Supplemental Appendix for "Whatever It Takes? The Impact of Conditional Policy Promises"

Valentin Haddad Alan Moreira Tyler Muir*

September 16, 2024

*Haddad: UCLA Anderson School of Management & NBER (valentin.haddad@anderson.ucla.edu); Moreira: University of Rochester Simon Graduate School of Business & NBER (alan.moreira@simon.rochester.edu); Muir: UCLA Anderson School of Management & NBER, (tyler.muir@anderson.ucla.edu).

Contents

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

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

B. Economic Model

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

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

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

Inelastic investors have W_I dollars of the risky asset at date 0 and are price inelastic. They can be thought of as insurance companies, pension funds, or other institutions that hold a large fraction of the bond market but do not trade frequently or are inattentive. In contrast, the arbitrageur should be thought of as a dealer bank, hedge fund, or other active trader. Inelastic investors have a stochastic demand shock at date 1 that leads them to sell $\tilde{B}$ dollars of the asset. It is convenient to define $\tilde{b} = \tilde{B}/W_1$ as the dollar sales made by the inelastic investors as a fraction of the arbitrageurs' date 1 wealth.[1] This fire sale shock is the only source of date 1 uncertainty. The fire sale shock

[1] This shock can also be interpreted as affecting total supply instead of sales by the inelastic investors. For example,

depresses prices but is independent of fundamentals of the asset payoff. While the COVID-19 episode fits well with this fire-sale interpretation, one could alternatively consider fundamental cash flow shock at date 1, that is a shock to date 1 cash-flow expectations. Such an approach might be more in line with other episodes with asset purchase announcements such as quantitative easing.

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

$$\alpha_1 = \frac{E_1[ln(X/P_1)] - log(R_f)}{Var_1(ln(X/P_1))} = \frac{\mu - p_1 - r_f}{\sigma^2} \tag{A1}$$

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

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

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

$$p_1 = \sigma^2(1 - \tilde{b} + q) + \mu - r_f \tag{A2}$$

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

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

$$p_1 - E[p_1|b] = \sigma^2(q - E[q|b]) \tag{A3}$$

Only purchases that deviate from what was expected given the announcement in date 0 have any effect, and when the policy maker simply fulfills their promises the effect is exactly zero. This zero effect does not mean that the date 1 intervention is ineffective, but simply that it was already reflected in the date-0 price response.[2]

it could be the consequence of firms' large debt issuance needs during COVID-19.

[2]While we consider here one announcement date paired with one purchase date, it is straightforward to extend the

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

$$\text{(A4)} \qquad m_2 = \left(\alpha_1 R_2 + (1-\alpha_1) r_f\right)^{-1} = \left((1+\tilde{b}-q)R_2 + (q-\tilde{b}) r_f\right)^{-1}$$

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

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

$$\text{(A5)} \qquad P_0 = \frac{1}{R_f}\frac{1}{E_0\left[\frac{1}{P_1}\right]}$$

It follows from the expression above that a date 0 announcement by the central bank to purchase a constant share of the asset at date 1 does not change the pricing kernel between dates 0 and 1. Because the intervention pushes up prices proportionally both at date 0 and date 1 it does not change asset risk between these dates. Thus, our framework recovers the correct price support function.[3]

Deterministic purchases do not affect the risk premium at date 0 for two reasons. The first is that purchases, whether deterministic or state-dependent, do not remove risk from the arbitrageur balance sheet until date 1. Removal of risk only impacts the pricing kernel from date 1 forward. The second is that deterministic purchases do not change the risk of the asset because it moves

model to have multiple announcement dates paired and multiple implementation dates dates. The key intuition is that the first announcement reveals the policy rule that the policy maker will adopt in the follow on announcements. In this way the response to the first announcement is driven not only by the immediate follow-on purchases but also by the follow-on announcements which themselves will appear ineffective.

[3]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 $g(p_1) = g_a$.

prices uniformly up. Stochastic purchases can impact the pricing kernel between date 0 and date 1 through the effect they have on the risk of the asset between dates 0 and 1. Plugging in the date-1 price of the asset gives the date-0 price as

$$p_0 = \mu + \sigma^2 - ln\left(E_0\left[exp\left(\sigma^2(\tilde{b} - q)\right)\right]\right) \tag{A6}$$

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

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

C. Relaxing the Stochastic Discount Factor Invariance

We first review some standard model results and facts about how changes in short-term stochastic discount factor affect asset prices. We then derive the results of Section I.C.4 which generalize our approach to recover the conditional price support.

1 The typical effect of changes in short-term stochastic discount factor

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

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

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

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

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

$$f^{\mathbb{Q}}(p_1) = \frac{1}{p_1}\frac{1}{\sqrt{2\pi\sigma^2}}e^{-\gamma(\log(p_1)-\mu)-\frac{1}{2}\gamma^2\sigma^2}e^{-(\log(p_1)-\mu)^2/2\sigma^2} \tag{A7}$$

$$= \frac{1}{p_1}\frac{1}{\sqrt{2\pi\sigma^2}}\exp\left[-\frac{1}{2\sigma^2}\left(\log(p_1)-\left(\mu-\gamma\sigma^2\right)\right)^2\right] \tag{A8}$$

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

$$g(p_1) = e^{(\gamma-\gamma')\sigma^2} - 1 \tag{A9}$$

We see that despite the change in discount factor disproportionately affecting the left tail of the distribution, the price support function is constant. Not surprisingly, in the real world, returns are not exactly log-normal, and the SDF need not follow this functional form, so we now discuss the unconditional data.[4]

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

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

Of course, our specific settings reflect somewhat unusual market conditions, which leads us to construct additional measures to assess the impact of variation in short-term risk premia. See Section I.C.4 for this discussion.

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

2 Testing a constant price support

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

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

$$m'(s) = \Theta\left(s, \frac{p_1'(s)}{p_0'}\right) \tag{A10}$$

$$= \Theta\left(s, \frac{p_1(s)(1+g)}{p_0(1+g)}\right) \tag{A11}$$

$$= \Theta\left(s, \frac{p_1(s)}{p_0}\right) \tag{A12}$$

$$= m(s). \tag{A13}$$

This confirms our guess and concludes the proof.

