announcement carried much larger effects than the second, consistent with investors forming expectations of future announcements from the initial announcement. The OMT featured conditional commitments to purchase government debt. Again, the strongest response comes from the initial OMT announcement consistent with the state-contingent view. No purchases were made during the OMT program. Finally, the LTRO extended loans to banks. The LTRO announcements feature the same declining pattern.

In sum, the ECB announcements that involved direct purchases of sovereign debt (SMP) feature declining multipliers. Other programs aimed at reducing sovereign yields had declining effectiveness after the initial announcement was made.

Overall, the promises view provides a consistent and simple way to interpret the variation in the announcement effects we observe. Initial announcements induce investors to form expectations of future and more aggressive interventions in adverse states, and as a result, are associated with large effects. Conversely, the often larger follow-on interventions tend to induce only a muted price response as they are already baked in.

K.2 Longer-Run Effects on Risk in Corporate Bond Markets

We investigate the long-run effects of the Fed's corporate bond intervention by looking at how the dynamics of corporate bond tail risk changes after the programs are implemented. While our main analysis documents how the Fed's introduction of the policy has the immediate effect of reducing this tail risk 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 price data. The challenge is to have a good benchmark for how tail risk would behave absent interventions. We look at a variety of approaches to deal with this challenge, though these results should be taken as suggestive and depend on how reasonable our benchmarks are.

Our first approach constructs a tail risk index for corporate bonds using the slope of the implied volatility curve. We take implied volatility for options with a delta of 90 and subtract the implied volatility for option with delta of 10. This difference is insensitive to parallel movements in the implied volatility curve and increases when the implied volatility of the left tail rises relative to the right tail. We then take the same tail risk measure using S&P500 index options as well as options on the investment grade CDX index.

Table OA.5 shows that tail risk sensitivity changed after the announcements: tail risk in corporate bond markets are usually positively related to tail risk in equity or CDS markets; after the interventions, this sensitivity disappears altogether. 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 9, 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.

The post interaction term is strongly negative and statistically significant, meaning corporate bond tail risk becomes much less sensitive to broader tail risk in the economy after the interventions. The sum of the two coefficients represents the total sensitivity in the post period and is, if anything, slightly negative. We find similar results using the CDS index in place of the stock market as a gauge of tail risk variation in corporate bonds. The CDS index is useful because it is a more direct measure of the cash-flow risk that corporate bonds are exposed to. Finally, the results in the pre-period are not driven by extreme behavior during the acute phase of COVID where all tail risk measures spike. To show this, we add a COVID dummy, equal to 1 for the period of February 1, 2020 to April 9, 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. Finally, consistent with the persistence of promises, we find no difference if we change the post period to June after the purchases occur.

This declining sensitivity is not only present in tail risk, but also in overall bond yields. In

Table OA.5: 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 9, 2020, the dummy "covid" equals 1 from February 1, 2020 to April 9, 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***	0.43***		
	(0.05)	(0.02)		
$Tail_t^{SP500} \times post$	-0.78***	-0.63***		
	(0.07)	(0.05)		
$Tail_t^{SP500} \times covid$		0.68***		
		(0.15)		
$Tail_t^{CDS}$			0.27***	0.14***
			(0.04)	(0.02)
$Tail_t^{CDS} \times post$			-0.37***	-0.24***
			(0.04)	(0.02)
$Tail_t^{CDS} \times covid$				0.90***
				(0.16)
post	0.16***	0.14***	-0.06***	-0.02*
	(0.01)	(0.01)	(0.01)	(0.01)
covid		-0.12***		0.35***
		(0.03)		(0.06)
Constant	-0.04***	-0.02***	0.11***	0.06***
	(0.01)	(0.00)	(0.01)	(0.01)
Observations	2,769	2,769	1,510	1,510
R-squared	0.25	0.29	0.26	0.44

Appendix Table OA.6, we find that corporate bond returns are usually negatively related to changes in the VIX. This sensitivity is divided in half in the post-intervention period.

