



BANK OF ENGLAND

Staff Working Paper No. 592

An agent-based model of dynamics in corporate bond trading

Karen Braun-Munzinger, Zijun Liu and Arthur Turrell

April 2016

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Authority Board.



BANK OF ENGLAND

Staff Working Paper No. 592

An agent-based model of dynamics in corporate bond trading

Karen Braun-Munzinger,⁽¹⁾ Zijun Liu⁽²⁾ and Arthur Turrell⁽³⁾

Abstract

We construct a heterogeneous agent-based model of the corporate bond market capturing the interaction of market maker behaviour, fund trading strategies, and cash allocation by investors in funds to study feedback effects and the impact of market changes. The model parameters are calibrated against empirical data on US corporate bond trading. Where available, inputs are taken from market data. Others are calibrated through matching statistical features of market returns such as auto-correlations, volatility and fat tails. We use the model to explore the impact of shocks. We find that the sensitivity of the market maker to demand and the degree to which momentum traders are active strongly influence the over and undershooting of yields in response to a shock. This suggests that correlation in funds' trading strategies can exacerbate extreme price movements and contribute to the procyclicality of financial markets. While the behaviour of investors in funds based on past experience plays a comparatively smaller role in model dynamics, it represents another source of amplification which could be particularly problematic if investors respond to a shock with greater risk aversion. Simple measures to reduce the speed with which investors can redeem investments can reduce the extent of yield dislocation. We also explore the impact of the growth in passive investment, and find that it increases the tail risk of big yield dislocations after shocks, though, on average, volatility may be reduced.

Key words: Agent-based model, corporate bond market, trading strategies.

JEL classification: C63, G11, G12, G17.

(1) Bank of England. Email: karen.braun-munzinger@bankofengland.co.uk

(2) Bank of England. Email: zijun.liu@bankofengland.co.uk

(3) Bank of England. Email: arthur.turrell@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful to Nicola Anderson, 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.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

1 Introduction

The ability of financial markets to absorb shocks depends on factors such as market micro-structure, vulnerabilities of the entities operating within them and behavioural choices of market participants. The importance of interaction of different market participants and evidence that their decision-making may not always be fully rational suggest that agent-based computational economics may be a useful modelling approach. This is reflected in a number of recent papers that use such methods to explore market dynamics in a stress (see, for example, Bookstaber, Foley, and Tivnan [2015]). While these papers have underlined the importance of leverage and risk management choices (see also Thurner, Farmer, and Geanakoplos [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. This is particularly problematic given 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 IMF [April 2015].

To fill that gap, we use a heterogeneous agent approach to model a fixed income market with mutual funds, incorporating payoff structures, reduced likelihood of short selling and slower speed of trading. This allows us to assess how mutual funds' trading strategies, the speed of yield adjustments by the market maker and the behaviour of end investors interact. We can 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. In doing so, we draw on small-scale heterogeneous agent trading models as summarised in Hommes [2006] and Tesfatsion [2006] and 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].

Our model differs from those papers in a number of ways. Here, 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 arguably more predictable with typically positive non-zero payoffs on indices. We reflect these aspects in 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, Shim, and Shin [2015], who demonstrate that discretionary sales follow forced sales for emerging market economies 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], Chen, Ferson, and Peters [2010b]. We add to the factors included in the market maker decision rule introduced by Farmer and Joshi [2002] to reflect decreasing willingness of market makers to bear risk in uncertainty by making the speed of adjustment dependent on past volatility.

Finally, as in Thurner et al. [2012], we include a stylised investor pool, but incorporate theoretical and empirical literature (see in particular Chen, Goldstein, and Jiang [2010a], Goldstein, Jiang, and Ng [2015], 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, ie investors respond more strongly to negative performance.

The model does not incorporate endogenously determined wider market factors such as interest rates, macroeconomic factors affecting the value of the bonds or alternative investment opportunities and is thus not suited well to forecasting future market movements. We impose shocks exogenously to investigate the role of market structure and practices (such as funds redemption policies) in dealing with such changes, however.

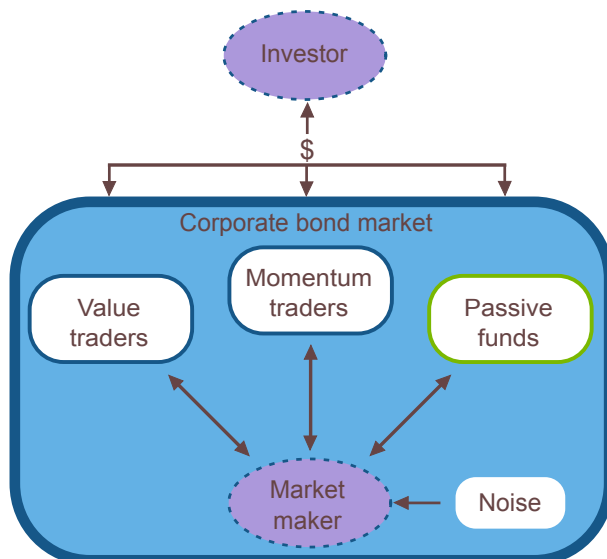


Figure 1: Stylised model setup.

2 The model

We propose a simple model of heterogeneous funds trading a corporate bond index via a stylised over-the-counter market maker. These funds are subject to cash in- and outflows from a stylised investor pool (Fig. 1). The index is made up of zero coupon bonds with constant average duration - this is the risky asset in the model. Funds can also hold cash. Payoffs from the asset are passed directly to the investor pool.

Each trading period t representing a trading day, 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 desired next period's returns divided by risk aversion γ and asset variance σ

$$\kappa_{j,t} = \frac{E_{j,t}(R_{t+1})}{\gamma\sigma} \quad (1)$$

In contrast to previous models, we do not assume that funds adjust their portfolio completely at every period. Instead, fund j adjusts with a reaction strength δ_j such that they plan to reach their desired portfolio in δ_j days based on their information at time t . For the fraction of holdings in the risky asset at time t for fund j denoted by $h_{j,t}$, this gives the change in the proportion of holdings in the risky asset as

$$\Delta h_{j,t} = h_{j,t} + \frac{\kappa_{j,t} - h_{j,t}}{\delta_j}$$

The demand of fund j in period t is then $z_{j,t} = \Delta h_{j,t} W_{j,t}$ where $W_{j,t}$ is the wealth held by each fund j at the beginning of period t . 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; $\delta_j \sim \mathcal{U}(L_R, U_R)$.

Expected returns $E_{j,t}(R_{t+1})$ 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. 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 considerable 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.

We assume that funds themselves do not switch strategies. This is motivated in part by the fact that funds set out their preferred strategies for their investors - frequent switching of approaches would run counter to that. In addition, we already capture growth of successful and decline of less successful strategies through the switching of cash between funds by investors.

Funds' expected returns are

$$E_{j,t}(R_{t+1}) = \begin{cases} (1 - L)(1 + Y_t + \alpha(Y_t - Y^*)) - 1 & \text{for value traders} \\ (1 - L)(1 + Y_t + \beta(\bar{Y} - Y_t)) - 1 & \text{for momentum traders} \end{cases} \quad (2)$$

where

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

is the average yield over the time period beginning at $t - t_{1w}$ and ending at time t , with Δt a single time step (single trading day). t_{1w} 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. 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. Y^* is the positive, non-zero fundamental yield of the bond, which is related to fundamental price P^* and face value F by

$$P^* = \frac{F}{(1 + Y^*)^D} \quad (3)$$

for bond duration D and assuming zero coupon. 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.

