To appear in *Quantitative Finance*, Vol. 00, No. 00, Month 20XX, 1–24

An agent-based model of corporate bond trading

K. Braun-Munzinger†, Z. Liu‡ and A. E. Turrell*§

†European Central Bank, 60640 Frankfurt am Main, Germany
‡Hong Kong Monetary Authority, 8 Finance Street, Hong Kong
§Bank of England, Threadneedle Street, United Kingdom

(*Received 00 Month 20XX; in final form 00 Month 20XX*)

We construct an heterogeneous agent-based model of the corporate bond market and calibrate it against US data. The model includes the interactions between a market maker, three types of fund, and cash investors. In general, the sensitivity of the market maker to demand and the degree to which momentum traders are active strongly influence the over- and under-shooting of yields in response to shocks, while investor behaviour plays a comparatively smaller role. Using the model, we simulate experiments of relevance to two topical issues in this market. Firstly, we show that measures to reduce the speed with which investors can redeem investments can reduce the extent of yield dislocation. Secondly, we find the unexpected result that a larger fraction of funds using passive investment strategies increases the tail risk of large yield dislocations after shocks.

Keywords: Agent-based modelling; Corporate bonds; Financial instability; Bond funds

JEL Classification: C63; G12; G23

*Corresponding author. Email:arthur.turrell@bankofengland.co.uk
The views expressed in this paper are solely those of the authors, and do not necessarily represent those of the Bank of England, the European Central Bank, or the Hong Kong Monetary Authority and should not be reported as such.

1. Introduction

The ability of financial markets to absorb shocks depends on factors such as market micro-structure, the vulnerabilities of the entities operating within them and the behavioural choices of market participants. Market maker inventory, for example, may well play a role in their willingness to contribute liquidity (Comerton-Forde *et al.* 2010). Equity markets have been the focus of much work examining these features (Coval and Stafford 2007, Christoffersen *et al.* 2014). Bond markets are less well examined, though there have been empirical studies such as the recent work by Goldstein *et al.* (2015) and a literature which identifies differences in the propensity of more or less well informed investors to withdraw investments in a shock.

We provide a simulation-based view of the market in corporate bonds which extends empirical studies in two ways. First, it allows the consequences of the limited rationality of different market participants to be assessed for the corporate bond market. Modelling of equity markets has shown that limited rationality can lead to large price swings (Hommes 2006). Second, it allows for the exploration of how outcomes in this market differ under the application of different policy tools.

These simulations are performed within an agent-based model, of the kind long used in the natural sciences (Turrell 2016). Agent-based models are being used to simulate the interactions of self-interested, heterogeneous and partially rational actors within economic systems (Farmer and Foley 2009). They have enjoyed particular success in finance where their ability to model agents without perfect rationality or perfect foresight in partial equilibrium have seen them reproduce a number of features of real-world financial markets (Ponta *et al.* 2011, Immonen 2017).

The agent-based model developed here allows us to assess how mutual funds' trading strategies, the speed of price adjustments by a market maker and the behaviour of end investors interacts. We draw on this simulated market to gain insights into the implications of market conditions such as the rise of passive investment strategies and to test the impact of public policies on shock propagation. The model is purposefully parsimonious so that it may be easily extended or adopted for other applications. It uses as few parameters as are required to reasonably reproduce stylised facts of the market. Where possible, the values of these parameters are based on empirical observations. For parameters that cannot be observed in the real world, some are fixed in the model if they are subsumed in other variables, or if model results are not sensitive to their values (within economically reasonable bounds). Out of the twenty parameters used in the model, nine are based on empirical observations, six are imposed, and only five are calibrated (see Table 5). This parsimony extends to the behaviours of the funds: despite the funds following relatively simple rules, the emergent behaviour of the price is reasonably close to that which is observed empirically.

We do not incorporate wider market factors such as interest rates, macroeconomic factors affecting the value of the bonds, or alternative investment opportunities and therefore the model does not fully capture some of the fat-tail behaviour in corporate bond prices observed in practice (see Fig. 7). Introducing extra parameters, such as non-Gaussian noise, might have resulted in a closer fit to fat-tailed distributions of returns but at the cost of making the model more complex. The model as presented with Gaussian noise has the property that departures from Gaussianity in steady state are generated endogenously by the interactions between the agents. We calibrate the model to match the US investment grade corporate bond index and the US open-ended mutual fund industry, based on data from Bank of America Merrill Lynch and Morningstar. As part of our analysis, we impose exogenous shocks to fundamentals in order to investigate the role of market structure and policy choices.

The results section shows that the sensitivity of the market maker to demand and the degree to which momentum traders are active are important for the response of the simulated bond market to shocks, especially with respect to the amplitude of yield (and price) dislocations to the long term steady state values. Two scenarios are explored. In one, it is demonstrated that a larger fraction of passive funds can increase the tail risk associated with shocks considerably. In another, we show that managing redemptions over a period longer than a single trading day could reduce the impact of shocks in stressed times when investor outflows are unusually large, as in the 2008 financial crisis.

The rest of the paper is organised as follows. Section 2 describes the model, and Section 3 describes the calibration approach. Model results are presented in Section 4 and Section 5 concludes.

Table 1. Agents in the model

Agent	Number
Representative investor	1
Representative market maker	1
Funds (three types)	1000

Table 2. Variables for the bond index

Variable	Variable name	Value
P^*	Fundamental price	91.57 units of cash
Y^*	Fundamental yield	1.282% per year
F	Face value	100 units of cash
D	Duration	6.917 years

2. The model

We propose a model of heterogeneous funds trading a corporate bond index via a stylised over-the-counter market maker.

Our model builds on the existing literature which simulates financial markets using heterogenous agents of limited rationality. For the behaviours of the agents, we draw on the heterogeneous agent trading models as summarised in Hommes (2006) and Tesfatsion (2006) which have been used to understand volatility and other aspects of financial market dynamics in, eg, Aymanns and Farmer (2015), Thurner *et al.* (2012), Franke and Westerhoff (2012) and Fischer and Riedler (2014).

However, the model presented is novel in a number of ways which are driven by the structure of the bond market. The risk-reward payoff in a fixed income environment differs substantially to that of equity trading in that potential gains and losses are capped and are arguably more predictable with typically positive non-zero payoffs on indices. We reflect these aspects in the trading strategies of investors, building on the typical value and momentum trading types. Those trading types are also motivated by empirical literature specific to mutual funds such as Lütje (2009), who finds evidence that a significant proportion of managers of mutual funds follow the market trend, and Shek *et al.* (2015), who demonstrate that discretionary sales follow forced sales for emerging market economy bonds, suggestive of trend following.

In addition, portfolio changes by funds in our model are gradual and of heterogeneous speed, reflecting trading costs and different timing ability as in Moneta (2015) and Chen *et al.* (2010b). Following Farmer and Joshi (2002), we also include the decreasing willingness of market makers to bear risk due to uncertainty by making the speed of adjustment of price with demand dependent on past volatility.

As in Thurner *et al.* (2012), we include a stylised investor pool, but incorporate insights from theoretical and empirical literature (see in particular Chen *et al.* (2010a), Goldstein *et al.* (2015) and Chen and Qin (2014)), which finds that investors respond to both overall market performance and to idiosyncratic fund performance in allocating funds. In contrast to investors in equity funds, corporate bond fund investors do not disproportionately reward high performers. Instead, fund flows vary linearly with performance or are even concave, i.e. investors respond more strongly to negative performance.

In designing this model, we also take into account a number of recent papers that use dynamic methods to explore stressed markets (see, for example, Bookstaber *et al.* (2015)). While these papers have underlined the importance of leverage and risk management choices (see also Thurner *et al.* (2012), Fischer and Riedler (2014), Aymanns and Farmer (2015)), there has been relatively less exploration of the role of unleveraged entities. In addition, empirical validation of models has typically been confined to equity and foreign exchange markets. Our model fills this gap in the literature, motivated by concerns over the strong growth in intermediation by mutual funds and structural changes in markets which may affect less liquid fixed income markets as discussed in Fender and Lewrick (2015) and the report of Navigating Monetary Policy Challenges and Managing Risks (2015).

There are five types of agents in the model, and these are shown in Table 1. There are two types of representative agent: investors and the market maker, and three types of agent which are funds: value traders, momentum traders, and passive funds.

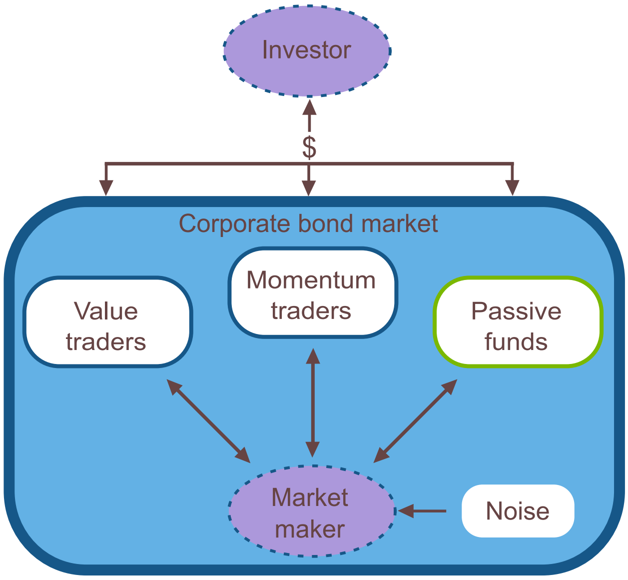

Figure 1. Stylised model setup.

There are two assets in the model; a risky asset, which is a corporate bond index, and a risk-free asset, cash. There is no return for holding cash. The bond index has a positive, non-zero fundamental yield Y^*, which is related to a fundamental price P^* and a face value F by

$$P^* = \frac{F}{(1+Y^*)^D} \tag{1}$$

for bond duration D and assuming zero coupon. P^* is the present fundamental face value of the bond. The parameters used are shown in Table 2.

With zero coupon, the annuity value is zero, and P^* is given by the maturity value alone. This is a convenient modelling assumption that has no impact on the dynamics of the model as we can replicate coupon-paying bonds with a portfolio of zero-coupon bonds. We do not model issuance. Note that P^*, D, F, and therefore Y^*, are all fixed unless stated otherwise and are taken from averages of the relevant variables in our data. The traded price of the bond index at time t is P_t, with P_t determined dynamically and endogenously.

A schematic of the model is shown in Fig. 1. Funds are subject to cash in- and outflows from a stylised investor pool. Funds invest in the index and hold cash. Payoffs from the risky asset are passed directly to the pool of investors. Funds buy and sell the index from a market maker, who sets prices based on the net demand of the funds, and on how volatile recent market conditions are.

The high-level summary of the sequence of events over a single simulated trading day can be seen in Fig. 2. Each trading day is represented by a value of t. We summarise the sequence of events in more detail here:

(i) A trading day, indexed by t, begins. The price has been set the previous day by the market maker at P_t.
(ii) Each fund is subject to inflows and outflows of cash from a single, representative external investor pool. Investment decisions are made each trading day on the basis of past return rates (both overall sectoral return rates and the past relative return rates of individual funds), as described in equations (2) and (3), and the flow equation, (4).
(iii) Each fund then decides on its desired fraction of risky versus risk-free assets by solving the mean-variance optimisation problem of equation (6). The nature of the expected return is determined according to the fund's type (fundamentalist, momentum trader, or passive fund).
(iv) Having decided on the desired fraction of risky assets, the fund moves to rebalance its portfolio

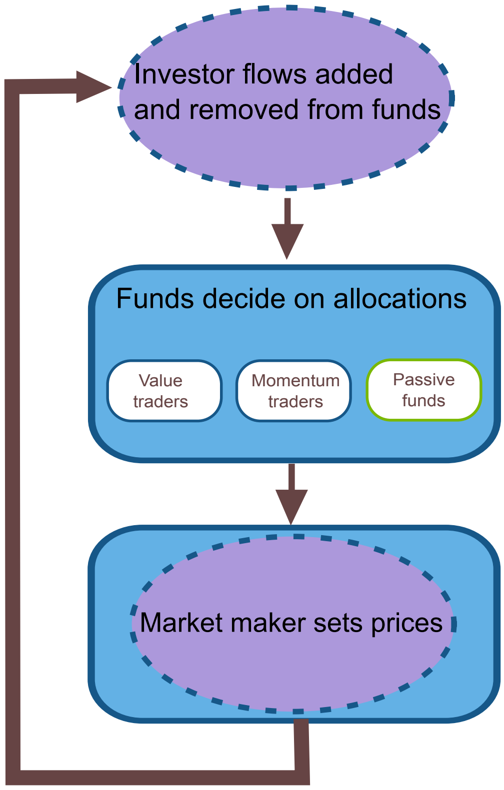

Figure 2. The sequence of events in a single simulated trading day.

towards the desired fraction by buying or selling the bond index at price P_t. Fund j only completes a portion $1/\delta_j$ of the portfolio rebalancing on day t.