Figure OA.13 instead uses monthly data on option-based pseudo credit spreads from Culp et al. (2018). These pseudo-spreads are constructed by equity options and Treasuries — not corporate bonds — and have been shown in earlier samples to exhibit remarkably similar properties than the actual spreads. We use the two-year maturity investment-grade pseudo credit spread from Culp et al. (2018).[42] We then compare this spread to the Bank of America investment grade option-adjusted credit spread index for maturities between one and three years taken from FRED. We plot both the actual credit spreads and pseudo spreads in Figure OA.13. From 2010 to 2020, the two spreads track each other quite well. In early 2020, when the COVID-19 crisis hits, actual spreads for investment-grade spike well beyond those implied by equity market options, consistent with investment grade bond prices becoming abnormally depressed in this episode. However, following the Fed's intervention, investment-grade spreads become quite low, and in fact reach their lowest point at any time over the 2010-2020 window. In contrast, equity markets still feature substantial volatility, implying higher than usual default risk on pseudo-bonds. This large gap is consistent with a market pricing of future interventions: a crash is still possible and priced in equity options, but the Fed would intervene in corporate bond markets and make it disappear.

Consistent with our evidence of abnormally low credit spreads after the intervention, Boyarchenko et al. (2020) and Becker and Benmelech (2021) find abnormally large issuance of

[42]We obtain the data from The Credit Risk Lab.

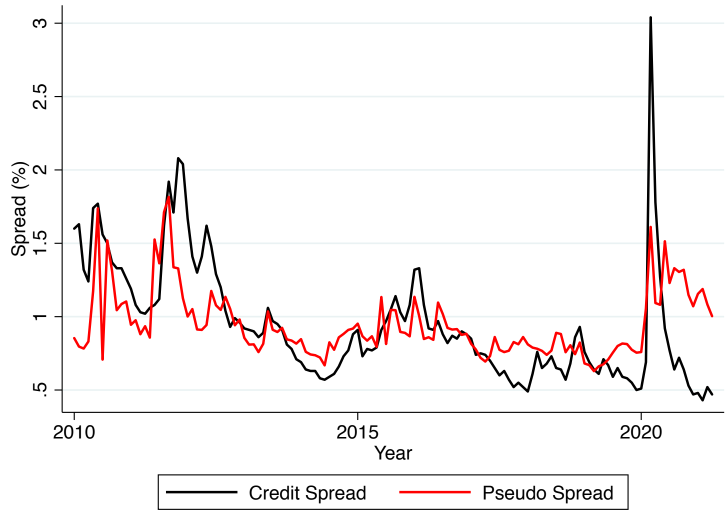

Figure OA.13: Spreads and Pseudo Spreads.
This figure plots actual credit spreads and pseudo spreads from Culp et al. (2018).

investment-grade bonds by firms after the interventions, and Balthrop and Bitting (2022) find this effect is persistent for firms eligible for Fed purchases under the original SMCCF facility.[43] Again, this increase in issuance fits with the narrative of an implicit subsidy of low spreads due to expectations of future Fed support.

[43] See also Acharya et al. (2022) who show empirical evidence that quantitative easing impacted firms bond issuance behavior

Table OA.6: 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 9, 2020, the dummy "covid" equals 1 from February 1, 2020 to April 9, 2020. Interaction effects capture whether this sensitivity is lower after Fed interventions. Robust standard errors given in parentheses.

Panel A: Corp Bond Returns

	(1)	(2)
	$r_t^{CorpBond,e}$	$r_t^{CorpBond,e}$
ΔVIX_t	-0.21***	-0.20***
	(0.02)	(0.02)
$\Delta VIX_t \times post$	0.10***	0.08***
	(0.03)	(0.03)
$\Delta VIX_t \times covid$		-0.05
		(0.04)
$post$	0.04	0.03
	(0.03)	(0.03)
$covid$		-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)	(2)
	$\Delta spread_t$	$\Delta spread_t$
$\Delta pseudo_t$	0.41**	0.16***
	(0.19)	(0.05)
$\Delta pseudo_t \times post$	-0.62***	-0.37***
	(0.23)	(0.13)
$\Delta pseudo_t \times covid$		1.74***
		(0.38)
$post$	-0.11**	-0.10**
	(0.05)	(0.05)
$covid$		0.08
		(0.23)
Constant	0.00	-0.01
	(0.02)	(0.01)
Observations	135	135
R-squared	0.22	0.69