

NBER WORKING PAPER SERIES

WHATEVER IT TAKES? THE IMPACT OF CONDITIONAL POLICY PROMISES

Valentin Haddad
Alan Moreira
Tyler Muir

Working Paper 31259
<http://www.nber.org/papers/w31259>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2023

We are grateful to Viral Acharya, Mike Chernov, Will Diamond, Darrel Duffie, Andrea Eisfeldt, Robert Goldstein, Zhiguo He, Anil Kashyap, Narayana Kocherlakota, Ralph Koijen, Arvind Krishnamurthy, Sydney Ludvigson, Emi Nakamura, Carolin Pflueger, Pierre-Olivier Weill, Stavros Panageas, and Jon Steinsson, as well as conference and seminar participants at UCLA, Chicago Booth, LBS, USC, MIT, Stockholm, Cornell, University of Utah, University of Florida, Georgia State University, the Federal Reserve Board, New York Fed, Boston Fed, UT Dallas, Temple, NYU, Stockholm, NBER Monetary, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Valentin Haddad, Alan Moreira, and Tyler Muir. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Whatever it Takes? The Impact of Conditional Policy Promises
Valentin Haddad, Alan Moreira, and Tyler Muir
NBER Working Paper No. 31259
May 2023
JEL No. E50,G0

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 persistently distort asset prices, risk, and the response to future announcements.

Valentin Haddad
University of California, Los Angeles
Anderson School of Management
Office C4.19
110 Westwood Plaza
Los Angeles, CA 90024
and NBER
vhaddad@ad.ucla.edu

Tyler Muir
University of California at Los Angeles
Anderson School of Management
110 Westwood Plaza
Los Angeles, CA 90024
and NBER
tyler.muir@anderson.ucla.edu

Alan Moreira
Simon Graduate School of Business
University of Rochester
Box 270100
Rochester, NY 14627
and NBER
alan.moreira@simon.rochester.edu

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 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, the price response to future announced purchases much weaker, and the scope for moral hazard concerns much larger, 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. 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 of long-term distortions in price and risk dynamics, as well as in 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 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 30% without intervention, the intervention raises them

by 35%.¹ However, in states where prices would have otherwise increased 10%, the intervention only increases them by an additional 1%. 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 much more aggressive intervention if economic conditions had worsened instead of improved.

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 this state-contingent behavior. 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 recovery 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,” $g(\cdot)$, that tells us 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 (Breedon 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 $g(\cdot)$ that ties these two distributions together. Which prices have been moved to which new levels to obtain the post-announcement distribution? This type of mathematical question is a transport problem, and we derive its solution. The price support function is unique as long as the policy is order-preserving: the Federal Reserve does not support prices so much in bad states of the world that they exceed prices in better states.

¹Prices are still dropping in these states, by $1 - (1 - 0.3) \times 1.35 = 5.5\%$.

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 30 times larger purchases if the price had fallen by 30% compared to the realized state, and around five times larger purchases compared to the median state. 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 per dollar of purchases would have to exceed \$200 in some states, which is orders of magnitude larger than estimates elsewhere in the literature.

A related way to interpret conditional policy actions that generate asymmetric price support is through multiple equilibria, where equilibrium selection is achieved by promising to backstop fire sale states. This is exemplified by Henry Paulson’s quote: “If you’ve got a bazooka, and people know you’ve got it, you may not have to take it out.”⁴ Under this view, promises of large enough policy actions in the case of a run or fire sale may be enough to prevent the fire sales themselves.

²See 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.

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

⁴For example, in Diamond and Dybvig (1983), the “run-free” equilibrium can be achieved by promising to make depositors whole in the event of a run. Making them whole requires the whatever-it-takes approach and is thus a state-contingent policy action in the bad state. However, once the market believes this there will be no runs in equilibrium.

This selection of equilibrium mechanism would show up as a sharp change in the probability of very bad states.

Another alternative is that our measurement framework fails to account for changes in the pricing of risk between the announcement and the purchases — changes in risk premia afterwards are not an issue — and that such changes drive the asymmetric price support.⁵ Our baseline calculation makes no assumption about risk-pricing in general, but assumes that risk-pricing does not change on announcement for payoffs at the horizon between the announcement of purchases and actual purchases. Broad changes in the pricing of risk, e.g. driven by variation in aggregate risk aversion, are unlikely to rationalize the data. Empirically, other assets such as stocks or high-yield bonds do not react to the announcement, despite these assets being at least as sensitive as investment-grade bonds to broad changes in the pricing of risk. Alternatively, specialized investors who only trade investment-grade corporate bonds might drive risk premia. In such theories, 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 affects 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.

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.

⁵Changes in risk-pricing after purchases are allowed by our framework and a likely justification for the price impact of purchases; see Appendix Section C. 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).

The price support function gives a sharp measurement of the short-term implications of state-contingent policy. But the announcement of purchases can also have long-term implications, which we study. For example, the market might believe that the Federal Reserve will now step into the corporate bond market whenever it gets distressed. Several pieces of evidence suggest such long-term effects. First, after the new policy, tail risk in corporate bond markets becomes far less sensitive to other measures of asset price tail risk, for example measured using equity index options. This suggests that the “policy put” is still present after the program has officially ended, dampening downside risk in this market. Further, corporate bond spreads become far less responsive to changes in pseudo-spreads implied by equity options constructed by Culp et al. (2018) and the level of credit spreads post-intervention appears far too low compared to the pseudo spreads. These results suggest that expectations of future interventions significantly impact asset price dynamics, consistent with the view that the possibility of future purchases in bad states can de-link bond prices from fundamentals. Such price distortions could potentially induce moral hazard by issuers and investors.

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., 2021).⁶ 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 finan-

⁶See also D’Amico and King (2013), Hamilton and Wu (2012).

cial 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. Our results 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). 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 beliefs.

Our results are related to broader work using information in options markets to 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.⁷

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: how much is the policy changing prices as a function of the state of the world.

⁷See also Grinblatt and Wan (2020), which discusses anticipated effects of announcements.

1.1 The Effect of Policy Promises on Asset Prices

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

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

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

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

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.

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 during

US 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 will scale up the policy by an additional amount Q^* if we are in a state at date 1 where the no-intervention price would fall below a cutoff value p^* . In this situation, the price at date 1 becomes $p'_1 = p_1(1 + \mathcal{M}(Q + 1_{p_1 \leq p^*}Q^*))$. The post-announcement price is

$$p'_0 = E[p_1] + E[p_1]\mathcal{M}Q + E[p_1 \times 1_{p_1 \leq p^*}]\mathcal{M}Q^*. \quad (2)$$

We see that both the baseline policy and the implicit promise shape the price response to the announcement. The promise provides an additional boost to the price equal to the product of the additional policy implemented, policy effectiveness, and the contribution of states where the promise is realized to the expected price.

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

$$\mathcal{M}_{\text{estimated}} = \mathcal{M} \left(1 + \underbrace{\frac{E[p_1 \times 1_{p_1 \leq p^*}]}{E[p_1]}}_{\text{contribution of promised states to the price}} \times \underbrace{\frac{Q^*}{Q}}_{\text{rel. size of the promise}} \right). \quad (3)$$

Because the promise provides additional price support, the effectiveness of the policy will be over-estimated. How large is the bias? First, the bias depends on how likely are the promises to be implemented, specifically how much the states where the promise is implemented contribute to the initial price. Of course this contribution is always less than 1, and can be small if a crash is unlikely. However, because new policy tools are often used in difficult and uncertain conditions — think of the midst of the COVID-19 crisis — these probabilities are likely non-negligible. Second, the bias depends on the size of the promised policy relative to the baseline amount Q^*/Q . This second term can be sizable, for example much larger than 1. Indeed, if the crash scenarios are

dramatic, the policy maker might expend significantly more resources. In the empirical evidence we study later on in this paper, we find that indeed, both the probability of additional support and the strength of additional price support are economically significant. To be concrete, the corporate bond market increased by about \$0.5-1 trillion in value when the Fed announced corporate bond purchases in March of 2020 though the Fed ended up only making purchases of around \$15 billion. Our estimates suggest that close to half of the response is due to a 30-fold increase in the size of the program in the lowest 20% of states.

How can one separate the promise from the baseline policy in this case? Option contracts on the asset offer a path forward. Intuitively, in the case with promises, put options with low strike prices pay off only in states where the promise is realized. In contrast, call options with high strike prices are not affected by the promise. Thus, the presence of policy promises can be detected by “more action” in out-of-the-money puts than out-of-the-money calls.

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. We introduce a flexible representation of conditional policies and the assumptions underlying its estimation. Then, we explain the two steps necessary to go from option data to a conditional policy.

1.2.1 The conditional price support function.

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

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

$$p'_1 = p_1 (1 + g(p_1)). \quad (4)$$

This assumption entertains the conditional nature of the announced policy in a flexible way. A fixed policy can result in a constant g , while our example with promises from the previous section corresponds to $g(p) = \mathcal{M}Q + \mathcal{M}Q^*1_{\{p \leq p^*\}}$. The price support $g(\cdot)$ summarizes the total impact of the policy, and on its own does not separate the 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.