(v) Having met all demand and supply requests at price P_t on day t, the market maker uses the net demand on day t, $\sum z_{j,t}$, to set a price for the next day, P_{t+1}, according to equation (11).

(vi) $t \rightarrow t+1$ and the next trading day begins.

Subsequent sections explain these steps in more detail.

2.1. *Investor flows*

Following these stages in detail, we begin with the investor flows. Investors are represented by a stylised investor pool who withdraw and invest based on the past excess returns of funds. There is no budget constraint for investors. Investors represent holders of financial assets other than US corporate bonds that may potentially invest in open-ended US corporate bond funds. According to Bank of England (2015), there are around \$227 trillion of financial assets globally, of which 11% is held by open-ended funds and less than 0.3% is held by open-ended US corporate bond funds (see Section 3 for the estimated size of open-ended US corporate bond funds). This suggests that the size of potential investors is much larger than the current size of open-ended US corporate bond funds, and provides justification for assuming no budget constraint for investors. In the simulations, introducing a budget constraint for investors that is relatively large, equivalent to ten times the total initial wealth of the funds, would make no difference to the results.

The flows into and out of the funds are determined in part by the aggregate performance of the bond index, and in part by the performance of individual funds. This is based on the empirical work by Chen and Qin (2014) and Goldstein *et al.* (2015). The former jointly estimate sensitivity of fund flows to macroeconomic conditions and individual fund performance and find significant relationships with both aggregate bond performance and performance rank. The latter estimate the relative flows between funds, controlling for aggregate in- and outflows, and find a significant concave relationship with performance. Motivated by that, investor flow, S, in our model is a linear function of both the rate of return on the index, R^S, and of each individual fund's rate of return, R^I. The rate of return on the index is taken relative to the price change over the last 20 trading days,

$$R_t^S = \frac{P_t - P_{t-20}}{20 \cdot P_{t-20}} \tag{2}$$

The rate of return based on individual funds is given by $R^I_{j,t} = R_{j,t} - R_t$ where

$$R_{j,t} = \frac{A_{j,t}P_t}{A_{j,t-20}P_{t-20}} - 1 \tag{3}$$

is the mean rate of return on the risky asset over the last month (20 days) for fund j, $A_{j,t}$ is the quantity of the risky asset held by fund j at time t and R_t is $R_{j,t}$ averaged over all funds. The total investor flow is then given by

$$S_{j,t} = sR_t^S + \left[I_- + H(R^I_{j,t})(I_+ - I_-)\right] R^I_{j,t} \tag{4}$$

with the first term the flow due to the sectoral rate of return, and the second term the flow due to the performance of individual funds. For the first term, a single coefficient s maps aggregate performance into an investor flow. For the fund specific rate of return, there are two possible coefficients, I_+ and I_-, depending on whether an individual fund has a positive or negative return relative to the other funds. $H(R^I_{j,t})$ is the step (or Heaviside) function, which is one for positive values of the argument and zero otherwise. The flow function then has two possible combinations of coefficients: (s, I_+), (s, I_-) depending on whether $R^I_{j,t} > 0$ for fund j at time t. The values of the coefficients for each term are set empirically.

2.2. *Fund portfolio decisions*

Funds meet the in- and out-flows (from investments and redemptions respectively) by keeping their fractional holdings of risky assets and non-risky assets the same. For example, this means that redemptions are met through both cash and the sale of the risky asset. In practice, bond mutual funds have a number of measures to meet redemptions, including cash balances, liquid assets and liquidity facilities. While we do include cash balances, we do not include liquid assets and liquidity facilities because (i) these measures are likely to have limited effects under stressed conditions (see Cetorelli *et al.* (2016)) and (ii) we do not have reliable data on these measures.

After investor flows $S_{j,t}$ are accounted for at price P_t, the remaining wealth of fund j at time t is given by

$$W_{j,t} \equiv A_{j,t}P_t + C_{j,t} \tag{5}$$

where $A_{j,t}$ is the stock of bonds (with price P_t), and $C_{j,t}$ is cash.

Fund j decides what proportion $\kappa_{j,t}$ of its wealth it would ideally like to allocate to the risky asset by myopic mean variance maximisation. That is, $\kappa_{j,t}$ is equal to fund j's expected next period return divided by risk aversion θ and asset volatility σ

$$\kappa_{j,t} = \max\left\{\frac{\mathrm{E}_{j,t}(R_{t+1}) - R_{\text{free}}}{\theta\sigma}, 1\right\} \tag{6}$$

where R_{free} is the risk-free rate of return on cash. Without loss of generality, we set $R_{\text{free}} = 0$. Funds have a budget constraint and cannot borrow, so they cannot demand more than their wealth in any given period.

In contrast to previous models, we do not assume that funds adjust their portfolio completely at every period. A distinction is made between a fund's desired fractional holding, $\kappa_{j,t}$, a fund's fractional holdings at time t, $h_{j,t}$, and a fund's desired next period holdings, $\hat{h}_{j,t+1}$. Fund j adjusts such that they plan to reach their desired portfolio, represented by $\kappa_{j,t}$, in δ_j days based on their information at time t. We call δ_j the reaction strength of fund j. The fraction of holdings in the risky asset at time t for fund j is denoted by $h_{j,t}$ and are given by

$$h_{j,t} \equiv \frac{A_{j,t}P_t}{W_{j,t}} = 1 - \frac{C_{j,t}}{W_{j,t}} \tag{7}$$

The desired proportion of holdings of the risky asset next period is

$$\hat{h}_{j,t+1} \equiv h_{j,t} + \frac{\kappa_{j,t} - h_{j,t}}{\delta_j} \tag{8}$$

Note that δ_j varies between funds but is constant over time and is uniformly distributed between a lower and upper reaction strength, L_R and U_R respectively so that $\delta_j \sim \mathcal{U}(L_R, U_R)$. The desired assets in the next period are given by

$$\hat{A}_{j,t+1} = \frac{W_{j,t}\hat{h}_{j,t+1}}{P_t}$$

The demand of fund j in period t is then given by the difference between the stock of risky assets that they currently hold and the desired next period holdings. Using the identities, this is proportionally equivalent to the difference between the current fraction of wealth held in the risky asset and a fund's optimal fraction of wealth held in the risky asset spread over the number of days given by the reaction strength:

$$z_{j,t} = \hat{A}_{j,t+1} - A_{j,t} = \frac{W_{j,t}}{P_t}\hat{h}_{j,t+1} - A_{j,t} = \frac{W_{j,t}}{\delta_j P_t}(\kappa_{j,t} - h_{j,t}) \tag{9}$$

where the last step uses equation (7).

Expected returns $\mathrm{E}_{j,t}(R_{t+1})$ differ by trading types. In line with the literature, there are two types of active trading. Value traders assume that yields will revert to some fixed value Y^* over time, and thus buy more of the risky asset when they believe it is undervalued and less when overvalued. Momentum traders believe short-term trends in yield will persist and so sell if yields are above a long-term average $\bar{Y}$ and buy if yields are below it. There is evidence of momentum trading in bond markets in the literature, for instance, in Durham (2016), which finds that a simple momentum strategy works in the US Treasury market, and in Jostova *et al.* (2013), which documents significant momentum in US corporate bond returns.

We assume that both trading types expect a positive non-zero payoff given the publicly known face value and expected loss rate of the corporate bond index. This simplification smoothes desired asset holding and means that we do not have to introduce short selling, which is not typically a strong component of corporate bond trading, in order to deal with discontinuities. Given this, funds' expected returns are

$$\mathrm{E}_{j,t}(R_{t+1}) = \begin{cases} (1-L)(1+Y_t+\alpha(Y_t-Y^*))-1 & \text{for value traders} \\ (1-L)(1+Y_t+\beta(\bar{Y}-Y_t))-1 & \text{for momentum traders} \end{cases} \tag{10}$$

where

$$\bar{Y} = \frac{1}{t_{\mathrm{lw}}} \sum_{t-t_{\mathrm{lw}}}^{t} Y_{t'}$$

is the average yield over the time period beginning at $t - t_{\mathrm{lw}}$ and ending at time t. Y_t is the market yield on the bond at time t, while α and β are constants controlling the degree to which traders are value or momentum oriented. t_{lw} is the 'window' over which momentum traders average when considering whether yields are likely to increase or decrease. L is the loss rate which is formed by the expectations of the firms in relation to both probability of default and the recovery rate in the event of default. All agents share the

Table 3. The static variables associated with agents

Agent type	Variable	Description	Values	Type
Representative market maker	λ	Market maker sensitivity to net demand	0.033	Calibrated
	v	Volatility coefficient	20	Calibrated
	ε	Noise demand	$\varepsilon \sim \mathcal{N}(0, N_{\rm ns})$	Calibrated
All active funds	L	Expected loss rate	0.04%	Empirical
	$\theta\sigma$	The product of risk aversion and return volatility	1	Imposed
	δ_j	Fund j reaction strength	$\delta_j \sim \mathcal{U}(L_{\rm R}, U_{\rm R})$ (days)	Imposed
-Value traders	α	Value trader strength	0.008	Calibrated
-Momentum traders	β	Momentum trader strength	1.65	Calibrated
	$t_{\rm lw}$	Window for momentum traders	100 days	Imposed
Representative investor	s	Systematic flow strength	0.25	Empirical
	I_+	Fund-specific positive return flow strength	0.621	Empirical
	I_-	Fund-specific negative return flow strength	1.128	Empirical
	τ	Sector-wide redemption window size	1 day	Empirical

same expectations regarding the loss rate. This is motivated by the fact that rating agencies publish data on loss rates associated with these bonds for a given rating and fund investors would be unlikely to have access to private information on the companies in question.

We assume that funds themselves do not switch between trading strategies. This is motivated in part by the fact that funds set out their preferred strategies for their investors and therefore do not engage in frequent switching of strategies. Additionally, the growth and decline of funds based on the success of their strategies is driven by the switching of cash between funds by investors. A third set of funds, passive investment funds, only trade in response to in- and outflows from investors rather than taking a view on price. In our model, passive funds include exchange-traded-funds (ETFs), given that most ETFs track an index like a passive fund.

2.3. *The market maker*

Prices are adjusted by the (stylised) market maker as a log-linear function of net demand $\sum_j z_{j,t}$ following the approach outlined in Farmer and Joshi (2002) and Kyle (1985), supported empirically by Kyle and Obizhaeva (2016);

$$\ln(P_{t+1}) = (\lambda + vV)\left(\sum z_{j,t} + \varepsilon\right) + \ln(P_t) \tag{11}$$

This stylised marker-marker is not explicitly balance sheet constrained, assuming a market in which market makers are required to continue to make markets and where substitutability of market making services is present. We incorporate a linear function of long-run volatility V into the speed with which the market maker changes prices in response to demand. This allows us to proxy increased risk aversion as a consequence of volatility, which is consistent with the empirical evidence that corporate bond dealer bid-ask spreads tend to increase during market stress (see Anderson *et al.* (2015)). Here, the market maker's yield adjustment in response to net demand is controlled by an explicit, calibrated parameter λ. The extent to which net demand is passed through into price is also affected by the volatility V observed in the market over the past 20 trading days. Volatility has a calibrated coefficient, v. Noise, ε, is drawn from a normal distribution, $\varepsilon \sim \mathcal{N}(0, N_{\rm ns}^2)$, with $N_{\rm ns}$ the standard deviation. This noisy demand represents the actions of market participants – for instance banks, hedge funds, insurance companies – and trading strategies that are not explicitly modelled.

