

Trading and Liquidity with Limited Cognition*

Bruno Biais,[†]Johan Hombert,[‡]and Pierre-Olivier Weill[§]

December 11, 2010

Abstract

We study the reaction of financial markets to aggregate liquidity shocks when traders face cognition limits. While each financial institution recovers from the shock at a random time, the trader representing the institution observes this recovery with a delay, reflecting the time it takes to collect and process information about positions, counterparties and risk exposure. Cognition limits lengthen the market price recovery. They also imply that traders who find that their institution has not yet recovered from the shock place market sell orders, and then progressively buy back at relatively low prices, while simultaneously placing limit orders to sell later when the price will have recovered. This generates round trip trades, which raise trading volume. We compare the case where algorithms enable traders to implement this strategy to that where traders can place orders only when they have completed their information processing task.

Keywords: Liquidity shock, limit-orders, asset pricing and liquidity, algorithmic trading, limited cognition, sticky plans

J.E.L. Codes: G12, D83.

*We are grateful to the editor Harald Uhlig, and four anonymous referees for insightful comments. Many thanks, for helpful discussions and suggestions, to Andy Atkeson, Dirk Bergemann, Darrell Duffie, Emmanuel Farhi, Thierry Foucault, Xavier Gabaix, Alfred Galichon, Christian Hellwig, Hugo Hopenhayn, Vivien Lévy-Garboua, Johannes Horner, Boyan Jovanovic, Ricardo Lagos, Albert Menkveld, John Moore, Stew Myers, Henri Pages, Thomas Philippon, Gary Richardson, Jean Charles Rochet, Guillaume Rocheteau, Ioanid Rosu, Larry Samuelson, Tom Sargent, Jean Tirole, Aleh Tsyvinski, Juuso Välimäki, Dimitri Vayanos, Adrien Verdelhan, and Glen Weyl; and to seminar participants at the Dauphine-NYSE-Euronext Market Microstructure Workshop, the European Summer Symposium in Economic Theory at Gerzensee, École Polytechnique, Stanford Graduate School of Business, New York University, Northwestern University, HEC Montreal, MIT, and UCI. Paulo Coutinho and Kei Kawakami provided excellent research assistance. Bruno Biais benefitted from the support of the “Financial Markets and Investment Banking Value Chain Chair” sponsored by the Fédération Bancaire Française and from the Europlace Institute of Finance, and Pierre-Olivier Weill from the support of the National Science Foundation, grant SES-0922338.

[†]Toulouse School of Economics (CNRS-CRM, IDEI), biais@cict.fr

[‡]HEC Paris, johan.hombert@hec.fr

[§]University of California Los Angeles and NBER, poweill@econ.ucla.edu

1 Introduction

“The perception of the intellect extends only to the few things that are accessible to it and is always very limited” (Descartes, 1641)

We analyze trades and price dynamics when investors face cognition limits and the market is hit by an aggregate liquidity shock. The shock induces a transient drop in the willingness and ability of financial institutions to hold assets such as stocks or bonds. It can be triggered by changes in the characteristics of assets, e.g., certain types of institutions, such as insurance companies or pension funds, are required to sell bonds that lose their investment grade status, or stocks that are delisted from exchanges or indices (see, e.g., Greenwood, 2005). Alternatively, an aggregate shock can reflect events affecting the overall financial situation of a category of institutions, e.g., funds experiencing large outflows or losses (see Coval and Stafford, 2007), banks incurring large losses (see Berndt, Douglas, Duffie, Ferguson, and Schranz, 2005), or specialists building extreme positions (see Comerton Forde, Hendershott, Jones, Moulton, and Seasholes, 2010).¹ To recover from a shock, institutions seek to unwind their positions in multiple markets (e.g., credit default swaps (CDS) corporate bonds or mortgage backed securities (MBS)). Institutions can also raise new capital, secure credit lines or structure derivative trades to hedge their positions. These processes are complex and take time. However, once an institution arranges enough deals, it recovers from the liquidity shock.

In this context, traders must collect and process a large flow of information about asset valuations, market conditions and the financial status of their own institution. They must obtain and aggregate information from several desks, markets and departments about gross and net positions, the resulting risk exposure, and compliance with regulations. When traders have limited cognition, completing these tasks is challenging and requires significant time and effort.

We address the following issues: How do traders and markets cope with negative liquidity shocks? What is the equilibrium price process after such shocks? How are trading and prices affected by cognition limits? Do the consequences of limited cognition vary with market mechanisms and technologies?

¹A striking example of a liquidity shock and its consequences on institutions and market pricing is analyzed by Khandani and Lo (2008). These observe that, during the week of August 6th 2007, quantitative funds subject to margin calls and losses in credit portfolios had to rapidly unwind equity positions. This resulted in a sharp but transient drop in the S&P 500. But, by August 10th 2007 prices had in large part reverted.

We consider an infinite-horizon, continuous-time market with a continuum of rational, risk-neutral competitive financial institutions, deriving a non-linear utility flow from holding divisible shares of an asset, as in Lagos and Rocheteau (2009) and Gârleanu (2009). To model the aggregate liquidity shock, we assume that at time 0 this utility flow drops for all institutions, as in Duffie, Gârleanu, and Pedersen (2007) and Weill (2007). Then, as time goes by, institutions progressively switch back to a high valuation. More precisely, each institution is associated with a Poisson process and switches back to high-valuation at the first jump in this process. Unconstrained efficiency would require that low-valuation institutions sell to high-valuation institutions. However, in our model, such asset reallocation is delayed because of cognition limits. In line with the rational inattention models of Reis (2006a and 2006b), Mankiw and Reis (2002) and Gabaix and Laibson (2002), we assume that each trader engages in information collection and processing for some time and, only when this task is completed, observes the current valuation of her institution for the asset. We refer to this observation as an “information event.” When this event occurs, the trader updates her optimal asset holding plan, based on rational expectations about future variables and decisions.² The corresponding demand, along with the market clearing condition, gives rise to equilibrium prices. Thus, while Mankiw and Reis (2002) and Gabaix and Laibson (2002) analyze how inattention affects consumption, we study how it affects equilibrium pricing during liquidity shocks.

In the spirit of Duffie, Gârleanu, and Pedersen (2005), we assume that information events, and correspondingly trading decisions, occur at Poisson arrival times.³ Thus in our model, each institution is exposed to two Poisson processes: one concerns changes in its valuation for the asset, the other the timing of its trader’s information events. For simplicity, we assume these two processes are independent. Further, for tractability, we assume that these processes are independent across institutions. Thus, by the law of large numbers, the aggregate state of the market changes deterministically with time.⁴ Correspondingly, the equilibrium price process is

²Thus, in the same spirit as in Mankiw and Reis (2002), traders have “sticky plans” but rationally take into account this stickiness.

³Note however that the interpretation is different. Duffie, Gârleanu, and Pedersen (2005) model the time it takes traders to find a counterparty, while we model the time it takes them to collect and process information. This difference results in different outcomes. In Duffie, Gârleanu, and Pedersen investors do not trade between two jumps of their Poisson process. In our model they do, but based on imperfect information about their valuation for the asset. In a sense our model can be viewed as the dual of Duffie, Gârleanu, and Pedersen: they assume that traders continuously observe their valuation but are infrequently in contact with the market, while we assume that traders are continuously in contact with the market but infrequently refresh their information about their valuation.