The traded price P is determined dynamically and endogenously within the model. All agents share the same expectations regarding loss rate and agree on the face value of the bond. The former 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.

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.

Prices are adjusted by the (stylised) market maker as a log-linear function of excess demand $\sum_j z_{j,t}$ following Farmer and Joshi [2002]. 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.

$$\ln(P_{t+1}) = (\lambda + vV) \left(\sum_j z_{j,t} + \epsilon \right) + \ln(P_t) \quad (4)$$

Here, the market-maker's yield adjustment in response to excess demand or supply is controlled by an explicit, calibrated parameter λ . Excess demand is also passed through into prices via the volatility V and its calibrated coefficient v . Noise, ϵ , is drawn from a normal distribution, $\epsilon \sim \mathcal{N}(0, N_{ns}^2)$, with N_{ns} the standard deviation. This noisy demand represents the actions of participants - eg banks, hedge funds, insurance companies - and trading strategies that are not explicitly modelled.

Investors are represented by a stylised investor pool who withdraw and invest based on past excess returns of funds. These flows in 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, based on the empirical work by Chen and Qin [2014] and Goldstein et al. [2015]. The former jointly estimate sensitivity of fund flows to

macro-economic conditions and individual fund performance and find significant relationships with both aggregate bond performance and performance rank. The latter estimates the relative flows between funds, controlling for aggregate in- and outflows, and finds 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 each individual fund's rate of return, R^I , where

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

and

$$R_{j,t}^I = R_{j,t} - R_t$$

where $R_{j,t} = \frac{A_{j,t}P_t}{A_{j,t-1}P_{t-1}} - 1$ 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 flow is given by

$$S_{j,t} = sR_t^S + [I_- + H(R_{j,t}^I)(I_+ - I_-)] R_{j,t}^I$$

which has different fund specific coefficients (+ and -) for positive and negative rates of return. $H(R_{j,t}^I)$ 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 $R_{j,t}^I > 0$ for fund j at time t .

Funds are not allowed to borrow, ie they cannot demand more than their wealth in any given period. $S_{j,t} < 0$ represents a withdrawal by investors from fund j at time t . Fund j meets this redemption through both the cash and the sale of the risky asset in the proportion that they have them in their current holdings $h_{j,t}$. So for fund j facing redemption $S_{j,t+1}$ demand reduces as

$$z_{j,t+1} = (1 - (S_{j,t+1}/W_{j,t+1}))z_{j,t} \text{ for } S_{j,t} < 0 \quad (5)$$

where j refers to any of the individual funds. Individual fund wealth W_j each period is then equal to

$$W_{j,t+1} = W_{j,t} + (P_{t+1} - P_t)z_{j,t} + S_{j,t+1} \quad (6)$$

where $z_{j,t}$ represents the demand of fund j in the last period.

3 Empirical foundation and calibration

To assess the success of our model in capturing properties of trading that we see 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. 3 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. 2 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 (eg during the financial crisis) and when there were changes in the risk-free rate (eg 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.

We focus on the following empirical properties:



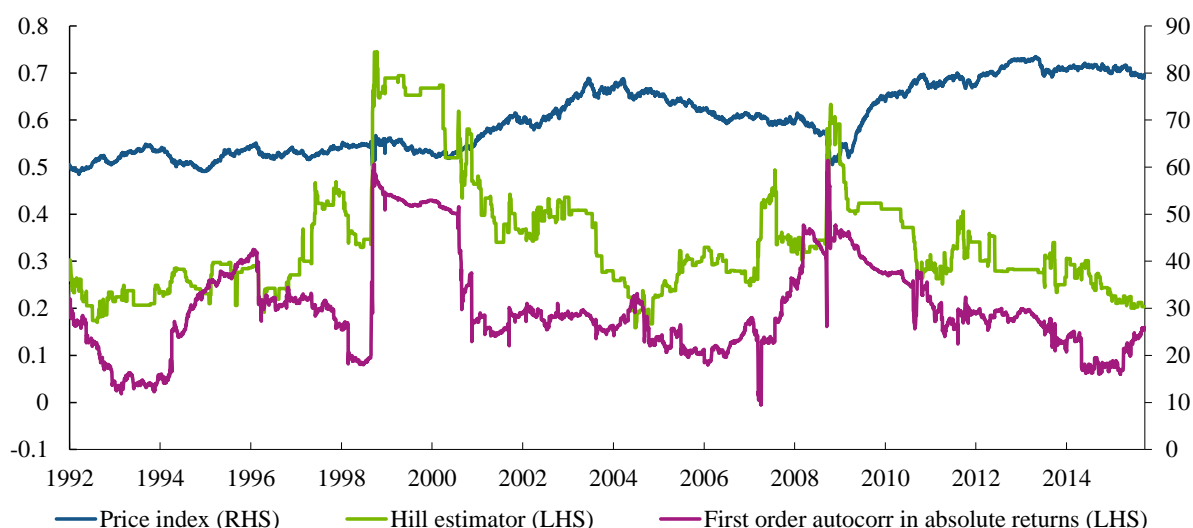


Figure 2: US IG corporate bond price index and rolling moments. Moments are calculated based on a two-year rolling window. Source: Bank of America Merrill Lynch and Bank calculations.

- First-order autocorrelation in returns (ρ_{-1}^r): 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.
- 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. A lower value represents higher kurtosis.
- 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.
- First, fifth- and tenth-order autocorrelation in absolute returns (ρ_{-1}^ν , ρ_{-5}^ν and ρ_{-10}^ν): as documented in Cont [2001], 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 technical reasons 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].

Model parameters are derived from empirical data where available.

- 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 (ie 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 2015Q2 is around 80, this implies that agents hold roughly 7.3 billion risky assets.
- Distribution of fund size: the distribution of size of agents in our model is calibrated to the empir-

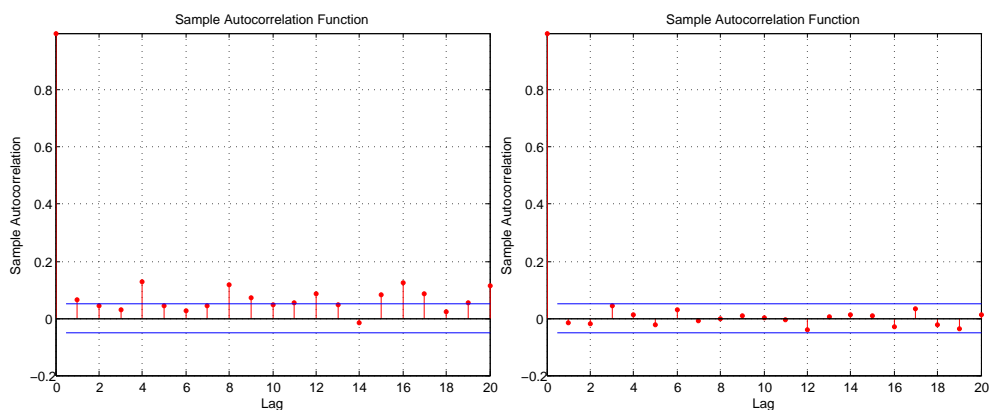


Figure 3: Autocorrelations of the empirical corporate bond return series. Source: Bank of America Merrill Lynch and Bank calculations.