Our goal is to recover the price support function g from data on option prices. Option contracts are useful because they allow us to zoom in on different parts of the state space. Consider for example a call option — a contract paying $\max(p_1 - K, 0)$ at date 1 — with a large strike price ($K^h > p^* + \mathcal{M}(Q + Q^*)$). Because this call only pays in states of the world where the promise is not realized, the change in option price at announcement is entirely driven by the baseline promise. Conversely, a put option with a low strike ($K^l < p^* - \mathcal{M}(Q + Q^*)$) only pays off in states in which the promise is realized. Then the change in option price at announcement reflects only interventions with the promise.⁸ Comparing the price of these contracts before and after the announcement reveals how values of the price are changed by the policy. However, one faces two challenges in implementing this idea. First, asset prices are affected not only by actual distributions of outcomes but also by risk adjustments. Said otherwise, option prices only reveal risk-neutral expectations. Second, multiple policy support functions can lead to the same change

⁸Formally, let $[x]^+ = \max(x, 0)$, then the price response of the high and low strike price option are respectively $E[[p_1(1 + \mathcal{M}Q) - K^h]^+] - E[[p_1 - K^h]^+]$ and $E[[K^l - p_1(1 + \mathcal{M}(Q + Q^*))]^+] - E[[K^l - p_1]^+]$.

in distribution. Next, we introduce two simple assumptions that overcome these issues, and argue they are plausible. We discuss generalizations of these assumptions in Section 1.3.

First, an assumption is necessary about the pricing of the asset and options on it at date 0. We assume pricing by the same risk-neutral distribution over underlying states of the world before and after the policy announcement.⁹

Assumption 2. Asset pricing. *The same risk-neutral distribution F_{p_1} over states of the world p_1 prices the asset and options before and after the policy announcement. That is:*

i) Before the announcement, for all functions h , a claim paying $h(p_1)$ at date 1 has price

$$\int h(p_1)dF_{p_1}(p_1) = E[h(p_1)].$$

ii) After the announcement, for all functions h , a claim paying $h(p'_1) = h(p_1(1 + g(p_1)))$ at date 1 has price

$$\int h(p_1(1 + g(p_1)))dF_{p_1}(p_1) = E[h(p_1(1 + g(p_1)))].$$

Underlying this assumption is the simple view that policies do not change the fundamental randomness of the world. Instead, they change what happens in various states of the world. This is standard in the setup of dynamic stochastic models: start with a primitive filtration and probability measure, and derive equilibrium outcomes. We go one step further and assume a constant risk-neutral probability measure. This assumption brings some flexibility: we do not impose coincidence of risk-neutral and historical probabilities, nor do we assume the ability to recover historical probabilities from prices. However it also has some bite: we are implicitly assuming that the stochastic discount factor between date 0 and date 1 is unaffected by the policy. When looking at the data, we explore ways to assess the plausibility of this assumption. For example, one can check whether a related asset for which the policy has no direct effect responded to the announcement. In the case of corporate bond purchases, we compare high-yield bonds, which were not initially targeted, to investment-grade bonds which were. A shift in pricing kernel perspective would imply large movements in the risk-neutral distribution of high-yield bonds while conditional policy targeted at investment-grade would not. Interpretations of the data with a more segmented view of financial markets would limit the usefulness of this comparison. We show how to address such situations in Section 1.3.

⁹To simplify notations, we are ignoring risk-free discounting between date 0 and 1. Equivalently, we are working with forward prices.

It is also worth pointing out that the framework does not put 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 we do not make assumptions about where \mathcal{M} comes from. An appealing mechanism is through effects of purchases on risk premia *from date 1 onwards*. Appendix C presents a model of this mechanism. This dimension is distinct from the properties of pricing between date 0 and 1 in Assumption 2.

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

Assumption 3. Order-preserving policy. *The post-policy price $p'_1 = p_1(1 + g(p_1))$ is increasing in the no-policy price p_1 .*

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

1.2.2 Estimation strategy

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

Step 1: Recovering the future price distribution. We follow the approach of Breeden and Litzenberger (1978) to recover the distribution of the future price of the asset. They show that

observation of option prices (calls or puts) across strikes 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

$$\frac{dPut(p_1, K)}{dK} = \lim_{h \rightarrow 0} \frac{Put(p_1, K + h/2) - Put(p_1, K - h/2)}{h} = 1_{\{p_1 < K\}} \quad (5)$$

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) $F_{p_1}(K)$.

The first step of our method is to apply this idea to the option curve (the relation between strike and put price) before and after the announcement. Doing so allows to recover the cumulative distribution function of the no-policy price p_1 and the post-policy price p'_1 , which we denote F_{p_1} and $F_{p'_1}$, respectively. In practice, one cannot observe option prices for all strikes, but instead for a finite number of specific strikes. Appendix A provides additional details on implementation.

Step 2: Solving the transport problem. Once we have the two distributions, the next task is to find the conditional price support $g(\cdot)$ that explains the change in distribution. This type of problem is known as a transport problem: how should we move all the values of a random variable to obtain a new distribution? The order-preserving property, Assumption 3, imposes that this transport is monotone. This feature is enough to guarantee existence and uniqueness (up to probability-0 events) of a solution $g(\cdot)$. There is a simple method to construct this mapping. Start from a value x and compute the corresponding CDF $F_{p_1}(x)$. Then, because the order of states of the world is unchanged, this value must map to another value y that falls at the same ranking, that is, the same CDF value. This corresponds to finding y such that $F_{p'_1}(y) = F_{p_1}(x)$. 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

¹⁰This is similar to constructing a Q-Q plot.

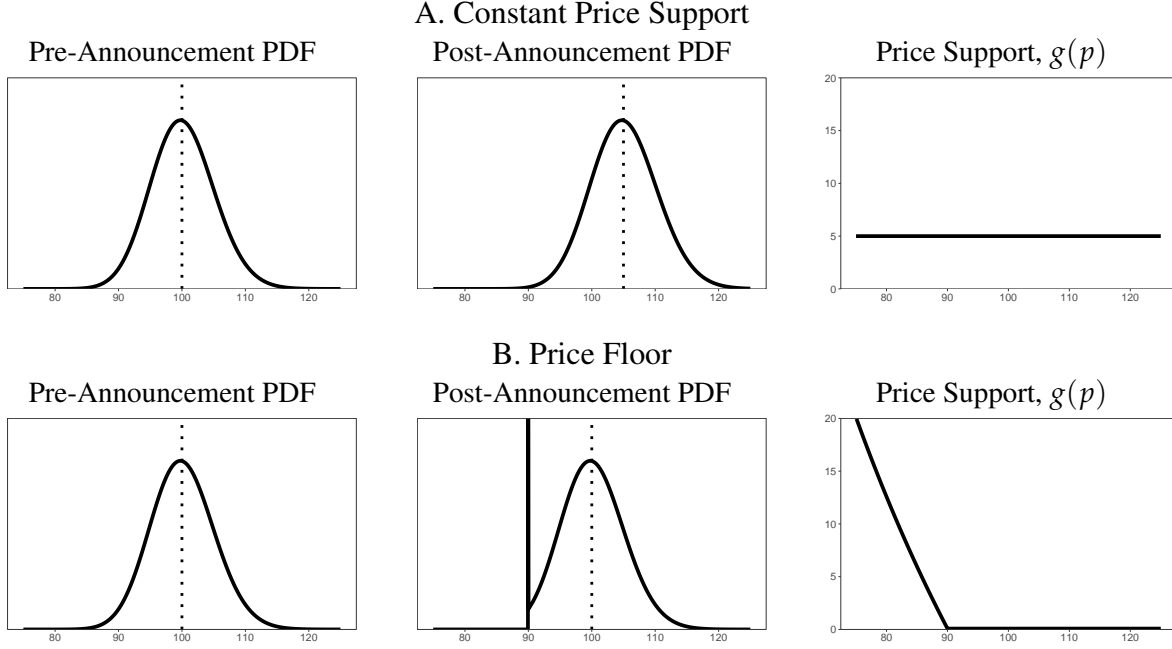


Figure 1: Examples of Distributions and Conditional Price Support Policies

The top row considers a constant price support: the price is increased by the same amount in all states. The bottom row considers a price floor: the price is forced to stay above a threshold 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 $g(\cdot)$.

change in price necessary to move from this initial value to the post-policy price. The following proposition summarizes this calculation.