2.4. *Model summary*

The variables which pertain to each individual entity are shown in Table 3. For the funds, Table 4 shows the variables which change as the model is run. A summary of model parameters, including global parameters such as the bond yield, may be found in Table 5.

Table 4. The simulated variables associated with all fund agent types

Variable	Description	Values	Units	Equation reference
$R_{j,t}$	Fund j return	$R_{j,t} \in \mathrm{R}$	-	(3)
$z_{j,t}$	Fund j demand	$z_{j,t} \in \mathrm{R}$	Bonds	(9)
$W_{j,t}$	Fund j wealth	$W_{j,t} \in \mathrm{R}$	Cash	(5)
$A_{j,t}$	Fund j assets	$A_{j,t} \in \mathrm{R}$	Bonds	
$C_{j,t}$	Fund j cash	$C_{j,t} \in \mathrm{R}_+$	Cash	
$\kappa_{j,t}$	Fund j desired allocation	$\kappa_{j,t} \in (0,1)$	-	(6)
$E_{j,t}(R_{t+1})$	Fund j expected return	$E_{j,t} \in \mathrm{R}$	-	(10)
$h_{j,t}$	Fund j holdings	$h_{j,t} \in (0,1)$	-	(7)
$\hat{h}_{j,t+1}$	Fund j expected next-period holdings	$\hat{h}_{j,t+1} \in (0,1)$	-	(8)
$S_{j,t}$	Fund j investor flow	$S_{j,t} \in \mathrm{R}$	Cash	(4)

3. Empirical foundation and calibration

To assess the success of our model in capturing the properties of trading in corporate bond markets, we compare statistics produced by the model to statistics from empirically observed bond prices. Given typical corporate bond indices assume that coupons and bond repayments are reinvested in the corporate bond market, but are not included explicitly in our model, we construct an empirical corporate bond price index (P) based on the average yield (Y) and duration (D) of US investment grade corporate bonds according to Eq. 1 and using data from the Bank of America Merrill Lynch (BAML).

The returns series appears very sensitive to movements in the risk free rate and shocks to fundamentals. Fig. 3 shows this for a corporate bond price index as constructed above: two of our statistics of interest, the Hill estimator (an estimator of fat tails) and the first order autocorrelation in absolute returns (a measure of volatility clustering), were volatile during the sample period and tend to jump when there were shocks to fundamentals (e.g. during the financial crisis) and when there were changes in the risk-free rate (e.g. during 1999–2001). To avoid these changes, we use empirical properties of the post-crisis sample (since end–2010) only. This comprises 1503 trading days for our data.

We focus on the following empirical properties:

- First-order autocorrelation in returns (ρ^r_{-1}): autocorrelations of the empirical corporate bond returns (based on the post-crisis sample) are not significantly different from zero (Fig. 3). Autocorrelation of returns at longer lags do not provide additional information about the return series and are therefore not included.
- Mean of the absolute returns (ν): following Franke and Westerhoff (2012), we use the mean of the absolute returns as a measure of the overall volatility of the return series.
- Hill estimator of the tail index of absolute returns (γ): following Franke and Westerhoff (2012), we use the Hill estimator to capture the fat tail property of the return series. Arranging the absolute returns in descending order, $\nu_1 \geq \nu_2 ... \geq \nu_T$, and setting $k = 0.05 \cdot T$ (rounded to an integer number), we have

$$\gamma = \frac{1}{k} \sum_{i=1}^{k} \ln \nu_i - \ln \nu_k \tag{12}$$

- First, fifth- and tenth-order autocorrelation in absolute returns (ρ^ν_{-1}, ρ^ν_{-5} and ρ^ν_{-10}): as documented in Cont (2001), the volatility of financial assets can display a positive autocorrelation over several days, suggesting that high-volatility events tend to cluster in time. As shown in Fig. 3, absolute returns of the empirical corporate bond price index can be positively correlated at multiple lags. We match autocorrelations at the first as well as longer lags to capture this persistence. The autocorrelation of the empirical series is particularly large at every four lags, which could be driven by regular coupon payments or outliers. In order to reduce the influence of such effects, our fifth- and tenth-order lag autocorrelations are computed as three-lag averages as in Franke and Westerhoff (2012). We do not look at autocorrelations with more than ten lags which have little marginal effect on our results.

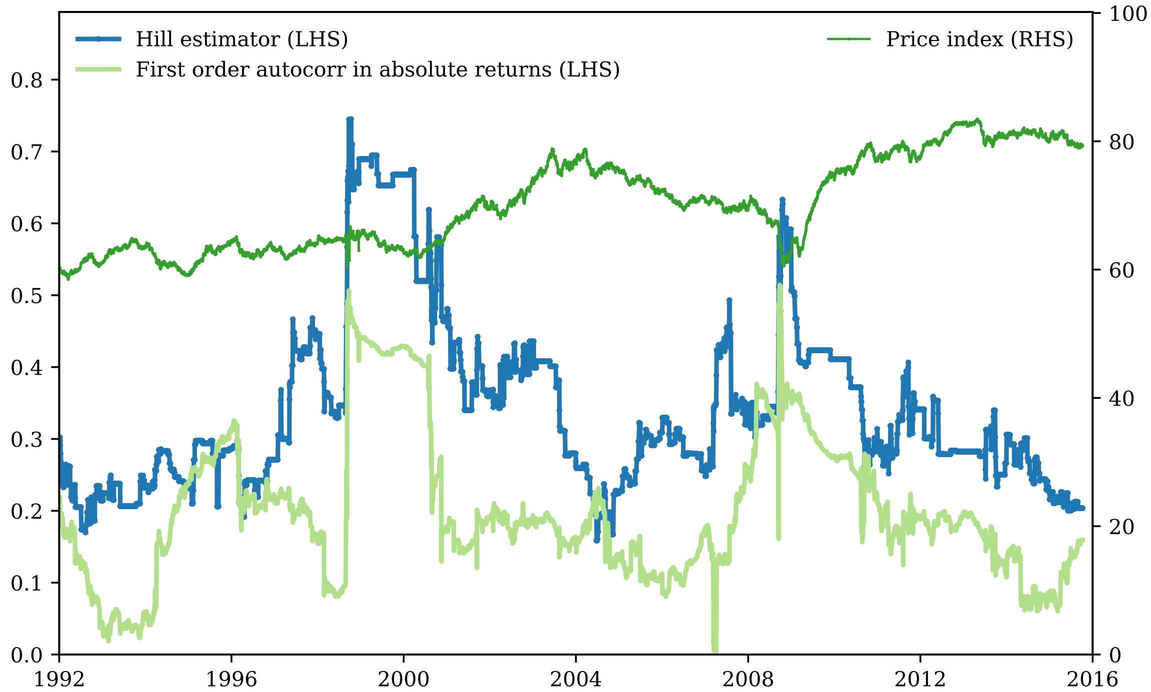

Figure 3. US IG corporate bond price index and rolling moments. Moments are calculated based on a two-year rolling window. Source: Bank of America Merrill Lynch and Bank calculations.

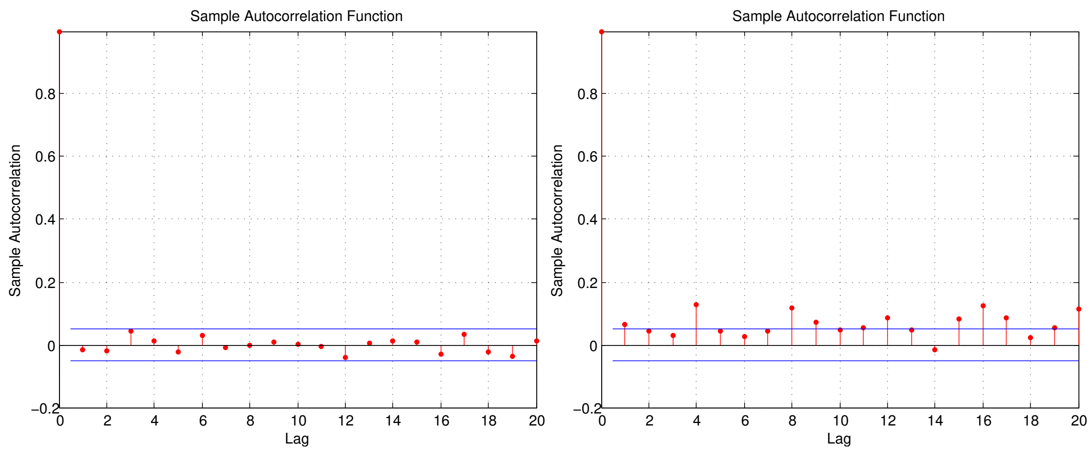

(a) Raw returns

(b) Absolute returns

Figure 4. Autocorrelations of the empirical corporate bond return series. Source: Bank of America Merrill Lynch and Bank calculations.

Model parameters are derived from empirical data where available, as follows:

- Number of agents: there are around 1000 open-ended mutual funds that invest in US corporate bonds (source: Morningstar) which we model with the same number of computational agents.
- Size of agents: based on Morningstar data, the total amount of US corporate bonds held by open-ended mutual funds (i.e. mutual funds that allow daily investment and redemptions, as in the case in our model) is around \$740 billion as at 2015Q2. According to Bank of America Merrill Lynch (BAML) data, there are around \$6.7 trillion of US corporate bonds outstanding as at 2015Q2, of which \$5.3 trillion were investment grade. This suggests that open-ended mutual funds hold about \$585 billion of investment grade corporate bonds. Given that the US corporate bond price index as at

Table 5. Summary of parameters

Empirically Determined Parameters	Value
Number of agents	1000
Size of agents	\$7.3 billion
Proportion of index funds	0.2
Systematic flow strength, s	0.25
Fund-specific positive flow strength, I_+	0.621
Fund-specific negative flow strength, I_-	1.128
Duration of risky asset, D	6.917 years
Fundamental yield, Y^*	1.282% per year
Expected loss rate, L	0.04% per year
Sector-wide redemption window size, τ	1 day
Imposed Parameters	**Value**
Proportion of value traders	0.4
Proportion of momentum traders	0.4
Reaction strength lower bound, L_R	10 days
Reaction strength upper bound, U_R	20 days
Momentum trader long window, t_{lw}	100 days
Risk aversion times return volatility, $\theta\sigma$	1
Face value of bond index, F	100
Calibrated Parameters	**Value**
Market maker sensitivity, λ	0.033
Volatility component, v	20
Value trader strength, α	0.008
Momentum trader strength, β	1.65
Noise level, N_{ns}	0.0266

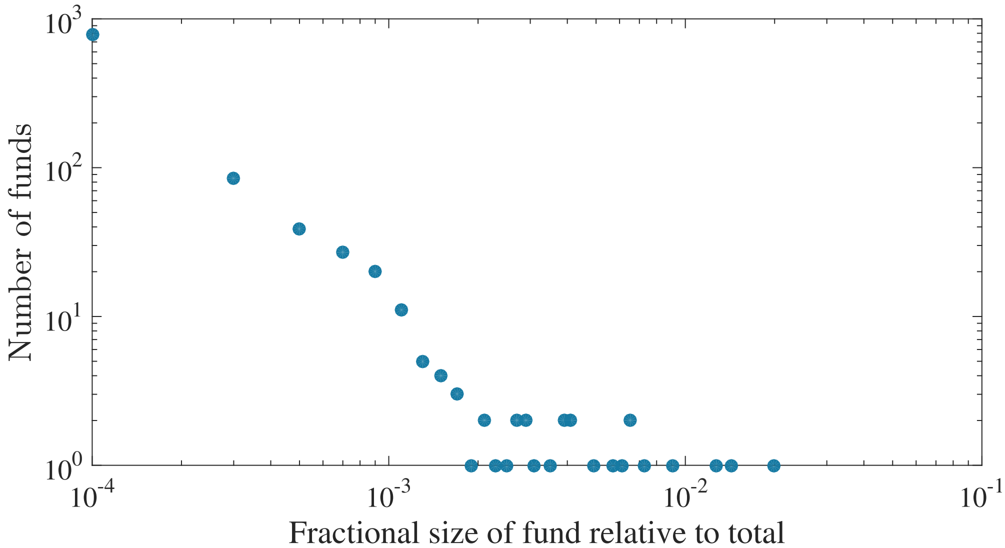

Figure 5. Distribution of fund sizes relative to total sector holdings.

