Whatever It Takes?
The Impact of Conditional Policy Promises

Valentin Haddad, Alan Moreira, and Tyler Muir*

January 25, 2023

Abstract

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

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

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 market infers that policymakers might do "whatever it takes" in that they will go to much greater lengths to backstop markets if necessary. Under this view, the strong asset price response is not only driven by the announced quantity itself, but also by the value of the perceived promise to support markets if the situation gets worse. Distinguishing these views is important. Under the whatever-it-takes view, the potential expansion of the central banks' 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. While asset purchases are an increasingly important tool for central banks globally, these same issues arise in 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.

In this paper, we propose and implement a method to measure the state contingent impact of policy actions. Our method uses changes in the distribution implied by option prices before and after policy announcements to characterize the impact of policy across many possible states. We find pervasive evidence of 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 during the 2008 crisis. These findings are consistent with the "whatever it takes" or "policy put" view that markets perceived larger interventions if conditions worsened. Our results shed light on the effectiveness of these policies, the "weakening

announcement effect" that effectiveness appears much smaller for later announcements of the same policy tool, and debates about how such policies can potentially create moral hazard concerns.

Our primary empirical setting is the Fed's announcement of corporate bond purchases in March of 2020, which dramatically boosted corporate bond prices despite a small amount of ultimate purchases. We find that the announcement increased corporate bond prices by 40% in states where they would have otherwise decreased 30% without intervention, and increased prices only 5% in states where they would have increased 10% in the counterfactual of no intervention – a striking asymmetric effect with a signature very similar to a put option. We quantify how much this state-contingent policy mattered for the overall announcement effect: at least 50% of the large price recovery from the announcement came from additional policy in the left tail. Our results explain why this particular announcement led to an immediate, dramatic increase in corporate bond prices of between $500 billion and $1 trillion, despite very small ultimate purchases of about $15 billion: the market expected much more aggressive intervention if economic conditions had worsened instead of improved.[1]

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. 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. Conversely, an out-of-the-money call that pays off only when prices are high reveals policy when conditions improve. This simple idea already leads to evidence supportive of more intervention in the left tail. Using options on corporate bond ETFs, we see a stronger recovery in the price of contracts targeting bad states of the world at the

[1]This accords with Jerome Powell's statement to the U.S. House of Representatives on June 16th, 2020 that "markets are functioning pretty well, so our purchases will be at the bottom end of the range that we have written down."

announcement – implied volatilities fall much more in the left tail than the right tail of bond prices. We show how to go further than this qualitative assessment and obtain the entire state-contingent policy.

We then estimate what we call a "price support function," $g(.)$, that tell us the state by state response of prices to the policy intervention. Specifically, in the state that the price were to move to p absent intervention, the central bank intervention will raise it by $g(p)$ so that we observe $p + g(p)$ in that state. First, the price of options across different strikes reveals the market perception of the distribution of the future price or risk-neutral distribution (Breeden and Litzenberger (1978)). We estimate 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 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.

The price support function we recover in the data is strongly asymmetric and resembles a "policy put:" while price support is low and relatively flat in good states of the world, it increases to much larger values as one moves towards worse states.[2]

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,

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

2011; Haddad et al., 2021).[3] Further, Hesse et al. (2018), Meaning and Zhu (2011), and Bernanke (2020) find a "weakening announcement effect:" initial stage announcements of asset purchases in the US and Europe have large effects on asset prices but later stage announcements have little to no effect. An interpretation consistent with our results is that early announcement responses are inflated by the value of state-contingent actions to do more if the situation worsens, but in later stages of asset purchase programs this state-contingent policy has already been reflected in prices. Thus, even a large announcement later on can appear to have zero effect.

We can go from measuring contingent price support to measuring contingent policy actions by making additional assumptions on their relation. For example, consider the simple view that the elasticity of bond prices to bond purchases is constant across states. In this case, the asymmetry that we document implies 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 Fed, and we argue it is more plausible than the view of a fixed small quantity but state-dependent elasticity – under that view the price impact per dollar of purchases would have to exceed $200 in some states.

We then bring evidence from additional assets to sharpen our inference. We explore in which states the Fed was likely to provide more support. Corporate bond prices can fall either 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. Using copula methods to model their joint distribution, we find that the distribution of synthetic bonds did not show as severe of a left tail initially. Moreover, this left tail did not shrink as substantially after the announcement. This suggests that the market expected large interventions in states where corporate bond markets were highly

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

dislocated – states with a much larger CDS-bond basis than at the time of announcement.

An important issue is whether changes in the pricing of risk due to the policy can drive the announcement effects instead of our inferred price support. In our baseline calculation we do not make any assumptions about risk-pricing in general, but we do assume that risk-pricing does not change on announcement for payoffs at the horizon between the announcement of purchases and actual purchases. We allow for arbitrary unmeasured changes in risk-pricing for periods beyond when the purchases occur. For example, we allow for risk-premia effects that operate though a change in the amount of risk that investors bear once the central bank buys the assets as in Vayanos and Vila (2021). Empirically, we see no response to the policy announcement in the closely related high-yield market, suggesting that the broad pricing of credit risk does not change on announcement. This cuts against the view of a broad change in the price of risk.

To further explore and clarify these issues we study a segmented markets model as in Vayanos and Vila (2021) with a specialized investor who only trades investment-grade corporate bonds. The investors' pricing kernel is given by $risk\ of\ asset \times quantity\ held$. With constant purchases, we show that the investors' pricing kernel does not change on announcement for payoffs that occur before purchases are made, but will after purchases are made as this is when risk is taken off the investors' balance sheet. Prices rise at the announcement through lower risk premiums from the purchase date onward. Importantly, recovering an asymmetric price support function only happens when purchases are state dependent and asymmetric themselves. When purchases are state dependent, the risk of the asset itself can change and this can change the specialized investors pricing kernel between announcement and purchases. We show how to adjust our price support function in this case and find that our main conclusions of strongly asymmetric price support still hold.

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 if the market believes that the Federal Reserve will now step into the corporate bond market whenever it gets distressed. We present three pieces of evidence suggesting long-term effects of the introduction of bond purchases. We focus on properties of tail risk specifically, similarly to our main approach. First, after the new policy, tail risk in corporate bond markets becomes far less sensitive to other measures of asset price tail risk, for example measured using equity index options. This result suggests that the "policy put" is still present, dampening downside risk in this market. Second, corporate bond returns also become far less sensitive to changes in the VIX, which shows the downside dampening is relevant for the level of prices. Third, corporate bond spreads become far less responsive to changes in the "pseudo spreads" implied by equity options constructed by Culp et al. (2018). This last result demonstrates that the change in downside risk properties is restricted to, or strongest in, the corporate bond market. These results suggest that expectations of future interventions impact asset price dynamics. This is consistent with the view that the possibility of future purchases in bad states can de-link bond prices from fundamentals which can induce moral hazard by issuers and investors. But it is also consistent with the more benign view that future interventions make the bond market less prone to dislocations during crises.

A related way to view our results 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 support in the case of a run or fire sale may be enough to prevent the fire sales themselves. This effect shows up in the left tail of the distribution.

While we conduct an in-depth study of bond purchases during the COVID-19 crisis,

[4]For 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.

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

Our results speak to macro finance models that assess the impact of policies that support financial market prices during crashes such as asset purchases or equity injections to the financial sector (e.g., He and Krishnamurthy (2013), Moreira and Savov (2017), Vayanos and Vila (2021)). Quantitative theoretical exercises typically treat the policy implementation as "one-off" events for simplicity and ignore expectations of future policy across states. In contrast, our findings suggest that state-contingent policy is first order to understanding the effectiveness of policy announcements. Our results also relate to the literature on forward guidance and central banking communication starting with the seminal work of Gürkaynak et al. (2004) which shows that central bank statements have powerful effects on long-term yields. Other examples of this literature include Swanson (2011), Hanson and Stein (2015), and Nakamura and Steinsson (2018).

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

announcements.[5]

1. Measuring Conditional Promises

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

1.1 The Effect of Policy Promises on Asset Prices

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

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

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

the present value of the product of the fiscal multiplier with the pass-through from GDP to corporate profits.

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

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

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

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

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

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

$$p_0' = E[p_1] + E[p_1]\mathcal{M}Q + E[p_1 \times 1_{p_1 \leq p^*}]\mathcal{M}Q^*. \tag{2}$$

We see that both the baseline policy and the implicit promise shape the price response to the announcement. The promise provides an additional boost to the price equal to the 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:

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

Because the promise provides additional price support, the effectiveness of the policy will be overestimated. How large is the bias? First, the bias depends on how likely are the promises to be implemented, 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 $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 $14 billion. Our estimates suggest that close to half of the response is due to a 5-fold increase in the size of the program in the lowest 20% of states.

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

$$p'_{0,call} - p_{0,call} = E[\max(p_1(1 + \mathcal{M}Q) - K, 0)] - E[\max(p_1 - K, 0)]. \tag{4}$$

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

$$p'_{0,put} - p_{0,put} = E[\max(K - p_1(1 + \mathcal{M}(Q + Q^*)), 0)] - E[\max(K - p_1, 0)]. \tag{5}$$

These two cases highlight that we can learn about the conditional nature of the policy by using options to zoom in on various parts of the state space. One has to go further than this simple intuition to get to quantitative statements. For example, comparing equa-

tions (4) and (5) suggests that the presence of policy promises can be detected by "more action" in out-of-the-money puts than out-of-the-money calls. Because these two contracts are very different to start with, more careful calculations are needed to implement such a comparisons. This is the focus of next section.

1.2 A Method to Estimate Conditional Policy Promises

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

1.2.1 The conditional price support function.

We maintain the timing assumptions of our motivating example. At date 0, a policy is unexpectedly announced to be implemented at date 1.[6] The policy announcement potentially contains conditional promises. That is, the policy implementation can depend on the realized state of the world at date 1. Our first assumption is that the state of the world at date 1 maps exactly to the date-1 price of the asset absent policy intervention p_1. Given this mapping, the effect of the policy on the price can be represented by a price support function $g(.)$. 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* $g(p_1)$*,*

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

[6] We relax this assumption of a fixed timing when considering extensions of our approach.

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

Our goal is to recover the price support function g from data on option prices. As we discussed earlier, option contracts zoom in on different parts of the state space. 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 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.[7]

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 com-

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

parison, and suggest risk-premium effects of asset purchases. We show how to address such situations in Section 1.3.