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

$$g(p_1) = \frac{F_{p'_1}^{-1}(F_{p_1}(p_1)) - p_1}{p_1}. \quad (6)$$

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

Figure 1 illustrates how changes in distribution map to the price support function. We con-

sider two extreme cases of conditional promises. Panel A reports the probability density functions (PDFs) before and after the announcement in the left and middle panels, respectively, and the price support function for a constant price support in the right panel. In this case, there is no conditional promise: the price is increased by the same relative amount no matter what happens. The whole distribution experiences a shift to the right. Panel B reports the same quantities but for a price floor, that is $p'_1 = \max(p_1, \underline{p})$ for some threshold \underline{p} . This is the most extreme case of conditional promise: if the price falls too low absent intervention, the policymaker does whatever it takes to ensure it stays up to the threshold. In terms of distribution we see no change above the threshold but the probability below the threshold becomes accumulated right at the threshold. This corresponds to a sharply decreasing function $g(p)$ below the threshold. For each unit that the no-policy price falls further down, the price gets supported by one more unit to stay at the threshold. This slope of -1 in this range is actually the largest permitted while maintaining the order-preserving property. Interestingly this price support function coincides exactly with a put option payoff, giving a formal way to measure the commonly used notion of a “policy put.”

1.3 Relaxing the Pricing Assumption

Before turning to the data, we show two ways to relax Assumption 2 and entertain a more flexible impact of the policy intervention on the pricing kernel between date 0 and 1. We first introduce a more general setting, and then state the two sets of results.

1.3.1 General pricing framework

First, assume that there is a potentially multidimensional underlying state of the world s at date 1, with actual distribution $F^{\mathbb{P}}(s)$, which does not change before and after the announcement. The state s captures for example how severe the COVID-19 pandemic is at that date. Then, assume that all financial contracts are priced by a pricing kernel $m(s)$ before the announcement and $m'(s)$ after the announcement. This implies that a claim with payoff $x(s)$ at date 1 has price $E^{\mathbb{P}}[mx] = \int m(s)x(s)dF^{\mathbb{P}}(s)$ before the announcement and price $E^{\mathbb{P}}[m'x]$ after the announcement. For simplicity, we assume that the risk-free discount rate between date 0 and 1 is 0, so $E^{\mathbb{P}}[m] =$

$$E^{\mathbb{P}}[m'] = 1.^{11}$$

With this new notation, the price of the asset at date 1 both before and after the announcements are functions of the underlying state s , which we denote $p_1(s)$ and $p'_1(s)$. We maintain Assumption 1: the effect of the policy on the price is still entirely determined by the value of the price absent policy:

$$p'_1(s) = p_1(s)(1 + g(p_1(s))). \quad (7)$$

We also maintain Assumption 3, that is p'_1 is increasing in p_1 .

Our baseline exercise under Assumption 2 corresponds to the case in which the pricing kernel is not affected by the policy, $m(s) = m'(s)$. In this situation, the risk-neutral distribution of p_1 is unaffected by the policy.¹² We now turn to two settings in which we relate this assumption of invariant pricing kernel.

1.3.2 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 Proposition 1 recovers the price support for a large family of pricing kernels under this null hypothesis.

Proposition 2. *If the true price support function $g(\cdot)$ is constant and the pricing kernel before and after the intervention can be written:*

$$\begin{aligned} m &= \Theta(s, p_1(s)/p_0), \\ m' &= \Theta(s, p'_1(s)/p'_0) \end{aligned}$$

for the same function Θ , then equation (6) correctly recovers the price support function.

The empirical content of this proposition is that finding of a non-constant g reveals the presence

¹¹It is straightforward to adjust all date-0 prices for risk-free discounting, that is considering forward prices.

¹²Note that the risk-neutral distribution is equal to:

$$dF_{p_1}(p_1) = E^{\mathbb{P}}[m(s)|p_1]dF^{\mathbb{P}}(p_1), \quad (8)$$

where in a slight abuse of notation we denote $dF^{\mathbb{P}}(p_1)$ the physical distribution of p_1 implied by the distribution of s .

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 fact that the asset return matters for 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. Many asset pricing models also feature pricing kernels determined by the 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 B.1 derives the result more formally.

1.3.3 Adjusting estimates for endogenous risk premia

By taking a stand on the dependency of the pricing kernel to the properties of returns, we can go further and provide estimates of the price support function that take into account this effect. Specifically, we replace Assumption 2 by the following.

Assumption 4. 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 through its effect on the distribution of the asset return between date 0 and 1. 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.¹³ While this is a somewhat specific case, we show in our empirical application that it implies a large risk premium, which is also very responsive to the properties of the asset. Therefore, it leads to conservative estimates of the price support $g(\cdot)$ accounting for risk premium effects.

The following proposition shows how to recover the price support function when replacing Assumption 2 by Assumption 4.

¹³Martin (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.

Proposition 3. *Under Assumptions 1, 3, and 4, the price support function $g(\cdot)$ is the unique solution to equation (6), where the risk-neutral distribution is replaced by the forward-neutral distribution $F^{\mathcal{N}}$, which can also be obtained from option prices:*

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

In words, instead of focusing on matching the invariant risk-neutral distribution before and after the announcement, we concentrate on a distribution affected only by the exogenous part of the pricing kernel.¹⁴ 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 they can be replicated very simply by combining the same calls and puts as in our baseline case. Appendix Section B.2 derives these results.

2. Corporate Bond Purchases in 2020

We now turn to our main empirical application, the March 2020 Federal Reserve announcement of corporate bond purchases.

2.1 Background and Effect on Prices

On March 23rd, 2020 the Federal Reserve announced purchases of investment grade corporate bonds and corporate bonds ETFs through the Secondary and Primary Market Corporate Credit Facility (SMCCF and PMCCF). The announcement immediately raised investment-grade corporate bond prices by about 7.5% or around \$500 billion in market value. Following Haddad et al. (2021), we use the iShares investment grade corporate bond ETF (LQD) to proxy for the daily return in this market and show formal event-study regressions in Appendix Section D.2 that include controls for high-yield bonds, Treasuries, and stocks.¹⁵ This large ETF captures the immediate price

¹⁴The forward-neutral measure and its economic interpretation are discussed in Geman et al. (1995).

¹⁵Using a longer three-day window increases the abnormal excess return to about 10%. The narrow one-day window provides better identification at the cost that it may take the market time to process the announcement. The

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.¹⁶

The announced capacity of corporate bond purchases on March 23rd was \$300 billion, but the Fed left unclear the total amount they would buy. Ultimately, purchases occurred around three months later and totaled only around \$15 billion, 0.2% of the market capitalization of investment-grade corporate bonds, as of June 2020. No further purchases occurred afterwards. See Haddad et al. (2021) and Boyarchenko et al. (2020) for an in-depth discussion of the Fed’s announcement and additional details on purchases.

It is initially surprising that the corporate bond market increased by \$0.5 trillion on announcement given the quantity of purchases was small. The promises view is that investors in part viewed the announcement as a “policy put” or market backstop. That is, the Fed’s willingness to intervene in this market for the first time carried the potential for additional support or larger than realized purchases if conditions deteriorated. The presence of conditional promises would disproportionately reduce the probability of left tail events in the distribution of corporate bond prices, which we explore next.

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 2 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. While implied volatility dropped notably, the drop was most pronounced in the left tail (deltas below 30%). This empirical finding implies that the risk-neutral probability of extreme low prices was particularly sensitive to the announcement.¹⁷ Appendix Figure OA.2 formalizes this by converting the implied volatility curves into the option-implied cumulative distribution function (CDF) for future values of investment-grade bonds and shows a disproportionately larger drop for left tail outcomes.

narrower window is desirable 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.

¹⁶Haddad 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.

¹⁷This left tail drop is even more pronounced over a longer three-day window.

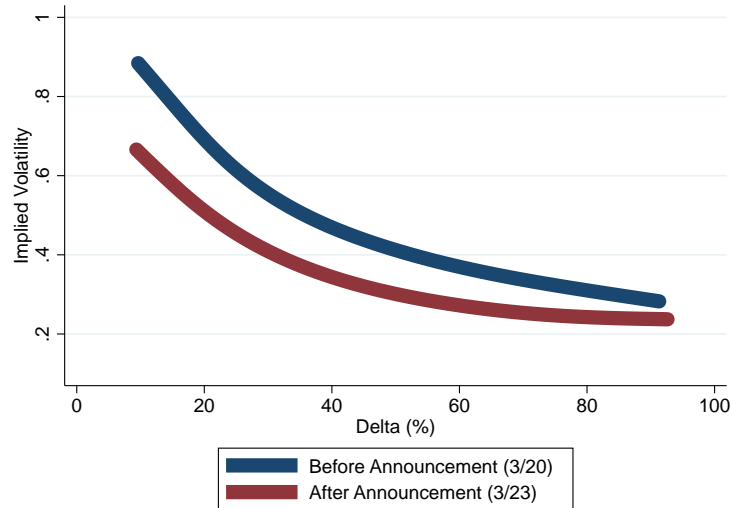


Figure 2: Implied volatility before and after announcement.

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

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

Figure 3 plots the function $g(p)$ expressed as a percentage of the no-policy price p . First, $g(p)$ is not flat as an unconditional price support would imply, but it is strongly downward sloping particularly for low values of the price. At values where the price drops 20-30%, the slope of the price support function is nearly -1 which suggests a policy close to a price floor (e.g., each dollar lost in this region is offset by conditional price support by the Fed). The price support 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 1%. If, prices would have declined by 30% instead, the Fed would push the price up by over 35%. Thus, the asymmetry is economically very large.