2015Q2 is around 80, this implies that agents hold roughly \$7.3 billion of risky assets.

- Distribution of fund size: the distribution of size of agents in our model is calibrated to the empirical distribution, as shown in Fig. 5
- Proportion of index funds: according to Morningstar data, around 20% of open-ended mutual funds that invest in US corporate bonds were index funds.
- Flow-performance relationship: as discussed, fund flows in our model depend on both market-wide factors, such as the return on the market index, and fund-specific factors, such as fund returns relative to a benchmark.
    - Market-wide factor, s: based on Morningstar data, we estimate a linear relationship between the aggregate monthly net flows of open-ended mutual funds that invest in US corporate bonds and the monthly returns on the BAML US corporate bond index with a one month lag. We find that

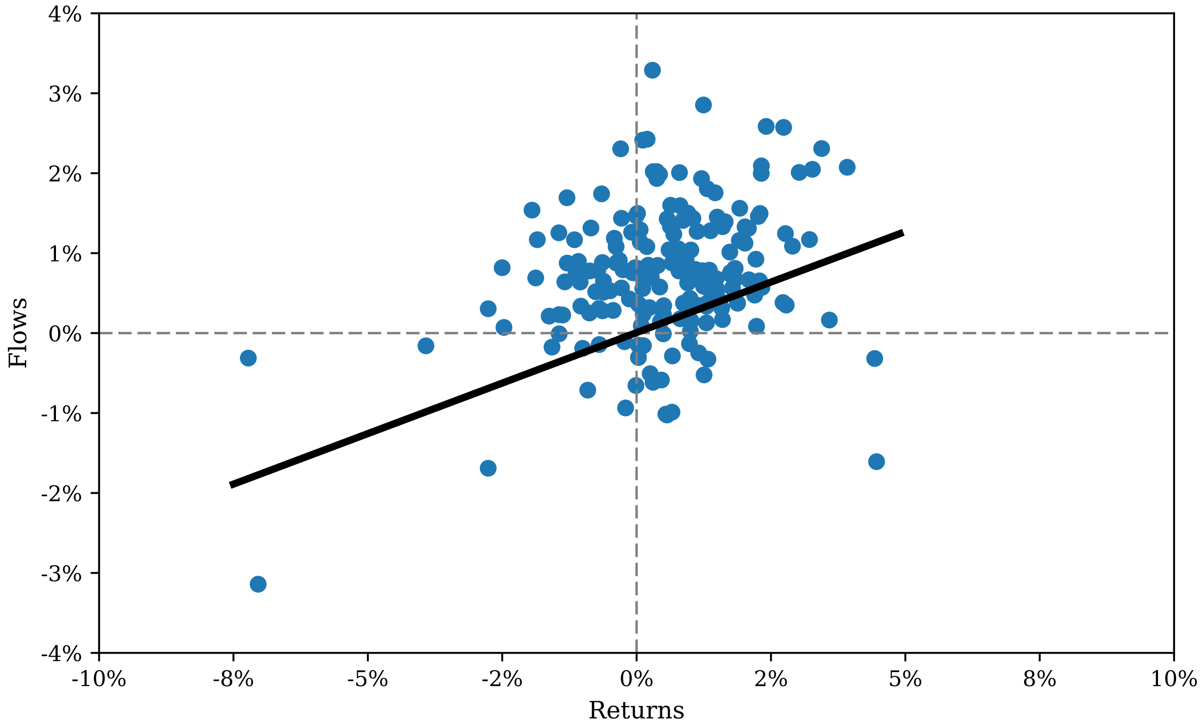

Figure 6. Aggregate net flows v.s. US corporate bond index returns. Source: Morningstar and Bank calculations. The linear relationship is statistically significant at the 1% level.

a 1% return on the index is associated with a 0.25% increase in aggregate net flow, as shown in Figure 6.

- ○ Fund-specific factors, I_+ and I_-: using the empirical findings of Goldstein *et al.* (2015), a fund will have an inflow of 0.621% for 1% excess returns (above the industry average), and an outflow of 1.128% for a -1% excess return.

- Duration, D: the duration of corporate bonds is set to 6.917 years, the average duration of the BAML US corporate bond index as at 2015Q2.
- Fundamental yield, Y^*: the fundamental yield is set to 1.282%, which is the difference between the average yield on BAML US corporate bond index and the 7-year US Treasury yield as at 2015Q2.
- Expected loss rate, L: the annual loss rate L is set to 0.04%, which was the annual credit loss rate of investment-grade corporate bonds in 2010 according to Moody's. When we increase the loss rate via a shock later in the paper, the long-term average yield of the bond also increases by a similar magnitude, which suggests that market price can adequately reflect changes in fundamental value in our model.

While there are qualitative data that suggest that market participants employ technical or trend-following trading strategies and differ in the speed with which they respond to new information, we are not aware of quantitative data on the relative likelihood of value versus momentum strategies and speed of change. We fix these in the model (see Table 5). The model is quite insensitive to the distribution of reaction strengths across funds but it is sensitive to the mean reaction strength, becoming unstable below a certain threshold, typically a few days. Given that funds are unlikely to re-balance their entire portfolio on this time scale we use a mean reaction strength of 15 days and a uniform distribution of reaction strength across funds.

The other parameters in the model (see Table 5) are estimated by matching model outputs to empirical properties using the Moment Coverage Ratio (MCR) proposed by Franke and Westerhoff (2012). The parameter values are optimised using a grid search which runs simulations with different sets of parameters many times iterating toward a set which maximise the MCR. Once the MCR is maximised and the parameters are fixed, we calculate the 95% confidence interval for each moment based on empirical data following the methodology in Franke and Westerhoff. The standard errors of the empirical moments (apart from the

Table 6. Moment confidence intervals

Moment	Name	Mean	Lower bound	Upper bound
ρ^r_{-1}	First-order autocorrelation in returns	-0.0142	-0.0726	0.0443
γ	Hill estimator of the tail index of absolute returns	0.2191	0.1699	0.2684
ν	Mean of the absolute returns	0.1510	0.1430	0.1591
ρ^ν_{-1}	First autocorrelation of absolute returns	0.0652	0.0093	0.1211
ρ^ν_{-5}	Fifth autocorrelation of absolute returns	0.0657	0.0098	0.1216
ρ^ν_{-10}	Tenth autocorrelation of absolute returns	0.0572	0.0087	0.1056

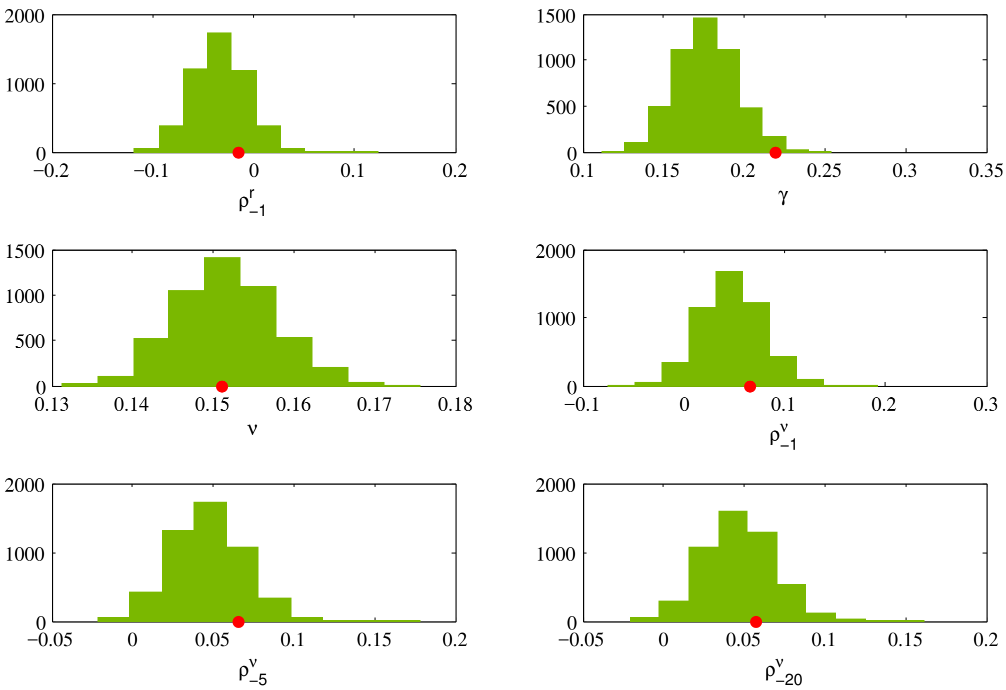

Figure 7. Simulated moments vs empirical moments based on 5,000 Monte Carlo runs. The length of each run is 1,500 trading days, the same as the length of the empirical sample used to estimate moment confidence intervals. The bars represent the distribution of simulated moments and the red dots represent the empirical moments.

Hill estimator which had a closed-form formula) are calculated using the delta method and shown in Table 6. We then run a large number of Monte Carlo simulations of the model and calculate the moments of each simulation. The joint MCR of the model is defined as the proportion of Monte Carlo runs that produce moments that jointly fall within the 95% confidence intervals of the empirical moments.

We plot the simulated moments against the empirical moments in Fig. 7. All empirical moments fall into the distribution of simulated moments. Table 7 shows the MCR of each moment as well as the joint MCR. The joint MCR of the calibrated parameters is 38.9%, which means that our model can generate simulated price paths that are not statistically different from the empirical path in roughly one out of every three runs. The joint MCR is much lower than the individual MCRs, similar to Franke and Westerhoff (2012). This demonstrates that the model and parameters are a reasonable platform for our sensitivity analysis and policy experiments in the following sections.

Moments are particularly sensitive to changes in the market maker's sensitivity to demand. Trading strategy formulation predominantly affects the degree of auto-correlation in raw and absolute returns. The model is fairly robust to changes in the allocation of investor flows.

To ensure that the conclusions are not unduly influenced by the specific choice of model, we also checked that the results were qualitatively similar with price-setting based on market clearing rather than on net demand. This meant replacing equation (11) with a price setting process which finds P_{t+1} such that $(\lambda + \nu V)\left(\sum z_{j,t} + \varepsilon\right) = 0$ given funds' desired allocations. We did not carry out a new calibration for the market-clearing based scheme but used the parameters as determined in the net demand scheme. We tested

Table 7. Moment coverage ratios

Moment	MCR (%)
ρ^{r}_{-1}	92.3
γ	62.4
ν	80.1
ρ^{ν}_{-1}	88.8
ρ^{ν}_{-5}	95.6
ρ^{ν}_{-10}	95.9
Joint	38.9

this with a loss rate shock, a loss rate shock combined with a change to β, and a loss rate shock combined with a change to λ. In each case the model responses were qualitatively the same and quantitatively very similar to the results presented in Section 4.

4. Simulation results

4.1. *Steady state*

The calibrated model, running in the steady state, reproduces stylised facts of the market, including those moments discussed in the previous section and shown in Fig. 7. We define the steady state as the simulation running with no changes to the parameters shown in Table 5 with the yield fluctuating around the long-run mean of Y^*. With noise disabled, the model quickly reaches, and remains at, the fundamental price and yield. While the model run shown in Fig. 8 is typical in that yields vary around this long run mean and do not significantly diverge from it, we also observe model runs where negative feedback loops cause more persistent long run yield increases, suggestive of a tail risk of more pronounced bubbles. This is acceptable as our calibration suggests it is a feature of the mechanics of the market and also because all of the 'experiments' with the model are performed many times so that minor variations in steady state conditions are aggregated over in our results.

We typically find repeating cycles of yield swings and corrections as shown in Fig. 8 for an illustrative single run. These swings reflect the interaction of trading strategies, market maker adjustments and investor flows with normally distributed noise. As the upper row of Fig. 8 demonstrates, such short to medium term yield changes are in line with empirical data.

As further corroboration that the model produces reasonable outcomes, Fig. 9 shows the distribution function of daily log-price returns, $\ln(\frac{P_{t+1}}{P_t})$, from pre-crisis (beginning January 1990), crisis (beginning May 2008), and post-crisis (beginning January 2010) periods as compared with simulated data. The empirical data have larger tails in the log-price return distribution function. This may also be seen in the Hill estimator of the tail index of absolute returns, γ, which is the worst performer of the variables included in the moment coverage ratio. The poor performance of the simulation in capturing tail behaviour is because real returns can be moved by macroeconomic news, while the simulation data shown is a 'steady state' run. As noted in the introduction, the use of non-Gaussian noise would be likely to improve this situation but at the cost of the parsimony of the model. What tails exist in the simulation data shown in Fig. 9 are endogenously generated.