3 Adjusting the estimates for a response of the stochastic discount factor

We prove Proposition 3.

The first part of the result is to notice that, as long as we find a distribution for $p_1(s)$ which is unaffected by the announcement, we can use the same idea as in our baseline of matching the quantiles of $p_1(s)$ and $p_1'(s)$. In our baseline setting, the risk-neutral distribution of $p_1(s)$ was invariant. This is not the case anymore under Assumption 5 because the pricing kernel is affected by the intervention. In contrast, the physical distribution remains unchanged, but we cannot recover it from option contracts. Instead, we focus on an intermediate distribution, only affected by the exogenous part of the pricing kernel $\theta(s)$, the forward-neutral distribution. We define this distribution in equation (11), which we repeat here:

$$dF^{\mathcal{N}}(p_1) = E[\theta(s)|p_1]dF(p_1). \tag{A14}$$

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

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

$$C_K(s) = \begin{cases} p_1(s)/p_0 & \text{if } p_1(s) \leq K \\ 0 & \text{if } p_1(s) \leq K. \end{cases} \tag{A15}$$

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

$$E[m(s)C_K(s)] = E[E[m(s)C_K(s)|p_1]] \tag{A16}$$

$$= \int_{-\infty}^{K} E[\theta(s)|p_1]\frac{p_0}{p_1}\frac{p_1}{p_0}dF(p_1) \tag{A17}$$

$$= F^{\mathcal{N}}(K) \tag{A18}$$

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

$$E[m'(s)C'_{K'}(s)] = E[E[m'(s)C'_{K'}(s)|p_1]] \tag{A19}$$

$$= \int_{-\infty}^{K} E^{\mathbb{P}}[\theta(s)|p_1]\frac{p'_0}{p'_1(p_1)}\frac{p'_1(p_1)}{p'_0}dF^{\mathbb{P}}(p_1) \tag{A20}$$

$$= F^{\mathcal{N}}(K) \tag{A21}$$

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

$$C_K(s) = \frac{K}{p_0}1_{\{p_1 \leq K\}} - \frac{1}{p_0}\max\left(K - p_1, 0\right). \tag{A22}$$

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

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

$$F^{\mathcal{N}}(K) = E\left[m(s)C_K(s)\right] \tag{A23}$$

$$= \frac{K}{p_0}\frac{\partial Put}{\partial K}(K) - \frac{1}{p_0}Put(K), \tag{A24}$$

where $Put(K)$ is the price of put options as a function of the strike K. Figure A1 illustrates the comparison of this replicating portfolio with that of the risk-neutral measure.

D. The Order-Preserving Condition

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

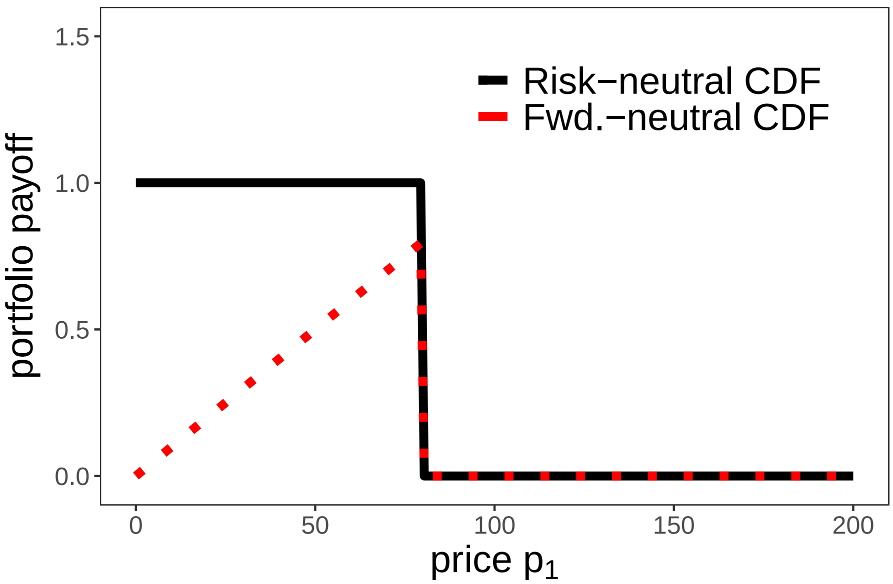

Figure A1: Estimating the risk-neutral and forward-neutral measures.
This figure reports the payoffs for contracts replicating the risk-neutral measure (solid black line) and forward-neutral measure (dotted red line) for the value $K = 80$.

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

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

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

A numerical example. We consider a variation of the whatever-it-takes example, where the true price support function has discrete jumps up at a series of cut-offs. Figure A2 reports the result of this exercise. The black line represents our assumption about the true price support function. With an assumption about the distribution of the underlying state (we assume a simple log-normal), we can compute an estimated price support function following our approach. This corresponds to the red line on the figure. The estimated price support function smoothes out the discrete increments, but otherwise exhibits a close quantitative behavior.

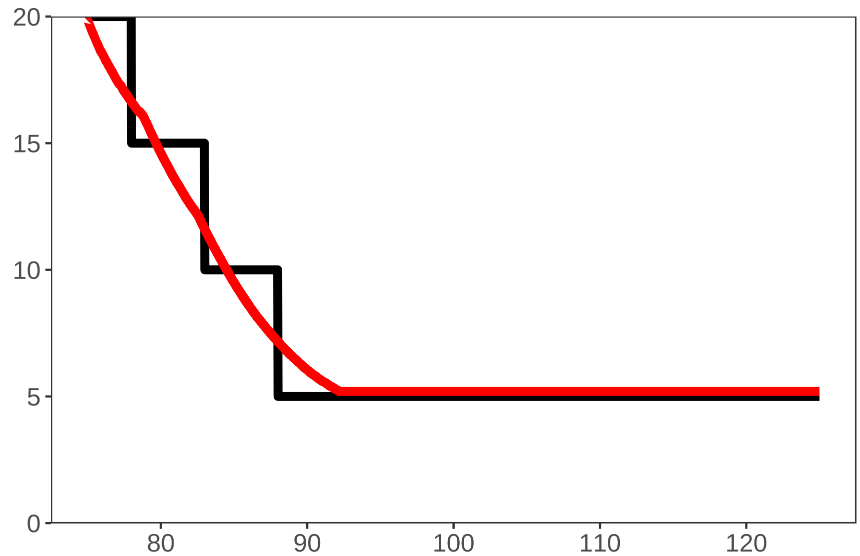

Figure A2: Price support function estimation with discrete jumps in policy.
The black line represents the actual price support function. We assume the underlying state is log-normal with mean level 100 and volatility of 7.5% and follow our method to estimate the price support function. The red line reports the estimate price support function.