⁴We also analyze an extension of our framework where the market is subject to recurring aggregate liquidity shocks, occurring at Poisson arrival times. While, in this more general framework, the price is stochastic, the

deterministic as well. We show equilibrium existence and uniqueness. In equilibrium the price increases with time, reflecting that the market progressively recovers from the shock. Limits to cognition lengthen the time it takes market prices to fully recover from the shock. Yet they do not necessarily amplify the initial price drop generated by that shock. Just after the shock, with perfect cognition, the marginal trader knows that her institution has a low valuation, while with limited cognition such trader is imperfectly informed about her institution's valuation, and realizes that, with some probability, it may have recovered. We also show that the equilibrium allocation is an information constrained Pareto optimum. This is because in our setup there are no externalities, as the holding constraints on holdings imposed by cognition limits on one trader, do not depend on the actions of other traders.

While we first characterize traders' optimal policies in terms of abstract holding plans, we then show how these plans can be implemented in a realistic market setting, featuring an electronic order book, limit and market orders, and trading algorithms. The latter enable traders to conduct programmed trades while devoting their cognitive resources to investigating the liquidity status of their institution. In this context, traders who find out their institution is still subject to the shock, and correspondingly has a low valuation for the asset, sell a lump of their holdings, with a market sell order. Such traders also program their trading algorithms to then gradually buy back the asset, as they expect their valuation to revert upward. Simultaneously, they submit limit orders to sell the asset, to be executed later when the equilibrium price will have recovered. To the extent that they buy in the early phase of the aggregate recovery, and then sell towards the end of the recovery, such traders act as market makers.⁵ The corresponding round-trip transactions reflects the traders' optimal response to cognition limits. These transactions can raise trading volume above the level it would reach under perfect cognition.

We also study the case where trading algorithms are not available and traders must implement their holding plans by placing limit and market orders when their information process jumps. With increasing prices, this prevents them from buying in between jumps of their information process. When the liquidity shock is large, this constraint binds and reduces the efficiency of the equilibrium allocation. It does not necessarily amplify the price pressure of the liquidity shock, however. Since traders anticipate they will not be able to buy back until

qualitative features of our equilibrium are upheld.

⁵In doing so they act similarly to the market makers analyzed by Grossman and Miller (1988). Note however that, while in Grossman and Miller agents are exogenously assigned market making or market taking roles, in our model, agents endogenously choose to supply or demand liquidity, depending on the realization of their own shocks.

their next information event, they sell less when they observe that their valuation is low. Such a reduction in supply limits the selling pressure on prices. Put differently, banning the use of algorithms could help alleviate the initial price pressure created by the liquidity shock. Yet, this policy would reduce welfare: indeed, in our model the equilibrium with algorithmic trading is information constrained Pareto optimal.⁶

The order placement policies generated by our model are in line with several stylized facts. Irrespective of whether algorithms are available, we find that successive traders place limit sell orders at lower and lower prices. Such undercutting is consistent with the empirical results of Biais, Hillion, and Spatt (1995), Griffiths, Smith, Turnbull, and White (2000) and Ellul, Holden, Jain, and Jennings (2007). Furthermore, our algorithmic traders both supply and consume liquidity, by placing market and limit orders, consistent with the empirical findings of Hendershott and Riordan (2010) and Brogaard (2010). Brogaard also finds that algorithms i) tend not to withdraw from the market after large liquidity shocks; ii) tend to provide liquidity by purchasing the asset after large price drops; and iii) in doing so profit from price reversals. Furthermore, Kirilenko, Kyle, Samadi, and Tuzun (2010) find that algorithms tend to buy after a rise in price and that these purchases tend to be followed by a further increase in price. They also find that algorithmic traders engage in high frequency round trip trades, buying the asset and then selling it to other algorithmic traders. All these findings are in line with the implications of our model.

Our analysis of the dynamics of markets in which traders choose whether to place limit or market orders is related to the insightful papers of Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2005), Rosu (2009), and Goettler, Parlour, and Rajan (2005, 2009). However, we focus on different market frictions than they do. While these authors study strategic behavior and/or asymmetric information under perfect cognition, we analyze competitive traders with symmetric information under limited cognition. This enables us to examine how the equilibrium interaction between the price process and order placement policies is affected by cognition limits and market instruments.

The next section describes the economic environment and the equilibrium prevailing un-

⁶In Section II of our supplementary appendix we analyze the case when traders can place only market orders when their information process jumps, i.e., limit orders and trading algorithms are ruled out. In such case the price reverts to its pre-shock level sooner. Indeed, when traders can place limit orders, the sell orders stored in the book exert a downward pressure on prices towards the end of the recovery. But, then again, the efficiency of the allocation is higher when traders can use limit orders, and even higher they can use algorithms. Indeed these market instruments enable traders to conduct mutually beneficial trades that would be infeasible with market orders only.

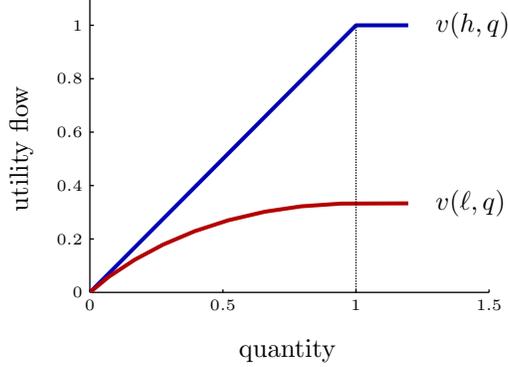


Figure 1: The utility flows of high- (in blue) and low-valuation (in red) investors, when $\sigma = 0.5$.

der unlimited cognition. Section 3 presents the equilibrium prevailing with limited cognition. Section 4 discusses the implementation of the abstract equilibrium holding plans with realistic market instruments such as limit and market orders and trading algorithms. Section 5 concludes. Proofs not given in the text are found in the appendix, and a supplementary appendix offers additional information about the model, along with proofs and analyses.

2 The economic environment

2.1 Assets and agents

Time is continuous and runs forever. A probability space (Ω, \mathcal{F}, P) is fixed, as well as an information filtration satisfying the usual conditions (Protter, 1990).⁷ There is an asset in positive supply $s \in (0, 1)$ and the economy is populated by a $[0, 1]$ -continuum of infinitely-lived agents that we call “financial institutions” (funds, banks, insurers, etc.) discounting the future at the same rate $r > 0$.

Each institution can be in one of two states. Either it derives a high utility flow (“ $\theta = h$ ”) from holding any quantity $q \geq 0$ of the asset, or it derives a low utility flow (“ $\theta = \ell$ ”), as illustrated in Figure 1. For high-valuation institutions, the utility flow per unit of time is $v(h, q) = q$, for all $q \leq 1$, and $v(h, q) = 1$, for all $q > 1$. For low-valuation institutions, it is $v(\ell, q) = q - \delta \frac{q^{1+\sigma}}{1+\sigma}$, for all $q \leq 1$, and $v(\ell, q) = 1 - \delta/(1 + \sigma)$, for all $q > 1$.⁸ The two parameters $\delta \in (0, 1]$ and $\sigma > 0$ capture the effect of the low state on utility flows. The

⁷To simplify the exposition, for most stated equalities or inequalities between stochastic processes, we suppress the “almost surely” qualifier as well as the corresponding product measure over times and events.