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 aboutwhere $\mathcal{M}$ comes from. An appealing mechanism is through effects of purchases on risk premia *from date* 1 *onwards*. Section 6 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 $g(.)$ 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 $g(.)$. 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

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

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

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

Our approach to recover the state price density is standard and follows Malz (2014). We obtain prices and use the standard Black-Scholes formula to translate prices into im-

plied volatilities. We then fit a cubic spline to the implied volatility curve. However, we are careful to not extrapolate the curves. 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 $g(.)$, only over a finite interval.

Step 2: Solving the transport problem. Once we have the two distributions, the next task is to find the conditional price support $g(.)$ 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 $g(.)$. 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'. The price support function is the change in price necessary to move from this initial value to the post-policy price. The following proposition summarizes this calculation.

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

$$g(p_1) = \frac{F_{p_1'}^{-1}(F_{p_1}(p_1)) - p_1}{p_1}. \tag{8}$$

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

Figure 1 illustrates how changes in distribution map to the price support function. We consider two extreme cases of conditional promises. Panel A and B report the probability density functions (PDFs) before and after policy announcement and the price support function for a constant price support. In this case, there is no conditional promise: the price is increased by the same amount no matter what happens. The whole distribution simply experiences a parallel shift to the right. Panel C and D report the same quantities but for a price floor, that is $p_1' = \max(p_1, \underline{p})$ for some threshold $\underline{p}$. This is the most extreme case of conditional promise: if the price falls too low absent intervention, the policymaker does whatever it takes to ensure it stays up to the threshold. In terms of distribution we see no change above the threshold but all the probability below the threshold becomes accumulated right at the threshold. This corresponds to a sharply decreasing function $g(p)$ below the threshold. For each unit that the no-policy price falls further down, the price gets supported by one more unit to stay at the threshold. This slope of -1 in this range is actually the largest permitted while maintaining the order-preserving property. Interestingly this price support function coincides exactly with a put option payoff, lending formal support for the commonly use notion of a policy put such as the "Greenspan put."

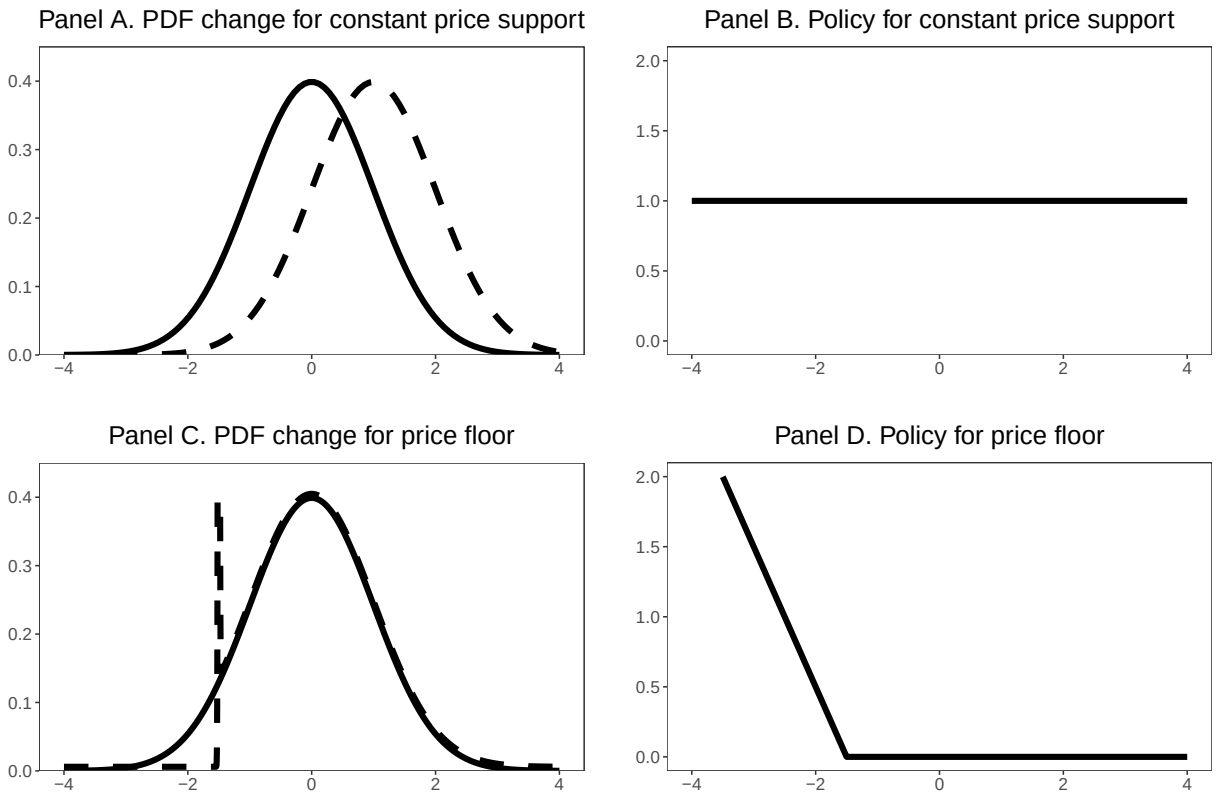


Figure 1: Examples of Distributions and Conditional Price Support Policies
The top row considers a constant price support: the price is increased by the same amount in all states. The bottom row considers a price floor: the price is forced to stay above a threshold $\underline{p}$. The left panels report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

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 dramatic 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 $\pi(s)$ at date 1 has price $E^{\mathbb{P}}[M\pi] = \int M(s)\pi(s)dF^{\mathbb{P}}(s)$ before the announcement and price $E^{\mathbb{P}}[M'\pi]$ after the announcement. For simplicity, we assume that the risk-free discount rate between date 0 and 1 is 0, so $E^{\mathbb{P}}[M] = E^{\mathbb{P}}[M'] = 1$.[8]

With this new notation, both the price of the asset at date 1 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:

$$p_1'(s) = p_1(s)(1 + g(p_1(s))). \tag{9}$$

We also maintain Assumption 3, that is p_1' in 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. This risk-neutral distribution is simply equal to:

$$dF_{p_1}(p_1) = E^{\mathbb{P}}[M(s)|p_1]dF^{\mathbb{P}}(p_1), \tag{10}$$

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

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

implied by the distribution of s. 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

One path to entertain more generality than Assumption 2 is to ask a more restrictive question than a full recovery of g. Specifically, we focus on the question of whether the price support function $g(.)$ is constant. This null hypothesis consists in supporting the price of the asset by the same proportion in all sets of the world. The next proposition shows that the approach of Proposition 1 correctly recovers the price support for a large family of pricing kernels under this hypothesis.

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

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

for the same function Θ, then equation (8) *correctly recovers the price support function.*

The empirical content of this proposition is that if the pricing kernel is within this family, then a finding of a non-constant g reveals the presence of state contingency in the price support. What is this family of pricing kernels? 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 between these dates 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 kernel 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 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 A.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. *Assume that 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.*

Under this assumption the pricing kernel can be affected by the announcement, because the distribution of the asset return between date 0 and 1 is changed. Specifically, it assumes that the part of the pricing kernel that is endogenous to returns behaves as for a log-utility investor with all her wealth invested in the asset. While this is a somewhat specific case, we will show that in our empirical application 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 $g(.)$ accounting for risk premium effects.

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

Proposition 3. *Under Assumptions 1, 2, and 4, the price support function* $g(.)$ *is the unique solution to* (8)*, where the risk-neutral distribution is replaced by the numeraire-equivalent distribution* $F^{\mathcal{N}}$*, which can also be obtained from option prices:*

$$dF^{\mathcal{N}}(p_1) = E^{\mathbb{P}}[\theta(s)|p_1]dF^{\mathbb{P}}(p_1). \tag{11}$$

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. As showed by equation (22), 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 migh seem unusual, we show they can be replicated very simply by combining the same calls and puts as in our baseline case. Appendix Section A.2 derives these results.

2. Data

We use a variety of financial instruments that have traded option contracts referenced to them and were the direct target of policy announcements.[9] We aim to use options of maturities close to three months which are frequently the most liquid.[10]

3. Corporate Bond Purchases in 2020

3.1 Background and Effect on Prices

On March 23rd, 2020 the Federal Reserve unveiled facilities that would purchase investment grade corporate bonds and corporate bonds ETFs through the Secondary and Primary Market Corporate Credit Facility (SMCCF and PMCCF). This announcement was important because it was the first time the Fed had directly targeted the corporate bond

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

[10]Ideally 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.

market. As Haddad et al. (2021) show, the announcement came at a time when corporate bond prices were depressed and dislocated (investment grade debt was trading at steep discounts to Treasuries even taking into account credit risk measured from the CDS market).[11] These facilities took time to set up so that no bonds were purchased for several months, and there was a total (announced) capacity of $300 billion.

The announcement of the SMCCF and PMCCF had a significant and immediate impact on corporate bond prices. Haddad et al. (2021) show this with the iShares investment grade corporate bond ETF (LQD) using a window of one to three days around the announcement.[12] 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 much less frequently. The one-day announcement window return is 7%, and the abnormal excess return is about the same (with controls for high-yield bonds and the stock market, see Appendix **??**). The 7% return translates into around a $500 billion increase in market value for investment grade corporate bonds.[13]

Most importantly for this paper, the announcement to purchase corporate bonds was unprecedented as the Fed had never purchased corporate bonds before. Because in an event study we can only cleanly measure the effect of the announcement surprise, this

[11] 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.

[12] 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 almost no movement. See also appendix **??** for a summary of the evidence

[13] Using a three-day window for the announcement increases the raw return and abnormal excess return to about 10%. A shorter one-day window provides better identification at the cost that it may take the market time to process the announcement. The narrower window is likely particularly 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.

setting is a particularly good laboratory to measure the proposition that these announcements work through state contingent promises.

While the Fed left unclear the total amount of bonds they would buy, the announced total capacity of purchases on March 23rd was $300 billion. Ultimately, only around $15 billion of purchases were ultimately made, or around 0.2% of the market capitalization of investment grade corporate bonds.

It is initially surprising that the corporate bond market increased by $0.5 trillion on announcement given so few purchases were made. The promises view is that investors understood that purchases would have been larger had the market fallen further and that this conditional promise is an important driver of the recovery. Thus, the presence of conditional promises will disproportionately reduce the probability of left tail events in corporate bond markets.

3.2 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 17 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 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.[14]

Figure 3 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 potential price for investment-grade bonds in three months using options with a three month maturity but with varying moneyness. We compare this distribution

[14]This left tail drop was even more pronounced over a longer three-day window.