Figure 3: 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.

A natural question is whether these movements in the implied volatility curve are typical and so could have happened by chance. 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 3 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. To construct these confidence intervals, we bootstrap daily pairs of implied volatility curves for options on the investment-grade bond ETF. From these, we construct a price support function $g(p)$ on each day. To account for the fact that implied volatility is much higher around the Fed announcement than on other days, we compute the price support in units of standard deviation based on at-the-money implied volatility on the initial day. We then scale up the bootstrapped values based on at the money implied volatility on the trading day before the announcement was made.

The Appendix shows that this finding is a robust feature of the data. Appendix Section D.8 shows our results are robust to bid-ask spreads in option prices and discusses the liquidity of the options we use. Appendix Section D.5 shows our main result is robust to using a longer window in our event study to allow more time for markets to react at the cost of tighter identification. While

magnitudes are slightly larger compared to the results in the one-day window, the asymmetric effect 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 by constructing 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 7.4% 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 no change in the no-policy price, $g(p_0)$. Formally, we define $\tilde{g}(p) = g(p)$ when $p > p_0$ and $\tilde{g}(p) = g(p_0)$ when $p \leq p_0$. Appendix Figure OA.3 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 implied probabilities of each state. Comparing this counterfactual number to the actual announcement return quantifies the effect of the policy put. We find that 53% 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 3.9% or about \$275 billion.

2.4 Interpreting the Asymmetric Price Support Function

Our main methodology delivers conditional price support, but does not specify the mechanism in terms of specific policy actions in different states. We explore several channels.

Conditional policy actions. First, a natural interpretation is that all price support comes from conditional quantities: the Fed will buy more in bad states than good states. To gauge implied magnitudes, we assume that the price impact or multiplier (\mathcal{M}) is the same across states so that $g(p) = \mathcal{M}q(p)$. 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 30% as a reference “bad state.” The Fed would buy 6 times as many bonds in this bad state compared to the state where prices don’t change ($p = 100$), and over 30 times as many bonds compared to the case where prices appreciate by 10%. These are large relative magnitudes. Gauging absolute quantities requires taking a stand on price impact, \mathcal{M} . The literature estimates values in the range of 0.3 to 5.¹⁸ For example, if $\mathcal{M} = 5$, then the Fed would purchase \$500 billion of bonds in the bad state,

¹⁸Gabaix and Koijen (2021) estimate $\mathcal{M} \approx 5$ for the stock market, Greenwood and Vayanos (2014) estimate 0.4

\$80 billion in the case of no change, and \$16 billion in the good state of 10% appreciation. Recall that the announced facility size on March 23rd was \$300 billion (later expanded on April 9th to \$900 billion) so the numbers in the bad state would imply a mild expansion of the facility. The relatively good state quantity of \$16 billion roughly matches the actual purchases the Fed made. This interpretation based on quantities fits the narrative provided by Jerome Powell in 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.” Using values for \mathcal{M} in the bottom of the range found in the literature (a more elastic bond market) would increase the implied quantities.

A related way to view these results is through a multiple equilibrium setting. By committing to intervene enough in the case of a fire sale or run, the Fed could potentially avoid this outcome. For example, in models like Diamond and Dybvig (1983), credibly promising to insure deposits conditional on the run state can prevent runs. This interpretation is closely related to the whatever-it-takes conditional quantity view as it requires credibly promising to purchase whatever is needed to support the price. In this interpretation, the quantity numbers do not necessarily need to be deployed but fiscal capacity from the government needs to ensure they could be deployed if needed. This view relates to work on corporate bond mutual fund fragility (Ma et al. (2020)) and work on runs in safe assets (Eisenbach and Phelan (2022)).

Conditional policy effectiveness. Another interpretation is that asymmetric price support comes entirely from a multiplier $\mathcal{M}(p)$ that varies across states. This interpretation 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. Second, the multiplier in the bad state would have to be around 200 (e.g., by purchasing 1% of the market cap of corporate bonds you raise the price by 200%), which again appears orders of magnitude above other estimates and intuitively implausible. This view also does not explain several other features of the data — for example the changes in asset price dynamics — that we return to in Section 4.

Changes in the pricing kernel. A third interpretation is that the asymmetric price support comes in part through changes in the pricing kernel that Assumption 2 ruled out. Intuitively, conditional purchases might change relative state prices.¹⁹ In our event studies, we find that high-yield bonds

for Treasuries using shocks to supply, and Bretscher et al. (2022) estimate around 0.3 for corporate bonds.

¹⁹For example, Pflueger and Rinaldi (2022) studies a model where policy announcements affect risk-aversion.

and the stock market do not respond to the announcement which cuts strongly against a broad pricing kernel change.²⁰ For example, a reduction in investors' risk-aversion would impact high-yield bonds more than investment grade due to their higher exposure to credit risk.

While the lack of response in high-yield bonds and the stock market suggests that broad macroeconomic risk premiums (or credit risk premiums) do not change over the event, it is possible that market segmentation in investment-grade bonds leads to specific or segmented pricing kernel effects. In section Section 1.3.2, we show that a constant price support pattern generates no change in the pricing kernel for a wide class of standard discount factors. Hence, pricing kernel adjustments are only likely to affect the magnitude of our estimates, not their asymmetry. We assess this view by using the extended framework of Section 1.3.3, and include a risk-premium adjustment in our calculations. Figure 4 shows the price support in the presence of endogenous pricing kernel effects following Proposition 3. The asymmetry, with a much stronger support the left tail is sharp in this case. The magnitude of the support is only dampened by about a third, despite an aggressive adjustment for risk premia: our calculation implies a 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.

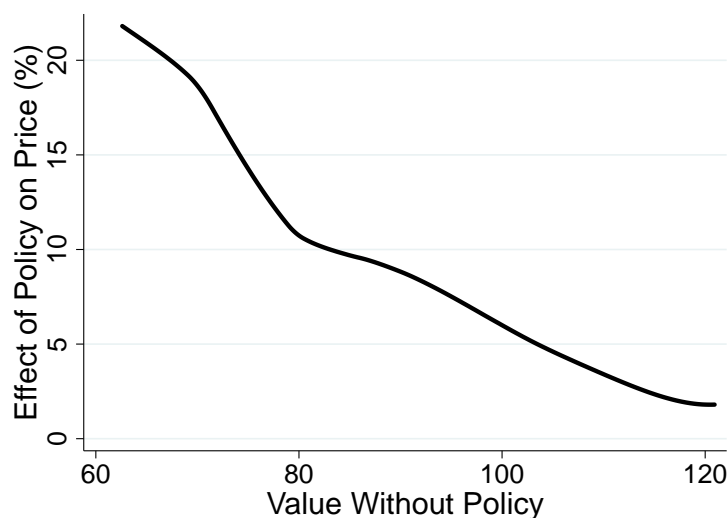


Figure 4: Risk-Premium Adjusted Price Support.

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, under Assumption 4.

²⁰See also Haddad et al. (2021) for more comparisons across asset classes.

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 CDX and Treasuries equal to the historical average and using copula functions following Haugh (2016).²¹ This enables us to decompose our earlier finding into effects driven by movements in the synthetic bonds and the dislocation component. Specifically, Figure 5 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 the statements from Powell emphasizing “markets [...] functioning.”

3. Conditional Promises Everywhere

We expand our analysis to a large set of announcements worldwide in which policymakers step into 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, and various dates associated with the implementation and unwinding of quantita-

²¹See Appendix D.7 for details on implementation and robustness to correlation assumptions.




Figure 5: Decomposition of announcement effects: basis vs synthetic.

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).

tive easing (QE) in the United States from 2008-2013. 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.

Table 1 lists all the announcement dates, and reports the fraction of the announcement report explained by downside support. Figure 6 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.

2008 US financial sector bailout. We first study the October 10th, 2008 announcement of large equity injections to the banking sector as well as guarantees on various forms of bank debt in an effort to “restore confidence in the financial system.” 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.²² 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 6 — reveals that the policy acted as a put option to the

²²See also Kelly et al. (2016a) who use options markets to evaluate government guarantees on the financial sector.

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.

Event	Fraction Explained by Left Tail
High-Yield April 9th 2020	9%
Oct 13th 2008 (Paulson Plan)	46%
BoJ Purchase Speech	11%
<i>US Quantitative Easing Events:</i>	
Nov 25th 2008	2%
Dec 16th	27%
March 19th	20%
June 19th, 2013 (Tantrum)	10%
<i>ECB Announcements:</i>	
May 10, 2010	21%
July 26, 2012	9%
Aug 2, 2012	37%
Sep 6, 2012	20%
Average	20%

financial sector. For example, we see price support of 40% if the equity of the financial sector were to fall by 50% and price support of less than 10% if the equity of the financial sector increased by 50%. This asymmetry confirms that the the policy was largely effective because of a willingness to intervene more in poor states of the world. Communicating this willingness was likely deliberate, with 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 6 studies Japan on April 4th, 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 6 looks at several US quantitative easing (QE) announcements using three-month options on the 10-year Treasury Note futures contract.