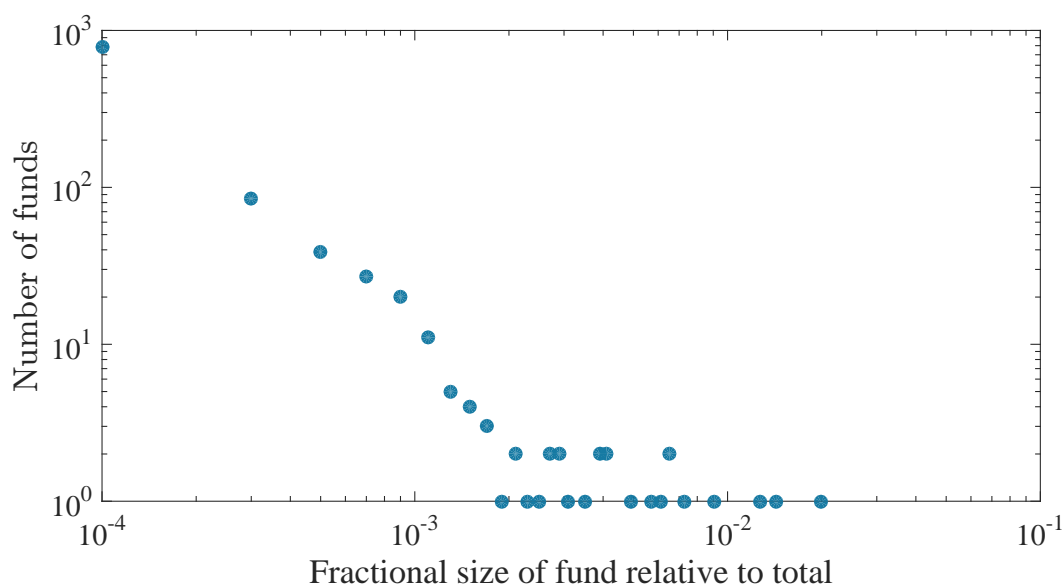


Figure 4: Distribution of fund sizes relative to total sector holdings.

ical distribution, as shown in Fig. 4

- 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. We find that a 1% return on the index is associated with a 0.25% increase in aggregate net flow, as shown in Figure 5.
 - 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.

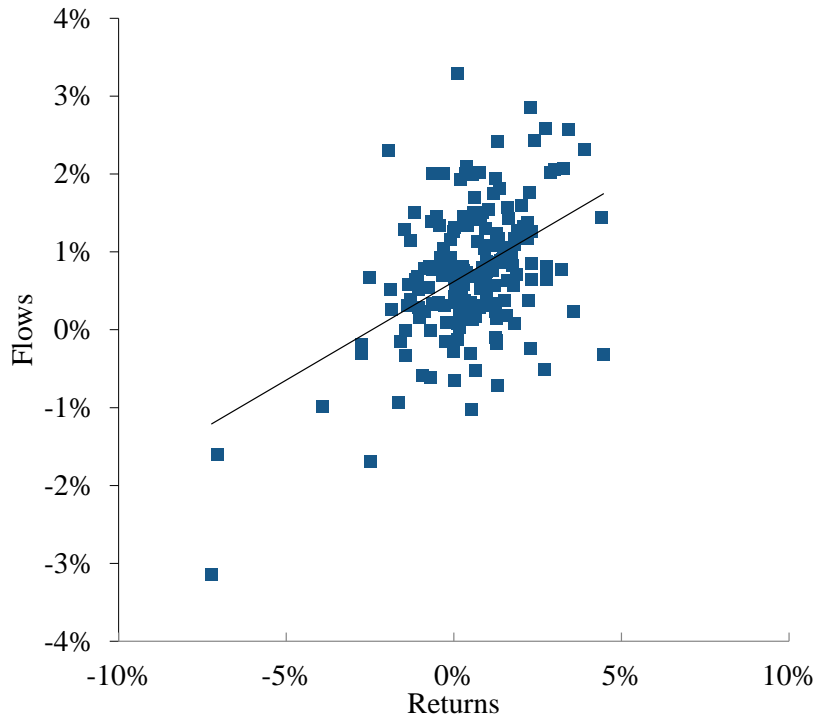


Figure 5: Aggregate net flows v.s. US corporate bond index returns. Source: Morningstar and Bank calculations. The linear relationship is statistically significant at the 1% level.

- Duration, D : the duration of corporate bonds is set to 6.917 years, based on 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 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 1). 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 across funds.

Table 1: 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
Fundamental yield, Y^*	1.282%
Expected loss rate, L	0.04%
Imposed Parameters	Value
Proportion of value traders	0.4
Proportion of momentum traders	0.4
Reaction strength lower bound, L_R	10
Reaction strength upper bound, U_R	20
Momentum trader long window, t_{lw}	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

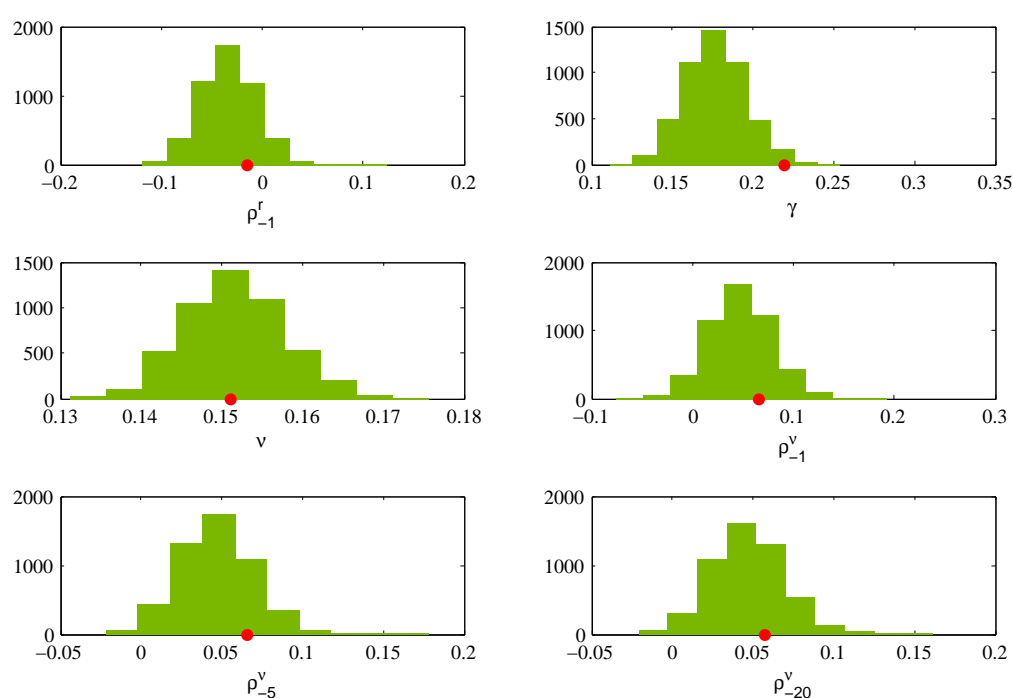
The other parameters in the model (see Table 1) are estimated by matching model outputs to empirical properties using the Moment Coverage Ratio (MCR) proposed by Franke and Westerhoff [2012]. The parameter values are optimised using a grid search which runs simulations with different sets of parameters 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 Hill estimator which had a closed-form formula, $\frac{\gamma^2}{k}$) are calculated using the delta method and shown in Table 2. 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. 6. All empirical moments fall into the distribution of simulated moments. Table 3 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. This demonstrates that the model and parameters are a reasonable platform for our sensitivity analysis and experiments in the following sections.