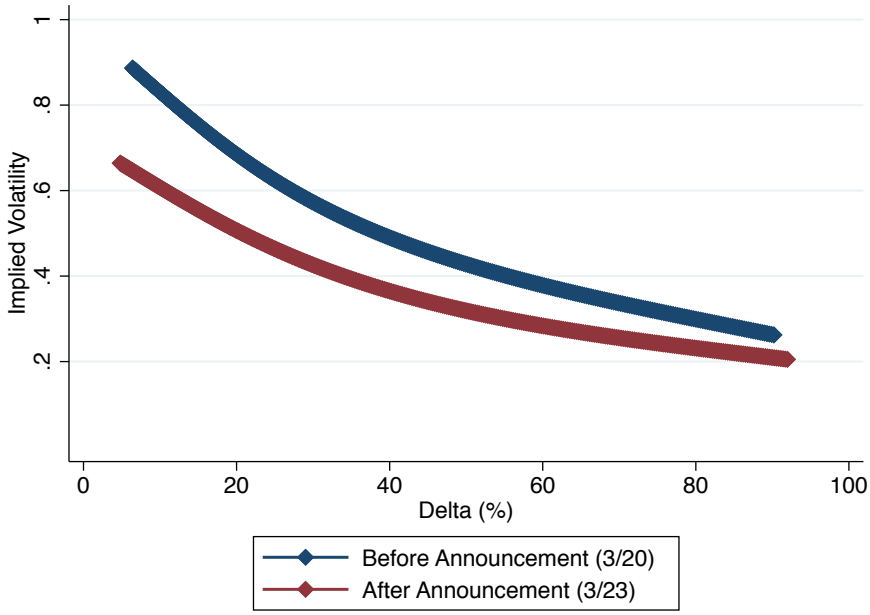


Figure 2: Implied volatility before and after announcement.
This figure provides implied volatility from options on an Investment-Grade corporate bond ETF (LQD) on March 20th and 23rd, 2020 as a function of the option delta. Time to maturity is 3 months.

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.

3.3 Conditional Price Support

We now use apply our measurement framework to recover the 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

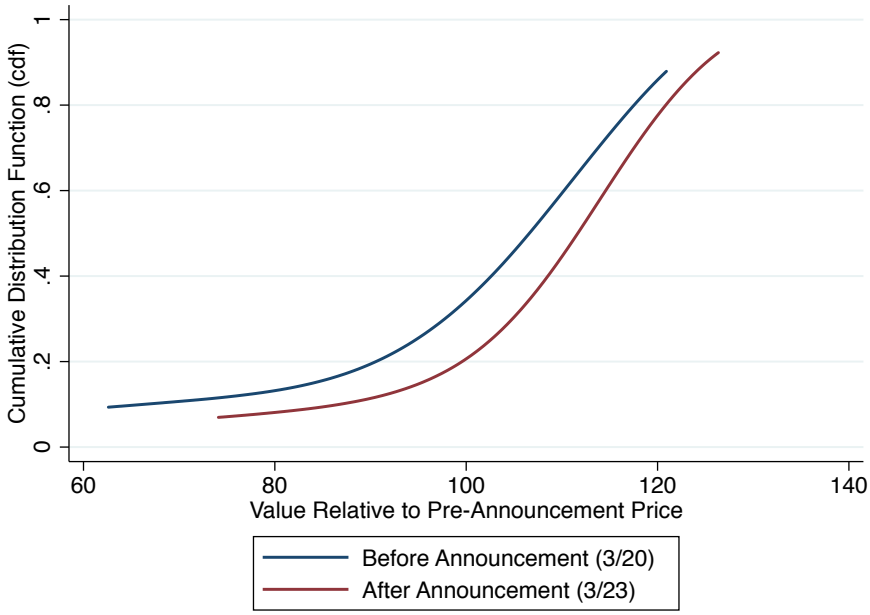


Figure 3: CDF based on option prices.
This figure shows the implied CDF of future returns on corporate bonds extracted from option prices.

fundamental state of the corporate bond market. We interpret $g(p)$ as the total price effects of any Fed intervention in state p. For example, if the Fed purchases more bonds in distress periods when the price p is low, then $g(p)$ will be declining in p. Notice in this case $g(p)$ will shrink the left tail of the distribution of bond prices. In contrast, an unconditional promise to purchase bonds with a constant price impact of purchases would mean $g(p)$ were a constant, which would simply shift the entire distribution in parallel to the right. The new CDF that includes the effects from conditional purchases is based on the post intervention price which we call p'. We then have $p' = p \times (1 + g(p))$. Thus, by using the behavior of the CDF pre and post announcement we can find $g(p)$ and hence assess the impact of the Fed as a function of the non-intervention price p.

Figure 4 shows the implied effect of the policies on prices by plotting the function $g(p)$ expressed as a percentage of the no-policy price p. We can use this plot to gauge

the magnitudes of conditional Fed intervention. First, $g(p)$ is not flat as an unconditional price support would imply, but 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 a put option, lending support to the view of a "policy put." This suggests that bond investors perceived stronger intervention would occur if the situation in the corporate bond market deteriorated. To gauge magnitudes it is worth picking two points on the figure. If, absent any policy intervention, prices would have increased by 20%, the Fed would push the price up by an additional 1.5%. 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.

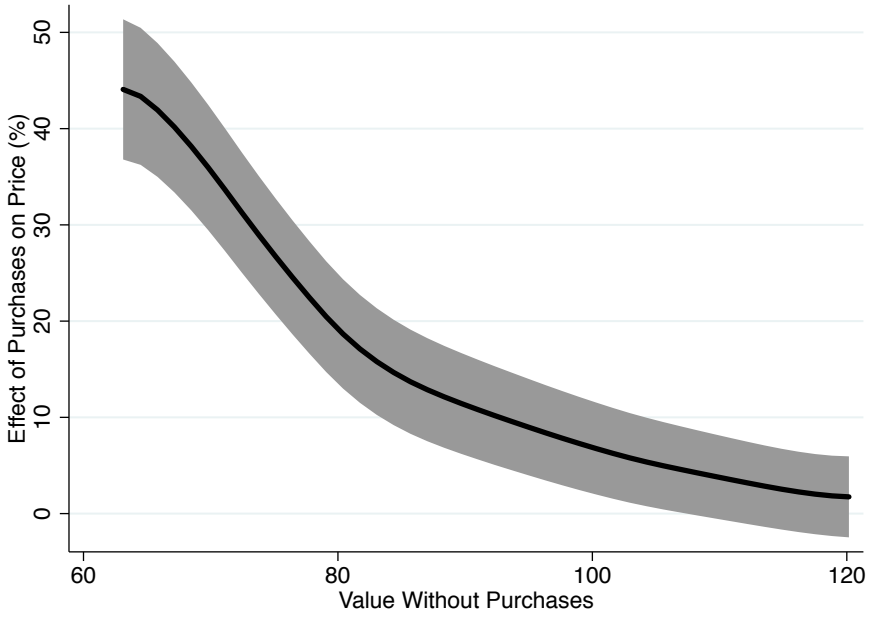


Figure 4: Effect on prices.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

A natural question is whether these movements in the implied volatility curve are typical and so could have happened just 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 4 is built exactly to measure how likely the pattern we document can be due to chance. We construct this by bootstrapping daily pairs of implied volatility curves for options on the invest-grade bond ETF. From these we can construct our price support function $g(p)$ at each day. We compute the price support in units of standard deviation based on at the money implied volatility on the initial day. This deals with the fact that implied volatility is much higher around the Fed announcement than on other days. We then scale these up based on at the money implied volatility on the trading day before the announcement was made. The confidence interval indicates statistically significant price support for the left tail of the distribution, and only insignificant estimates of price support at the upper end. Thus the confidence bounds tell us that this pattern is extremely unlikely to emerge by chance alone.

Revisiting Announcement Effects

We next compute the fraction of the initial announcement return stemming from conditional promises. Recall that $E[g(p)]$ is roughly equal to the one-day announcement return of 7.4% shown earlier. To gauge the effects of this return coming from state-contingent policy in the left tail, suppose in the left tail the Fed simply supported prices by a constant amount equal to $g(100)$, the price support expected absent any change in the no-policy price. Let $\tilde{g}(p) = g(p)$ when $p > 100$ and $\tilde{g}(p) = g(100)$ when $p \leq 100$. We compute $E[\tilde{g}(p)]$ using the implied probabilities of each state which gives the change in price with no policy put, and thus the effect due to the policy put is $E[g(p)] - E[\tilde{g}(p)]$. We find that about 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.

From Price Support to Quantities

Our main methodology delivers conditional price support, but does not specify how this price support is achieved in terms of specific policy actions in different states. One natural view is that price support in different states is achieved through conditional quantities of purchases. To explore this, we make additional assumptions to interpret our numbers in terms of quantities purchased. Assume the price impact of purchases in the bond market is a constant, $\mathcal{M}$, so that the total change in bond market value of purchases are $g(p) = \mathcal{M} \times q(p)$ where $q(p)$ represents the total quantity of bonds or ETFs purchased by the Fed as a function of the no-policy price p. We express the quantity purchased $q(p)$ as a fraction of the total market value of bonds outstanding (because we express price moves in percent rather than total changes in bond market value, the purchases are then also in percent of total value of bonds outstanding). First, we focus on relative statements that don't take a stand on $\mathcal{M}$. As a baseline, $g(p)$ is around 7% when the price without purchases remains at its current level of 100 (e.g. no change in underlying state). But the effect on prices $g(p)$ is 40% if prices were to drop by 30% absent the announcement. With a constant multiplier $\mathcal{M}$, this implies over 5 times as many bonds being purchased in these bad states of the world compared to the baseline case of no change. Comparing to good states where corporate bonds appreciate by 10% or more, we see purchases up to 30 times more than the bad state. Hence, there is a very strong expectation for significantly more intervention if the bond market were to deteriorate.

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

Fed had, directly or indirectly, only purchased about $15 billion of corporate bonds. Since $g(p)$ in this state is about 1.7%, this implies a purchase of $15 billion had a 1.7% effect on bond prices. The total size of the corporate bond market is about $ 7 trillion, thus implying a multiplier $\mathcal{M} = g(p)/purchase(p) = 1.7\%/(15\ billion/7\ trillion) \approx 8$. This implies a large multiplier, but note that using this number then implies smaller required purchases to achieve large price effects in the left tail. Applying this number to the state where the no-policy price fell by 30% would imply purchases of more than $300 billion in these states, exactly the size of the facility.