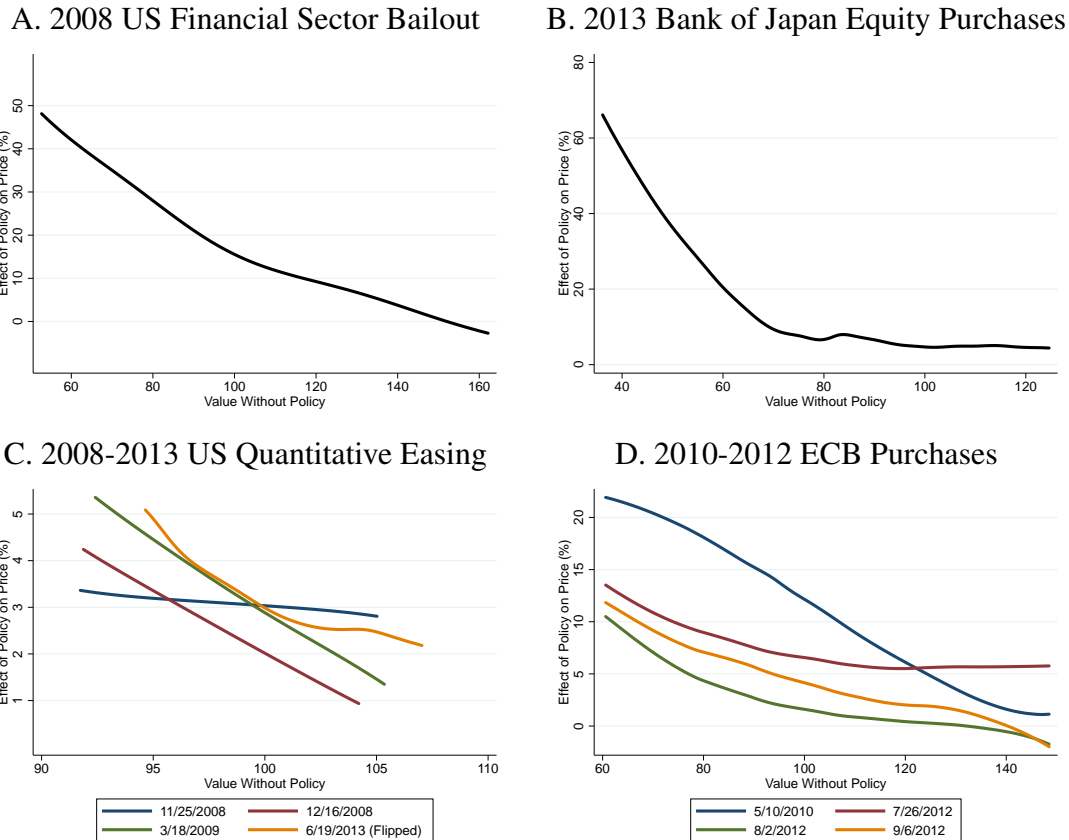


Figure 6: Price Support Function Across Policy Announcements

This figure reports price support as a function of the pre-policy price, normalized to 100 before announcement across policy announcements.

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 25th, 2008 and the FOMC statements of December 16th, 2008 and March 18th, 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. Interestingly, the initial policy announcement on November 25th 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 19th, 2013 where the Fed announced that purchases would decline and markets had a strong negative reaction. The sign is flipped in the plot to be consistent with the other announcements. We see an overall decline in prices (sharp increase in yields) with an upward, rather than downward slope. This announcement is thus associated with not only a tapering of purchases but an associated decline in 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.²³ For the events the set of events we have, where the stock market tracks developments in debt markets, we can 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 6 shows strong effects of conditional policy from these announcements.

²³This 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.

Across all the 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 “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 26th, 2012 speech by Mario Draghi has relatively weaker asymmetry in the price support function compared to the follow on announcements on August 2nd and September 6th. The July 26th 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 2nd announcement), the market priced in stronger conditional price support. The August 2nd and September 6th 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 9th, 2020 high-yield announcement. The announcement of corporate bond purchases by the Federal Reserve on March 23rd focused on investment grade bonds. However, the Fed made an additional announcement on April 9th, 2020 that expanded the facilities to include high-yield bonds. If this announcement contains implicit promises, we would expect them to show up particularly in high-yield bonds. This announcement is useful because we should expect the opposite patterns for high-yield versus investment-grade bonds compared to March 23rd. This is exactly what we see. Appendix Figure OA.6 plots the price support from this announcement following our same methodology applied to options on the high-yield ETF (HYG). We see the same effects of asymmetry: price support is very high for low prices, peaking at over 10%, but is much lower at around 5% for higher levels of prices. This provides strong support for implicit promises 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 the Promises View for Market Dynamics

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 will now step into the corporate bond market whenever it gets distressed which would affect the response of bond prices to bad economic news. Understanding whether price dynamics are affected is important for concerns about whether these policies generate moral hazard from market participations — investors might be less careful in the bonds that they hold if they expect support any time the market crashes. We study these longer run effects on market dynamics through two different lenses.

Section 4.1 studies changing asset price dynamics in the presence of promises. Specifically, if future interventions in crash states are expected, then ex-ante measures of crash risk will be lower and less sensitive to adverse shocks. We show this in the context of corporate bonds after the Fed’s intervention in March of 2020. The results are consistent with a belief of future intervention in the case of a crash.

Section 4.2 considers implications for future purchase announcements. If an initial announcement contains perceived promises of future interventions in bad states of the world, then the observed asset price response to future announcements when bad states happen will be weaker because market participants anticipate them. Ignoring promises would lead to a mistaken interpretation that asset purchase policies are running out of effectiveness. We show patterns of response to announcements consistent with the promises view.

4.1 Hidden Risks and Long-Run Effects in 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.

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 2 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 9th, 2020 when the Fed had already announced the expansion of its corporate bond facilities. Notably, in the period prior to this, corporate bond tail risk and equity market tail risk co-move strongly so that tail risk in corporate bonds was highly sensitive to tail risk in equity markets.

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 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 1st, 2020 to April 9th, 2020. Including an interaction with this dummy doesn’t change our conclusions, and in this case the non-interacted coefficient measures the sensitivity of corporate bond tail risk to other tail risk excluding the COVID episode. 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 Appendix Table OA.2, 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 7 instead uses monthly data on option-based pseudo credit spreads from Culp et al.

Table 2: Long Term Effects on Corporate Bond Tail Risk

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

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

(2018). These pseudo-spreads are constructed by equity options and Treasuries — not corporate bonds — and have been showed 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).²⁴ 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 7. 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,

²⁴We obtain the data from The Credit Risk Lab.




Figure 7: Spreads vs pseudo spreads.

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

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 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.²⁵ Again, this increase in issuance fits with the narrative of an implicit subsidy of low spreads due to expectations of future Fed support.

4.2 Is the Effect of Asset Purchases Getting Weaker?

Another long-term manifestation of policy promises is in how financial market respond to subsequent announcements. If investors already expect these announcements based on the state of the economy, the estimated announcement effects of these interventions should be zero or greatly attenuated.

Hesse et al. (2018), Meaning and Zhu (2011), and Bernanke (2020) find that it is indeed the case that announcements of asset purchases appear to reflect a weakening in the effectiveness of asset purchases. That is, in both the US and Europe, initial announcements of asset purchases

²⁵See also Acharya et al. (2022) who show empirical evidence that quantitative easing impacted firms bond issuance behavior

appear to have powerful effects, but later announcements of the same asset being purchased appear to have much weaker or even negligible effects. 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. Figure 8 shows the impact of these announcements on yields over time for the US, UK, and Europe. When the announcement comes with specific quantities, we divide price responses by this quantity to obtain a multiplier of price movement per dollar purchased. The multiplier, while imprecise, offers a more consistent way to compare announcements, as the interventions vary in the amount of bonds purchased both over time and across countries. Appendix D.11 details the estimates of central bank announcements used in this figure in the US, Europe, and the UK using the work of Joyce and Tong (2012), Meaning and Zhu (2011), Gagnon et al. (2018), Vissing-Jorgensen (2021), and Krishnamurthy et al. (2018). For simplicity, we focus on a subset of announcements, and focus on the programs that have appeared the most powerful.²⁶

Figure 8 reports the announcement responses (black bars) and the multipliers (gray bars). Strikingly, both the total yield changes and multipliers quickly decay to near zero in all three settings.