⁸The short-selling constraint is without loss of generality in the following sense. If we extend the utility functions to $q < 0$ in any way such that they remain concave, then the equilibrium outcomes we characterize are unaffected. In particular, $q < 0$ never arises.

parameter δ controls the level of utility: the greater is δ , the lower is the marginal utility flow of low-valuation institutions. The parameter σ , on the other hand, controls the curvature of low-valuation institutions utility flows. The greater is σ , the less willing such institutions are to hold extreme asset positions.⁹ Because of this concavity, it is efficient to spread holdings among low-valuation institutions. This is similar to risk-sharing between risk-averse agents, and as shown below will imply that equilibrium holdings take a rich set of values.¹⁰ This is in line with Lagos and Rocheteau (2009) and Gârleanu (2009). Note that, even in the $\sigma \rightarrow 0$ limit, low-valuation investors' utility flow is reduced, by a factor $1 - \delta$, but in that case the utility flow is piecewise linear.¹¹

The difference between the high and the low states can be interpreted as a holding cost, or a capital charge or shadow cost associated with positions. As discussed in the introduction and in Duffie, Gârleanu, and Pedersen (2007) a variety of institutional factors can generate such costs, e.g., regulatory constraints on holdings (see, e.g., Greenwood, 2005), need for cash (see Coval and Stafford, 2007, Berndt et al., 2005), position limits (see Comerton Forde et al., 2010, or Hendershott and Seasholes, 2007), or tax considerations.

Lastly, in addition to deriving utility from the asset, institutions can produce (or consume) a non-storable numéraire good at constant marginal cost (utility) normalized to 1.

2.2 Liquidity shock

To model liquidity shocks we follow Duffie, Gârleanu, and Pedersen (2007) and Weill (2007). Before the shock, each institution is in the high-valuation state, $\theta = h$, and holds s shares of the asset. But, at time zero, the liquidity shock hits all the institutions, and they make a switch to low-valuation, $\theta = \ell$. Note, however, that the shock is transient. As discussed in the introduction, to cope with the shock, institutions seek to unwind positions, raise capital, secure credit lines, or hedge positions. All this process is complex and takes time. However, once the institution is able to arrange enough deals, it recovers from the liquidity shock. To capture the recovery process we assume that, for each institution, there is a random time at which it reverts

⁹The curvature of low-valuation utilities contrasts with the constant positive marginal utility high-valuation institutions have for $q < 1$. One could have introduced such curvatures for high-valuations as well, as in Lagos and Rocheteau (2009) or Gârleanu (2009) at the cost of reduced tractability, without qualitatively altering our results.

¹⁰Note also that the holding costs of low-valuation institutions are homothetic. This results in homogenous asset demand and, as will become clear later, this facilitates aggregation.

¹¹For the $\sigma \rightarrow 0$ limit, see our supplementary appendix, (Biais, Hombert, and Weill, 2010b), Section III.

to the high-valuation state, $\theta = h$, and then remains there forever. For simplicity, we assume that recovery times are exponentially distributed, with parameter γ , and independent across investors. Hence, by the law of large numbers, the measure μ_{ht} of high-valuation investors at time t is equal to the probability of high-utility at that time conditional on low-utility at time zero.¹² Thus

$$\mu_{ht} = 1 - e^{-\gamma t}, \tag{1}$$

and we denote by T_s the time at which the mass of high-utility institutions equals the supply of the asset, i.e.,

$$\mu_{hT_s} = s. \tag{2}$$

2.3 Equilibrium without cognition limits

Consider the benchmark case where institutions continuously observe their θ_t . To find the competitive equilibrium, it is convenient to solve first for efficient asset allocations, and then find the price path that decentralizes these efficient allocations in a competitive equilibrium.¹³

In the efficient allocation, for $t > T_s$, all assets are held by high-valuation institutions, and all marginal utilities are equalized. Indeed, with an (average) asset holding equal to $s/\mu_{ht} < 1$, the marginal utility is 1 for high-valuation institutions, while with zero asset holdings marginal utility is $v_q(\ell, 0) = 1$ for low-valuation institutions. In contrast, for $t \leq T_s$, we have $\mu_{ht} \leq s$, and each high-valuation institution holds one unit of the asset while the residual supply, $s - \mu_{ht}$, is held by low-valuation institutions. The asset holding per low-valuation institution is thus:

$$q_t = \frac{s - \mu_{ht}}{1 - \mu_{ht}}. \tag{3}$$

This is an optimal allocation because all high-valuation institutions are at the corner of their utility function: reducing their holdings would create a utility loss of 1, while increasing their holdings would create zero utility. Low-valuation institutions, on the other hand, have holdings

¹²For simplicity and brevity, we do not formally prove how the law of large numbers applies to our context. To establish the result precisely, one would have to follow Sun (2006), who relies on constructing an appropriate measure for the product of the agent space and the event space.

¹³Note that, with quasi-linear utilities and unlimited cognition, in all Pareto efficient allocation of assets and numéraire goods, the asset allocation maximizes, at each time, the equally weighted sum of the institutions' utility flows for the asset, subject to feasibility.

in $[0, 1)$, so their marginal utility is strictly positive and less than 1.

For $t \leq T_s$, as soon as an institution switches from $\theta = \ell$ to $\theta = h$, its holdings jump from q_t to 1, while as long as its valuation remains low, it holds q_t , given in (3), which smoothly declines with time. This decline reflects that, as time goes by, more and more institutions recover from the shock, switch to $\theta = h$ and increase their holdings. As a result, the remaining low-valuation institutions are left with less shares to hold.

Equilibrium prices reflect the cross-section of valuations across institutions. In our setting, by the law of large numbers, there is no aggregate uncertainty and this cross-section is deterministic. Hence, the price also is deterministic. For $t \leq T_s$, it is equal to the present value of a low-valuation institution's marginal utility flow:

$$p_t = \int_t^\infty e^{-r(z-t)} v_q(\ell, q_z) dz,$$

where q_z is given in (3). Taking derivatives with respect to t , we find that the price solves the Ordinary Differential Equation (ODE):

$$v_q(\ell, q_t) = rp_t - \dot{p}_t \equiv \xi_t. \tag{4}$$

The left-hand side of (4) is the institution's marginal utility flow over $[t, t + dt]$. The right-hand side is the opportunity cost of holding the asset: the cost of buying a share of the asset at time t and reselling it at $t + dt$, i.e., the time value of money, rp_t , minus the capital gain, \dot{p}_t . Finally, when $t \geq T_s$, $v_q(\ell, q_z) = v_q(\ell, 0) = 1$ and the price is $p_t = 1/r$.

Thus, the price increases deterministically towards $1/r$, as the holdings of low valuation institutions go to zero and their marginal utility increases. Institutions do not immediately bid up this predictable price increase because the demand for the asset builds up slowly: on the intensive margin, high-valuation institutions derive no utility flow if they hold more than one unit; and, on the extensive margin, the recovery from the aggregate liquidity shock occurs progressively as institutions switch back to high utility flows. Thus, there are “limits to arbitrage” in our model, in line with the empirical evidence on the predictable patterns of price drops and reversals around liquidity shocks.¹⁴

Throughout this paper we illustrate our results with numerical computations based on the parameter values shown in Table 1. We take the discount rate to be $r = 0.05$, in line with Duffie,

¹⁴ See, e.g., for short-lived shocks the empirical findings of Hendershott and Seasholes (2007), Hendershott and Menkveld (2010) and Khandani and Lo (2008).