We find in sensitivity testing that the level of volatility is driven strongly by market maker characteristics, while trading strategies predominantly affect the persistence of price dislocations. To ensure that effects are not an artefact of small numbers, we initialise the simulations with a realistic number of firms (1000) and the empirical size distribution (shown in Fig. 5). The steady state is not strongly sensitive to either agent numbers or the shape of the distribution.

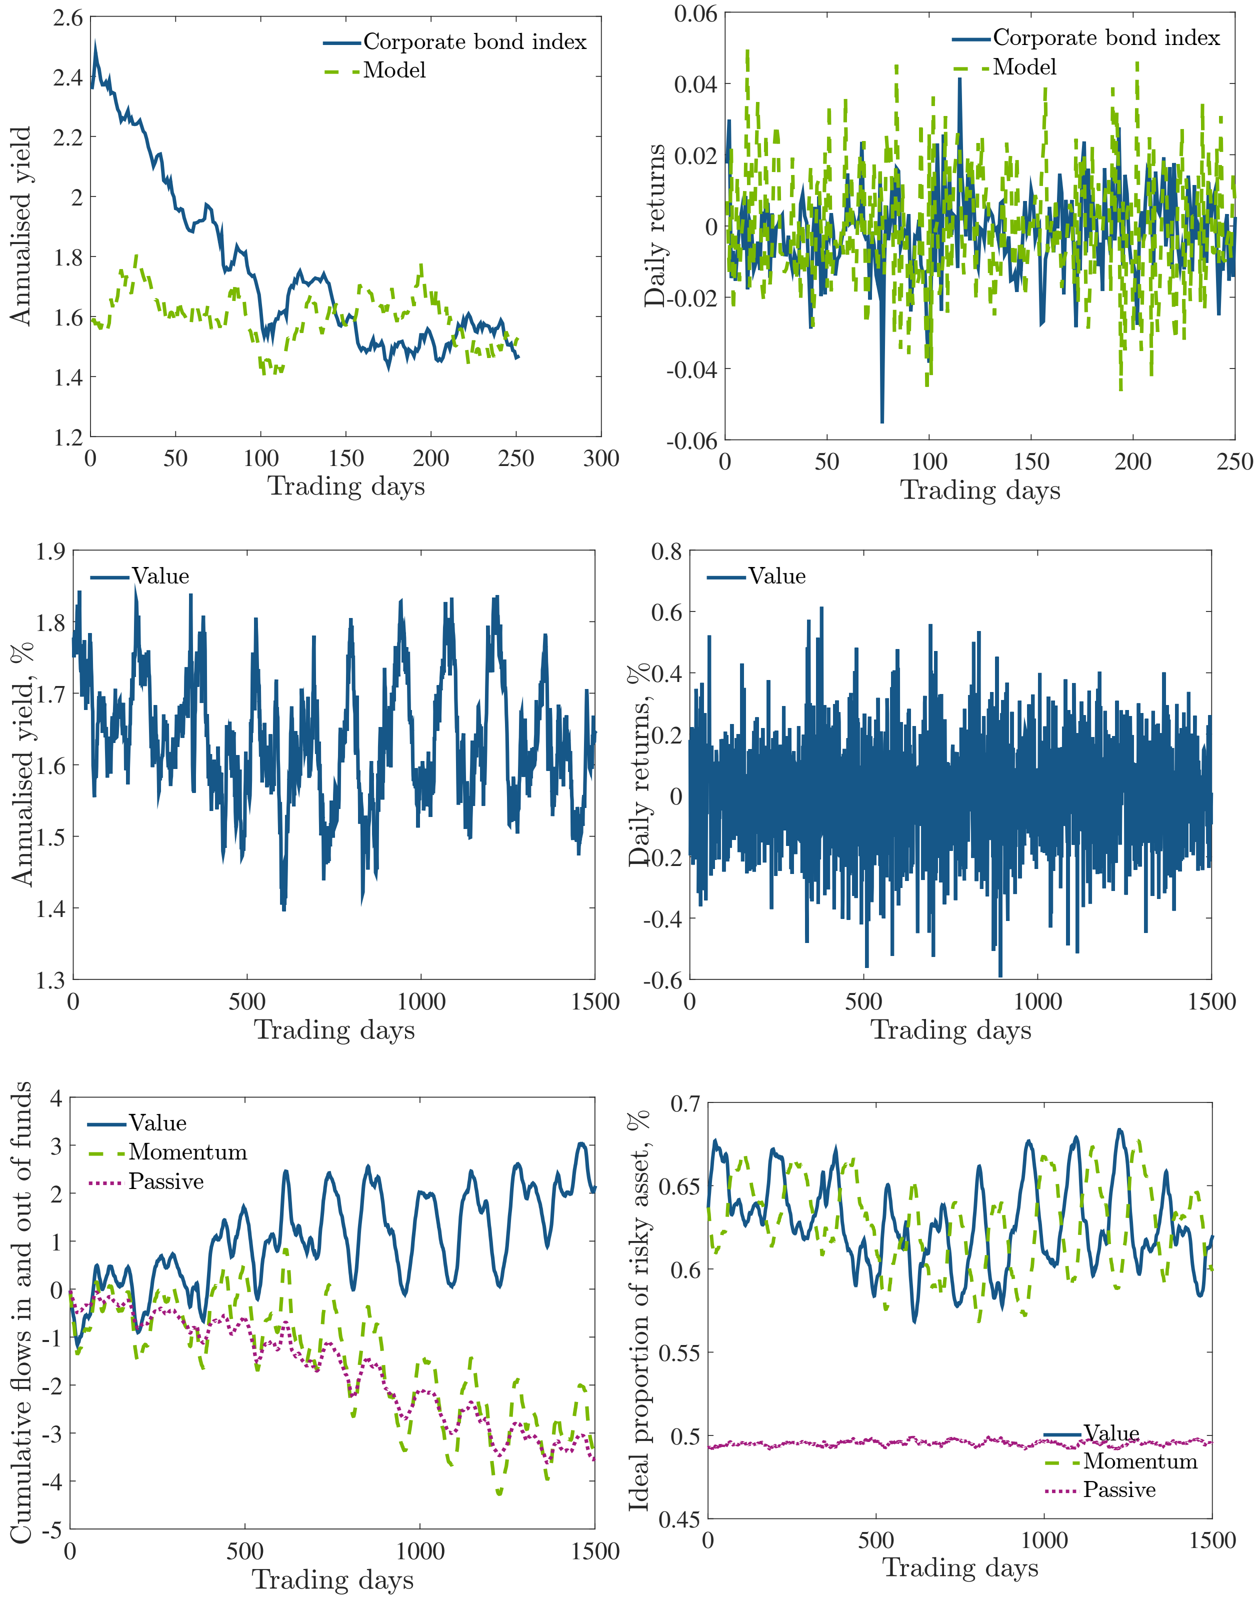

Figure 8. Output from a single model run over 250 trading days; clockwise starting from the upper left panel are shown: yield (annualised) over a trading year, returns over a trading year, returns over a six year period, funds' ideal portfolio composition, cumulative net flows from/to the investor pool and yield (annualised) over a six year period. We contrast results with empirical data for the first two charts, plotting data from an US Investment Grade corporate bond index alongside the model output. Source: Bank of America/Merrill Lynch and Bank calculations.

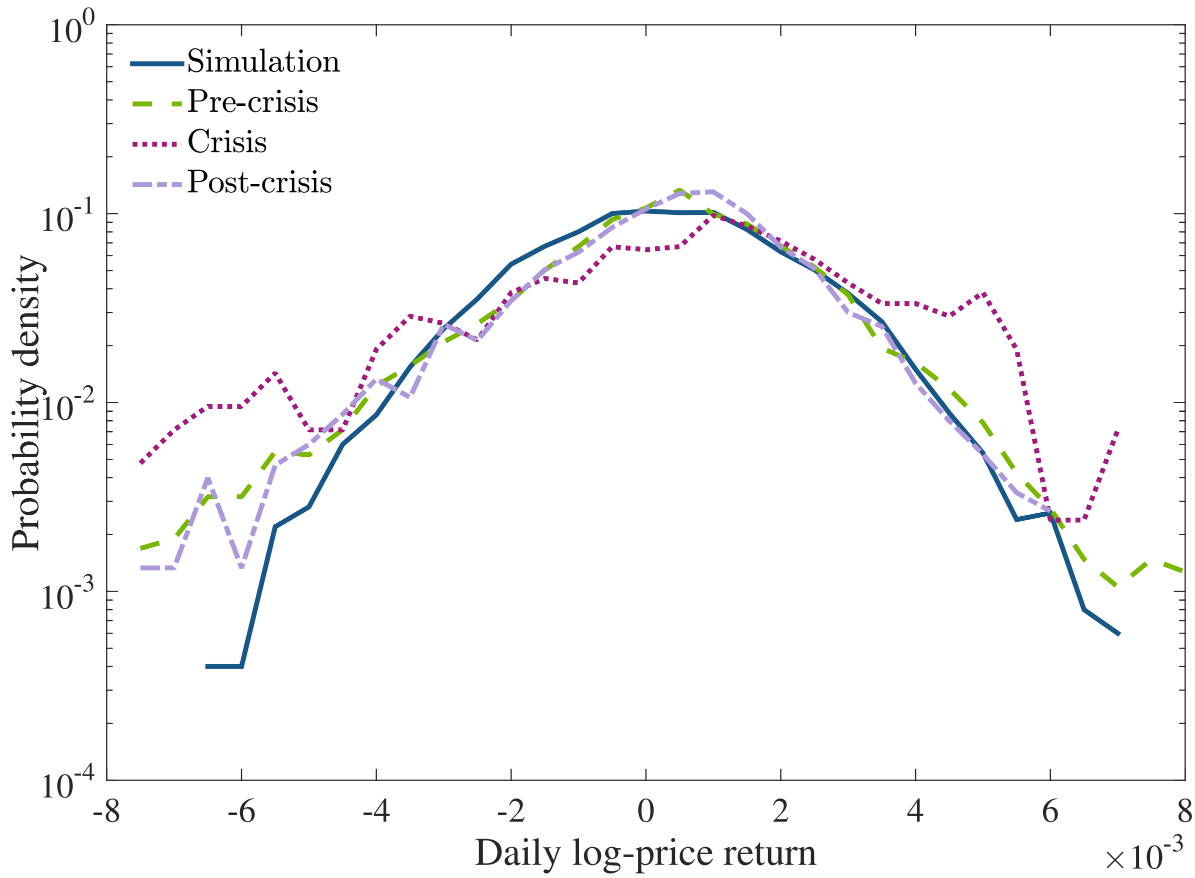

Figure 9. The distribution function of daily log-price returns of three periods shown against data generated by a single model run.

Table 8. Summary of maximum changes in yield under the most extreme parameter value changes.