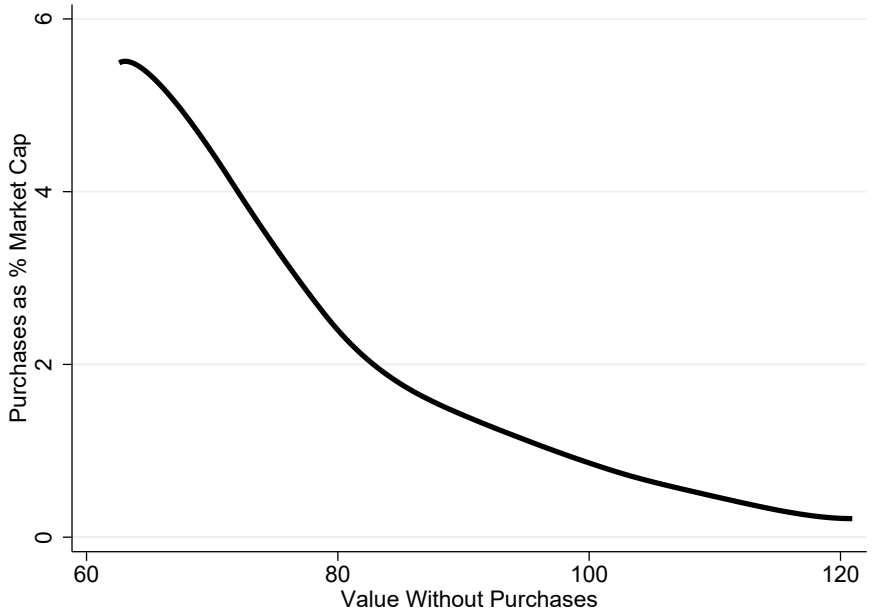


Figure 5: Implied Purchases Based on Multiplier.
This figure plots implied purchases (in billions of USD) as a function of the future distribution of prices before the intervention (normalized to be 100 before the announcement). The multiplier is chosen to match the observed price increase in the state where the Fed actually made purchases of $15 billion.

We note that this value of the multiplier $\mathcal{M}$ is at the upper end of estimates taken elsewhere in the literature. Gabaix and Koijen (2021) estimate $\mathcal{M} \approx 5$ for the stock market. If we use this value instead, it would imply higher purchases in these bad states. That is,

in the state where bond prices fell by 30% purchases would need to be around $500 billion which would be close to 7% of the total corporate bond market and a substantial increase in the size of the Fed programs.

Figure 5 plots implied purchases in billions of dollars using a variety of alternative multipliers. Calibrating to smaller values for $\mathcal{M}$ (a more elastic bond market) would imply even larger numbers. For example, $\mathcal{M} = 1$ would imply purchases close to 35% of the total corporate bond market in bad states of the world. We acknowledge that a specific value for $\mathcal{M}$ is potentially controversial. But by providing the numbers of the total response in various states, readers can easily infer how much conditional purchases they would need to believe for a given multiplier. Overall, we believe the most natural view that comes out from our analysis is for a fairly inelastic bond market combined with an expectation that the Fed would dramatically increase purchases if market conditions were to get worse.

An alternative extreme is to set the quantity of purchases Q to a constant, but allow the multiplier $\mathcal{M}(p)$ to be state-dependent, with higher multipliers in bad states. This view would imply a multiplier at the median state of about 35 and a multiplier in extreme bad states of over 200 (e.g., purchasing 1 of bonds in these states would raise their price by 200). We note these numbers are orders of magnitude larger than found elsewhere in the literature, suggesting the conditional quantity view is a much more plausible explanation.

3.4 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 purchases in 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 known as the CDS bond basis. For example, in the model we outline in Section 6 a high CDS bond basis happens when the arbitrageur has a low risk-bearing capacity and pension funds need to sell the bond for liquidity reasons. 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) by assuming a correlation between the CDX and Treasuries equal to the historical average and using copula functions following Haugh (2016). Our main finding is that the intervention massively reduces the basis risk – the risk that corporate bond prices would get sharply lower then what priced by interest rate and credit risk fundamentals. This pattern of buying more in states of high dislocations is in line with with statements from the Fed. In Senate testimony in June 2020, Jerome Powell stated "markets are functioning pretty well, so our purchases will be at the bottom end of the range that we have written down." The direct implication is that purchases would have been significantly larger had these markets not been functioning well, presumably coinciding with states where bond prices were substantially lower than fundamentals would warrant.

Figure 6 plots the distribution of the basis, defined as the dislocation between the cash market and the synthetic. Specifically, let $basis = p_{cash} - p_{Treas} - p_{CDS}$ where p_{cash} is the price of the corporate bond ETF, p_{Treas} is the price of a duration matched Treasury, and p_{CDS} is the price of the credit risk component of investment grade bonds inferred from CDS prices. The main finding highlighted in the Figure is that the intervention dramatically shifts the tail of the basis distribution, while also shifting the whole distribution to

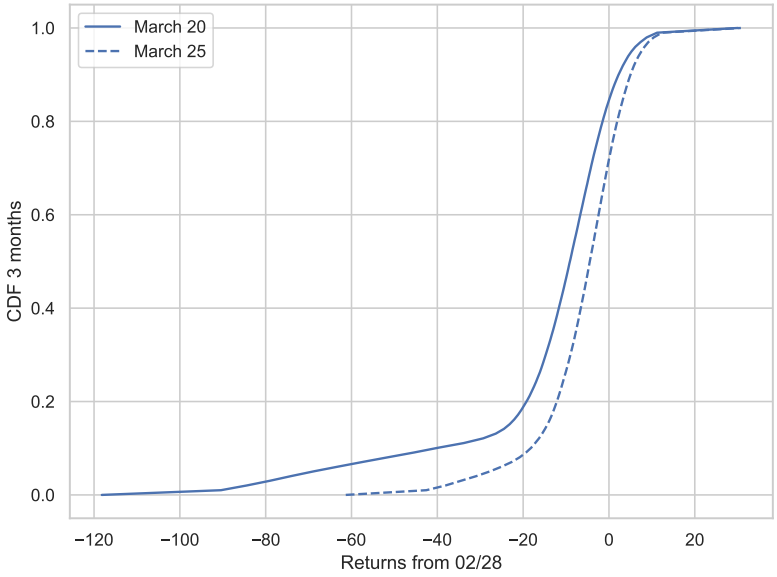


Figure 6: Distribution of the CDS bond basis.
This figure plots the CDF of corporate bond prices minus the price of a synthetic corporate bond constructed from CDS and Treasury options.

the right (shrinking the overall basis). This is consistent with stronger Fed intervention in states where the corporate bond market is highly dislocated so that a large basis opens up. The 10th percentile of the distribution shifts from a -60% to -20% drop in corporate bond prices relative to the synthetic price of the corporate bond. These results point to states where interventions would be larger, but do not necessarily mean the promised intervention is a free lunch. Even if prices in these states were low because of a financial market dislocations, it is still the case that the Fed bears credit risk by purchasing these bonds.

3.5 Robustness

High-yield comparison

Importantly, in Figure 7 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 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.

Segmented SDF Effects

While the lack of response in high-yield bonds and the stock market suggests that broad macroeconomic risk premiums (or broad 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. Figure 8 shows the price support in the presence of endogenous pricing kernel effects following Proposition 3. We note that the asymmetric effects in the left tail remain in this case, even though the risk premium adjustment is very large. Specifically, in the adjustment the implied risk premium for investment grade

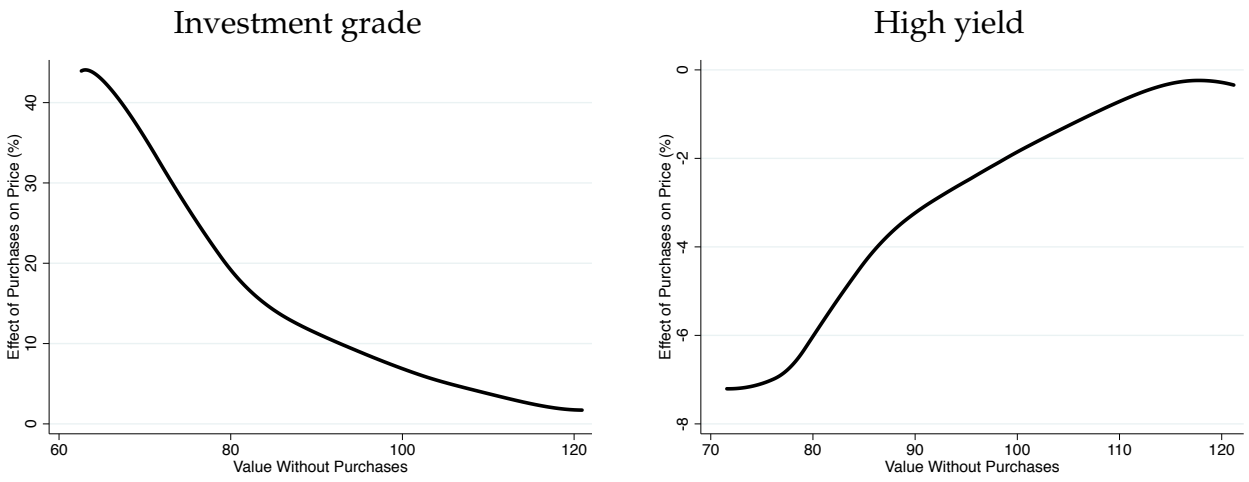


Figure 7: Investment Grade vs High Yield.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

corporate bonds is 16.5%, which is about 20 times the long term average of 0.8% reported in Giesecke et al. (2011).

Liquidity, event windows, and other issues

In the Appendix, we 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. While magnitudes are slightly larger compared to the results in the one-day window, the asymmetric effect is similar.

4. Conditional Price Support Everywhere

Here we provide select additional evidence of key policy announcements to study state-dependent policy. Our list is non-exhaustive but seeks to illustrate both the role of state-

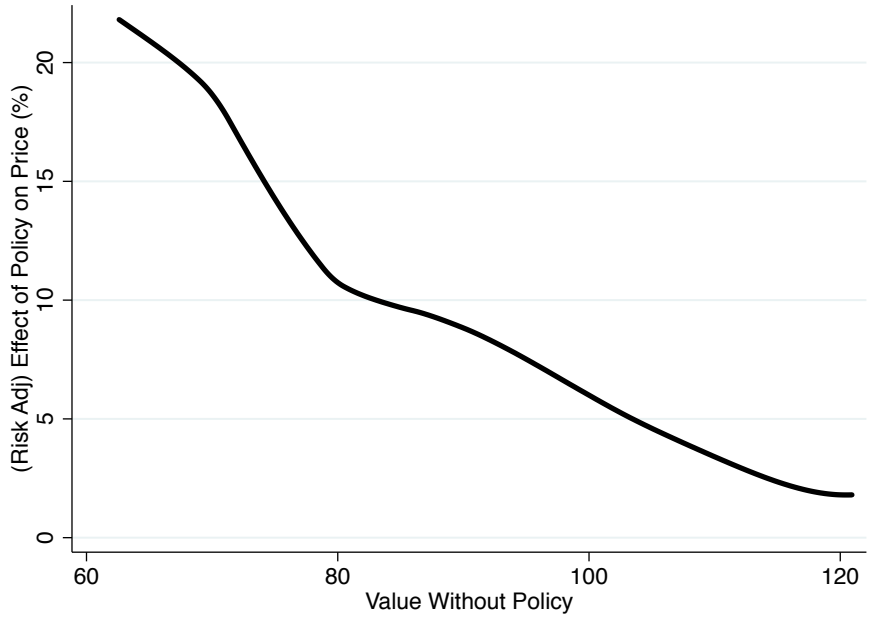


Figure 8: Risk-adjusted Price Support.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

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. We study equity injections in the US financial sector during 2008, an announcement of large asset purchases by the Bank of Japan in 2013, and various dates associated with the implementation and unwinding of quantitative easing (QE) in the United States from 2008-2013.