E. Effect of Anticipation

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

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

Then we have:

$$F^{-} = \left(1-\theta^{-}\right) F_{p_1} + \theta^{-} F_{p_1'} \tag{A25}$$

$$F^{+} = \left(1-\theta^{+}\right) F_{p_1} + \theta^{+} F_{p_1'}. \tag{A26}$$

Rearranging these equations allow to inverse the two distributions:

$$F_{p_1} = \frac{\theta^{+}F^{-} - \theta^{-}F^{+}}{\theta^{+} - \theta^{-}} \tag{A27}$$

$$F_{p_1'} = \frac{1}{\theta^{+} - \theta^{-}} \left[(1-\theta^{-})F^{+} - (1-\theta^{+})F^{-}\right]. \tag{A28}$$

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

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

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

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

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

Application to the 2020 corporate bond purchases. Figure A3 reports the price support function implied by a few alternative combinations of θ^{-} and θ^{+} for our main empirical example (the March 23 announcement). The results are as one would expect. Anticipation makes the effect of the policy larger. Residual uncertainty about the implementation has a similar effect. However, both features do not strongly change the estimated asymmetry in price support across states.

F. Data Sources

We use a variety of financial instruments that have traded option contracts referenced to them and were the direct target of policy announcements.These include options on the iShares investment grade corporate bond ETF (LQD), the iShares high yield corporate bond ETF (HYG), the future on the S&P500 index, the future on the ten-year Treasury bond, the financial sector ETF (XLF), the future on the Nikkei index, the future on the European stock index (STOXX), and the CDX investment grade credit basket spread. We aim to use options of maturities close to three months which are frequently the most liquid. Data from exchange traded instruments come from OptionMetrics

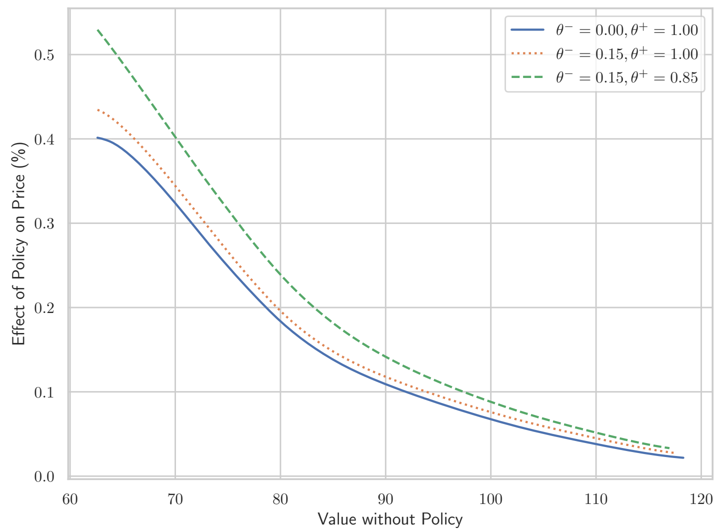

Figure A3: Effects of anticipation on price support. This figure shows the price support function under alternative assumptions about the probability of policy intervention before (θ^-) and after the announcement (θ^+).

(2022). Data on CDX and their options come form Bloomberg (2022) and IHS Markit (2022).[5] Data on ETF prices are from The Center for Research in Security Prices (2022).

G. Details of Corporate Bond Purchase Analysis

1 Event study for corporate bond purchases

The announcement of the SMCCF and PMCCF had a significant and immediate impact on corporate bond prices. Table A1 shows the return response for the iShares investment grade corporate bond ETF (LQD) using a window of one to three days around the announcement. Data are from Yahoo Finance (2022). 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

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

Table A1: Bond Price Response to the March 23 Announcementt

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

	(1)	(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)
N	2,988	2,988	2,988	2,988
R^2	0.11	0.87	0.09	0.87

identification at the cost that it may take the market time to process the announcement.[6] 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.

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

Figure A4 plots the implied price support function in our main exercise along with a dashed line for flat price support below the median price absent policy. We compute the expectations of each of these two price support functions to assess how much the extra support in the left tail increased the price at the time of announcement.

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

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

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

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

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

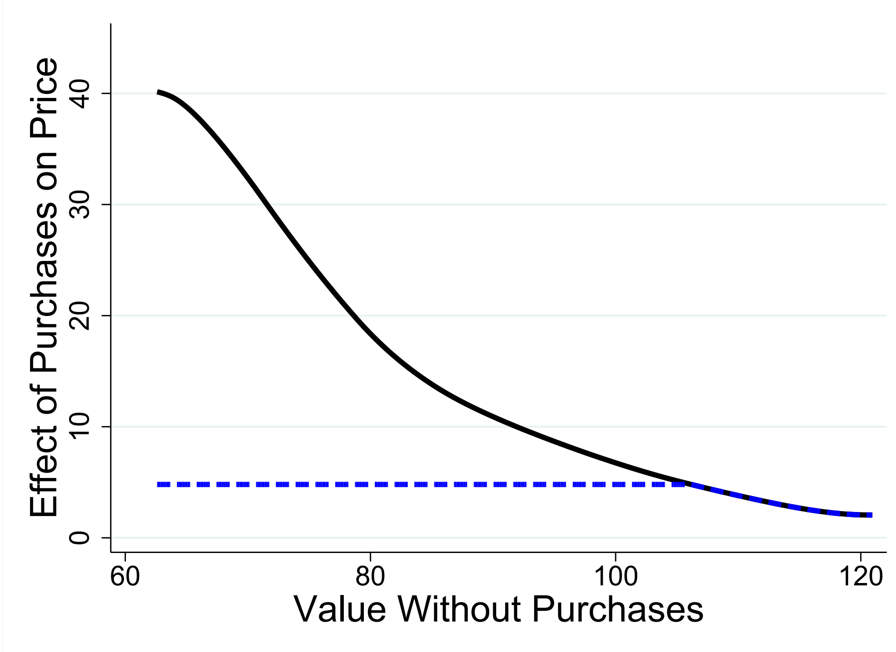

Figure A4: Counterfactual Price Support Function Without Downside Support.
This figure shows the implied price support (expressed in percent) as a function of the pre-policy price. The pre-policy price is normalized to 100 before announcement. The dashed line indicates a flat price support function below the median price absent policy.