Parameter	Description	Maximum value simulated (relative to baseline)	Change in yield subsequent to expected loss rate shock for max. parameter value simulated
λ	Market maker sensitivity	1.5×	58%
β	Momentum trader strength	1.0×	54%
s	Systematic flow strength	3.0×	63%
v	Volatility component	1.2×	57%

4.2. *Response to shocks to the expected loss rate under different conditions*

The extent to which endogenous market features either aggravate or dampen downside shocks is particularly interesting from a financial stability perspective. To explore this, we take a given shock to the market and then explore how the market response to that shock changes as other structural parameters of the model are changed from their values as described in Section 3. The variable which we shock is the expected loss rate, L. This loss rate reflects the funds' assumptions about the likelihood of bonds to default. L incorporates information about the probability of default and the rate of recovery into a single number. As a parameter, it appears in the portfolio allocation decisions of funds in Eq. (10) as the factor in $(1-L)$: a larger L, all else being equal, implies that funds will decide to hold fewer bonds. A sudden shock to L means, in this model, that the value of L changes in a single period at time t_s according to

$$L_t = L + H(t-t_s)\cdot(L'-L) \quad (13)$$

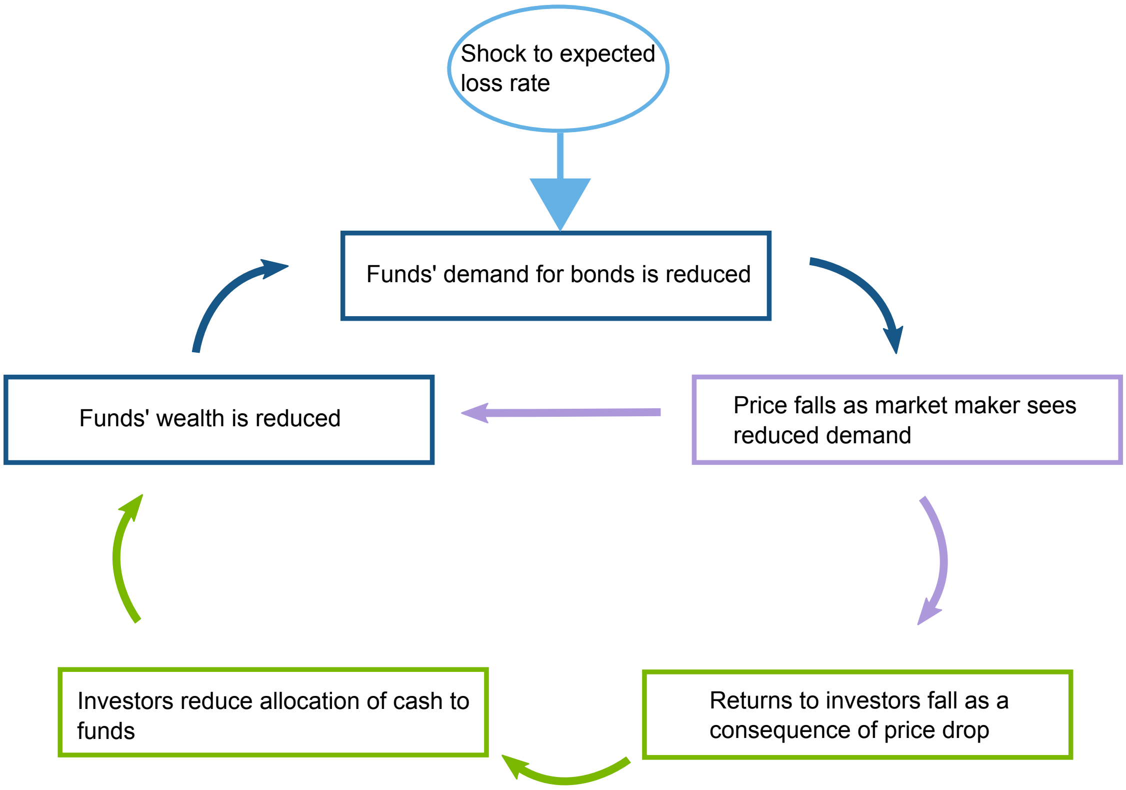

Figure 10. A schematic which shows the feedback loops following a shock to the value of the expected loss rate. The colours in the feedback loop indicate the different market players; funds, the market maker, and the investor pool.

where L' is the post-shock value of the loss rate, t_s is the time at which the shock transition occurs and $H(t)$ is the Heaviside step function defined by

$$H(x) \equiv \frac{\mathrm{d}}{\mathrm{d}x} \max\{x, 0\}$$

The shock transition to L' is in contrast to the steady state situation of section 4.1, in which L remains constant at a value of $L = 0.04\%$ per annum (p.a.) based on data from 2010 (see section 3). We choose as our baseline loss rate shock a final value of $L' = 0.35\%$ per annum.

However, before describing the results of the model with this shock and the other parameters varied in isolation, we describe the general effect of a loss rate shock of varying size. A schematic of what occurs in an arbitrary but positive (ie value-increasing) shock to the expected loss rate value is shown in Fig. 10. An increase in the loss rate to L' changes the funds' desired proportion of risky assets, and so reduces their net demand. In response to this, the market maker reduces prices according to equation (11). A consequence of the price fall is that funds' wealth is reduced, and the price return to investors is lower, potentially triggering investor outflows. This then causes a further reduction in fund wealth. Asset sales by funds prompt further price cuts, and the feedback loop continues. The feedback loop can be broken by the value traders, as their expected return on the asset increases as it becomes more undervalued according to the value traders' equation (10).

Fig. 11 shows the median response over a 100 model runs for simulated loss rate shocks of varying strength. The final loss rates, L', are given by 0.08%, 0.22%, and 0.33% per annum. The median response to a shock is as expected according to the schematic in Fig. 10. Initially, active funds sell the index in order to re-balance their portfolios according to (10) and the market maker drops prices, with fund outflow an exacerbating factor. Eventually, the value traders see the asset as undervalued enough to demand more, rather than less, of it and their demand slows the price fall. Once the price fall is arrested and reversed, momentum traders also reverse their expectations and begin to buy, but this overshoots the new equilibrium yield and

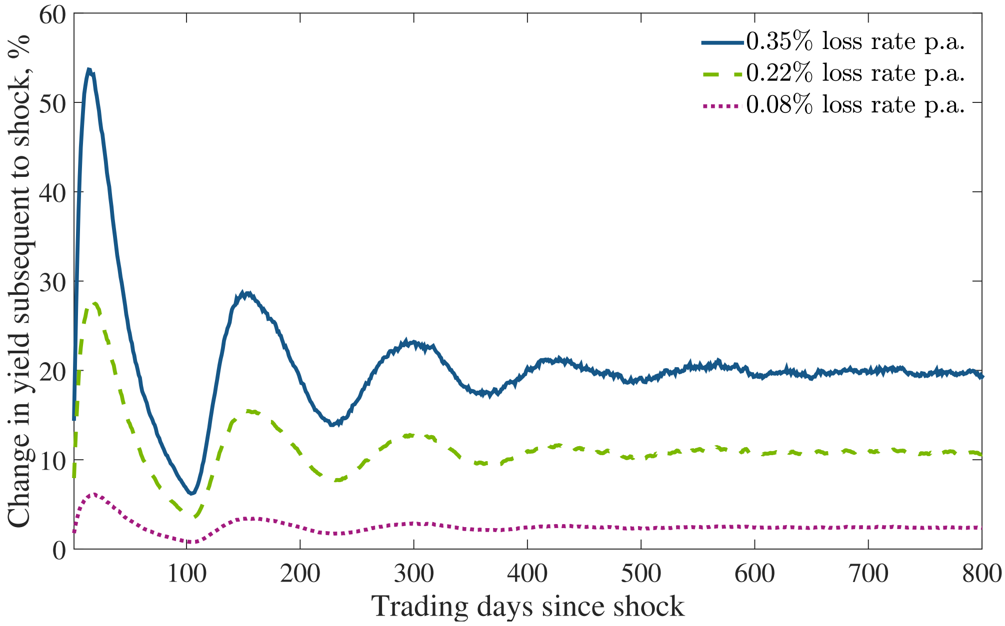

Figure 11. Expected loss rate shocks; a sudden change in firms' expected loss rate, L, causes both short-term fluctuations in yield, and a new, higher equilibrium yield. Results presented are the median over 100 individual simulations runs. The final loss rate is shown in the figure.

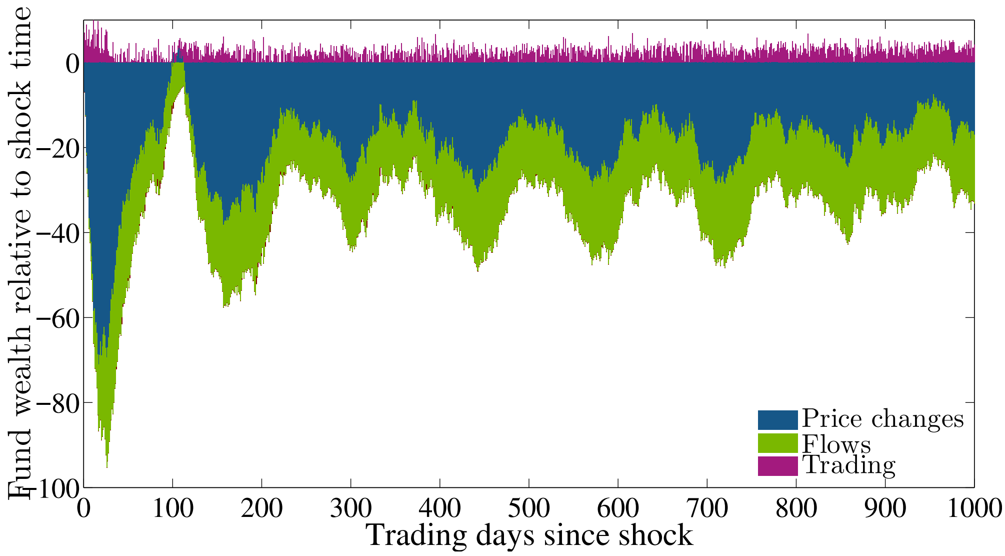

Figure 12. The effect of an expected loss rate shock on wealth averaged over all funds. The decomposition shows how price changes, investor flows and trading all contribute to the post-shock changes in wealth. Results presented are the median over 100 individual simulations runs.

value traders provide stability. Over several cycles, the fluctuations stabilise and the new equilibrium yield is reached.

In Fig. 11, the factor increases in the loss rate (ie L'/L) lead to increases in equilibrium yields of similar magnitude. These results suggest that the pricing of default risk in our model is close to risk-neutral pricing in the long run. Note that the loss rates shown in the figure are the post shock values of L per annum, not the factor increase. The factor increases are $\times 2$, $\times 5.5$ and $\times 8.75$, in ascending order in L'. We also find that the strength of these responses are similar in order of magnitude to empirical data – for the largest loss rate, the response is around one third of the change in yield during the 2008 financial crisis.

As an illustration of the interaction of dynamics, Fig. 12 breaks down the impact of changes on wealth of funds following this shock into the different components, ie mechanical wealth changes as a consequence of price changes, investor flows and trading. This shows, for example, the delay in onset of investor flows: as investors base their investment decisions on the previous month's returns, outflows are delayed and fluctuate over time, prolonging the adjustment process. Funds benefit from active trading in aggregate, as the profits and losses of value traders and momentum traders partially offset each other. We can capture this as the residual after price movements and investor flows have been taken into account. These feedback effects are broadly similar to what occurs in response to a shock to funds' wealth, which could be triggered, for example, by a reassessment by investors of the liquidity risk inherent in holding fund assets or news on fund management. Such a wealth shock is applied by removing a fraction f of each fund's wealth as a proportion of total sector asset holdings so that $W_t = W_{t-1}(1-f)$. Funds meet these outflows through a combination of asset sales and cash in the same proportion as their holdings. An example where a wealth shock has been applied is shown in Fig. 18. The fall in demand by funds prompts an increase in yield exacerbated by momentum traders and further investor outflows responding to associated poor fund performance. As with the expected loss rate shock, overshoots correct over time, but it is a prolonged process with several cycles of overshooting in both directions.

We now turn to scenarios in which the loss rate shock such that $L = 0.04\%$ p.a. changes to to $L = 0.35\%$ p.a. is applied, but where other parameters of the model, such as λ, are different from their calibrated values throughout the simulation. This gives an indication as to how important the parameters of the model are in the market's response to the onset of a crisis. In particular, we investigate changes to the market maker sensitivity, to trading strategies, and to the strength of investor response. The maximum change in yield subsequent to the expected loss rate shock under the most extreme values of the parameters simulated are summarised in Table 8. In the simulations, the same noise is used for the different values of the parameters so that the results are genuinely comparable. In many real world scenarios, it might be the case that two or more parameters could change concurrently (for example changes in the market maker's sensitivity to demand and the fraction of passive funds might well be correlated). These concurrent changes in structural model parameters are not analysed but would give rise to extra interaction effects. Note that the increase from $L = 0.04\%$ p.a. to $L = 0.35\%$ p.a. is of a comparable order of magnitude to the annual loss rate changes between 2007 and 2008.

Fig. 13 shows the impact of different starting values for one component of the market maker's sensitivity to demand, λ. This affects both the amplitude of the post-shock yield fluctuations and the longer-run level of average post-shock yields. The change in the amplitude of fluctuations is because λ amplifies the effects of net demand on price levels. Longer run increases of average yield reflect fund losses due to increased volatility, which affects their demand persistently.