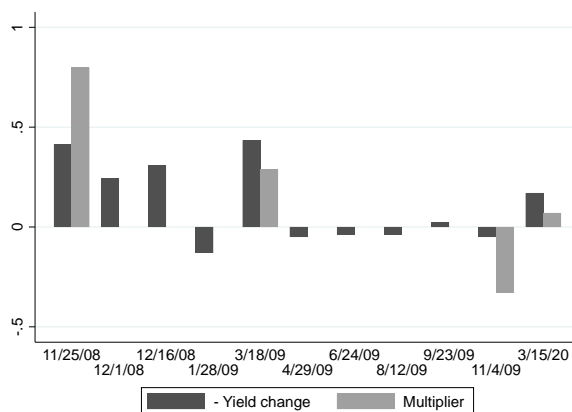
At the same time, the literature has found that the effects of the *original* announcements appear very large relative to estimates of the multiplier \mathcal{M} from other sources of variation. For example, Greenwood and Vayanos (2014) estimate the supply effect on bond yields in the time-series and state: “our estimate that a unit decrease in maturity-weighted debt to GDP lowers the long-term yield by 40 bps is somewhat smaller than the QE estimates.”²⁷ They estimate QE effects from the initial QE1 program of at least 130bps for the same change in supply.

The view that promises are important predicts both of these findings and in fact relates them tightly. Early announcements, which convey not only immediate purchases but the potential for additional purchases in the future if conditions were to get worse, should move prices “too much” relative to the headline purchase amounts. This is the substantial value of the policy put which we have quantified in Section 3. Similarly, later announcements should only matter to the extent they are a surprise relative to the state-contingent policy plan investors perceive. Because investors


²⁶We omit some announcements during QE2 and QE3, for example, but these effects have been shown to be near zero. Similarly, for the UK, Meaning and Zhu (2011) show that the second round of purchase announcements had a negligible effect on yields.

²⁷See also Krishnamurthy and Vissing-Jorgensen (2012).

A. US QE1 and COVID-19



B. ECB Purchases



C. Bank of England QE

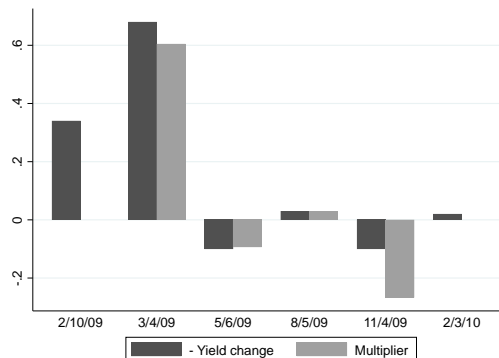


Figure 8: Weakening Announcement Effects of Asset Purchases

This figure plots announcement effects to asset purchase announcements made by the US Federal Reserve, Bank of England, and the ECB. See appendix 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)

already perceive further purchases if the state of the economy is poor, the later announcements effects will be very weak relative to the numbers announced. The promises lens predicts that using the headline quantity announcement to measure the effect of the policy leads to price impact (price movements per dollar purchased) that are volatile, initially too large, and then too small for later announcements. Importantly, the weakening announcement effect does not imply that the effectiveness of purchases themselves has weakened.

An alternative explanation for the declining multiplier documented above is that the multiplier is higher when economic conditions are worse. Since the initial QE announcements occurred in periods when the economy was in particularly bad shape, this view could perhaps explain the “weakening effect.” Several pieces of evidence cut against this view as fully explaining the patterns

in the data.

The price response to the interventions in the Treasury market at the outset of the COVID shock in 2020 are particularly informative to distinguish the promises view from the economic conditions view as explanations for the time-variation in multiplier. These announcements are studied extensively in Vissing-Jorgensen (2021), who find very small multipliers looking at both announcements and actual purchases.²⁸ Vissing-Jorgensen (2021) reports that “an increase of 0.1 (buying 10% of supply) leads to a 5.35 bps larger decline in yields.” This is at least an order of magnitude lower than what is observed in either QE1 in 2008-2009, or relative to the work relating bond supply and yields in Greenwood and Vayanos (2014), despite occurring at a time when the Treasury market and economy were under extreme stress. Our interpretation is that the bond market expected large purchases of Treasuries given the prior experience of QE and the state of the economy. Under this view, it is not that purchases were not effective, but that the market already expected them to occur so that the announcement effect is not informative about effectiveness. The results for Treasuries during COVID contrast sharply with what we have shown 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.

Finally, this experience also contrasts sharply with that of the Bank of Canada (Arora et al., 2021) during the same time period. The Bank of Canada announced purchases of government bonds on March 27th, 2020. Government bond yields declined 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, contrasting with the US experience where such purchases were made in the global financial crisis.

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

²⁸These announcements were very large and quickly translated into actual purchases: within three weeks of the initial March 15th announcement, the Fed had purchased over \$1 trillion in Treasuries.

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 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 several other policy announcements and find that perceptions of downside support are pervasive.

The prevalence of state-contingent policy can lead to important long-term implications. We find evidence suggestive of these implications. For example, prices of corporate bonds appear distorted after promises are made, and this might lead to too much issuance. The state-contingent policy view also helps rationalize jointly the 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.
- 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, 943–83.
- Bianchi, Francesco, Martin Lettau, and Sydney C Ludvigson, 2022, Monetary policy and asset valuation, *The Journal of Finance* 77, 967–1017.
- 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* .
- Charoenwong, Ben, Randall Morck, and Yupana Wiwattanakantang, 2021, Bank of japan equity purchases: The (non-) effects of extreme quantitative easing, *Review of Finance* 25, 713–743.

- Cieslak, Anna, and Annette Vissing-Jorgensen, 2021, The economics of the fed put, *The Review of Financial Studies* 34, 4045–4089.
- 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* .
- Diamond, Douglas W, and Philip H Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of political economy* 91, 401–419.
- 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.
- 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.
- 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? .
- 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.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, 2012, The aggregate demand for treasury debt, *Journal of Political Economy* 120, 233–267.
- 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 .
- 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.
- 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.

Our 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 then compute first and second differences to recover the implied cumulative distribution function or state-price density. 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 will see next that we are still able to recover exactly the function $g(\cdot)$, only over a finite interval.

B. Price Support Function with a Flexible Pricing Kernel.

We derive the results of Sections 1.3.2 and 1.3.3 which generalize our approach to recover the conditional price support.

B.1 Testing a constant support, Section 1.3.2

We prove Proposition 2. First, notice that if, in equilibrium, the pricing kernel does not change, then we are back to the setting of Assumption 2. 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:

$$m'(s) = \Theta \left(s, \frac{p'_1(s)}{p'_0} \right) \tag{OA.1}$$

$$= \Theta \left(s, \frac{p_1(s)(1 + g)}{p_0(1 + g)} \right) \tag{OA.2}$$

$$= \Theta \left(s, \frac{p_1(s)}{p_0} \right) \tag{OA.3}$$

$$= m(s). \tag{OA.4}$$

This confirms our guess and concludes the proof.

B.2 Adjusting the estimates, Section 1.3.3

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 4 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:

$$dF^{\mathcal{N}}(p_1) = E^{\mathbb{P}}[\theta(s)|p_1]dF^{\mathbb{P}}(p_1). \quad (\text{OA.5})$$

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

$$C_K(s) = \begin{cases} p_1(s)/p_0 & \text{if } p_1(s) \leq K \\ 0 & \text{if } p_1(s) > K. \end{cases} \quad (\text{OA.6})$$

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

$$E[m(s)C_K(s)] = E[E[m(s)C_K(s)|p_1]] \quad (\text{OA.7})$$

$$= \int_{-\infty}^K E^{\mathbb{P}}[\theta(s)|p_1] \frac{p_0}{p_1} \frac{p_1}{p_0} dF^{\mathbb{P}}(p_1) \quad (\text{OA.8})$$

$$= F^{\mathcal{N}}(K) \quad (\text{OA.9})$$

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

$$E[m'(s)C'_{K'}(s)] = E[E[m'(s)C'_{K'}(s)|p_1]] \quad (\text{OA.10})$$

$$= \int_{-\infty}^{K'} E^{\mathbb{P}}[\theta(s)|p_1] \frac{p'_0}{p'_1(p_1)} \frac{p'_1(p_1)}{p'_0} dF^{\mathbb{P}}(p_1) \quad (\text{OA.11})$$

$$= F^{\mathcal{N}}(K) \quad (\text{OA.12})$$

Replicating the contract C_K . Finally, the last question 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:

$$C_K(s) = \frac{K}{p_0} 1_{\{p_1 \leq K\}} - \frac{1}{p_0} \max(K - p_1, 0). \quad (\text{OA.13})$$

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 . Remember 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:

$$F^N(K) = E[m(s)C_K(s)] \tag{OA.14}$$

$$= \frac{K}{p_0} \frac{\partial Put}{\partial K}(K) - \frac{1}{p_0} Put(K), \tag{OA.15}$$

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.




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 $K = 80$.

C. 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 policymaker (e.g., a central bank). The policy maker announces asset purchases at date 0 which are made at date 1.

The specialized arbitrageur has log utility over final wealth and chooses his 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 the gross return R_f and 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 who 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 face 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.²⁹ 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 episode fits primarily with this fire-sale interpretation, we could easily have a fundamental cash flow shock at date 1, i.e. a shock to date 1 cash-flow expectations, and this may be a better interpretation of other episodes with asset purchase announcements (e.g., 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

$$\alpha_1 = \frac{E_1[\ln(X/P_1)] - \log(R_f)}{\text{Var}_1(\ln(X/P_1))} = \frac{\mu - p_1 - r_f}{\sigma^2} \quad (\text{OA.16})$$

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 $E_1[\cdot]$ 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

$$p_1 = \sigma^2(1 - \tilde{b} + q) + \mu - r_f \quad (\text{OA.17})$$

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.