$$p_0 = \int p_1(s) dF^{\mathbb{Q}}(s), \tag{A29}$$

$$p'_0 = \int p_1(s)(1+g(s)) dF^{\mathbb{Q}}(s), \tag{A30}$$

(A31)

where p_0 and p'_0 are the date-0 prices of the asset before and after the policy is announced. The after-announcement price if a counterfactual policy $\tilde{g}(s)$ was announced follows immediately:

$$\tilde{p}_0 = \int p_1(s)(1+\tilde{g}(s)) dF^{\mathbb{Q}}(s) \tag{A32}$$

We compute this price for a price support function truncated at is median as in Figure A4. Table 1 reports

$$1 - \frac{\tilde{p}_0 - p_0}{p'_0 - p_0}, \tag{A33}$$

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

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

$$m(s) = \theta(s)\frac{p_0}{p_1(s)} \tag{A34}$$

$$dF^{\mathcal{N}}(s) = \theta(s)dF(s), \tag{A35}$$

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

$$dF^{\mathcal{N}}(s) = \frac{p_1(s)}{p_0}dF^{\mathbb{Q}}(s), \tag{A36}$$

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

$$1 = \int dF^{\mathbb{Q}}(s) \tag{A37}$$

$$1 = \int \frac{p_0}{p_1(s)}\theta(s)dF(s) \tag{A38}$$

$$\frac{1}{p_0} = \int \frac{1}{p_1(s)}\theta(s)dF(s) \tag{A39}$$

$$\frac{1}{p_0} = \int \frac{1}{p_1(s)}dF^{\mathcal{N}}(s). \tag{A40}$$

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

$$\frac{1}{p_0'} = \int \frac{1}{p_1(s)(1+g(s))}dF^{\mathcal{N}}(s), \tag{A41}$$

$$\frac{1}{\tilde{p}_0^{\mathcal{N}}} = \int \frac{1}{p_1(s)(1+\tilde{g}(s))}dF^{\mathcal{N}}(s). \tag{A42}$$

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

$$1 - \frac{\tilde{p}_0^{\mathcal{N}} - p_0}{p_0' - p_0}. \tag{A43}$$

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

4 The size of the investment grade corporate bond market

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

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

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

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

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

5 Longer event window

We demonstrate robustness to using a longer window in our event study. Our main results use one day which tightens identification. However, it could also be reasonable to allow more time for markets to react at the cost of less tight identification since a longer period means that other shocks could be affecting markets. Figure A5 shows that the results are similar if we expand the event window to three days. While magnitudes are slightly larger compared to the results in the one-day window, the asymmetric effect remains pronounced. Our main analysis attributes 50% of the announcement effect to additional promises in the left tail. Using a three-day window this number falls to about 30% because the right tail remains elevated. However, the dollar value from additional left tail promises increases from about $250 billion using the one-day window to about $300 billion when we expand to the three-day window. This comes from the overall return on corporate bonds being larger over three days compared to one day. This shows that the choice of event-window length does not have a substantial effect on these results.

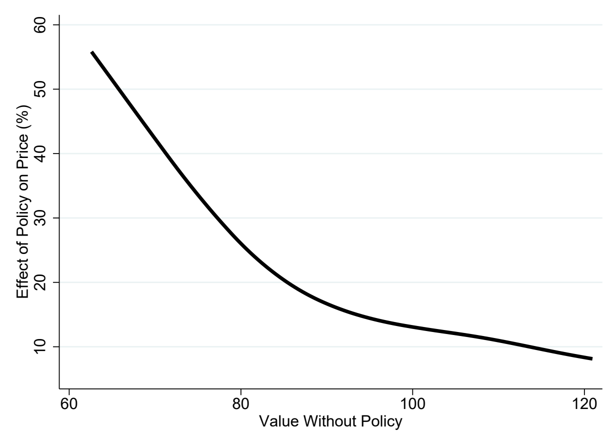

Figure A5: Price Support Function for the March 23 Announcement: 3-Day Window.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

6 Comparison to high yield bonds

Figure A6 contrasts the effects on investment grade bonds with those for high yield, using options on the largest and most liquid high yield bond ETF (HYG). Noticing the vastly different y-axes, the overall returns for high yield are basically flat. This result suggests that the announcement did not coincide with other macroeconomic news affecting corporate bond markets, since the effects are strongly concentrated in investment-grade bonds which were the target of the purchases. Second, and more importantly, they speak against the possibility that changes in the pricing kernel are driving our results as a lowering of the price of credit risk should show up disproportionately in high-yield bonds, which is not the case.