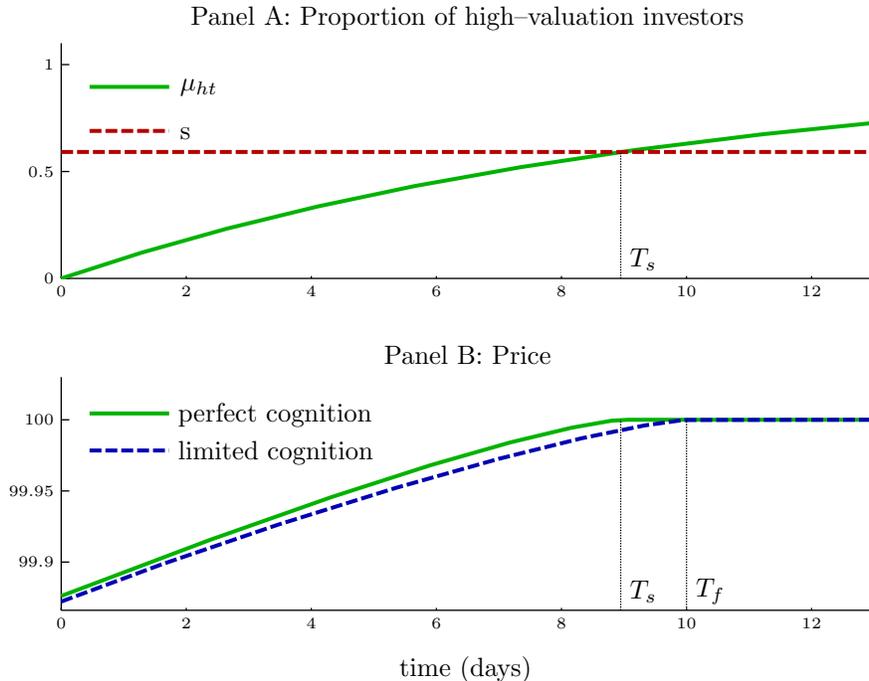


Figure 2: Population of high-valuation investors (Panel A) and price dynamics when $\sigma = 0.3$ (Panel B).

Gârleanu, and Pedersen (2007). We select the liquidity shock parameters to match empirical observations from large equity markets. Hendershott and Seasholes (2007) and Hendershott and Menkveld (2010) find liquidity price pressure effects of the order of 10 to 20 basis points, with duration ranging from 5 to 20 days. During the liquidity event described in Khandani and Lo (2008), the price pressure subsided in about 4 trading days. Adopting the convention that there are 250 trading days per year, setting γ to 25 means that an institution takes on average 10 days to switch to high valuation. Setting the asset supply s to 0.59 then implies that with unlimited cognition the time it takes the market to recover from the liquidity shock (T_s) is approximately 9 days, as illustrated in Figure 2, Panel A. Lastly, for these parameter values, setting $\delta = 1$ implies that the initial price pressure generated by the liquidity shock is between 10 and 20 basis points, as illustrated in Figure 2, Panel B.¹⁵

¹⁵Duffie, Gârleanu and Pedersen (2007) provide a numerical analysis of liquidity shocks in over-the-counter markets. They choose parameters to match stylized facts from illiquid corporate bond markets. Because we focus on more liquid electronic exchanges, we chose parameter values different from theirs. For example in their analysis the price takes one year to recover while in ours it takes less than two weeks.

Table 1: Parameter values

Parameter		Value
Intensity of information event	ρ	250
Asset supply	s	0.59
Recovery intensity	γ	25
Discount rate	r	0.05
Utility cost	δ	1
Curvature of utility flow	σ	{0.3, 0.5, 1.5}

3 Equilibrium with limited cognition

We now turn to the case where agents have limited cognition. In the first subsection below, we present our assumptions on cognition limits. In the second subsection we solve for equilibrium. To do so, we follow several standard steps: we first compute the value function of the traders, then we maximize this function to pin down demands, finally we substitute demands into the market clearing condition and we obtain the equilibrium price. In the third subsection we present the properties of the equilibrium, regarding welfare, holding plans, trading volume, and prices. In the fourth subsection we show that our qualitative results are upheld in an extension of the model where recurring preference shocks lead to stochastic equilibrium prices.

3.1 Assumptions

3.1.1 Limited cognition

Each institution is represented in the market by one trader.¹⁶ To determine optimal asset holdings, the trader must analyze the liquidity status of her institution, θ_t . This task is cognitively challenging. As mentioned in the previous section, to recover from the shock, the institution engages in several financial transactions in a variety of markets, some of them complex, opaque and not computerized. Evaluating the liquidity status of the institution requires collecting, analyzing and aggregating information about the resulting positions. Our key assumption is that, because of limited cognition and information processing constraints, the trader cannot continuously and immediately observe the liquidity status of her institution.¹⁷ Instead, we assume there is a counting process N_t such that the trader observes θ_t at each jump of N_t (and only

¹⁶For simplicity we abstract from agency issues and assume that the trader maximizes the inter-temporal expected utility of the institution.

¹⁷Regulators have recently emphasized the difficulty of devising an integrated measurement of all relevant risk exposures within a financial institution (see Basel Committee on Banking Supervision, 2009). Academic research has also underscored the difficulties associated with the aggregation of information dispersed in multiple departments of a financial institution (see Vayanos, 2003).

then).¹⁸ At the jumps of her information process N_t the trader submits a new optimal trading plan, based on rational expectations about $\{N_u, \theta_u : u \geq t\}$, and her future decisions. This is in line with the rational inattention model of Mankiw and Reis (2002). For simplicity, the traders' information event processes, N_t , are assumed to be Poisson distributed, with intensity ρ , and independent from each other as well as from the times at which institutions emerge from the liquidity shock.¹⁹

3.1.2 Conditions on asset holding plans and prices

When an information event occurs at time $t > 0$, a trader designs an updated asset holding plan, $q_{t,u}$, for all subsequent times $u \geq t$ until the next information event. Formally, denoting $D = \{(t, u) \in \mathbb{R}_+^2 : t \leq u\}$, we let an *asset holding plan* be a bounded and measurable stochastic process

$$\begin{aligned} q : D \times \Omega &\rightarrow \mathbb{R}_+ \\ (t, u, \omega) &\mapsto q_{t,u}(\omega), \end{aligned}$$

satisfying the following two conditions:

Condition 1. *For each $u \geq t$, the stochastic process $(t, \omega) \mapsto q_{t,u}(\omega)$ is \mathcal{F}_t -predictable, where $\{\mathcal{F}_t\}_{t \geq 0}$ is the filtration generated by N_t and θ_t .*

Condition 2. *For each (t, ω) , the function $u \mapsto q_{t,u}(\omega)$ has bounded variations.*

Condition 1 means that the plan designed at time t , $q_{t,u}$, can depend only on the trader's time- t information about her institution: the history of her information-event counting process, and of her institution utility status process up to, but not including, time t .²⁰ Condition 2 is a technical regularity condition ensuring that the present value of payments associated with $q_{t,u}$

¹⁸The time between jumps creates delays in obtaining fresh information about θ_t , which can be interpreted as the time it takes the risk management unit or head of strategy to aggregate all relevant information and disseminate it to the traders.

¹⁹For simplicity, we don't index the information processes of the different traders by subscripts specific to each trader. Rather we use the same generic notation, " N_t ", for all traders.

²⁰We add the "not including" qualifier because the asset holding plans are assumed to be \mathcal{F}_t -predictable instead of \mathcal{F}_t -measurable. This predictability assumption is standard for dynamic optimization problems involving decisions at Poisson arrival times (see Chapter VII of Brémaud, 1981). For much of this paper, however, we need not distinguish between \mathcal{F}_t -predictability and \mathcal{F}_t -measurability. This is because the probability that the trader type switches exactly at the same time an information event occurs is of second order. Therefore, adding or removing the type information accruing exactly at information events leads to almost surely identical optimal trading decisions.

is well defined. To simplify notations, from this point on we suppress the explicit dependence of asset holding plans on ω .