We summarize our results for these announcements in Table 1. The subsections below discuss each set of announcement dates.

Table 1: Summary of Additional Events

This table applies our methodology to many other announcement events. 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 subsections below.

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

4.1 Evidence from April 9th High-Yield Announcement

The announcement of corporate bond purchases on March 23rd focused on investment grade bonds. However, the Fed made an additional announcement on April 9th, 2020, that expanded the facilities to include high-yield bonds. If this announcement contains implicit promises, we would expect them to show up particularly in high-yield bonds. Figure 9 plots the price support from this announcement following our same methodology applied to options on the high-yield ETF (HYG). We see the same effects of asymmetry: price support is very high for low prices, peaking at over 10%, but is much lower at around 5% for higher levels of prices. This provides strong support for implicit promises boosting the value of high-yield bonds. In contrast, investment-grade is now much flatter, consistent with this announcement not reflecting any additional promises to investment-grade.

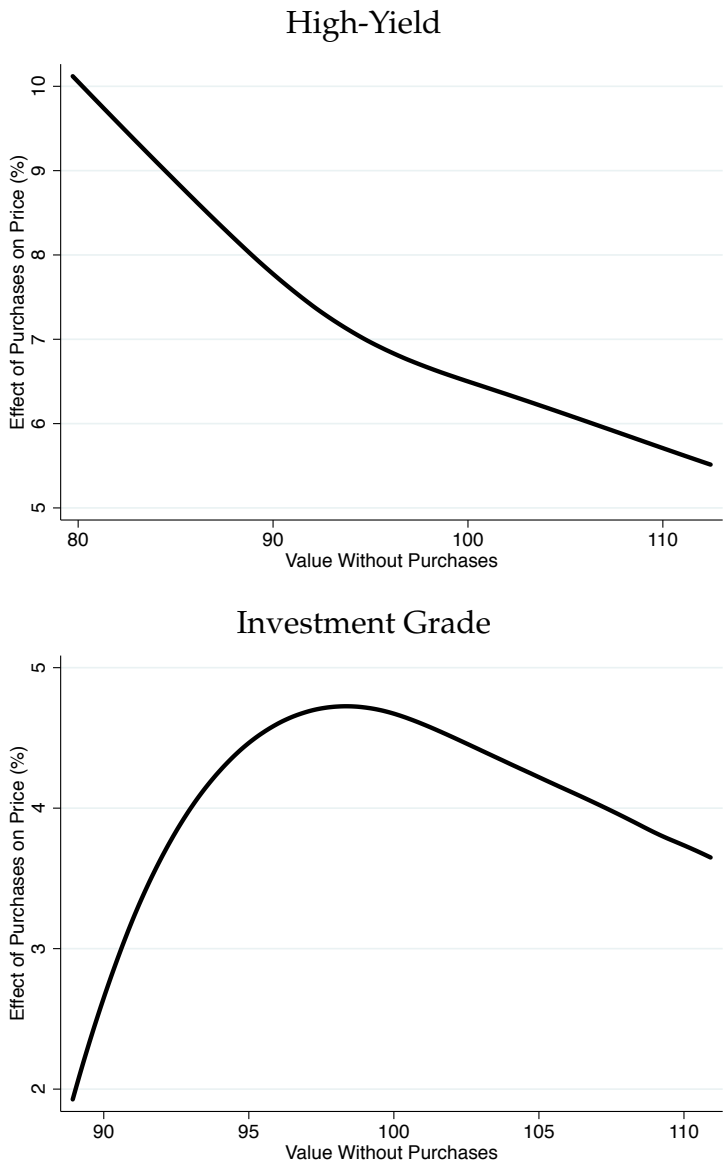


Figure 9: High-Yield Announcement, April 9th, 2020. Effect on prices.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

4.2 2008 Financial Sector Bailout

Figure 10 studies the "Paulson Gift" (Veronesi and Zingales (2010)) on October 10th 2008 which announced large equity injections to the banking sector, as well as guarantees on various forms of bank debt, in an effort to "restore confidence in the financial system." Veronesi and Zingales (2010) find large effects of this intervention on bank equity and bank debt.[15] To analyze this case, we use option prices of the Financial Select Sector SPDR Fund which focuses exposure on the financial sector. Our method reaches similar conclusions but maps out the price support of this policy as a function of the underlying equity value. For example, we see price support of 40% if the equity of the financial sector were to fall by 50% and price support less than 10% if the equity of the financial sector increased by 50%. This provides strong evidence of a type of policy put to the financial sector. This effect was likely a goal of the policy itself – by providing strong conditional promises that the US government would do whatever it takes to keep the financial sector solvent, both debt and equity prices rose substantially. Paulson himself famously stated "if you've got a bazooka, and people know you've got it, you may not have to take it out."

4.3 Bank of Japan Asset Purchases, 2013

Figure 11 studies Japan on April 4th 2013 after a speech given by Haruhiko Kuroda, the head of the central bank, in which he outlined large purchases of government bonds and equities to drive up asset prices and inflation. Charoenwong et al. (2021) systematically study equity purchases by the Bank of Japan and find that they increase equity prices. We use three-month options on the Nikkei index in a three day window around the announcement. This episode points to conditional promises that provide price support of around 10% for adverse states vs about 6% for good states. This was very much intended

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

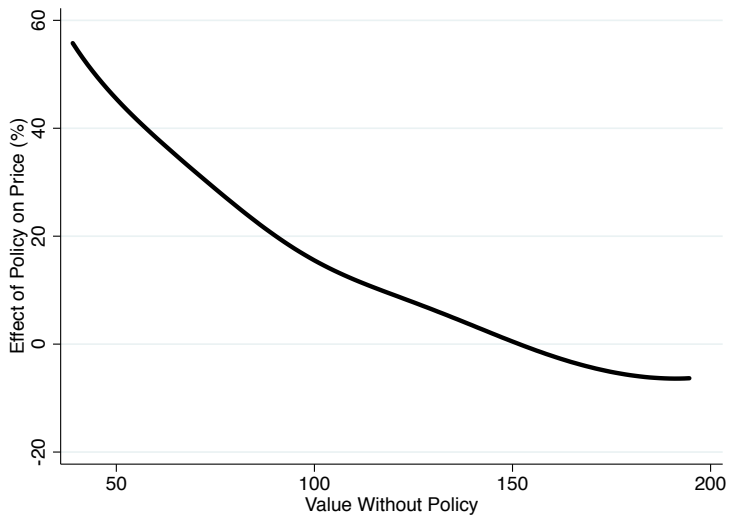


Figure 10: "Paulson's Gift:" Banking Sector Bailout in October 2008. Effect on prices.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

from the announcement by Kuroda as the speech outlined a "whatever it takes" type approach to support prices and achieve growth.

4.4 Quantitative Easing, 2008-2013

Figure 12 looks at US quantitative easing (QE) announcements on four separate dates using three-month options on the 10 year Treasury Note futures contract. Treasuries are likely an imperfect asset to study the conditional impact of QE, because future conditional purchases may depend on many other state variables besides the price of the 10 year Treasury, which our measurement will not capture. We focus our analysis on four QE announcements. The first is the initial announcement of large scale asset purchases (LSAPs) on November 25th, 2008. The second is the FOMC statement on December 16th, 2008. The third is the FOMC statement on March 18th, 2009. The last is the "Taper Tantrum" on June 19th, 2013. The first three policy announcements see an increase in Treasury

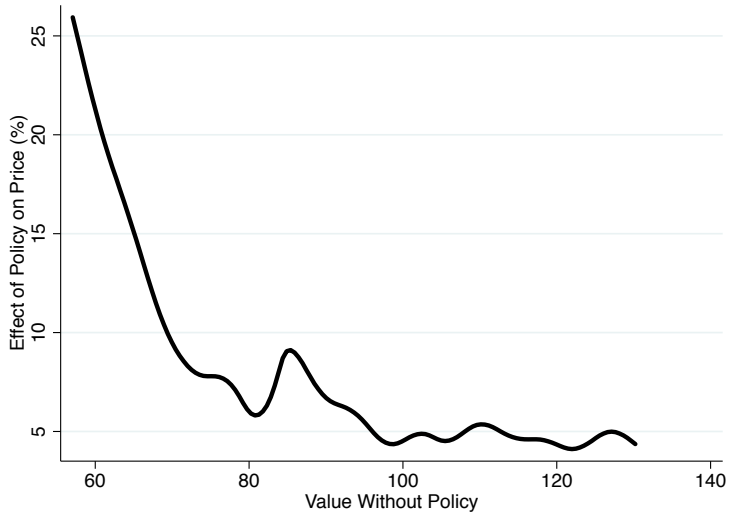


Figure 11: Announcement of Purchases by Bank of Japan in April, 2013. Effect on prices.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

prices (fall in yields) and the price support function is strongly downward sloping. All of these announcements are associated with increased purchases. The magnitudes are fairly similar with about 4-6% price support in cases where prices fall. The downward slope indicates a market belief for more purchases or price support should prices fall by 15% or more, roughly 3% price support at current prices, and roughly 2-3% for a 10% increase in prices. The last announcement (lower right panel) is the "Taper Tantrum" where it was announced purchases would decline. We see an overall decline in prices (sharp increase in yields) with an upward, rather than downward slope. This announcement is thus associated with not only a tapering of purchases but an associated decline in conditional price support from future purchases. Overall, the events of QE are associated with stronger action in the left tail, though the effects are more mild than other events. These results speak to a broader literature that estimates the channels through which QE operates by using event studies (e.g., Vissing-Jorgensen and Krishnamurthy (2011)). Our results suggest

that part of the large price response on announcement comes from stronger interventions in the left tail.