Figure A7 plots the price support from the April 9, 2020 announcement that expanded the facilities to include high-yield bonds. We follow the same methodology applied to options on

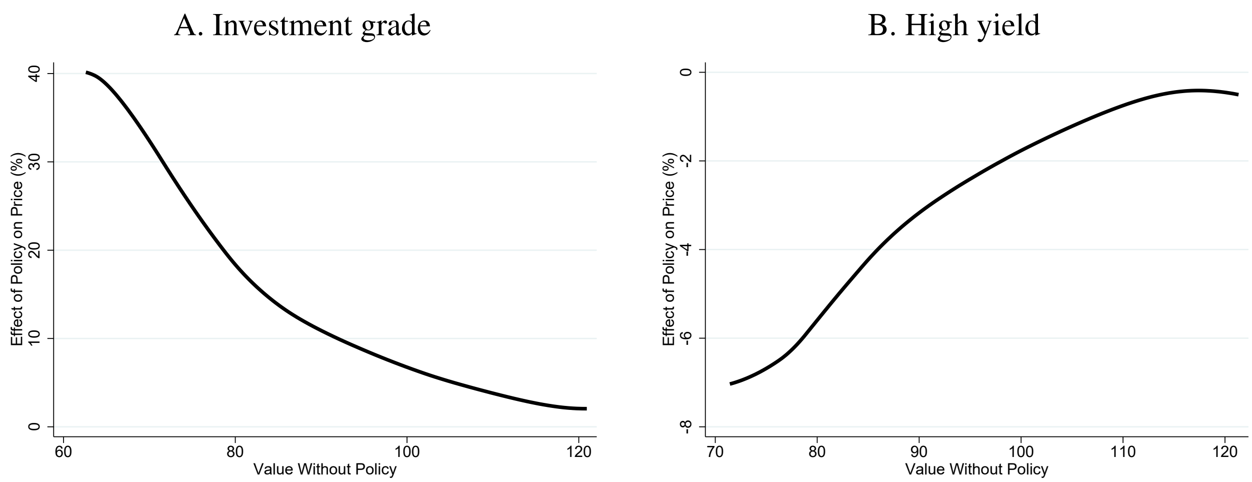

Figure A6: Comparison of Imvestment Grade and High Yield for March 23.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

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

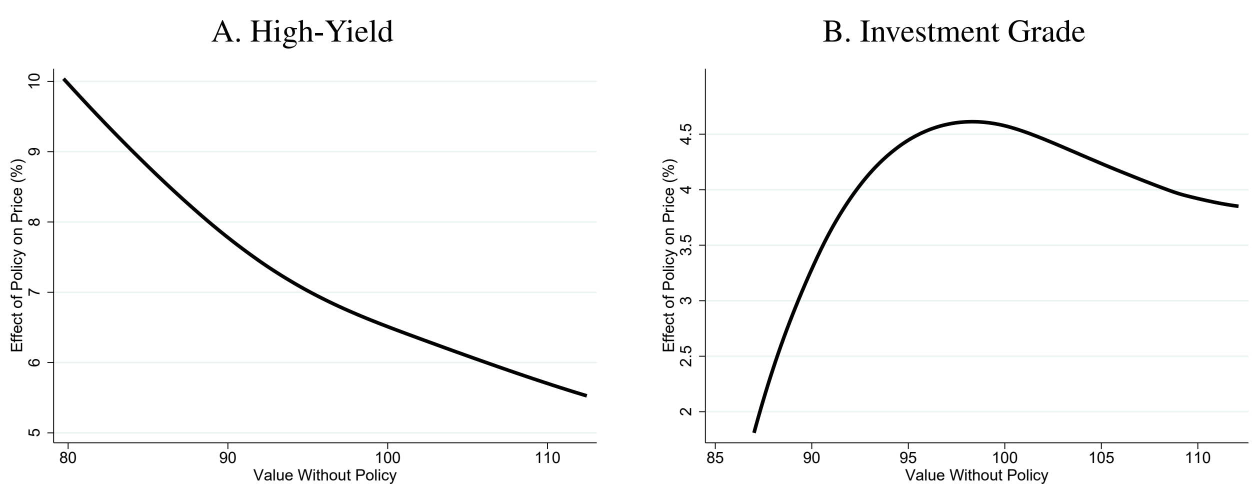

Figure A7: High-Yield Announcement on April 9, 2020: Price Support Functions.
This figure shows the implied price support (expressed in percentage as a return) as a function of the pre-policy price, normalized to 100 before announcement.

7 Liquidity

It is well known that liquidity in option markets—and derivatives markets more broadly—is heavily skewed. Trading volumes in the instruments we consider are not extremely high. For example, the options on the investment grade corporate bond fund we investigate in Section II.C have an overall trading volume in 2020 (in terms of contract notionals) around \$100 billion. Bid-ask spreads are about 3% but grow considerably for low strike prices, reaching values as high as 30%. Thus, it is natural to investigate the robustness of the patterns we document to liquidity costs. In Section II.C we show that the negatively slopped support function we recover is very unlikely to have happened by chance. It therefore cannot be driven the overall level of liquidity in this market, since if this pattern was liquidity-driven we should expect it to show up recurrently in the data. Still, one could be concerned that liquidity disappeared exactly around the announcement since those were unprecedented times. To evaluate this possibility we replicate our recovery procedure but using bid and ask quotes. Such quotes are firm offers to buy and sell and therefore less subject to liquidity concerns. In Figure A8 below we report the recovered price support function with all four possible combinations. Of particular interest is the line that depicts the price support function implied by the bid-ask pair since it reflects the prices at which investors could have bought (in small quantities) options before the announcement and sold after the announcement. Thus, the implied price support function corresponds to actual returns for an investor even if we account for the illiquidity implied by a wide bid-ask spread. While accounting for these price differences has a visible impact on the price support function, the broad magnitude and shape are mostly unaltered.

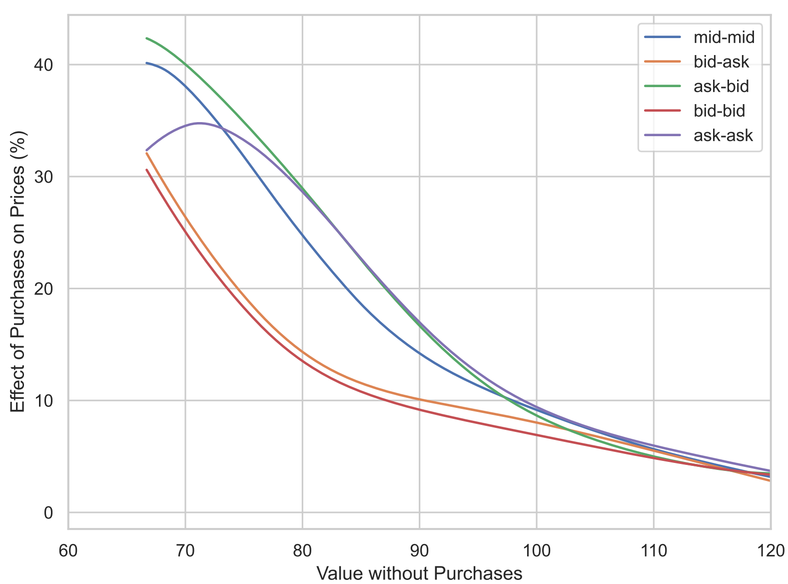

Figure A8: Price Support Function Using Bid and Ask Quotes.