At this stage of the analysis, we assume that traders have access to a sufficiently rich menu of market instruments to implement any holding plan satisfying Conditions 1 and 2. We examine implementation in Section 4, where we analyze which market instruments are needed to implement equilibrium holding plans, and the equilibrium that arises when the menu of market instrument is not sufficiently rich.

The last technical condition concerns the price path:

Condition 3 (Well-behaved price path). *The price path is bounded, deterministic and continuously differentiable (C^1).*

As in the unbounded cognition case, because there is no aggregate uncertainty, it is natural to focus on deterministic price paths. Further, in the environment that we consider the equilibrium price must be continuous, as formally shown in our supplementary appendix (see Biais, Hombert, and Weill, 2010b, Section VI). The economic intuition is as follows. If the price jumps at time t , all traders who experience an information event shortly before t would want to “arbitrage” the jump: they would find it optimal to buy an infinite quantity of asset and re-sell these assets just after the jump. This would contradict market-clearing. Finally, the condition that the price be bounded is imposed to rule out bubbles (see Lagos, Rocheteau, and Weill, 2007, for a proof that bubbles cannot arise in a closely related environment).

3.2 Equilibrium

3.2.1 Intertemporal payoffs

For any time $u \geq 0$, let τ_u denote the time of the last information event before u , with the convention that $\tau_u = 0$ if no information event occurred. Correspondingly $q_{0,u}$ represents the holdings of a trader who had no information event by time u and thus no opportunity to update her holding plan. Given that all traders start with the same holdings at time zero, we have $q_{0,u} = q_{0,0} = s$ for all $u \geq 0$.

The trader’s objective is to maximize the inter-temporal expected value of utility flows, net of the cost of buying and selling assets. With the above notations in mind, this can be written

as:

$$\mathbb{E} \left[\int_0^\infty e^{-ru} \left(v(\theta_u, q_{\tau_u, u}) du - p_u dq_{\tau_u, u} \right) \right], \quad (5)$$

where $v(\theta_u, q_{\tau_u, u}) du$ is the utility enjoyed, and $p_u dq_{\tau_u, u}$ is the cost of asset purchases during $[u, u + du]$, given the holding plan chosen at τ_u , the last information event before u . Given our distributional assumptions for the type and information processes, and given technical Conditions 1 to 3 we can rewrite this objective equivalently as:

Lemma 1. *The inter-temporal payoffs associated with the holding plan $q_{t,u}$ is, up to a constant:*

$$V(q) = \mathbb{E} \left[\int_0^\infty e^{-rt} \int_t^\infty e^{-(r+\rho)(u-t)} \left\{ \mathbb{E}_t [v(\theta_u, q_{t,u})] - \xi_u q_{t,u} \right\} du \rho dt \right], \quad (6)$$

where $\xi_u = rp_u - \dot{p}_u$.

The interpretation of equation (6) is as follows. The outer expectation sign takes expectation over all time- t histories. The “ ρdt ” term in the outer integral is the probability that an information event occurs during $[t, t + dt]$. Conditional on the time- t history and on an information event occurring during $[t, t + dt]$, the inner integral is the discounted expected utility of the holding plan until the next information event. At each point in time this involves the difference between a trader’s time- t expected valuation for the asset, $\mathbb{E}_t [v(\theta_u, q_{t,u})]$, and the opportunity cost of holding that asset, ξ_u . This is similar to the result in Lagos and Rocheteau (2009) that an investor’s demand depends on his current marginal utility from holding the asset as well as his expected marginal utility in the future. Finally, the discount factor applied to time u is adjusted by the probability $e^{-\rho(u-t)}$ that the next information event occurs after u .

3.2.2 Market clearing

In all what follows we focus on the case where all traders choose the same holding plan, which is natural given that traders are *ex-ante* identical.²¹ Of course, while traders choose *ex-ante* the same holding plan, *ex-post* they realize different histories of N_t and θ_t , and hence different asset holdings.

The market clearing condition requires that, at each date $u \geq 0$, the cross-sectional average asset holding be equal to s , the per-capita asset supply. By the law of large numbers, and given

²¹By *ex-ante* identical we mean that traders start with the same asset holdings and have identically distributed processes for information event and utility status.

ex-ante identical traders, the cross-sectional average asset holding is equal to the expected asset holding of a representative trader. Hence, the market clearing condition can be written:

$$\mathbb{E}[q_{\tau_u, u}] = s. \tag{7}$$

for all $u \geq 0$. Integrating against the distribution of τ_u , and keeping in mind that $q_{0, u} = s$, leads to our next lemma:

Lemma 2. *The time- u market clearing condition, (7), writes:*

$$\int_0^u \rho e^{-\rho(u-t)} \left\{ (1 - \mu_{ht}) \mathbb{E}[q_{t, u} | \theta_t = \ell] + \mu_{ht} \mathbb{E}[q_{t, u} | \theta_t = h] - s \right\} dt = 0. \tag{8}$$

This lemma states that the aggregate net demand of traders who experienced at least one information event is equal to zero. The first multiplicative term in the integrand of (8), $\rho e^{-\rho(u-t)}$, is the density of time- t traders, i.e., traders whose last information event occurred at time $t \in (0, u]$. The first and second terms in the curly bracket are the gross demands of time- t low- and high-valuation traders respectively. The last term in the curly bracket is their gross supply. It is equal to s because information events arrive at random, which implies that the average holding of time- t traders just prior to their information event equals the population average.

3.2.3 Equilibrium existence and uniqueness

We define an equilibrium to be a pair (q, p) subject to Conditions 1, 2 and 3 and such that: i) given the price path, the asset holding plans maximize $V(q)$ given in (6), and ii) the holding plans are such that the market clearing condition (8) holds at all times. In this subsection we first present, in Lemmas 4 to 6, properties of holding plans implied by i) and ii). Then, based on these properties we obtain our first proposition, which states the uniqueness and existence of equilibrium and gives the equation for the corresponding price.

Going back to the value $V(q)$, in equation (6), and bearing in mind that a trader can choose any function $u \mapsto q_{t, u}$ subject to Conditions 1 and 2, it is clear that the trader inter-temporal problem reduces to point-wise optimization. That is, a trader whose last information event occurred at time t chooses her asset holding at time u , $q_{t, u}$, in order to maximize the difference

between her expected valuation for the asset and the corresponding holding cost:

$$\mathbb{E}_t [v(\theta_u, q_{t,u})] - \xi_u q_{t,u}. \quad (9)$$

Now, for all traders, utilities are strictly increasing for $q_{t,u} < 1$ and constant for $q_{t,u} \geq 1$. So, if one trader finds it optimal to hold strictly more than one unit at time u , then it must be that $\xi_u \leq 0$, implying that all other traders find it optimal to hold more than one unit. Inspecting equation (8), one sees that in that case the market cannot clear since $s < 1$. We conclude that:

Lemma 3. *In equilibrium, $\xi_u > 0$ and $q_{t,u} \in [0, 1]$ for all traders.*

To obtain further insights on holding plans, consider first a time- t high-valuation trader, i.e., a trader who finds out at time t that $\theta_t = h$. Such trader knows that her valuation for the asset will stay high forever. Hence

$$\mathbb{E}_t [v(\theta_u, q_{t,u})] = v(h, q_{t,u}), \forall u \geq t. \quad (10)$$

Next, consider a time- t low-valuation trader, i.e. a trader who finds out at time t that $\theta_t = \ell$. This trader anticipates that her utility status will remain low by time u with probability $(1 - \mu_{hu})/(1 - \mu_{ht})$. Hence:

$$\mathbb{E}_t [v(\theta_u, q_{t,u})] = q_{t,u} - \delta \frac{1 - \mu_{hu}}{1 - \mu_{ht}} \frac{q_{t,u}^{1+\sigma}}{1 + \sigma}, \forall q_{t,u} \in [0, 1] \quad (11)$$

Comparing (10) and (11), one sees that, for all asset holdings in $(0, 1)$, high-valuation traders have a uniformly higher marginal utility than low-valuation traders. Now let

$$S_u \equiv \int_0^u \rho e^{-\rho(u-t)} (s - \mu_{ht}) dt, \quad (12)$$

the gross asset supply brought by all traders minus the maximum (unit) demand of high-valuation traders, integrating across all traders with at least one information event.