²⁹It is also possible to interpret this shock as coming from firms' potentially large debt issuance needs during COVID, rather than sales by the inelastic investors.

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,

$$p_1 - E[p_1|b] = \sigma^2(q - E[q|b]) \quad (\text{OA.18})$$

Only purchases that deviate from what was expected given the announcement in date 0 have any effect, and when the policy maker simply fulfill their promises the effect is exactly zero. This zero effect does not mean that the date 1 intervention was ineffective, but simply that it was already reflected in the date-0 price response.³⁰

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

$$m_2 = (\alpha_1 R_2 + (1 - \alpha_1)r_f)^{-1} = ((1 + \tilde{b} - q)R_2 + (q - \tilde{b})r_f)^{-1} \quad (\text{OA.19})$$

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 $E_0 \left[\frac{W_0 P_1}{W_1 P_0} \right] = 1$ 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 $E_0 \left[\frac{W_0 P_1}{W_1 P_0} \right] = 1$. Using $E_0 \left[\frac{W_0}{W_1} R_f \right] = 1$, we have

$$P_0 = \frac{1}{R_f} \frac{1}{E_0 \left[\frac{1}{P_1} \right]} \quad (\text{OA.20})$$

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

³⁰While here we have one announcement date paired with one purchase date, it is immediate to extend the model to speak to the evidence more explicitly by having two announcement dates paired with two purchase dates. The key is that the first announcement reveals the policy rule the policy maker will adopt in the follow on announcement(s). In this way the effect of the first announcement is driven not only by the immediate follow on purchases but also by the follow on announcements which themselves will appear to be ineffective.

not change asset risk between these dates. Thus, our framework recovers the correct price support function.³¹

Deterministic purchases do not affect risk premiums at date 0 for two reasons. The first is that purchases, whether deterministic or state-dependent, do not remove risk from the arbitrageur 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 date 0 to 1 pricing kernel only 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

$$p_0 = \mu + \sigma^2 - \ln \left(E_0 \left[\exp \left(\sigma^2 (\tilde{b} - q) \right) \right] \right) \quad (\text{OA.21})$$

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 affect on the risk of the asset. In the last case one needs to adjust our methodology to account for changes in the risk of the asset, but no other effects.

D. Empirical Appendix

D.1 Data

We use a variety of financial instruments that have traded option contracts referenced to them and were the direct target of policy announcements.³² We aim to use options of maturities close to three months which are frequently the most liquid.³³

³¹Note that $g(p_1) = (F_{p'_1}^{-1}(F_{p_1}(p_1)) - 1)$, 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 $F_{p_1}(y) = F^P(y) \frac{1}{R_{fy} E_0[\frac{1}{y}]}$ and $F_{p'_1}(y) = F_{p'_1}^P(y) \frac{1}{R_{fy} E[\frac{1}{y}]}$ 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 $g(p_1) = g_a$.

³²These include options on the iShares investment grade corporate bond ETF (LQD), the iShares high yield corporate bond ETF (HYG), the future on the S&P500 index, the future on the ten year maturity Treasury bond, the financial sector ETF (XLF), the future on the Nikkei index, and the CDX investment grade credit basket spread.

³³Ideally one would like longer maturity options as well to study whether implicit promises are longer-term in nature, but in practice the liquidity in the vast majority of these markets is heavily concentrated around or below three months.

Table OA.1: Announcement Effect

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

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

D.2 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 the market time to process the announcement.³⁴ 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.

D.3 Option-Implied CDFs

Figure OA.2 shows the option-implied cumulative distribution function (CDF) for future values of investment-grade bonds. More specifically, we capture the (risk-neutral) distribution of the poten-

³⁴We think for this event it is particularly desirable to have a narrow window given that volatility was very high and also the fact that the CARES acts was signed into law four days after this announcement.

tial price for investment-grade bonds in three months using options with a three month maturity but with varying moneyness. We compare this distribution before and after the announcement was made. The figure reveals a clear rightward shift in this overall distribution, but again most notably there is significantly more action in the left tail of the distribution. Before the event there was about a 15% chance that the value of investment-grade bonds would drop by 30% or more. This state of the world is vastly reduced after the policy is announced.




Figure OA.2: CDF based on option prices.

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

D.4 How much did promises contribute to the overall price movement?

Figure OA.3 plots the implied price support function in our main exercise along with a dashed line for flat price support below the current value. 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.

D.5 Longer Event Window

We next show robustness to using a longer window in our event study. Our main results use one day which tightens identification. However, it could also be reasonable to allow more time for markets to react at the cost of tighter identification since a longer period means that other shocks could be affecting markets. The lower left panel of Figure OA.4 shows our results are similar if we expand our announcement effects to a three-day window. While magnitudes are slightly larger compared to the results in the one-day window (reproduced in the upper left panel), the asymmetric effect is similar. In our earlier analysis, we attributed about 53% of the announcement effect to additional

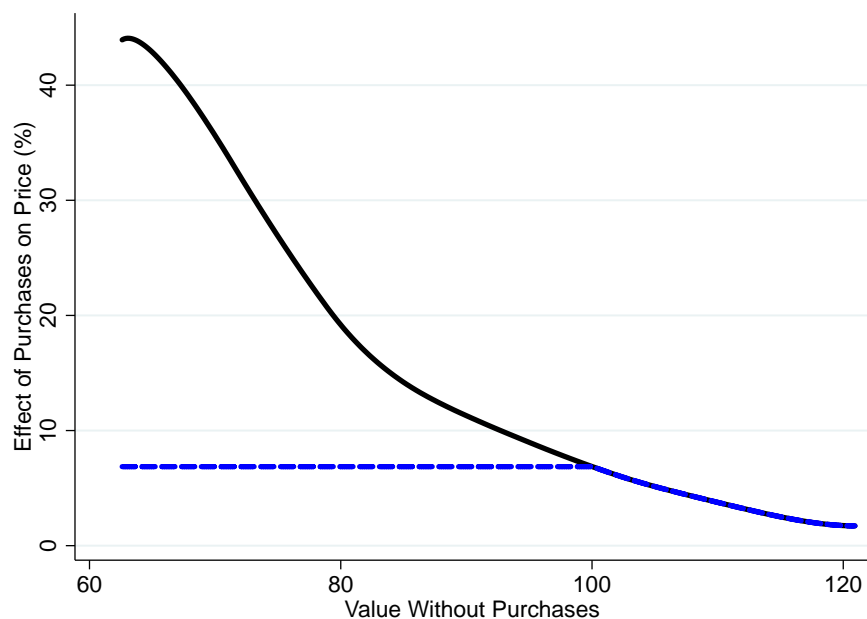


Figure OA.3: Reinterpreting Announcement Effects.

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 current price.

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

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

D.6 Comparison to High Yield Bonds

Importantly, in Figure OA.5 we contrast the effects on investment grade bonds with those for high yield, using options on the largest and most liquid high yield bond ETF (HYG). The upper right panel shows that, over a one-day window, the overall returns for high yield are actually slightly negative. If anything, the pattern is also upward sloping, meaning less price support at low prices and vice versa. These results are useful for two reasons. First, they suggest that the

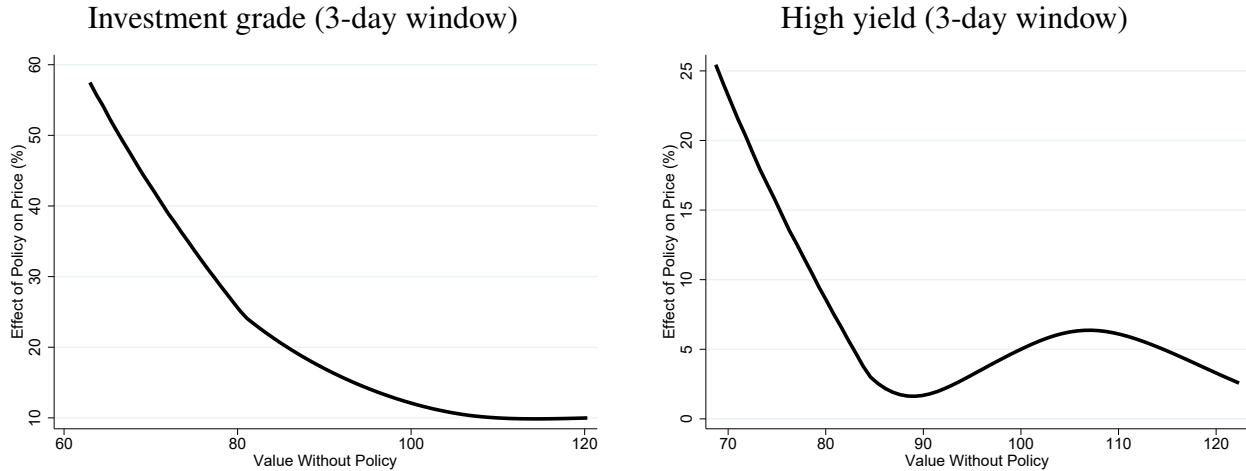


Figure OA.4: Investment Grade vs High Yield.