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

H. Confidence Intervals

1 Bootstrap procedure

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

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

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

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

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

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

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

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

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

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

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

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

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

Figure A9 reports the results. We now plot the 95% confidence bands in their initial position. These bands are not centered around 0: by construction we select large value of returns, so the average value of g is also large. The average level of g obtained by selecting these outliers (the dotted red line) is very close to the realized return for our event (the horizontal black line). This indicates that we are looking at broadly the same intensity of outliers. The blue band reveals that a strongly negative slope of the price support function is not a feature of the simulated data. Said otherwise, we reject the hypothesis that the negative slope of g we see for the announcement is explained by the fact that there was a large return.

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

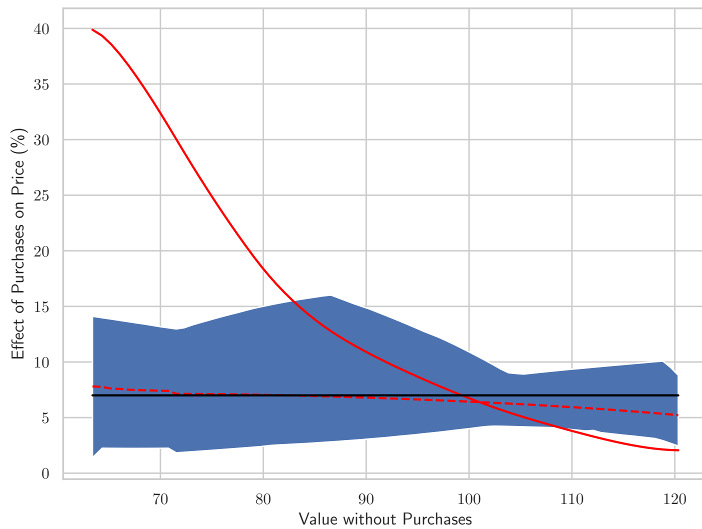

Figure A9: Standard Errors: Only Sampling Large Returns. The figure shows out baseline result with 95% confidence bands that now adjusts the baseline procedure to sample only from extreme (both positive and negative) 1% (implied-volatility) normalized return days. See text above for detailed description.

magnitude of the event day return is so large. To get at larger magnitudes of returns, we extend the time window to construct the "usual" behavior.[7]

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

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

$$Slope_{t,h} = \frac{g_{t,h}(x_t)|_{F(x_t)=\underline{F}} - g_{t,h}(x_t)|_{F(x_t)=0.5}}{x_t|_{F(x_t)=\underline{F}} - x_t|_{F(x_t)=0.5}}, \quad \text{(A44)}$$

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

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

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

Figure A10 reports a scatter plot of all slope and realized return pairs, i.e. all date-horizon observations. Our event is the blue star in the bottom right corner with a slope close to -1 associated with a return of about 7%. The different colors show the different horizons. The range of return realizations grows as the horizon increases, getting progressively closer to the 7% return.

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

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

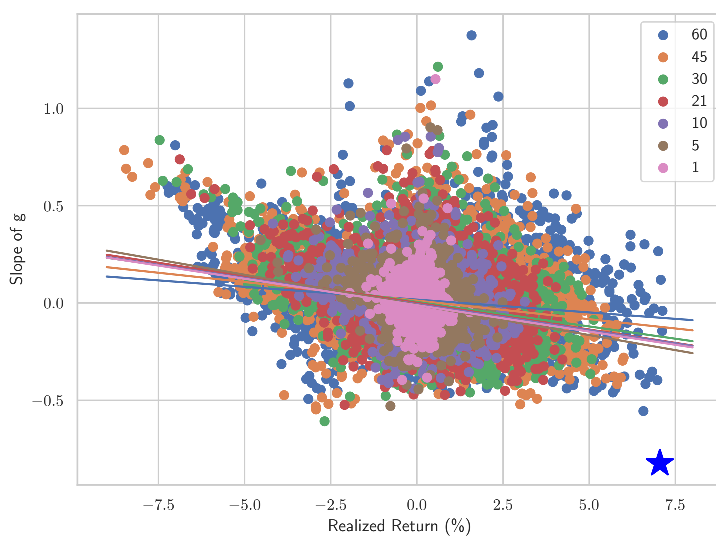

Figure A10: Relation Between Realized Return and Slope of the Price Support Function. Each dot represents a date-horizon pair. Different colors are different horizons ranging from 1 up to 60 trading days. The different lines are ordinary least-square predictions for each horizon. The March 23 event is shown with a star marker in the bottom right of the plot. See text for details.

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

I. State-Dependent Multiplier Calculations

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

To address this directly, we calibrate the Greenwood and Vayanos (2014) results to our setting based on the standard deviation of returns on investment-grade bonds implied by option prices. Specifically, we define $\mathcal{M}(p) = \mathcal{M} \times \left(1 - \frac{1}{4}R(p)/\sigma\right)$. That is, we take our baseline constant value of $\mathcal{M}$ and multiply by $\left(1 - \frac{1}{4}R(p)/\sigma\right)$ where $R(p)$ is the (net) return on corporate bonds in a given state p and σ is the standard deviation of returns using option implied volatility. This gives us quantitative variation in multipliers consistent with what Greenwood and Vayanos (2014) find empirically.[8] We then plot the associated quantities $Q(p)$ that we back out from using $g(p) = \mathcal{M}(p) \times Q(p)$. Figure A11 below plots the quantities when we use a multiplier that depends on the realized state.