Given the definition of S_u , and based on the above ranking of marginal utilities, one sees that the economy can be in one of two regimes. The first regime arises if $S_u > 0$: in that case, since $q_{t,u} \leq 1$ for high-valuation traders, market-clearing implies that $q_{t,u} > 0$ for some low-valuation trader. But then all high-valuation traders must hold one unit, since they have uniformly higher marginal utility for holdings in $[0, 1]$. The second regime arises if $S_u < 0$. In this case market clearing implies that some high-valuation trader must find it optimal to hold strictly less than

one share: such trader either strictly prefer to hold zero share, or is indifferent between any holding in $[0, 1]$. But since high-valuation traders have uniformly higher marginal utility for holdings in $[0, 1]$, this implies that all low-valuation traders hold zero shares. Summarizing:

Lemma 4. *Let T_f be the unique strictly positive solution of $S_u = 0$. Then:*

- if $u \in (0, T_f)$ then, for all $t \in (0, u]$, $\theta_t = h$ implies $q_{t,u} = 1$;
- if $u \in [T_f, \infty)$ then, for all $t \in (0, u]$, $\theta_t = \ell$ implies $q_{t,u} = 0$.

Next, consider the demand of high-valuation traders when $u > T_f$. We know from the previous lemma that low-valuation traders hold no asset. Thus, high-valuation traders must hold the entire asset supply. Moreover, since $S_u < 0$, the market-clearing condition implies that some high-valuation traders must hold strictly less than one share. Keeping in mind that high-valuation traders have the same linear utility flow over $[0, 1]$, this implies they must be indifferent between any holding in $[0, 1]$. Thus we can state the following lemma.

Lemma 5. *For all $u > T_f$, the average asset holding of a high-valuation trader is*

$$\frac{\int_0^u \rho e^{-\rho(u-t)} s dt}{\int_0^u \rho e^{-\rho(u-t)} \mu_{ht} dt},$$

but the distribution of asset holdings across high-valuation traders is indeterminate.

Now turn to the demand of low-valuation traders when $u < T_f$. Taking first-order conditions when $\theta_t = \ell$ in (9), we obtain, given $q_{t,u} \in [0, 1]$:

$$q_{t,u} = 0 \quad \text{if } \xi_u \geq 1 \tag{13}$$

$$q_{t,u} = 1 \quad \text{if } \xi_u \leq 1 - \delta \frac{1 - \mu_{hu}}{1 - \mu_{ht}} \tag{14}$$

$$q_{t,u} = (1 - \mu_{ht})^{1/\sigma} Q_u \quad \text{if } \xi_u \in \left(1 - \delta \frac{1 - \mu_{hu}}{1 - \mu_{ht}}, 1 \right), \quad \text{where } Q_u \equiv \left[\frac{1 - \xi_u}{\delta(1 - \mu_{hu})} \right]^{1/\sigma}. \tag{15}$$

Equation (13) states that low-valuation traders hold zero unit if the opportunity cost of holding the asset is greater than 1, their highest possible marginal utility, which arises when $q = 0$. Equation (14) states that low-valuation traders hold one unit if the opportunity cost of holding the asset is below the lowest possible marginal utility, which arises when $q = 1$. Lastly, equation (15) pins down a low-valuation trader's holdings in the intermediate interior case by equating to 0 the derivative of (11) with respect to $q_{t,u}$.

As discussed above, prior to time T_f the holdings of some low-valuation traders must be strictly greater than zero: thus, holdings are determined by either (14) or (15) and $\xi_u > 0$. By the definition of Q_u , $\xi_u \leq 1 - \delta(1 - \mu_{hu})/(1 - \mu_{ht})$ if and only if $(1 - \mu_{ht})^{1/\sigma} Q_u \geq 1$. Hence, the asset demand defined by (14) and (15) can be written as

$$q_{t,u} = \min\{(1 - \mu_{ht})^{1/\sigma} Q_u, 1\}. \quad (16)$$

Substituting the demand from (16) into the market-clearing condition (8) and using the definition of S_u in (12), the following lemma obtains.

Lemma 6. *If $u \in (0, T_f)$, then for all $t \in (0, u]$, $\theta_t = \ell$ implies $q_{t,u} = \min\{(1 - \mu_{ht})^{1/\sigma} Q_u, 1\}$ where:*

$$\int_0^u (1 - \mu_{ht}) \min\{(1 - \mu_{ht})^{1/\sigma} Q_u, 1\} \rho e^{-\rho(u-t)} dt = S_u. \quad (17)$$

Equation (17) is a one-equation-in-one-unknown for Q_u that is shown in the proof appendix to have a unique solution. Taken together, Lemmas 4 through 6 imply:

Proposition 1. *There exists an equilibrium. The equilibrium asset allocation is unique up to the distribution of asset holdings across high-valuation traders after T_f , and is characterized by Lemma 4-6. The equilibrium price path is unique, is increasing until T_f , constant thereafter, and solves the following ODE:*

$$u \in (0, T_f) : \quad rp_u - \dot{p}_u = 1 - \delta(1 - \mu_{hu})Q_u^\sigma \quad (18)$$

$$u \in [T_f, \infty) : \quad p_u = \frac{1}{r}. \quad (19)$$

As in the perfect cognition case, the price deterministically increases until it reaches $1/r$. One difference is that, while under perfect cognition this recovery occurs at time T_s (defined in equation (2)), with limited cognition it occurs at the later time $T_f > T_s$. For $u < T_f$, the time- u low-valuation traders are the marginal investors, and the equilibrium price is such that their marginal valuation is equal to the opportunity cost of holding the asset, as stated by (18). For $u > T_f$, the entire supply is held by high-valuation investors. Thus the equilibrium price is equal to the present value of their utility flow, as stated by (19).²² This proposition is illustrated

²²They must be indifferent between trading or not. This indifference condition implies that $1 - rp_u + \dot{p}_u = 0$. And, $p_u = 1/r$ is the only bounded and positive solution of this ODE.

in Figure 2, Panel B, which plots the equilibrium price under limited cognition. Note that for this numerical analysis we set the intensity of information events ρ to 250, which means that traders observe refreshed information on θ on average once a day. This is a plausible frequency, given the time it takes to collect and aggregate information across desks, departments and markets in a financial institution.

3.3 Equilibrium properties

In this subsection we discuss the properties of the equilibrium price and trades and compare them to their unbounded cognition counterparts.