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

announcement didn't coincide with other macroeconomic news affecting corporate bond markets, since the effects are strongly concentrated in investment-grade bonds which were the target of the purchases. Second, and more importantly, they speak to the possibility that changes in the pricing kernel are driving our results.

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

Figure OA.6 plots the price support from the April 9th, 2020 announcement that expanded the facilities to include high-yield bonds. We follow our same methodology applied to options on the high-yield ETF (HYG).

D.7 In Which States was the Fed Expected to Buy? Details on the Copula method

Here we discuss in more detail how we are able to pinpoint what exactly drove the price response. We start by using options on the ten year treasury futures (formally the name is Ultra 10-year T-Note futures) and option on the CDX north America investment grade index. This CDX index tracks the CDS spreads of 125 most liquid investment grade corporates. With these options we can recover the risk-neutral density for the distributions of credit spreads and interest rates. The options on the corporate bond ETF gives 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 tree prices (and how this distribution




Figure OA.5: Investment Grade vs High Yield.

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

changed). The options only give us the marginal distribution of each one. To recover the joint distribution we rely on Copula method.

Here how it works. Say 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 and we are willing to assume that these correlations are constant. Let this correlation matrix be given by C . We then sample $\tilde{x}_i, \tilde{y}_i, \tilde{z}_i \sim N(0, C)$ where $N(0, C)$ is a multivariate normal distribution with zero mean and univariate standard deviation and i index the draws. We then compute for each realization quantile in the standard normal distribution. For example $f x_i = F(\tilde{x}_i)$ where F is the cdf of a standard normal distribution. We then have $\{f x_i, f y_i, f z_i\}_{i=1}^N$ where N is the number of draws. Finally we use the original marginal densities to invert back the realization, i.e., $\{x_i, y_i, z_i\} = \{F_x^{-1}(f x_i), F_z^{-1}(f z_i), F_z^{-1}(f z_i)\}$. This procedure allow us to simulate from the joint distribution in a way that is consistent with the marginal distributions recovered from options prices. And therefore allow 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. So 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 were to push the correlations to -0.75 or 0.25 (see Figure). 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 Figure). 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.

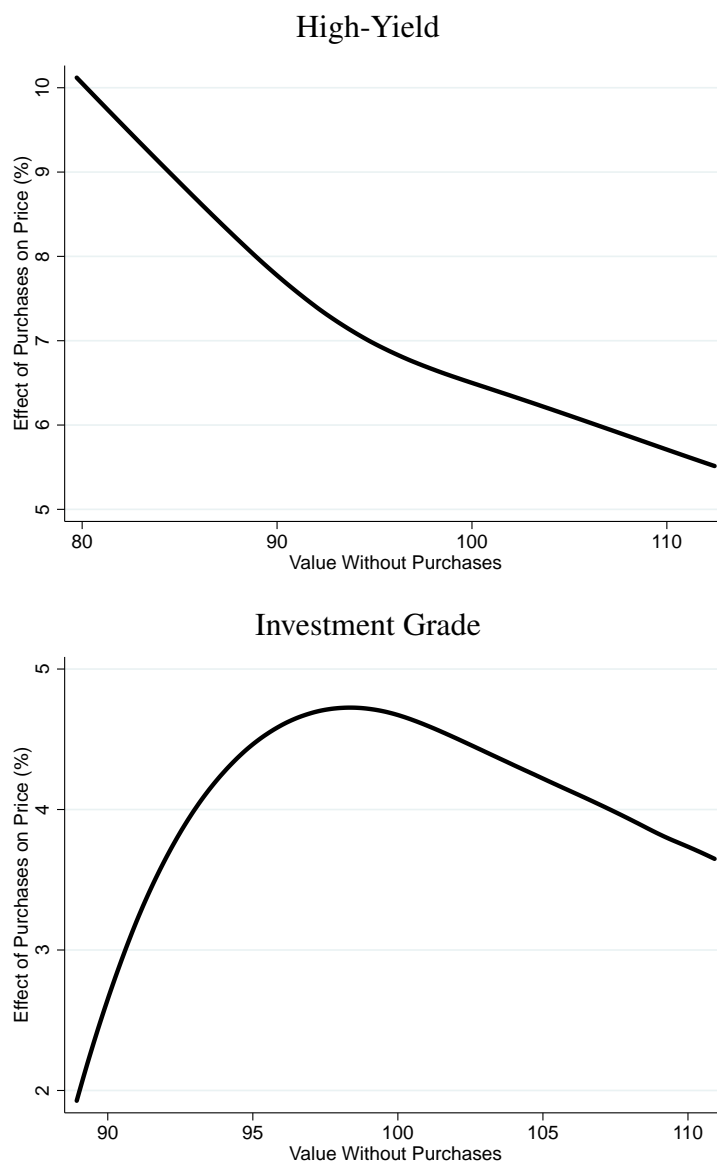



Figure OA.6: High-Yield Announcement, April 9th, 2020. Effect on prices.

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

D.8 Robustness: Liquidity of Options

It is well known that liquidity in option markets—and derivatives markets more broadly—is very heavily skewed. And it is certainly the case that trading volumes in these instruments are far from very high. For example, for the options on the investment grade corporate bond fund we investigate in section 2.3 the overall trading volume in 2020 (in terms of contract notional) were around one hundred billion dollars. Bid-ask spreads were about 3% but grew considerably for low strike prices, reaching values as high as 30% for options we use in our analysis. Thus, it is natural to ask how robust are the patterns we document. In Section 2.3 we show that the negatively

Correlation: interest rates and credit spreads



Correlation: synthetic and cash bond




Figure OA.7: Decomposition of announcement effects: robustness

In this figure we look at how the decomposition in Figure 5 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.

sloped support function we recover is very unlikely to have happen just 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. But 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 now using bid and ask prices. In Figure OA.8 below we report the recovered price support function with all the four pairs. 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 tells you the actual returns of an investor even if we account for the illiquidity implied by a wide bid-ask spread.

D.9 Additional Plots of Specific Announcements

D.10 Corporate Bond Returns and VIX

D.11 Details on Announcement Effects

This subsection provides more detail on data sources used in Section 4.2.

D.11.1 United Kingdom

Table OA.3 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 contained a fairly narrow and specific quantity range. First, note that only the first 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 10th, 2009, did not contain concrete information but suggested that purchases were likely. On March 4th, purchases of £75 billion lead to a decline of 70bps in the 10 year Gilt yield. In contrast, the next three announcements fea-




Figure OA.8: Effect on prices 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 (which is the mid point between these quotes which we use in our baseline analysis)

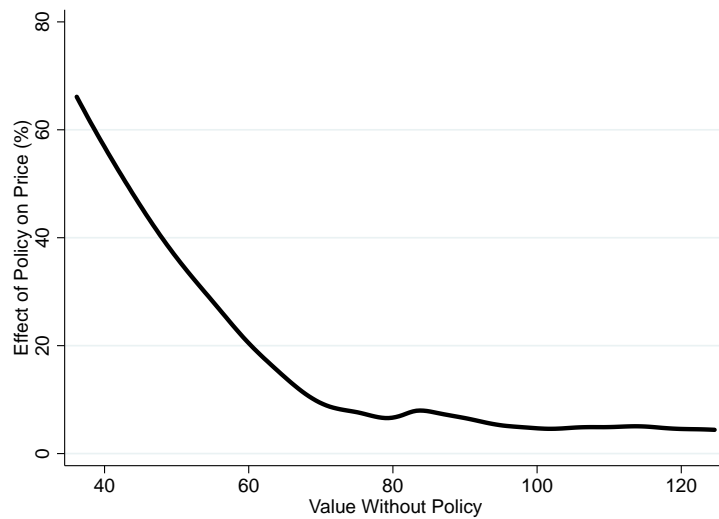


Figure OA.9: Announcement of Purchases by Bank of Japan in April, 2013. Effect on prices.

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

Table OA.2: Long Term Effects on Corporate Bond Prices

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

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

tured 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

Table OA.3: 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]

(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.

D.11.2 United States

Table OA.4 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.³⁵ 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 lead 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 option pricing results in 3 for the initial announcements which indicate implied policy puts from the early 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

³⁵We find similar results using a weighted average of the yield responses where weights are given by the relative supply of each.

Table OA.4: 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.

Date	Yield Responses					Quantities				Multiplier
	Treas	Agy	MBS	Avg		Treas	Agy	MBS	Total	
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

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 15th 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 23rd when the purchase amounts shifted to “unlimited” and the Fed continued to purchase large quantities. These announcements quickly translated into actual purchases – within three weeks of the initial March 15th 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 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 also contrast sharply with what we document for cor-

Table OA.5: 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	

porate 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 27th, 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.

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.

D.11.3 Eurozone

Table OA.5 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.