Table 2: Moment confidence intervals

Moment	Mean	Lower bound	Upper bound
ρ_{-1}^r	-0.0142	-0.0726	0.0443
γ	0.2191	0.1699	0.2684
ν	0.1510	0.1430	0.1591
ρ_{-1}^v	0.0652	0.0093	0.1211
ρ_{-5}^v	0.0657	0.0098	0.1216
ρ_{-10}^v	0.0572	0.0087	0.1056

Figure 6: Simulated moments v.s. empirical moments



Note: Based on 5,000 Monte Carlo runs. The length of each run is 1,500 simulation 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.

Table 3: Moment coverage ratios

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

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.

4 Simulation results

4.1 Steady state

When running the calibrated model, we typically find repeating cycles of yield swings and corrections as shown in Fig. 7 for an illustrative single run. These reflect the interaction of trading strategies, market maker adjustments and investor flows with normally distributed random noise. As the upper row of Fig. 7 demonstrates, such short to medium term yield changes are in line with the empirical data. As further corroboration that the model produces reasonable outcomes, Fig. 8 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 compared with simulated data. The empirical data has larger tails because returns can be moved by macroeconomic news, while the simulation data shown is from a ‘steady state’ run in which no changes to fundamentals occur.

We define the steady state as the simulation running with no changes to the parameters shown in Table 1 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. 7 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 reinforce each other to the point of 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 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 artifact of small numbers, we initialise the simulations with a realistic number of firms (1000) and the empirical size distribution (shown in Fig. 4). Results are not particularly sensitive to either agent numbers or the shape of the distribution.

4.2 Response to shocks

As the previous section illustrates, mutual fund trading calibrated to the post-crisis period on its own is typically fairly stable, in line with empirical data. That may no longer hold in a shock. From a financial stability perspective, we are particularly interested in the extent to which endogenous market features either aggravate or dampen downside shocks.

We explore system responses to shocks to the value of the risky asset as well as to market conditions such as fund wealth. In discussing results, for brevity, we focus primarily on a shock to 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. (2) 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