3.3.1 Welfare

To study welfare we define an asset holding plan to be *feasible* if it satisfies Conditions 1 and 2 as well as the resource constraint, which is equivalent to the market-clearing condition (7). Furthermore, we say that an asset holding plan q *Pareto dominates* some other holding plan q' if it is possible to generate a Pareto improvement by switching from q' to q while making time zero transfers among traders. Because utilities are quasi linear, q Pareto dominates q' if and only if $W(q) > W(q')$, where

$$W(q) = \mathbb{E} \left[\int_0^\infty e^{-ru} v(\theta_{\tau_u}, q_{\tau_u, u}) du \right]. \quad (20)$$

The next proposition states that in our model the first welfare theorem holds:

Proposition 2. *The holding plan arising in the equilibrium characterized in Proposition 1 maximizes $W(q)$ among all feasible holding plans.*

This proposition reflects that, in our setup, there are no externalities, in that the holdings constraints imposed by limited cognition for one agent (and expressed in conditions 1 and 2) do not depend on the actions of other agents. These constraints translate into simple restrictions on the commodity space (conditions 1 and 2), allowing us to apply the standard proof of the first welfare theorem (see Mas-Colell, Whinston, and Green, 1995, Chapter 16)

3.3.2 Holdings

As set forth in equation (16), for a trader observing at t that her valuation is low, the optimal holdings at time $u > t$ are (weakly) increasing in Q_u . Relying on the market clearing condition,

the next proposition spells out the properties of Q_u .

Proposition 3. *The function Q_u is continuous, such that $Q_{0^+} = s$ and $Q_{T_f} = 0$. Moreover, if*

$$s \leq \frac{\sigma}{1 + \sigma} \tag{21}$$

Q_u is strictly decreasing with time. Otherwise, it is hump-shaped.

The economic intuition is as follows. At time 0^+ the mass of traders with high-valuation is negligible. Therefore low-valuation traders must absorb the entire supply. Hence, $Q_{0^+} = s$. At time T_f high valuation traders absorb the entire supply. Hence, $Q_{T_f} = 0$.

When the per-capita supply of assets affected by the shock s is low, so that (21) holds, the incoming flow of high-valuation traders reaching a decision at a given point in time is always large enough to accommodate the supply from low-valuation traders. Correspondingly, in equilibrium low valuation traders sell a lump of their assets when they reach a decision, then smoothly unwind their inventory until the next information event.

In contrast, when s is so large that (21) fails to hold, the liquidity shock is more severe. Hence, shortly after the initial aggregate shock, the inflow of high-valuation traders is not large enough to absorb the sales of low-valuation traders who currently reach a decision. In equilibrium, some of these sales are absorbed by “early” low-valuation traders who reached a decision at time $t < u$ and have not had another information event. Indeed, these “early” low-valuation traders anticipate that, as time goes by, their institution is more likely to have recovered. Thus, their expected valuation (in the absence of an information event) increases with time and they find it optimal to buy if their utility is not too concave, i.e., if σ is not too high. Correspondingly, near time zero, Q_u is increasing, as depicted in Figure 3 for $\sigma = 0.3$ and 0.5.

Combining Lemma 4, Lemma 5, Lemma 6 and Proposition 3, one obtains a full characterization of the equilibrium holdings process, which can be compared to its counterpart in the unlimited cognition case. When cognition is not limited, as long as an institution has not recovered from the shock, its holdings decline smoothly, and, as soon as it recovers, its holdings jump to 1. Trading histories are quite different with limited cognition. First an institution’s holdings remain constant until its trader’s first information event. Then, if at her first information event the trader learns that her institution has a low valuation, the trader sells a lump. After that, if (21) does not hold, the trader progressively buys back, then eventually sells out

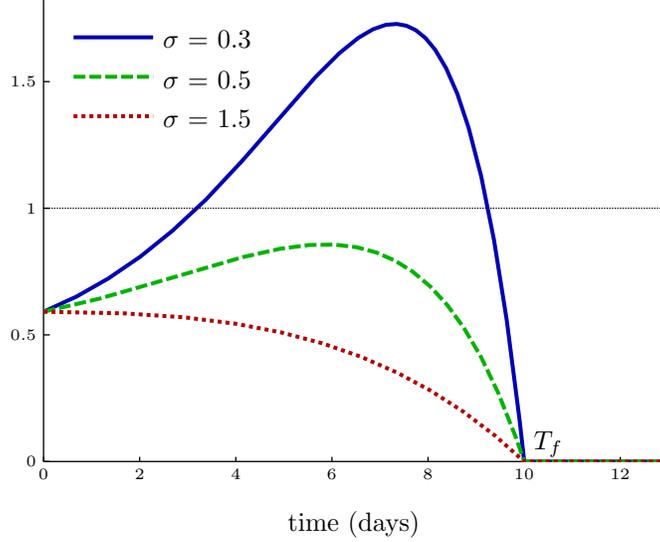


Figure 3: The function Q_u for various values of σ

at the next jump of her information process. This process continues until the trader finds out her valuation has recovered, at which point her holdings jump to 1. Such round-trip trades do not arise in the unbounded cognition case.

3.3.3 Trading volume

Because they result in round-trip trades, hump-shaped asset holding plans generate extra trading volume relative to the unlimited cognition case. Specifically, consider a trader who, at two consecutive information events t_1 and t_2 , discovers that she has a low-valuation. During the time period $(t_1, t_2]$ she trades an amount of asset equal to

$$\int_{t_1}^{t_2} \left| \frac{\partial q_{t_1, u}}{\partial u} \right| du + |q_{t_2, t_2} - q_{t_1, t_2}|. \quad (22)$$

The first term in (22) is the flow of trading between time t_1 and time t_2 dictated by $q_{t_1, u}$, the time- t_1 holding plan. The second term is the lumpy adjustment at time t_2 .

Note that (16) implies that, whenever $|\partial q_{t_1, u} / \partial u|$ is not 0, it has the same sign as Q'_u . Note also that, because at time t_2 the trader observes that the institution has still not recovered, $q_{t_2, t_2} < q_{t_1, t_2}$. Hence, if Q_u is decreasing, (22) is equal to $q_{t_1, t_1} - q_{t_2, t_2}$. In contrast, if Q_u is increasing, the amount traded between t_1 and t_2 is

$$q_{t_1, t_2} - q_{t_1, t_1} + q_{t_1, t_2} - q_{t_2, t_2} = \underbrace{2(q_{t_1, t_2} - q_{t_1, t_1})}_{\text{round-trip trade}} + q_{t_1, t_1} - q_{t_2, t_2} \quad (23)$$

Since, $q_{t_1, t_2} > q_{t_1, t_1}$, (23) is greater than $q_{t_1, t_1} - q_{t_2, t_2}$. The first term in the equation is the extra volume created by the round-trip trade: the purchase of $q_{t_1, t_2} - q_{t_1, t_1}$ during (t_1, t_2) followed by a sale of the same quantity at time t_2 .

The comparison between trading volume with unbounded cognition and with limited cognition depends on two effects: i) On one hand, the above discussion shows that cognition limits generate round-trip trades that tend to raise trading volume. ii) On the other hand, since information events are infrequent, agents' ability to realize mutually beneficial trades is limited (e.g., if $\rho = 0$ there are no trades.) For the parameter values we consider, Figure 4 shows that, for early times, effect i) dominates so that trading volume is greater with limited cognition. However, as shown in the figure, for later times, effect ii) dominates so that trading volume is lower with limited cognition. Indeed, round-trip trades become less prevalent, and stop after the time T_ψ when Q_u reaches its maximum. When ρ goes to infinity, there is only effect i) so that, as stated in the next proposition, trading volume is at least as large with limited cognition.

Proposition 4. *When ρ goes to infinity, the equilibrium price and allocation converge to their unlimited cognition counterparts. Relative to unlimited cognition, the extra volume at time t with limited cognition converges to*

$$\gamma \max \left\{ \frac{s - \mu_{ht}}{\sigma} - (1 - s), 0 \right\}.$$

When condition (21) does not hold, this asymptotic extra volume is strictly positive for small t .