Market data suggest that market maker responses to demand changes are not uniform over time. Rather, bid-ask spreads widen substantially in the face of increased uncertainty. As our model does not include a bid-ask spread, we proxy this through making the market maker response also linearly dependent on volatility over the previous 100 trading days (see Eq. (11)). When we increase the factor with which volatility is included, we observe that it increases both the extent of overshooting and the speed with which yield trends reverse (Fig. 16). Similarly to λ, the model is relatively sensitive to changes in this parameter.

Changes to α, the value trader strength, have very limited impact on market dynamics close to the baseline value. For significant increases (i.e. for values tens of times higher than the baseline model), model overshoots are dampened somewhat, though the long run outcome is not affected.

In contrast, we find that the strength of sensitivity of momentum traders, β, strongly affects the model response. As Fig. 14 shows, there is no overshoot without momentum trading, reducing short- to medium-term yield dislocations significantly. The contribution to price dislocation caused by momentum traders is well-known from other models, and has been discussed in De Long *et al.* (1990), Scharfstein and Stein (1990), Hong and Stein (1999) and Hommes (2006). Here, we find the same effect when momentum traders operate in the bond market and use yield as the basis for their decisions.

We also considered the impact of fund flows. At the levels suggested by empirical data, investor flows contribute to the impact of shocks through additional wealth losses, leading to larger swings and longer-term yield increases, but are relatively weak compared to both trading decisions and market maker action. Even tripling fund outflows, as shown in Fig. 15, only increases maximum yield by around ten percent, and longer run yields by around one to two percent. In our model, market-wide fund flows play a similar

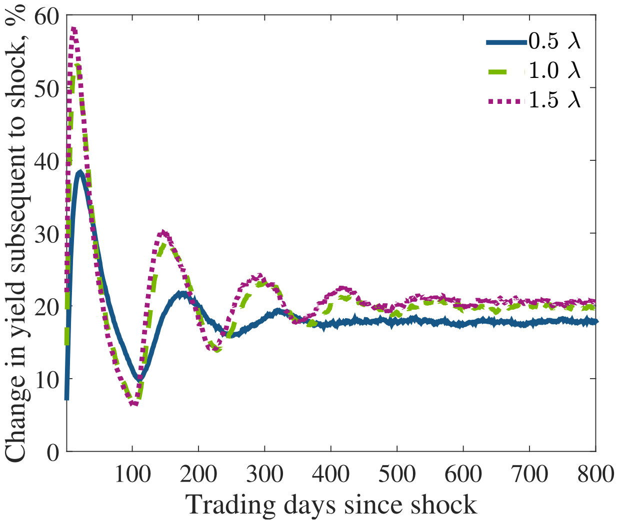

Figure 13. Yield change with a 0.35% p.a. loss rate shock and different market maker sensitivities to net demand (defined by equation (11)) relative to the base line value of λ.

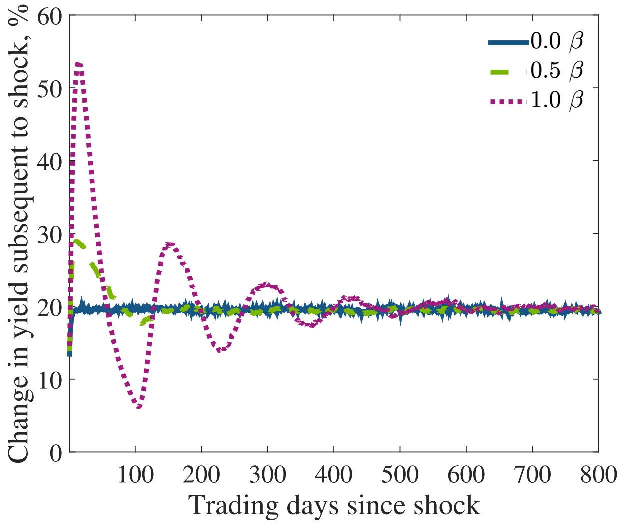

Figure 14. Yield change with a 0.35% p.a. loss rate shock and different momentum trader strengths (defined by equation (10)) relative to the base line value of β.

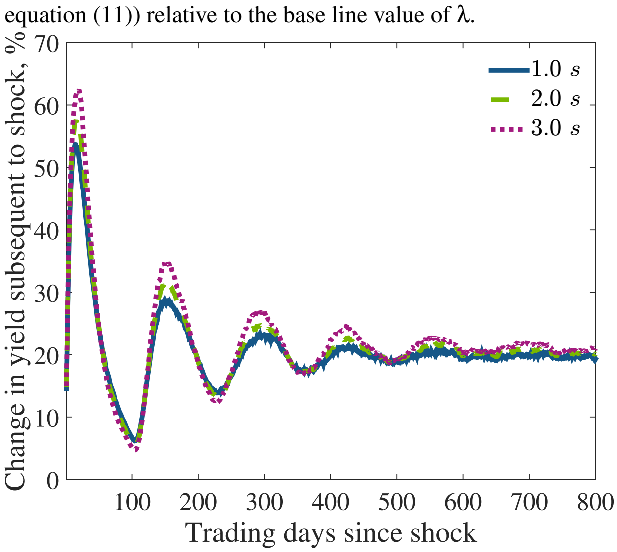

Figure 15. Yield change with a 0.35% p.a. loss rate shock and different systematic flow strengths (defined by equation (4)) relative to the base line value of s.

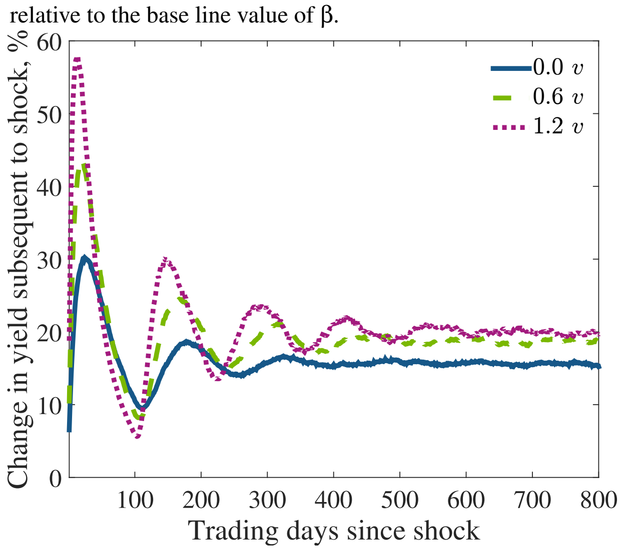

Figure 16. Yield change with a 0.35% p.a. loss rate shock and different market maker volatility coefficients (defined by equation (11)) relative to the base line value of v.

role to momentum traders, because they both result in greater holdings (in absolute terms) of the risky asset following price rises, and vice versa. A key difference is that market-wide fund flows also go to value traders, who may trade against the trend.

4.3. *Scenario: increased fraction of passive funds under stressed conditions*

Motivated by the trend in which funds increasingly use passive investment strategies, we assessed what occurs under the established shocks to loss rate, as defined in equation (13), and a shock to wealth as described in section 4.2, when the fraction of funds using passive investment strategies is higher than in the baseline calibration. The wealth shock used is a sudden need to reduce by 5% the total wealth of each fund for all funds.

Quantitatively, the impact of the increased presence of passive funds is dependent on the shock in question. Figs. 17 and 18 show this via the median change in yield (where the median is taken over the 100 trading days following the shock). For each shock, the scenario is repeated 250 times.

In the case of the loss rate shock, the presence of the passive funds is, overall, stabilising and, in most

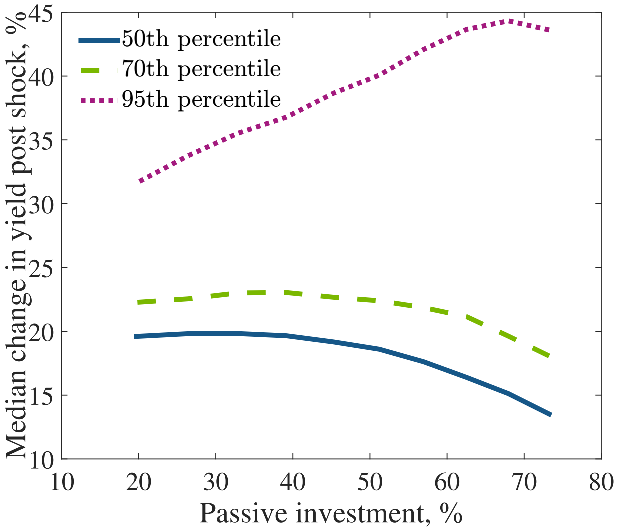

Figure 17. Distribution of outcomes for median yield over 100 trading days after a 0.36% increase in loss rate shock: percentiles indicate the value taken from 250 runs.

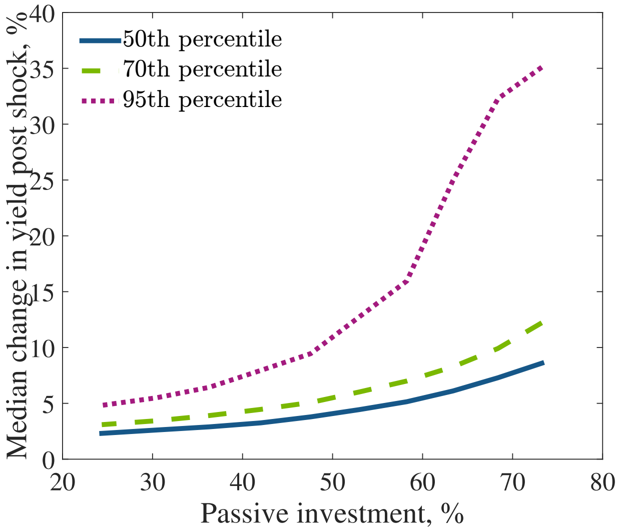

Figure 18. Distribution of outcomes for median yield over 100 trading days after a 5% fund wealth shock: percentiles indicate the value taken from 250 runs.

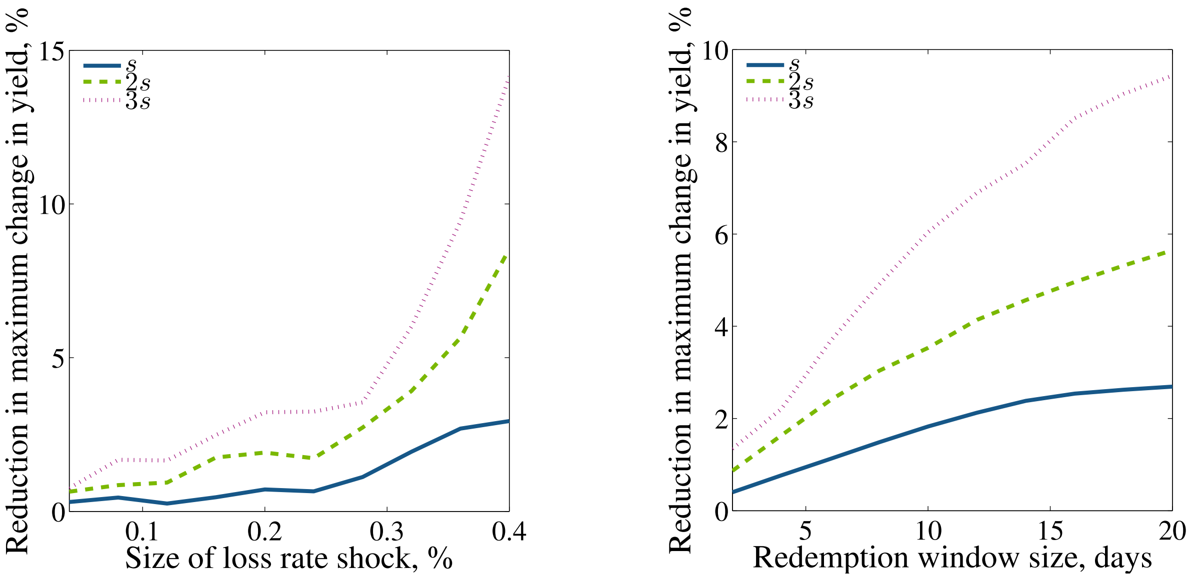

Figure 19. When flow strength, s, is large, redemption management can significantly reduce the magnitude of the maximum change in yield following a large loss rate shock. The redemption management policy splits any redemption over 20 trading days in the first panel, ie $\tau = 20$. The loss rate shock is 0.35% in the second panel.

case, lowers the post-shock median yield as shown in Fig. 17. This is the expected result: in contrast to momentum traders and value traders, passive investors do not trade based on the news about the loss rate and so the market impact of loss rate changes is smaller with the increased fraction of passive funds.