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

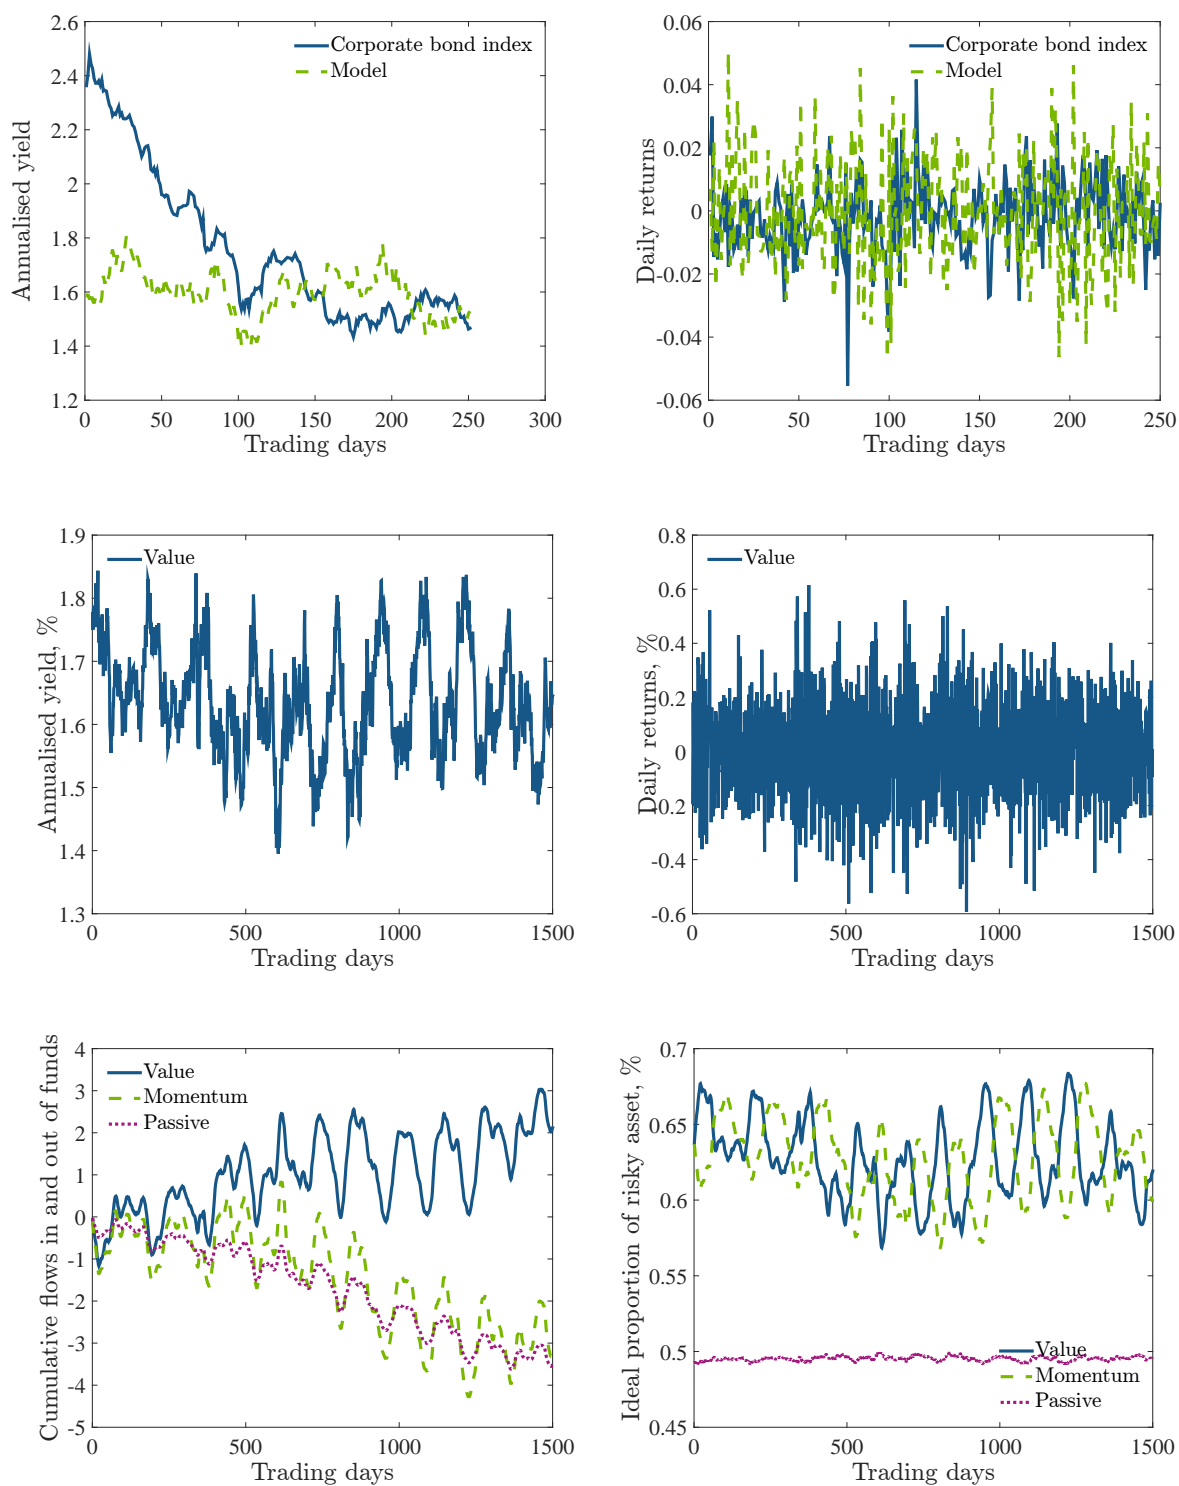


Figure 7: Output from a single model run over 250 trading days; clockwise starting from the upper left panel are shown: yield (annualised) over a trading year, returns over a trading year, returns over a six year period, funds' ideal portfolio composition, cumulative net flows from/to the investor pool 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.

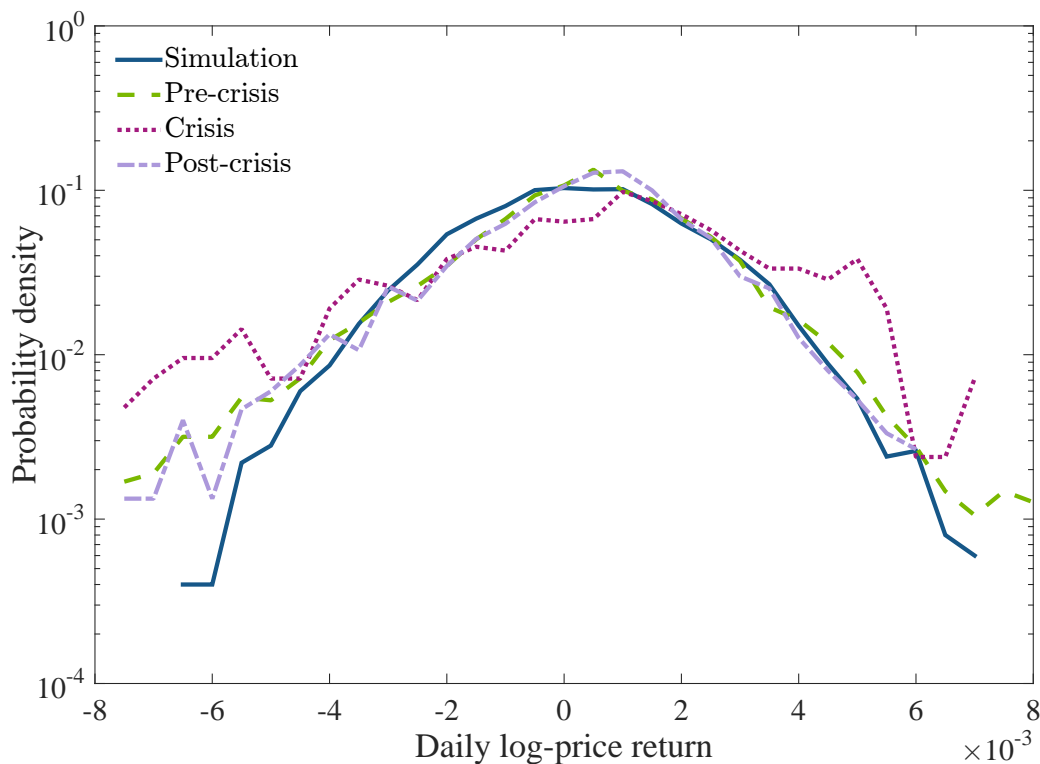


Figure 8: The distribution function of daily log-price returns of three periods shown against data generated by a single model run.

where $H(t)$ is the step function. This is in contrast to the steady state situation of section 4.1, in which L remains constant. As stated in section 3, this parameter takes the value $L = 0.04\%$ per annum (p.a.) for 2010 and is our base line value before any shocks arrive in the form of the change to L .

A schematic of what occurs in a shock to the expected loss rate value is shown in Fig. 9. The change in L feeds into the proportion of risky assets that funds are willing to hold, thus changing excess demand. This changes price and yield via the market maker, who responds to excess demand according to equation (4). One consequence of this is that funds' wealth is reduced. Furthermore, the fall in price means that the returns to investors are lower, triggering investor outflows. This then causes a further reduction in fund wealth. Reductions in fund wealth must be met by both selling the risky asset, and through the risk-less asset. Asset sales by funds prompt further price cuts, and the feedback loop continues.

Fig. 10 shows the median response over a 100 model runs for simulated loss rate shocks of varying strength. As these demonstrate, shocks to loss rates lead to increases in equilibrium yields of similar magnitude, which suggests that the pricing of default risk in our model is close to risk-neutral pricing in the long run. 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. 11 breaks down the impact of changes on wealth of funds 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.

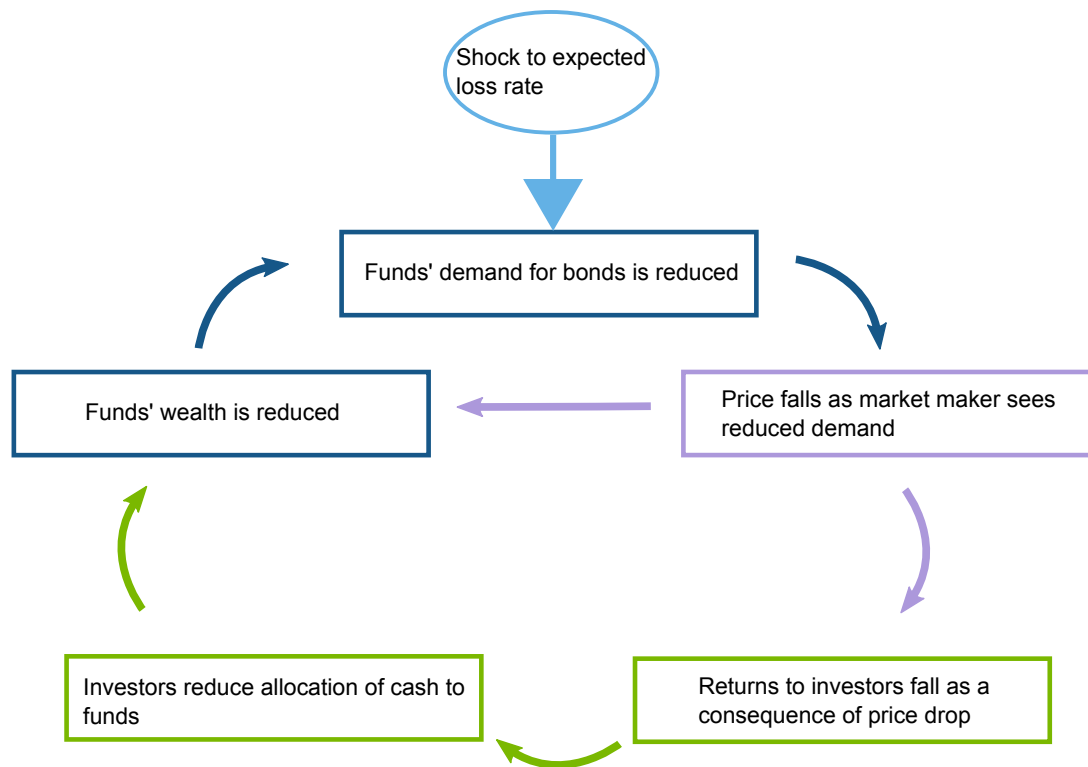


Figure 9: A schematic which shows the feedback loops following a shock to the value of the expected loss rate. The colours in the feedback loop indicate the different market players; funds, the market maker, and the investor pool.