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

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

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

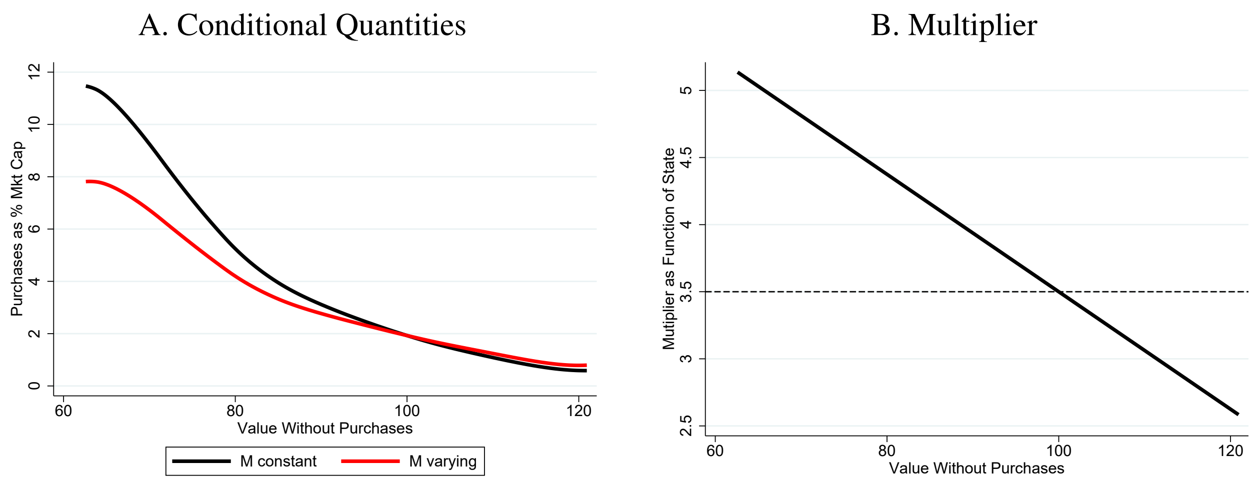

Figure A11: Price Support Function with a State-Dependent Multiplier

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

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

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

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

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

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

More specifically, we apply the method in two steps because it is more intuitive to think about the correlation between the synthetic and the cash instrument then to think about the correlation of the cash instrument and the different pieces of the synthetic. We first apply the copula method to the CDX and treasury options. We use a correlation of -0.25 which is consistent with the historical data. The results are quantitatively similar if we change the correlations to -0.75 or 0.25 (see Figure A12 Panel A). We set the correlation between the synthetic and the cash instrument to 0.8 which is consistent with the historical average. This correlation tends to go down during crisis (for example a 30-day moving average estimator has the lowest realization of 0.25 in our 9-year sample) so we show results with 0.4 and 0 as well (see Figure A12 Panel B). The key result is unchanged.

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

K. Implications of Promises for Market Dynamics

1 Are the effects of asset purchases getting weaker: empirical details

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

1.1 United Kingdom

Table A2 contains the response of 10 year Gilt yields to six announcements of purchases between February 2009 and February 2010. The Bank of England is unique in that most of these announcements contain a fairly narrow and specific quantity range. First, note that only the first two

A. Varying the correlation between interest rates and credit spreads

B. Varying the correlation between synthetic and cash bond

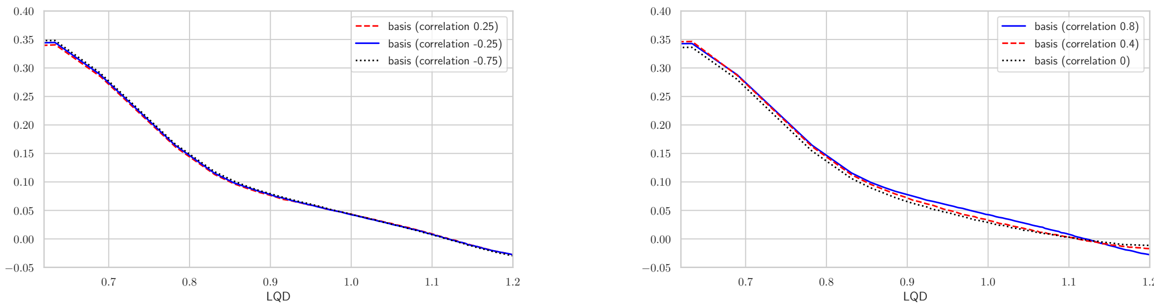

Figure A12: Decomposition of Announcement Effects: Robustness

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

Table A2: 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]

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

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

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

1.2 United States

Table A3 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.[9] The initial announcement, which stated the Fed would purchase "up to" \$600 billion across these categories, led to an average decline in yields of about 40 bps. This equates to a multiplier of about 0.8 (e.g., for a purchase sized at 1% of market cap, prices would increase by 0.8%). The next significant announcement in QE1 came in March 2009, where the Fed expanded quantities. While yields moved by about the same amount, the quantities were larger. This leads to a lower multiplier. Later announcements, for example dropping the "up to" language and effectively confirming the Fed would purchase the maximum stated amount, had no effect on yields. These patterns also fit the support function results we present in Section III for the initial announcements, which indicate the presence of implied policy puts in these initial announcements.

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

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

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

Table A3: US Announcement Effects

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

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

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

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

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

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

Table A4: ECB Announcement Effects

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

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

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

1.3 Eurozone

Table A4 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.

2 Longer-run effects on risk in corporate bond markets

We investigate the long-run effects of the Fed's corporate bond intervention by looking at how the dynamics of corporate bond tail risk changes after the programs are implemented. While our main analysis documents how the Fed's introduction of the policy has the immediate effect of reducing this tail risk when the Fed initially announces it will intervene, we now assess whether tail risk is less sensitive to economic conditions going forward. These longer-term effects are not easily captured in our earlier framework which focuses on conditional promises over shorter maturities at which we have option price data. The challenge is to have a good benchmark for how tail risk would behave absent interventions. We look at a variety of approaches to deal with this challenge, though these results should be taken as suggestive and depend on how reasonable our benchmarks are.