The situation is different for the wealth shock, as shown in Fig. 18. In this case, all funds must sell in order to achieve their preferred portfolio positions. However, in contrast to the loss rate shock, the median outcome for yield with an the increased fraction of passive investors is a higher post-shock median yield.

Under both types of shock, there is an unexpected effect: the extent of tail risk, here given by the possibility of larger post-shock median yield, is increased substantially by the increased presence of passive funds. This is shown in Figs. 17 and 18 by the lines representing the 70th and 95th percentile of outcomes from the 250 runs. In the case of the loss rate shock, which does not change passive investors behaviour, this result is particularly surprising.

For both types of shock, the distribution of outcomes broadens significantly as the level of passive investment increases. In both cases, the driver of this effect is the net demand in the market.

Clearly, this will be more negative on average in the wealth shock case, as all funds are forced to sell. For the wealth shock, part of the increase in tail risk at higher levels of passive investment is because the active funds seek to rebalance their portfolio (meeting some of their short fall with cash), while passive funds seek to divest their lost wealth completely out of the risky asset.

In the median outcome to a loss rate shock, the weaker response of yield is due to there being fewer active traders who need to rebalance their portfolios. But because the sizes of the active funds are drawn from the fat-tailed empirical distribution shown in Fig. 5, there is a chance that the sum of demand that the market maker observes from the remaining active funds attains a high value. The result is the increased tail risk seen in Fig. 17.

At current levels of passive fund representation (as noted previously, passive funds hold around 20% of assets), these yield swings are less severe. But there are indications that a substantial proportion of nominally actively managed funds adopt investment strategies that are passive in practice. The fraction of passive funds may therefore be considerably higher than the headline figure suggests.

4.4. *Policy experiment: redemption management*

Using the scenario of heightened investor sensitivity in the light of a change to expected loss rates, we can use our model to explore the impact of action to reduce the speed with which investor redemption requests are fulfilled. In our model, and reflecting common practice for open-ended funds, investors can redeem their holdings daily such that the redemption window, $\tau = 1$. This reduces liquidity risk for investors, enabling them to realise their investment at its market value at short notice, but can be costly, if it forces funds to sell into an adverse market environment and prompts the type of feedback loops discussed above.

Funds can address this in a number of ways, dependent on the jurisdiction within which they operate. They can negotiate informally with investors requesting a large redemption to spread payment over several days or weeks. They could implement redemption constraints or gates when faced with large outflows. They could raise liquidity for anticipated redemptions ahead of time. And they could change the speed of redemption offered more generally, providing repayment on a slower time table.

We introduce a policy to manage the risk outflows in the model as a simple linear change to the speed of payout. Rather than receiving the full amount on day 1, payments to investors are spread equally over a number of days so that $\tau > 1$. A fund j meets its redemption in the usual way, splitting it between the risky asset and cash in the proportion given by its current proportion of holdings of the risky asset $h_{j,t}$. However, the fund now spreads the sale of bonds over all of the days it is paying out rather than selling them on a single day, thus reducing its own impact on the price on the first day. In steady state, this has no effect on model behaviour, reflecting relatively slow responses by investors. While we see yield fluctuations (see Fig. 8), these tend to happen over a matter of weeks or months – longer than we might expect funds to spread outflows.

This changes, however, if we focus on a stressed system. While, even then, a lower payout speed has a fairly marginal impact in the base case, the policy is effective in reducing yield overshoots when the shock and the change in investor outflows triggered by the shock are both large. The first panel in Fig. 19 shows that the redemption management policy can reduce the maximum change in yield following a loss rate shock by more than 10% when the size of the shock is large (above 0.35% p.a.) and the strength of investor flows is three times larger. The second panel shows that the policy becomes more effective when redemption is spread across a longer window, but that the marginal benefit of an extra day decreases as the window size becomes larger. This is because the beneficial effects of the redemptions must ultimately be limited: in the long-run, a change to the loss rate will necessarily change the price which the risky asset is traded at. The policy can only reduce the extent of price dislocation following a shock but this entirely what it is designed to achieve and it should not interfere with the fundamental value.

5. Conclusion

We have described the first agent-based and empirically calibrated model of the corporate bond market with both active and passive traders. Statistics (or stylised facts) produced by the model match data from the US investment grade corporate bond index well. We have shown that the sensitivity of the market maker to demand and the degree to which momentum traders are active are important for the response of the simulated bond market to shocks, especially with respect to the amplitude of yield (and price) dislocations to the long term steady state values. The effect on price dislocation of the presence of momentum traders is well-established in the literature but here it is directly applied to a new market and a new, though related, metric – yield.

Combining these well-known elements of agent-based models of trade allowed us to analyse two scenarios of genuine importance for the market for bonds.

In one of the scenarios, we demonstrated that a higher fraction of passive funds has the intuitive effect of reducing the median extent of price and yield dislocations in response to a shock to the loss rate. However, we also find that it can increase the tail risk of larger dislocations considerably, for both shocks to the loss rate on bonds and more general shocks to the wealth of funds. This is important as the fraction of passive funds has been increasing in this market. Additionally, some market participants that are not explicitly modelled are likely to behave like passive investors. Finally, the rate of change of the tail risk of high dislocation is high at the current level of passive investment.

The second scenario looked at how the exacerbating effect of investor outflow (due to a first mover advantage) on yield overshoots could be reduced. Simulations showed that managing redemptions by spreading them over longer time periods can reduce the maximum extent of yield overshoots when investor outflows are unusually large, as in the 2008 financial crisis.

These findings highlight the importance of the interaction between trading behaviour and market maker response, suggesting that measures to reduce the risk of momentum trading in a shock may help mitigate the risk of overshoots during price adjustments. These could include ensuring the availability of information on assets to strengthen trading based on fundamentals, or mechanisms to control the speed of trading such as temporary suspensions.

More generally, this approach to modelling fixed income markets may prove useful in assessing systemic risk and mitigating policy measures more broadly. It offers a straightforward way of integrating data on market participant behaviour and assessing its interaction with market microstructure. Explicitly incorporating typical buy and hold investors in the insurance and pension fund sector and leveraged entities such as hedge funds would be natural extensions.

6. Acknowledgements

We are grateful to anonymous referees, attendees of seminars at the University of Oxford, University College London, Scuola Superiore Sant'Anna and the Bank of England, and Nicola Anderson, Tomaso Aste, David Bholat, Matthew Everitt, Doyne Farmer, Marco Galbiati, Andy Haldane, Michael Hume, Sujit Kapadia, Pippa Lowe, Joseph Noss, David Pugh, Gayathri Sundar and Nicholas Vause for help and comments.

References

Anderson, N., Webber, L., Noss, J., Beale, D. and Crowley-Reidy, L., The resilience of financial market liquidity. *Bank of England Financial Stability Paper*, 2015.

Aymanns, C. and Farmer, J.D., The dynamics of the leverage cycle. *Journal of Economic Dynamics and Control*, 2015, **50**, 155–179.

Bank of England, Financial Stability Report December. Technical report, Bank of England, 2015.

Bookstaber, R., Foley, M.D. and Tivnan, B., Market Liquidity and Heterogeneity in the Investor Decision Cycle. *Office of Financial Research Working Paper Series*, 2015.

Cetorelli, N., Duarte, F. and Eisenbach, T., Are Asset Managers Vulnerable to Fire Sales? [blog post]. Liberty Street Economics Retrieved 29th December 2016, 2016.

Chen, Q., Goldstein, I. and Jiang, W., Payoff complementarities and financial fragility: Evidence from mutual fund outflows. *Journal of Financial Economics*, 2010a, **97**, 239–262.

Chen, Y., Ferson, W. and Peters, H., Measuring the timing ability and performance of bond mutual funds. *Journal of Financial Economics*, 2010b, **98**, 72–89.

Chen, Y. and Qin, N., The Behavior of Investor Flows in Corporate Bond Mutual Funds. *Available at SSRN 2022059*, 2014.

Christoffersen, S.E., Musto, D.K. and Wermers, R., Investor flows to asset managers: Causes and consequences. *Annu. Rev. Financ. Econ.*, 2014, **6**, 289–310.

Comerton-Forde, C., Hendershott, T., Jones, C.M., Moulton, P.C. and Seasholes, M.S., Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues. *The Journal of Finance*, 2010, **65**, 295–331.

Cont, R., Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 2001, **1**, 223236.

Coval, J. and Stafford, E., Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 2007, **86**, 479–512.

De Long, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J., Noise trader risk in financial markets. *Journal of political Economy*, 1990, pp. 703–738.

Durham, J.B., Can Investors Use Momentum to Beat the U.S. Treasury Market? [blog post]. Liberty Street Economics Retrieved 29th December 2016, 2016.

Farmer, J.D. and Foley, D., The economy needs agent-based modelling. *Nature*, 2009, **460**, 685–686.

Farmer, J. and Joshi, S., The price dynamics of common trading strategies. *Journal of Economic Behavior & Organization*, 2002, **49**, 149–171.

Fender, I. and Lewrick, U., Shifting Tides - Market Liquidity and Market-Making in Fixed Income Instruments. *BIS Quarterly Review March 2015*, 2015.

Fischer, T. and Riedler, J., Prices, debt and market structure in an agent-based model of the financial market. *Journal of Economic Dynamics and Control*, 2014, **48**, 95–120.

Franke, R. and Westerhoff, F., Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics and Control*, 2012, **36**, 1193–1211.

Goldstein, I., Jiang, H. and Ng, D., Investor flows and Fragility in Corporate Bond Funds. *unp*, 2015.

Hommes, C.H., Chapter 23 Heterogeneous Agent Models in Economics and Finance. *Handbook of Computational Economics*, 2006, pp. 1109–1186.

Hong, H. and Stein, J.C., A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, 1999, **54**, 2143–2184.

Immonen, E., Simple agent-based dynamical system models for efficient financial markets: Theory and examples. *Journal of Mathematical Economics*, 2017.

Jostova, G., Nikolova, S., Philipov, A. and Stahel, C.W., Momentum in corporate bond returns. *Review of Financial Studies*, 2013, p. hht022.

Kyle, A.S., Continuous Auctions and Insider Trading. *Econometrica*, 1985, **53**, 1315–35.

Kyle, A.S. and Obizhaeva, A.A., Market Microstructure Invariance: Empirical Hypotheses. *Econometrica*, 2016, **84**, 1345–1404.

Lütje, T., To be good or to be better: asset managers' attitudes towards herding. *Applied Financial Economics*, 2009, **19**, 825–839.

Moneta, F., Measuring bond mutual fund performance with portfolio characteristics. *Journal of Empirical Finance*, 2015, **33**, 223–242.

Navigating Monetary Policy Challenges and Managing Risks, Global Financial Stability Report. Technical report, International Monetary Fund, 2015.

Ponta, L., Pastore, S. and Cincotti, S., Information-based multi-assets artificial stock market with heterogeneous agents. *Nonlinear Analysis: Real World Applications*, 2011, **12**, 1235–1242.

Scharfstein, D.S. and Stein, J.C., Herd behavior and investment. *The American Economic Review*, 1990, pp. 465–479.

Shek, J., Shim, I. and Shin, H.S., Investor redemptions and fund manager sales of emerging market bonds: how are they related?. *BIS Working Papers*, 2015, **509**.

Tesfatsion, L., Chapter 16 Agent-Based Computational Economics: A Constructive Approach to Economic Theory. *Handbook of Computational Economics*, 2006, pp. 831–880.

Thurner, S., Farmer, J.D. and Geanakoplos, J., Leverage causes fat tails and clustered volatility. *Quantitative Finance*, 2012, **12**, 695–707.

Turrell, A., Agent-based models: understanding the economy from the bottom up. *Bank of England Quarterly Bulletin*, 2016, **56**, 173–188.