We find broadly similar feedback effects 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. This 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. 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 previous shock, overshoots correct over time, but it is a prolonged process with several cycles of overshooting in both directions. For brevity, we focus on a loss rate shock in the remaining discussion, only bringing in other shocks where there are differences of particular interest.

To dig deeper into the role of different aspects in propagating shocks, we can repeat simulations under different steady state assumptions, but for the same set of noise runs. In particular, we investigate changes to market maker sensitivities, to trading strategies, and to the strength of investor response. For ease of presentation, we only show results for a single shock to the expected loss rate, increasing it from $L = 0.04\%$ p.a. to $L = 0.35\%$ p.a. which is of a comparable order of magnitude as the annual loss rate changes between 2007 and 2008.

Fig. 12 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 excess demand on price levels. Longer run increases of average yield reflects 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

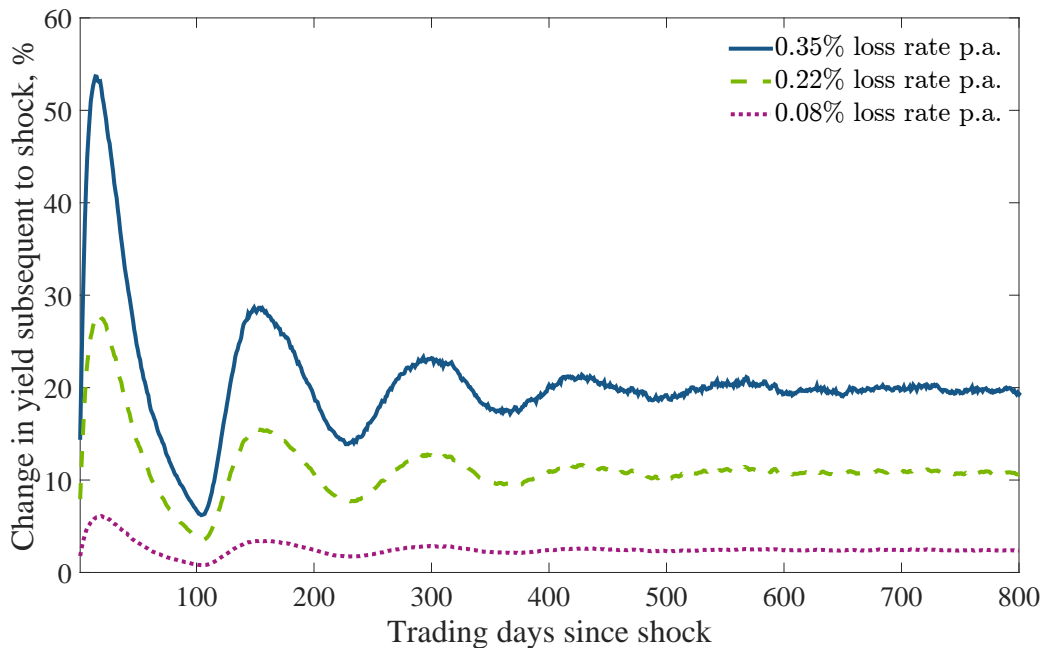


Figure 10: Expected loss rate shocks; a sudden change in firms' expected loss rate, L , causes both short-term fluctuations in yield, and a new, higher equilibrium yield. Results presented are the median over 100 individual simulations runs.

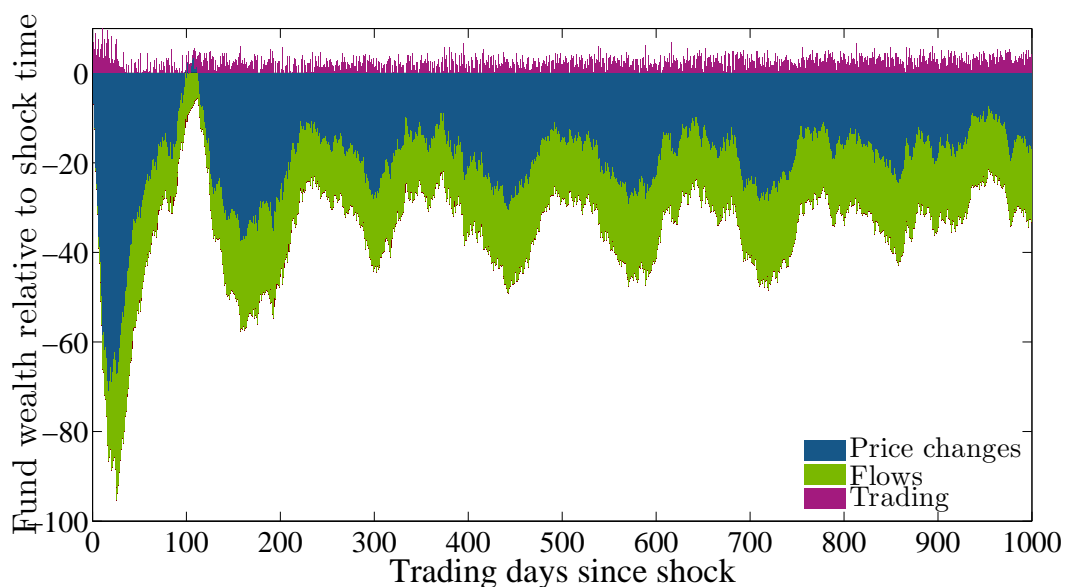


Figure 11: The effect of an expected loss rate shock on wealth averaged over all funds. The decomposition shows how price changes, investor flows and trading all contribute to the post-shock changes in wealth. Results presented are the median over 100 individual simulations runs.

volatility over the previous 100 trading days (see Eq. (4)). 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. 15). 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 (ie 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. 13 shows, there is no overshoot without momentum trading, reducing short- to medium-term yield dislocations significantly.

The median impact of an increased presence of passive funds is dependent on the shock in question. Their presence can reduce overshooting due to new information as, in contrast to momentum traders and value traders, they do not trade based on news (see Fig. 16). But in exogenous shocks such as reduced demand due to investor withdrawals, median yield dislocations can increase (see Fig. 17). This reflects increased tail risk of large yield swings in the absence of active investors which we show by plotting the 70th and 95th percentile of outcomes (for 250 runs).