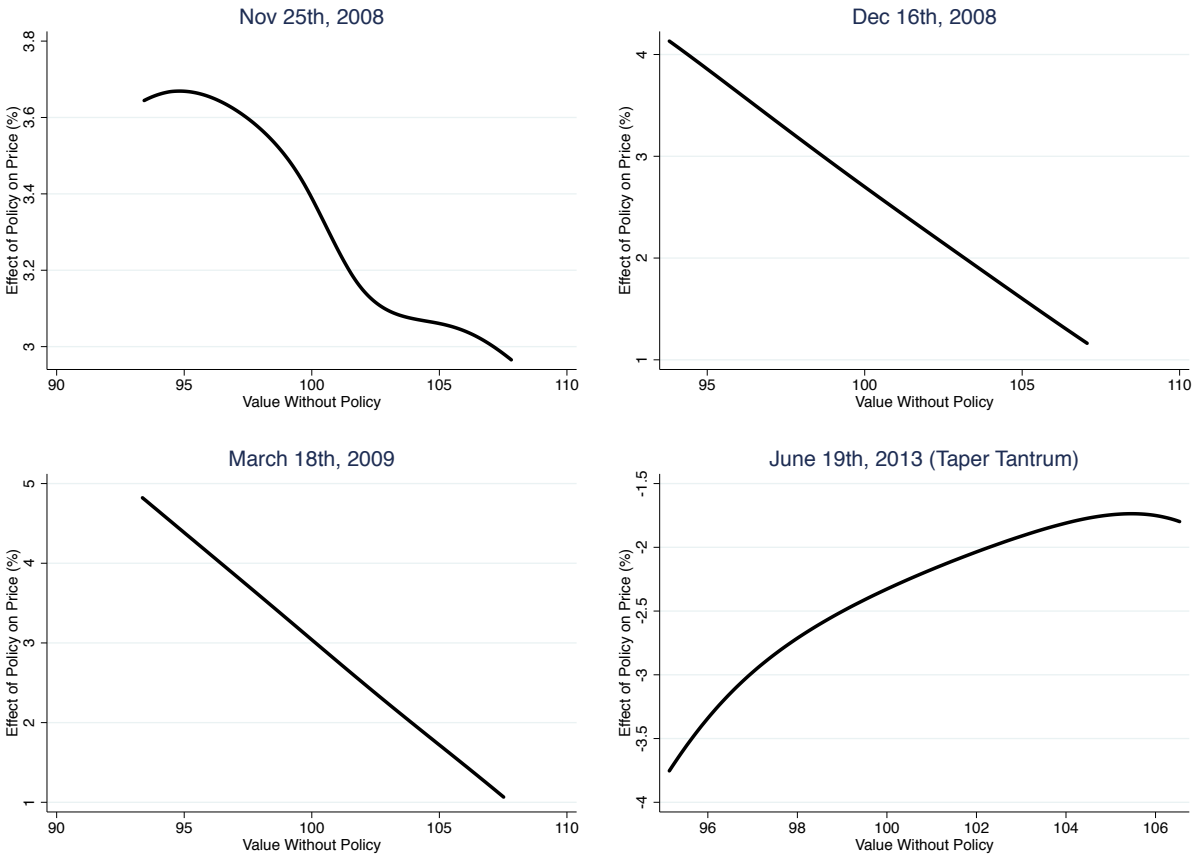


Figure 12: Quantitative Easing. Effect on prices.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

4.5 ECB, 2010-2012

Figure 13 looks at announcements of asset purchases by the European Central Bank. We use five announcement dates found by Krishnamurthy et al. (2018) to have substantial asset price effects. Ideally, for these announcements, we would like to have options on

sovereign bonds for high risk countries in the Eurozone (e.g. Italy). This is most directly where the asset purchase announcements were aimed. However, Krishnamurthy et al. (2018) show a very broad asset price response to the announcements. In fact, they estimate most of the increase in total value from these announcements came from the reaction of stock markets. This is likely due to different channels than what we saw for investment-grade bonds during COVID. One interpretation is for Europe there was a feared "doom loop" that sovereigns would be unable to pay creditors or roll over debt and this would lead to substantial declines in economic activity. This could be coming from higher taxation, strong fiscal adjustments, or losses born by holders of sovereign bonds which included a large portion of the banking sector in Europe. These losses could lead to a substantial credit crunch that would result in a decline in economic activity. The fact that the asset price effects were broad also means we can use options on the Euro Stoxx index to assess conditional purchases. This assumes that any conditional purchases will be correlated with the stock index value and that the stock index will respond to purchases in those states.

Figure 13 shows strong effects of conditional policy from these announcements. Across all the five announcements, the extra support in the left tail explains about 20% of the overall reaction on average. This is likely not surprising as part of the goal of the announcements was to promise to do "whatever it takes." While the exact quantities implied by this promise were vague, the intention to commit to promises in this case was explicit.

5. Implications of the Promises View for Market Dynamics

While these intervention can be powerful in stopping a financial crisis in the making, there is a widespread concern that these policies introduce moral hazard in market participants. For example, firms my leverage up more or investors might be less careful in

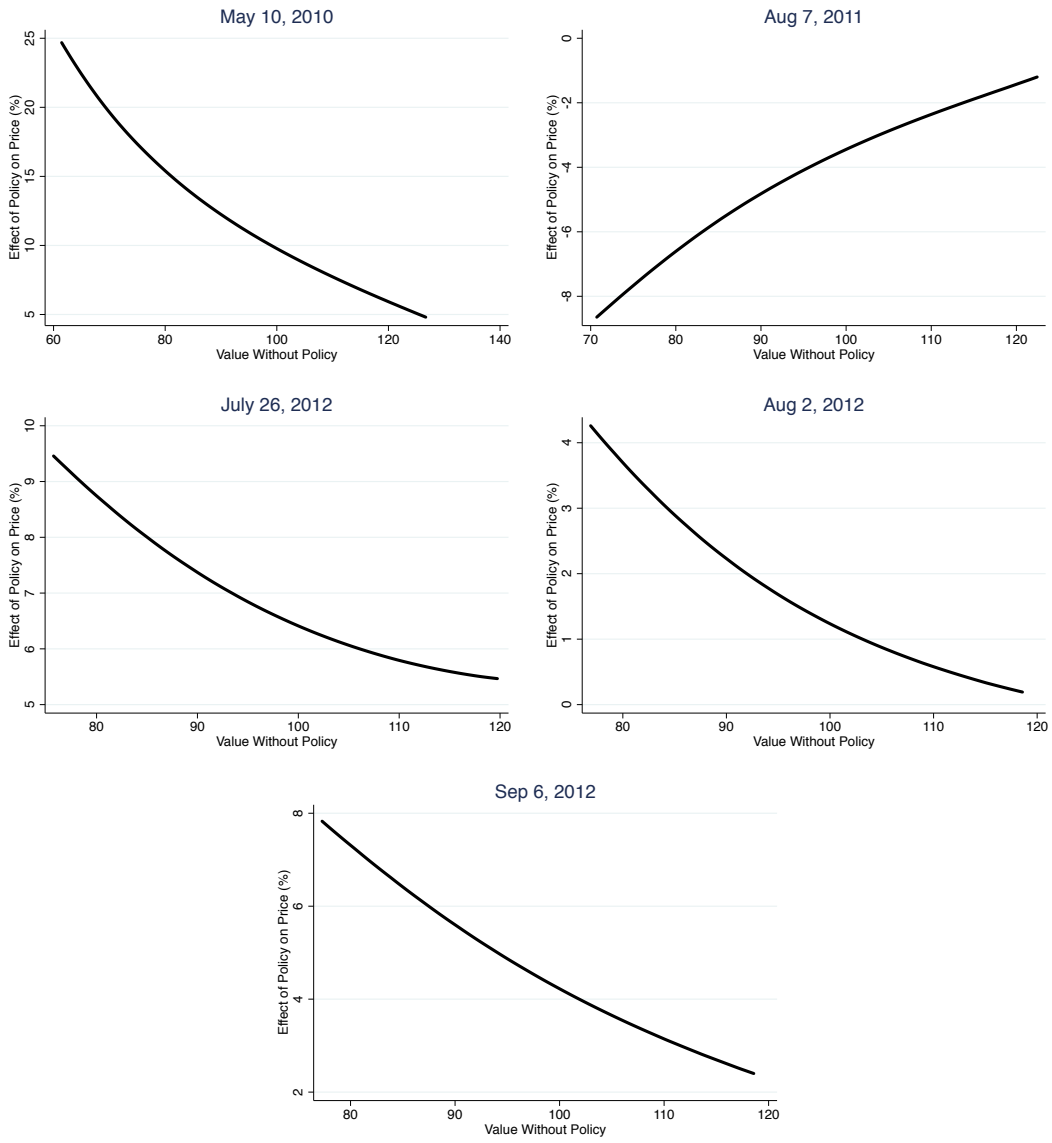


Figure 13: ECB Announcements. Effect on prices.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

the bonds that they hold. A pre-condition for such economic mechanisms to be operating is that market prices price in expectation of future not-yet announced interventions. We look at such effects through two different lenses.

Section 5.1 studies changing asset price dynamics in the presence of promises. Specifically, if future interventions in crash states are expected, then 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 5.2 considers the 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 will be weaker. This fits many facts about asset purchase programs and the "weakening announcement effect."

5.1 Hidden Risks and the 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 changed after the programs are implemented. While Our main analysis documents how the Fed announcement has the immediate effect of reducing this tail risk falls when the Fed initially announces it will intervene, we now assess whether tail risk is less sensitive to economic conditions going forward. The key idea here is that expectations of not yet announced future interventions decouple tail risk in the bond market from fundamentals. These longer-term effects are not easily captured in our earlier framework which focuses on conditional promises over shorter maturities at which we have option data. The challenge here 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 starts by constructing 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 (in the money options) minus the implied volatility for option with delta of 10

(out of the money options) as a measure of tail risk. The idea here is that this difference is insensitive to parallel movements in the implied volatility curve and goes up when the implied volatility of the left tail becomes higher then the implied volatility of the right tail, i.e. when left-tail events become more likely. We then take the same tail risk measure but using S&P500 index options and options on the investment grade CDX index. The idea here is that if tail risk increases in equities we should expect it increasing in bonds as well since they are ultimately claims on the same cash-flows. Table 2 shows that tail risk changed after the announcements and that corporate bond markets appeared less sensitive to tail risk in equity or CDS markets. Specifically, we regress the corporate bond tail risk on tail risk in equities and CDS markets using daily data from 2010 onward (for CDS, we only have data from 2015 onward). We then include a "post" dummy interaction term for the period after April 9th, 2020 when the Fed had already announced the expansion of its corporate bond facilities. Notably, in the period prior to this, corporate bond tail risk and equity market tail risk co-move strongly so that tail risk in corporate bonds was highly sensitive to tail risk in equity markets.