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

Table A5 shows that tail risk sensitivity changed after the announcements: tail risk in corporate bond markets are usually positively related to tail risk in equity or CDS markets; after the interventions, this sensitivity disappears altogether. Specifically, we regress the corporate bond tail risk on tail risk in equities and CDS markets using daily data from 2010 onward (for CDS, we only have data from 2015 onward). We then include a "post" dummy interaction term for the period after April 9, 2020 when the Fed had already announced the expansion of its corporate bond facilities. Notably, in the period prior to this, corporate bond tail risk and equity market tail risk co-move strongly so that tail risk in corporate bonds was highly sensitive to tail risk in equity markets.

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

This declining sensitivity is not only present in tail risk, but also in overall bond yields. In Table A6, we find that corporate bond returns are usually negatively related to changes in the VIX.

Table A5: Long Term Effects on Corporate Bond Tail Risk

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

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

This sensitivity is divided in half in the post-intervention period.

Figure A13 instead uses monthly data on option-based pseudo credit spreads from Culp et al. (2018). These pseudo-spreads are constructed by equity options and Treasuries — not corporate bonds — and have been shown in earlier samples to exhibit remarkably similar properties than the actual spreads. We use the two-year maturity investment-grade pseudo credit spread from Culp et al. (2018).[10] We then compare this spread to the Bank of America investment grade option-adjusted credit spread index for maturities between one and three years taken from FRED. We plot both the actual credit spreads and pseudo spreads in Figure A13. From 2010 to 2020, the two spreads track each other quite well. In early 2020, when the COVID-19 crisis hits, actual spreads for investment-grade spike well beyond those implied by equity market options, consistent with investment grade bond prices becoming abnormally depressed in this episode. However, following the Fed's intervention, investment-grade spreads become quite low, and in fact reach their lowest point at any time over the 2010-2020 window. In contrast, equity markets still feature substantial volatility, implying higher than usual default risk on pseudo-bonds. This large gap is consistent with a market pricing of future interventions: a crash is still possible and priced in equity options, but the Fed would intervene in corporate bond markets and make it disappear.

Consistent with our evidence of abnormally low credit spreads after the intervention, Boyarchenko et al. (2020) and Becker and Benmelech (2021) find abnormally large issuance of investment-grade bonds by firms after the interventions, and Balthrop and Bitting (2022) find this

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

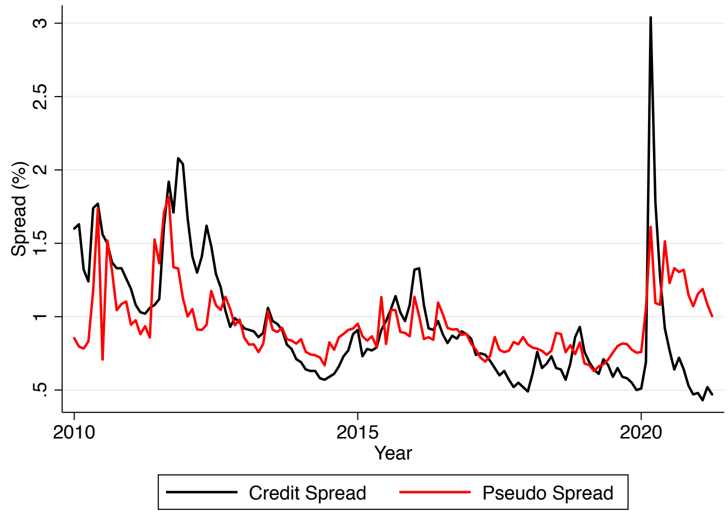

Figure A13: Spreads and Pseudo Spreads.
This figure plots actual credit spreads and pseudo spreads from Culp et al. (2018).

effect is persistent for firms eligible for Fed purchases under the original SMCCF facility.[11] Again, this increase in issuance fits with the narrative of an implicit subsidy of low spreads due to expectations of future Fed support.

References

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

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

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

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

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

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

Bloomberg, 2022, Option implied volatilities and futures, Accessed through Bloomberg Terminal, OVDV (April 2022).

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

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

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

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

Table A6: Long Term Effects on Corporate Bond Prices

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

Panel A: Corp Bond Returns

	(1)	(2)
	$r_t^{CorpBond,e}$	$r_t^{CorpBond,e}$
ΔVIX_t	-0.21	-0.20
	(0.02)	(0.02)
$\Delta VIX_t \times post$	0.10	0.08
	(0.03)	(0.03)
$\Delta VIX_t \times covid$		-0.05
		(0.04)
$post$	0.04	0.03
	(0.03)	(0.03)
$covid$		-0.12
		(0.34)
Constant	-0.01	-0.01
	(0.01)	(0.01)
Observations	2,987	2,987
R-squared	0.26	0.26

Panel B: Credit Spreads and Option-Based Pseudo Spreads

	(1)	(2)
	$\Delta spread_t$	$\Delta spread_t$
$\Delta pseudo_t$	0.41	0.16
	(0.19)	(0.05)
$\Delta pseudo_t \times post$	-0.62	-0.37
	(0.23)	(0.13)
$\Delta pseudo_t \times covid$		1.74
		(0.38)
$post$	-0.11	-0.10
	(0.05)	(0.05)
$covid$		0.08
		(0.23)
Constant	0.00	-0.01
	(0.02)	(0.01)
Observations	135	135
R-squared	0.22	0.69

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

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

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

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

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* .

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

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.

IHS Markit, 2022, Credit index options, Available through https://www.markit.com (accessed September 2021).

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

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

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

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

OptionMetrics, 2022, Ivydb, Available through Wharton Research Data Services (WRDS). https://wrds-www.wharton.upenn.edu/pages/about/data-vendors/optionmetrics/ (accessed November 2021).

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

The Center for Research in Security Prices, 2022, Crsp daily stock file, 1999-2022, Available through Wharton Research Data Services (WRDS). https://wrds-www.wharton.upenn.edu/pages/get-data/center-research-security-prices-crsp/annual-update/stock-security-files/daily-stock-file/ (accessed May 2021).

Vayanos, Dimitri, and Jean-Luc Vila, 2021, A preferred-habitat model of the term structure of interest rates, *Econometrica* 89, 77–112.

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

Yahoo Finance, 2022, Adjusted daily close prices, Available through https://finance.yahoo.com/ (accessed November 2021).