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 outflow sensitivity as shown in Fig. 14 only increases maximum yield increases by around ten percent, and longer run yields by around one to two percent.

4.3 Experiment on funds' redemption policy

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 reducing 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. 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 fund redemption 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. A fund j meets its redemption in the usual way, splitting it between the risky asset and the riskless asset 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. 7), 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 redemption management 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. 18 shows that the redemption management policy can reduce the maximum change in

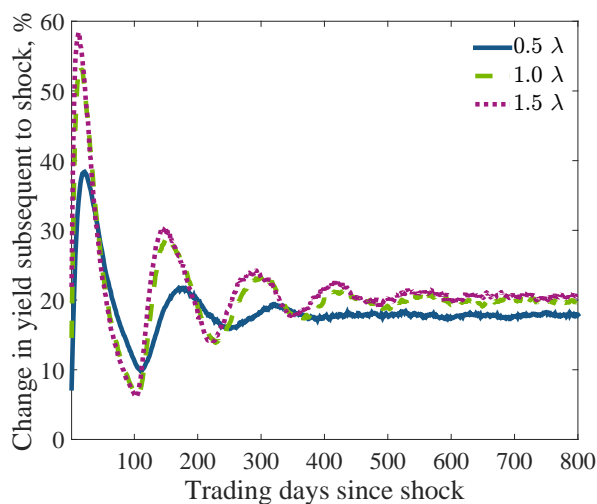


Figure 12: Yield change with a 0.35% p.a. loss rate shock and different market maker sensitivities relative to the base line value of λ .

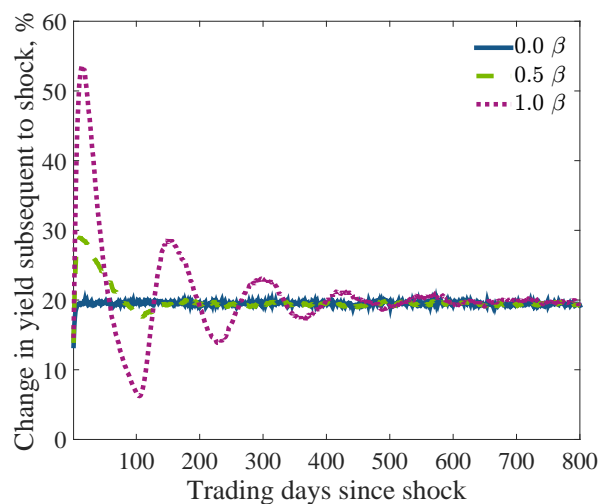


Figure 13: Yield change with a 0.35% p.a. loss rate shock and different momentum trader strengths relative to the base line value of β .

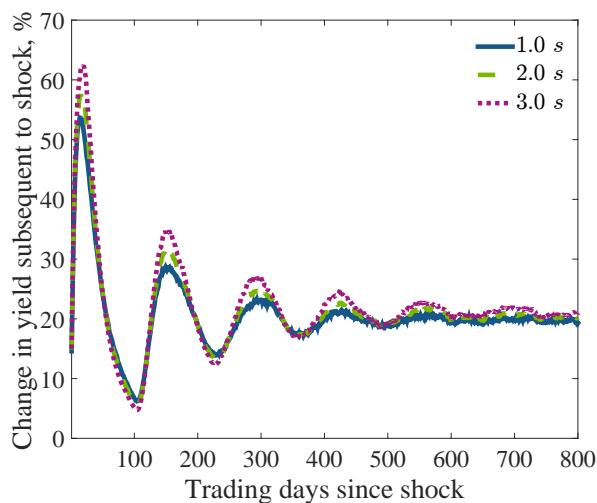


Figure 14: Yield change with a 0.35% p.a. loss rate shock and different systemic flow strengths relative to the base line value of s .

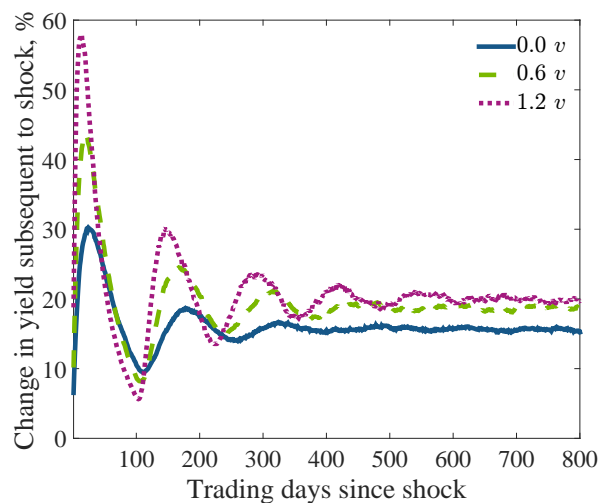


Figure 15: Yield change with a 0.35% p.a. loss rate shock and different market maker volatility coefficients relative to the base line value of v .

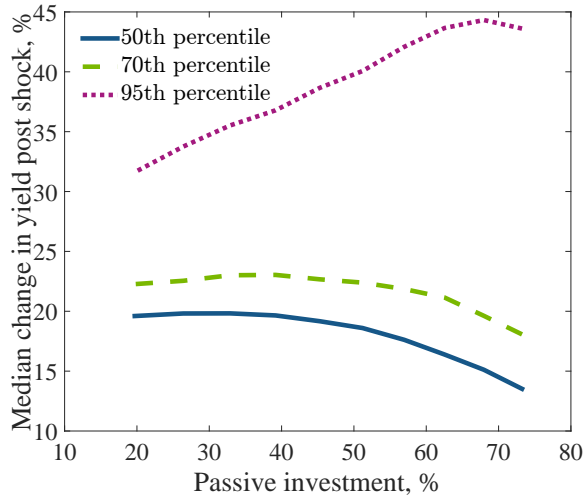


Figure 16: Distribution of outcomes for median yield over 100 trading days after a 0.36% increase in loss rate shock: percentiles indicate the value taken from 250 runs.

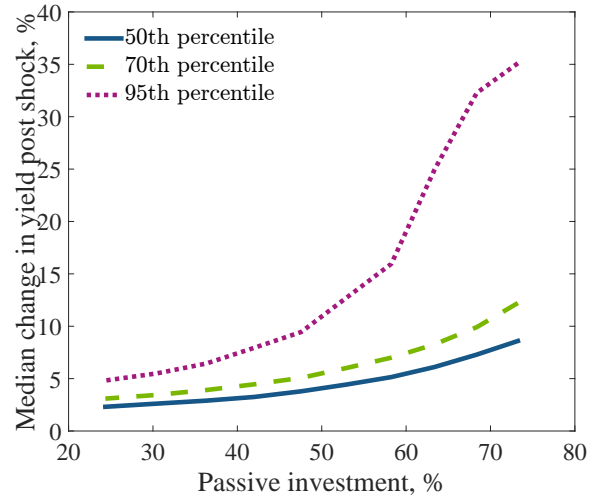


Figure 17: Distribution of outcomes for median yield over 100 trading days after a 5% fund wealth shock: percentiles indicate the value taken from 250 runs.