After the interventions occurred, this sensitivity dropped dramatically. The sum of the two coefficients represents the total sensitivity in the post period and actually goes slightly negative. But the main effect we focus on is that the post interaction term is strongly negative and statistically significant so that corporate bond tail risk becomes much less sensitive to broader tail risk in the economy. This is consistent with the view that financial markets are pricing in not-yet-announced future interventions to support bond prices if financial and/or economic conditions were to deteriorate. 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 here because it is a more more targeted measure of the cash-flow risk that corporate bonds are exposed to. For example, the fact that the results is similar tell us that change in financial leverage does not explain this decoupling as such change would be reflected in the CDS index as well. Finally, the results

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

in the pre-period are not driven by extreme behavior during the acute phase of COVID on financial markets where all tail risk measures spike. To show this, we add a COVID dummy, equal to 1 for the period of February 1st, 2020 to April 9th, 2020. Including an interaction with this dummy doesn't change our conclusions, and in this case the non-interacted coefficient measures the sensitivity of corporate bond tail risk to other tail risk excluding the COVID episode (e.g., from 2010-early 2020).

Figure Figure 14 instead uses monthly data on option-based pseudo credit spreads from Culp et al. (2018). These are credit spreads directly implied by equity options and Treasuries. We use the two year maturity investment grade pseudo credit spread from Culp et al. (2018) (see The Credit Risk Lab). We then compare this to the Bank of America

investment grade option adjusted credit spread index for maturities between one and three years taken from Fred. We plot the difference in actual credit spreads vs pseudo spreads in Figure 14. From 2010-2020 the two spreads track each other quite well. During early 2020 when the COVID-19 crisis hit, actual spreads for investment grade spiked well beyond those implied by equity market options. This is consistent with investment grade bond prices becoming abnormally depressed in this episode. However, following the Fed's intervention and the recovery, investment-grade spreads became quite low, and in fact reached their lowest point at any time over to 2010-2020 window. In contrast, equity markets still featured substantial volatility, implying higher than usual default risk on pseudo bonds. This keeps pseudo spreads elevated even after the Fed intervention. This large gap is consistent with the market pricing in future interventions in the case of a crash.

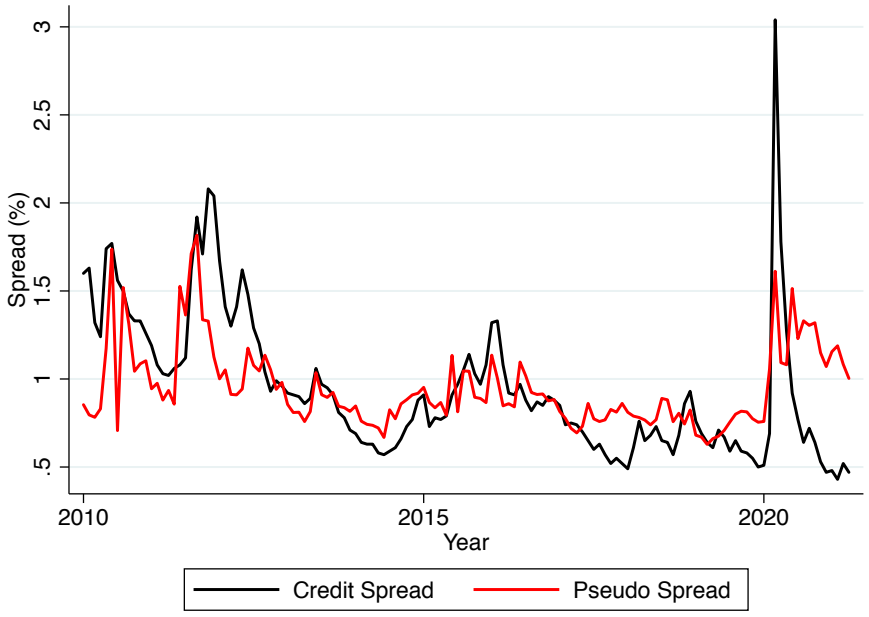


Figure 14: Spreads vs pseudo spreads.
This figure plots actual credit spreads vs pseudo spreads from Culp et al. (2018).

5.2 Weakening Patterns in Announcement Effects

An alternative way of looking for long-term pricing effects is to study how financial market respond to future announcements. If it is true that investors already expect these announcements as we argue above, the estimate announcement effects of these announcements 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 have a "weakening effect." That is, in both the US and Europe, initial announcements of asset purchases 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) that the effects of QE vary significantly in the literature in part due to which announcements are used. Figure 15 shows the impact of these announcements on yields (blue bars) over time for the US, UK, and Europe. The figure is a non-exhaustive list of all announcements. 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 is a more consistent way to compare announcements (both over time and across countries) as the interventions vary in the amount of bonds purchased. Appendix B.3 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).

At the same time, the literature has found the effects of the original announcements to be very large relative to other estimates. 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."[16] They estimate QE effects from the initial

[16]See also Krishnamurthy and Vissing-Jorgensen (2012).

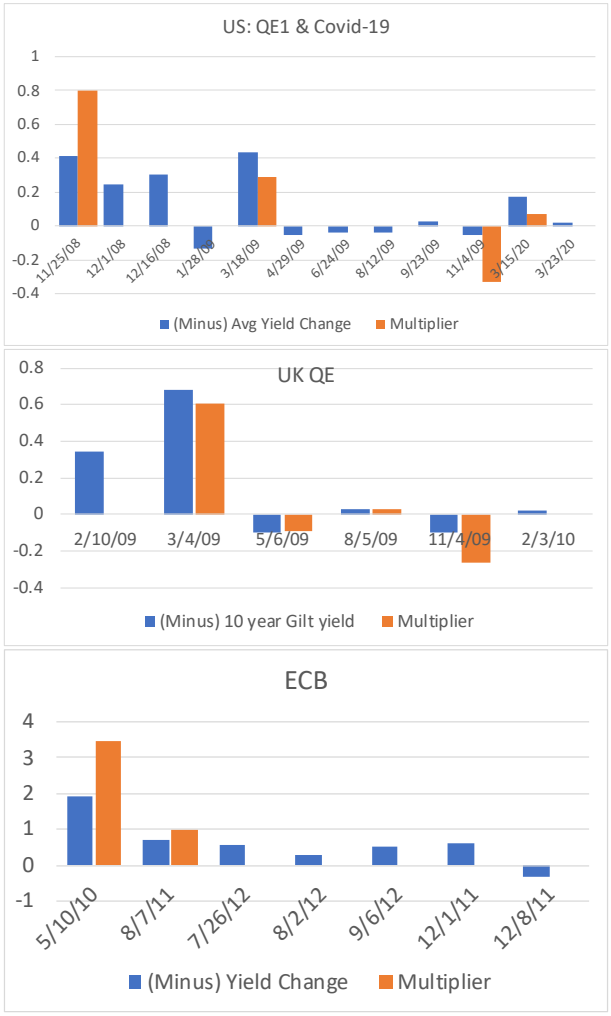


Figure 15: Weakening Announcement Effects in the Literature

This figure announcement effects to quantitative easing announcements made by the US Federal Reserve, Bank of England, and 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)

QE1 program of over 100bps for the same change in supply.

The promises lens 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 things were to get worse, should move prices "too much"

relative to the headline purchase amounts. This is because earlier announcements contain a policy put option that has substantial value, as the evidence in Section 4 shows. Similarly, later announcements will only matter to the extent they are a surprise relative to the policy plan investors perceive. Because investors 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. Thus, the weakening announcement effect does not imply that the effectiveness of purchases themselves has weakened.

here

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 "weakning 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 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.[17] Vissing-Jorgensen (2021) states "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

[17]These 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.

were under extreme stress. Our 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, 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. The key difference is that the Fed had never before purchased corporate bonds and thus the announcement was a surprise.

Finally, this experience also contrasts sharply 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.

6. 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 our 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 and these purchases are made at date 1.

The specialized arbitrageur has log utility over 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 the same at both dates but this is for simplicity and 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.[18] 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 indeed be a better economic 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

$$\alpha_1 = \frac{E_1[ln(X/P_1)] - log(R_f)}{Var_1(ln(X/P_1))} = \frac{\mu - p_1 - r_f}{\sigma^2} \tag{12}$$

where α_1 is the arbitrageur's portfolio share in the risky asset, X/P_1 denotes the gross

[18] It 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.

return on the asset from date 1 to date 2, $p_1 = ln(P_1)$, and $E_1[.]$ denotes the conditional expectation taken at time 1.

The central bank purchases q of the asset at date 1, where we denote q as a fraction of the arbitrageur's date 1 wealth. We allow this amount q to potentially 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. In a setting which fundamental

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

$$p_1 \quad = \quad \sigma^2(1 - \tilde{b} + q) + \mu - r_f \tag{13}$$

This equation gives a multiplier σ^2 for the effect of asset purchases q on the (log) price p_1 (where q here is as a share of the arbitrageur wealth). For example, higher purchases q remove the asset from the arbitrageur's balance sheet and raise prices by $q\sigma^2$ to compensate for bearing the additional risk, 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 market capitalization. That is, 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 total wealth invested in the risky asset, the multiplier will be large. This is because the purchases are absorbed by a relatively small amount of active agents, so they have a very large effect on active agents' positions. We also note that the multiplier is constant and does not depend on the realization of the state b at date 1.

One can use the date 1 pricing equation to show that this framework naturally explains the "weakening" effect of follow on purchases. 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,

$$p_1 - E[p_1|b] \quad = \quad \sigma^2(q - E[q|b]) \tag{14}$$

We see that only purchases that deviate from what was expected given the announcement in date 0 will have any effect, and when the policy maker simply fulfill their promises the effect is exactly zero. Of course, here the 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.[19]

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

$$m_2 \quad = \quad \left(\alpha_1 R_2 + (1-\alpha_1) r_f\right)^{-1} = \left((1+\tilde{b}-q)R_2 + (q-\tilde{b})r_f\right)^{-1} \tag{15}$$

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 and the discounting between 0 and 1 remains unchanged. Since the time 1 price

[19] While 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. All that our framework requires is that the first announcement reveals the policy rule the policy maker will adopt in the follow on announcement. 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.

has risen, the time 0 price will also rise. This make it transparent how asset purchases can impact prices by moving future risk-premium even if they do not impact the pricing kernel between the announcement date and the asset purchase date as we assume in our baseline analysis.