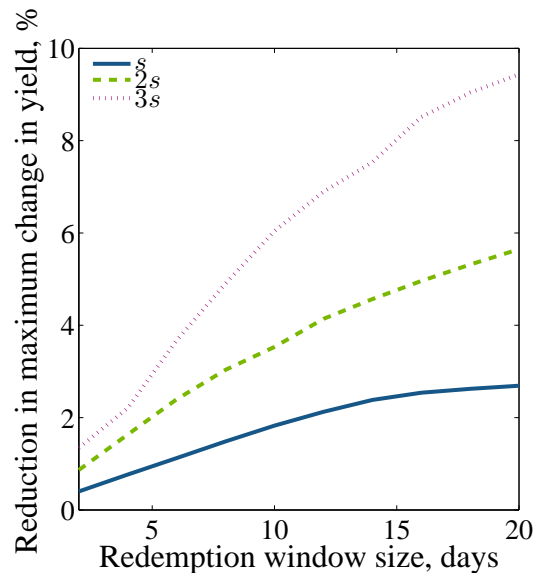
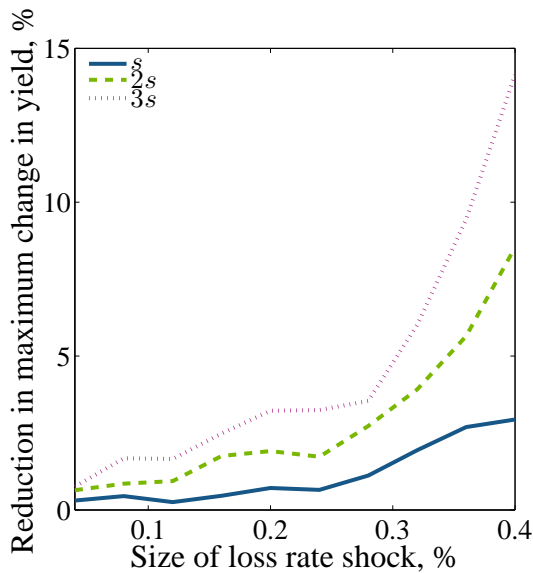


Figure 18: When flow strength, s , is large, redemption management can significantly reduce the magnitude of the maximum change in yield following a large loss rate shock. The redemption management policy splits any redemption over 20 trading days in the first panel. The loss rate shock is 0.35% in the second panel.

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 redemption management 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 redemption management 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 an empirically calibrated model of the corporate bond market with both active and passive traders. Statistics (or stylised facts) matched data from a US IG 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. Additionally, it was demonstrated that a larger fraction of passive funds can increase tail risk associated with shocks considerably. In contrast, managing redemptions could reduce the impact of shocks in stressed times when investor outflows are unusually large, as in the 2008 financial crisis.

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

References

- C. Aymanns and J. D. Farmer. The dynamics of the leverage cycle. *Journal of Economic Dynamics and Control*, 50:155–179, Jan 2015. ISSN 0165-1889. doi: 10.1016/j.jedc.2014.09.015. URL <http://dx.doi.org/10.1016/j.jedc.2014.09.015>.
- R. Bookstaber, M. D. Foley, and B. Tivnan. Market liquidity and heterogeneity in the investor decision cycle. *Office of Financial Research Working Paper Series*, 2015.
- Q. Chen, I. Goldstein, and W. Jiang. Payoff complementarities and financial fragility: Evidence from mutual fund outflows. *Journal of Financial Economics*, 97(2):239–262, Aug 2010a. ISSN 0304-405X. doi: 10.1016/j.jfineco.2010.03.016. URL <http://dx.doi.org/10.1016/j.jfineco.2010.03.016>.
- Y. Chen and N. Qin. The behavior of investor flows in corporate bond mutual funds. *Available at SSRN 2022059*, 2014.
- Y. Chen, W. Ferson, and H. Peters. Measuring the timing ability and performance of bond mutual funds. *Journal of Financial Economics*, 98(1):72–89, Oct 2010b. ISSN 0304-405X. doi: 10.1016/j.jfineco.2010.05.009. URL <http://dx.doi.org/10.1016/j.jfineco.2010.05.009>.
- R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1:223236, 2001.



- J. Farmer and S. Joshi. The price dynamics of common trading strategies. *Journal of Economic Behavior & Organization*, 49(2):149–171, Oct 2002. ISSN 0167-2681. doi: 10.1016/s0167-2681(02)00065-3. URL [http://dx.doi.org/10.1016/S0167-2681\(02\)00065-3](http://dx.doi.org/10.1016/S0167-2681(02)00065-3).
- I. Fender and U. Lewrick. Shifting tides - market liquidity and market-making in fixed income instruments. *BIS Quarterly Review March 2015*, 2015.
- T. Fischer and J. Riedler. Prices, debt and market structure in an agent-based model of the financial market. *Journal of Economic Dynamics and Control*, 48:95–120, Nov 2014. ISSN 0165-1889. doi: 10.1016/j.jedc.2014.08.013. URL <http://dx.doi.org/10.1016/j.jedc.2014.08.013>.
- R. Franke and F. Westerhoff. Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics and Control*, 36(8):1193–1211, Aug 2012. ISSN 0165-1889. doi: 10.1016/j.jedc.2011.10.004. URL <http://dx.doi.org/10.1016/j.jedc.2011.10.004>.
- I. Goldstein, H. Jiang, and D. Ng. Investor flows and fragility in corporate bond funds. *unp*, 2015.
- C. H. Hommes. Chapter 23 heterogeneous agent models in economics and finance. *Handbook of Computational Economics*, pages 1109–1186, 2006. ISSN 1574-0021. doi: 10.1016/s1574-0021(05)02023-x. URL [http://dx.doi.org/10.1016/S1574-0021\(05\)02023-X](http://dx.doi.org/10.1016/S1574-0021(05)02023-X).
- IMF. Global financial stability report. April 2015.
- T. Lütje. To be good or to be better: asset managers’ attitudes towards herding. *Applied Financial Economics*, 19(10):825–839, May 2009. ISSN 1466-4305. doi: 10.1080/09603100801964404. URL <http://dx.doi.org/10.1080/09603100801964404>.
- F. Moneta. Measuring bond mutual fund performance with portfolio characteristics. *Journal of Empirical Finance*, 33:223–242, Sep 2015. ISSN 0927-5398. doi: 10.1016/j.jempfin.2015.03.012. URL <http://dx.doi.org/10.1016/j.jempfin.2015.03.012>.
- J. Shek, I. Shim, and H. S. Shin. Investor redemptions and fund manager sales of emerging market bonds: how are they related? *BIS Working Papers*, 509, 2015.
- L. Tesfatsion. Chapter 16 agent-based computational economics: A constructive approach to economic theory. *Handbook of Computational Economics*, pages 831–880, 2006. ISSN 1574-0021. doi: 10.1016/s1574-0021(05)02016-2. URL [http://dx.doi.org/10.1016/S1574-0021\(05\)02016-2](http://dx.doi.org/10.1016/S1574-0021(05)02016-2).
- S. Thurner, J. D. Farmer, and J. Geanakoplos. Leverage causes fat tails and clustered volatility. *Quantitative Finance*, 12(5):695–707, May 2012. ISSN 1469-7696. doi: 10.1080/14697688.2012.674301. URL <http://dx.doi.org/10.1080/14697688.2012.674301>.