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

$$P_0 = \frac{1}{R_f}\frac{1}{E_0\left[\frac{1}{P_1}\right]} \tag{16}$$

It follows from the expression above that a date 0 announcement by the central bank to purchase a constant share of the asset at date 1 does not change the pricing kernel between dates 0 and 1. Because the intervention pushes up prices proportionally both at date 0 and date 1 it does not change asset risk between these dates. This means that in such a model our framework recovers the correct price support function.[20]

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. Thus, risk removal only impact the pricing kernel from date 1 forward. The second is that deterministic purchases do not change

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

the risk of the asset because it moves prices uniformly up. Because the arbitrageur prices the risky asset, the pricing kernel is, roughly speaking, $(arb\ asset\ exposure) \times (asset\ risk)$. That is, stochastic purchases might 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

$$p_0 = \mu + \sigma^2 - ln\left(E_0\left[exp\left(\sigma^2(\tilde{b} - q)\right)\right]\right) \tag{17}$$

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.

7. Conclusion

We provide a framework and methodology to evaluate the state-contingent impact of policy, which we apply to several policy announcements. We find a large role for state-contingent policy that indicates more intervention if conditions worsen. In our main empirical setting, the announcement of corporate bond purchases during the COVID-19 cri-

sis, 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 option to significantly expand the size of the intervention in bad states explains a large share of the market response to the announcement. Our methodology focuses on shorter horizons as this is the maturity that relevant options expire. We also find longer-run effects of state-contingent interventions in corporate bond markets after announcements of corporate bond purchases, consistent with expectations of future intervention if a future crash should occur. We extend our analysis to several other policy announcements as well and find support for conditional intervention. The state-contingent policy view helps explain many empirical phenomena including large multipliers from initial policy announcements and weakening or disappearing announcement effects for later announcements of the same policy.

References

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.

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.

Bernanke, Ben S, 2020, The new tools of monetary policy, *American Economic Review* 110, 943–83.

Boyarchenko, Nina, Anna Kovner, and Or Shachar, 2020, It's what you say and what you buy: A holistic evaluation of the corporate credit facilities, Working paper, Federal Reserve Bank of New York.

Breeden, Douglas T, and Robert H Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of business* 621–651.

Charoenwong, Ben, Randall Morck, and Yupana Wiwattanakantang, 2021, Bank of japan equity purchases: The (non-) effects of extreme quantitative easing, *Review of Finance* 25, 713–743.

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

Collin-Dufresne, Pierre, Benjamin Junge, and Anders B Trolle, 2021, How integrated are credit and equity markets? evidence from index options, *Evidence From Index Options (February 16, 2021). Swiss Finance Institute Research Paper* .

Culp, Christopher L, Yoshio Nozawa, and Pietro Veronesi, 2018, Option-based credit spreads, *American Economic Review* 108, 454–88.

D'Amico, Stefania, Vamsidhar Kurakula, and Stephen Lee, 2020, Impacts of the fed corporate credit facilities through the lenses of etfs and cdx, *Available at SSRN 3604744* .

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.

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.

Figlewski, Stephen, 2010, Estimating the implied risk-neutral density for the us market portfolio, *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle, Edited by Tim Bollerslev, Jeffrey Russell, and Mark Watson* .

Gabaix, Xavier, and Ralph SJ Koijen, 2021, In search of the origins of financial fluctuations: The inelastic markets hypothesis, Working paper, National Bureau of Economic Research.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack, 2018, The financial market effects of the federal reserve's large-scale asset purchases, *24th Issue (Mar 2011) of the International Journal of Central Banking* .

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.

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.

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.

Malz, Allan M, 2014, A simple and reliable way to compute option-based risk-neutral distributions, *FRB of New York Staff Report* .

Meaning, Jack, and Feng Zhu, 2011, The impact of recent central bank asset purchase programmes, *BIS Quarterly Review, December* .

Moreira, Alan, and Alexi Savov, 2017, The macroeconomics of shadow banking, *Journal of Finance* .

Nakamura, Emi, and Jón Steinsson, 2018, High-frequency identification of monetary non-neutrality: the information effect, *The Quarterly Journal of Economics* 133, 1283–1330.

O'Hara, Maureen, and Xing Alex Zhou, 2020, Anatomy of a liquidity crisis: Corporate bonds in the covid-19 crisis, *Available at SSRN 3615155* .

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.

A. 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.

A.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, take a guess-and-verify approach. Denoting $g(s) = \bar{g}$, 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:

$$M'(s) = \Theta\left(s, \frac{p_1'(s)}{p_0'}\right) \tag{18}$$

$$= \Theta\left(s, \frac{p_1(s)(1+g)}{p_0(1+g)}\right) \tag{19}$$

$$= \Theta\left(s, \frac{p_1(s)}{p_0}\right) \tag{20}$$

$$= M(s). \tag{21}$$

This confirms our guess and concludes the proof.

A.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 $\theta(s)$, which we call the numeraire-equivalent distribution. We define this distribution in equation (22), which we repeat here:

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

Let us show how to measure this distribution using options.
Consider a contract that pays off $C_K(s)$ defined by

$$C_K(s) = \begin{cases} p_1(s)/p_0 & \text{if } p_1(s) \leq K \\ 0 & \text{if } p_1(s) \leq K. \end{cases} \tag{22}$$

We define $C'_K(s)$ similarly after the announcement:

B. Empirical Appendix

B.1 Event Study for Corporate Bond Purchases

The announcement of the SMCCF and PMCCF had a significant and immediate impact on corporate bond prices. **??** 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.[21] 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.

B.2 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 16 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

[21] We 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.

Table 3: 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***	10.27**	7.37***	6.63***
	(3.78)	(1.25)	(0.01)	(0.07)
$r_t^{HighYield}$		0.54***		0.55***
		(0.04)		(0.04)
r_t^{SP500}		0.03		0.03
		(0.02)		(0.02)
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

(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 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).

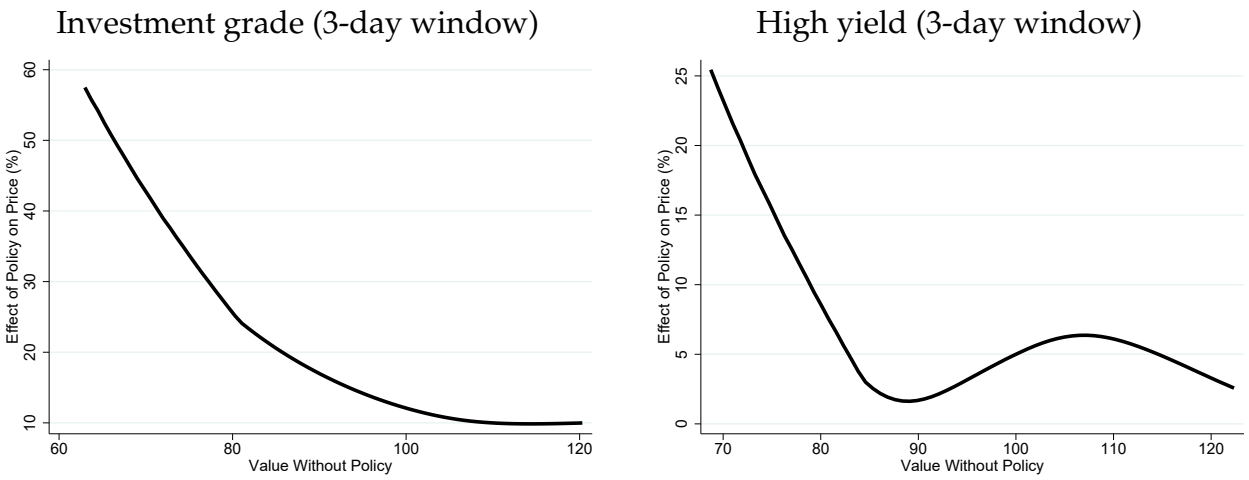


Figure 16: Investment Grade vs High Yield.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

B.3 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 3.3 the overall trading volume in 2020 (in terms of contract notionals) 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 3.3 we show that the negatively slopped 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 **??** 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.

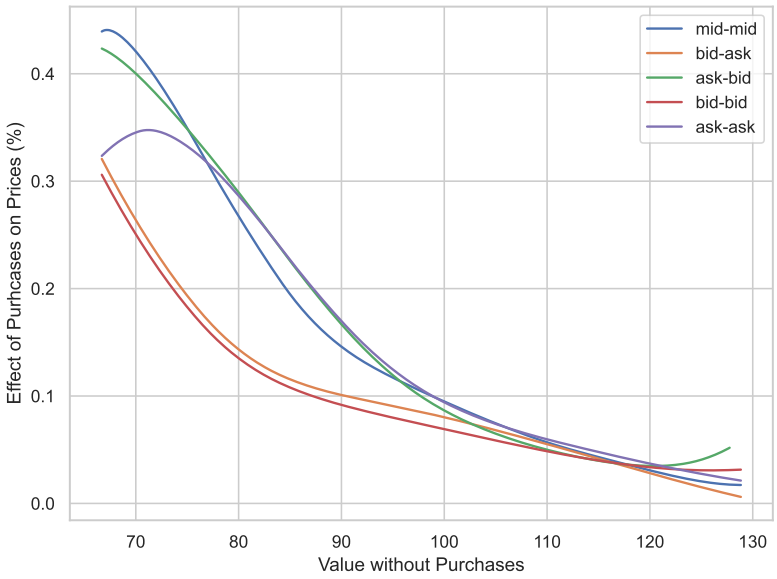


Figure 17: Effect on prices using Bid and Ask quotes.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement. Here we construct the price support function using both the bid (quote at which investors 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)

B.4 Details on Announcement Effects

This subsection provides more detail on data sources used in Section 5.2.

B.4.1 United Kingdom

Table 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 featured no changes in yields at all despite similar magnitudes of purchases. We can convert the yield changes and quantities into a

Table 4: 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]

price elasticity that gives the price impact, which we provide in the column "multiplier." The multiplier for the first announcement of 0.6 says that by purchasing 10% of the supply of Gilts the price of Gilts would fall by 6% (for a security with a duration of 10, this means a decline in yields of 60bps). When the announcement comes with a quantity range, we provide the range for the multiplier as well. The main finding is that the multiplier is much higher in the early announcement and is then quickly goes to zero.

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

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

B.4.2 United States

Table 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.[22] 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 4 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 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

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

Table 5: US Announcement Effects

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

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

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

This experience also contrasts with the Bank of Canada (Arora et al., 2021) during the same time period. The Bank of Canada announced purchases of government bonds on March 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, con-

Table 6: 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	

trasting 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.

B.4.3 Eurozone

Table 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.