The proposition reveals the crucial role of the curvature parameter, σ : as illustrated in Figure 3, when σ is small and utility is close to linear, the hump-shaped pattern in holdings is quite pronounced. This leads to very large extra trading volume

3.3.4 Prices

Proposition 1 implies that, from time T_s to time T_f , the price under limited cognition is lower than its counterpart with unlimited cognition, $1/r$. By continuity, the price is lower under cognition limits just before T_s . The next proposition states that this ranking of prices holds at all times if $s < \sigma/(\sigma + 1)$, but not necessarily otherwise.

Proposition 5. *If (21) holds, then at all time the price is strictly lower with limited cognition. But if s is close to 1 and σ is close to 0, then at time 0 the price is strictly higher with limited cognition.*

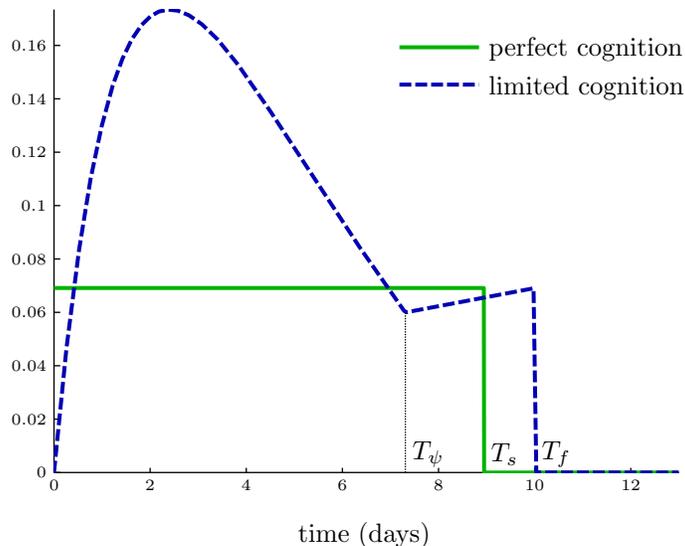


Figure 4: The trading volume under unlimited (green) and limited cognition (blue), when $\sigma = 0.3$. In the figure, T_ψ denotes the argument maximum of Q_u .

That the price would be higher without cognition limits sounds intuitive. Unbounded cognition enables traders to continuously allocate the asset to those who value it the most. Such an efficient allocation could be expected to raise the price, and this is indeed what happens when (21) holds. But, as stated in Proposition 5, there are cases where the price can be higher when cognition is limited than when it is unbounded. The intuition is as follows. Around time zero, low-valuation traders are marginal. With limited cognition, low-valuation traders have a higher expected marginal utility because they take into account the possibility that they may have switched to high-valuation. Consequently, they demand more assets, which tends to push up prices. This effect is stronger when low-valuation traders are marginal for a longer period, that is, when the shock is more severe (s close to one) and when their utility flow is not too concave (σ close to zero).

3.4 Robustness to recurring liquidity shocks

So far we focused on a simple model generating a deterministic equilibrium price process. In this subsection, we show that our qualitative results are upheld in a more general model, with stochastic shocks and stochastic prices. Thus, our baseline one-shock analysis can be viewed, at least qualitatively, as the impulse response of a more general model with recurring shocks.

In the spirit of Duffie, Gârleanu, and Pedersen (2007) we assume that aggregate liquidity shocks occur at Poisson arrival times with intensity $\kappa > 0$. As in the baseline model, when

a shock hits, all investors switch to the low-valuation state and recover later at independent exponential times with intensity γ . We also assume that, at the time of a shock, all traders withdraw from the market and halt trading, to re-design their holding plans. Traders resume trading and revise their holding plan at the first subsequent jump of their information process.²³ We focus on stationary equilibria in which the price is stochastic, but can be written as function of one random variable: the time elapsed since the last aggregate shock.

In Section IV, page 18 of the supplementary appendix, we show that the structure of the equilibrium remains qualitatively the same as in our baseline model. In particular, the market clearing condition is the same as in Lemma 2. Also, a trader's objective function is

$$\int_t^\infty e^{-(r+\rho+\kappa)(u-t)} \left\{ \mathbb{E}_t [v(\theta_u, q_{t,u})] + \kappa W(q_{t,u}) - \xi_u q_{t,u} \right\} du$$

where $\xi_u \equiv rp_u - \dot{p}_u - \kappa(p_0 - p_u)$,

and the time index, t or u , denote the time elapsed since the last liquidity shock. This is similar to the value function in Lemma 1, but with two changes. First, the definition of the opportunity cost of holding asset is modified so as to reflect the possibility that a new shock might hit the market with intensity κ , bringing down the price to p_0 . Second, the utility flow is augmented by $\kappa W(q_{t,u})$, where $W(q)$ is the value of holding q units of the asset at the time of a new aggregate liquidity shock, until the first subsequent information event. In Section IV, page 18 of the supplementary appendix, we solve for $W(q)$, compute the optimal holdings and, substituting them in the market clearing condition, obtain the following proposition:

Proposition 6. *With recurrent random aggregate liquidity shocks, there exists an equilibrium in which i) the price jumps each time the market is hit by a shock, ii) in between jumps the price is the function of the time u elapsed since the last shock which is the unique bounded solution of the ODE:*

$$(r + \kappa)p_u - \dot{p}_u = 1 + \frac{\kappa}{r} - \frac{\delta\kappa(r + \rho + \kappa)}{(r + \rho)(r + \rho + \kappa + \gamma)} Q_{h,u}^\sigma - \frac{\delta\kappa^2(r + \kappa)(r + \rho + \kappa)}{r(r + \rho)(r + \rho + \kappa + \gamma)} \int_0^\infty (e^{-(r+\kappa)z} - e^{-(r+\rho+\kappa)z}) Q_{h,z}^\sigma dz,$$

²³This is in line with evidence from the “flash crash” of May 6th, 2010: the Security and Exchanges Commission reports that automated trading systems paused in reaction to the sudden price decline in order to allow traders and risk managers to fully assess the risks before trading was resumed. See www.sec.gov/news/studies2010marketevents-report.pdf online.

where $Q_{h,u}$ is defined in equation (IV.6), page 21 of the supplement.

In the supplement we provide various numerical examples based on our analytical characterization. We consider various values of κ , set $\sigma = 0.3$, and otherwise keep the same parameter values as in Table 1. In the example where aggregate liquidity shocks occur every quarter on average ($\kappa = 4$) we find that holding policies are hump-shaped and that the price is increasing, just as in our baseline model. However, since traders anticipate the occurrence of future liquidity shocks, the price level is about 9.5% lower.

4 Implementation with realistic market instruments

So far, our characterization of equilibrium has been cast in abstract terms, such as holding plans and market clearing. We now study how these holding plans can be implemented with realistic market instruments. In so doing, we focus on electronic order driven markets. Such venues are the major trading mechanism for stocks around the world (e.g., in the US NASDAQ and the NYSE, and in Europe Euronext, the London Stock Exchange and the Deutsche Börse.) In these markets, traders can place limit sell (resp. buy) orders requesting execution at prices at least as large (resp. low) as their limit price. These orders are stored in the book, until they are executed, cancelled or modified. Traders can also place market orders, requesting immediate execution. A limit sell order standing in the book is executed, at its limit price, when hit by an incoming buy order (either a market order or a limit buy order with a higher price limit), if there are no unexecuted sell orders in the book at lower prices, or at the same price but placed at an earlier point in time. The case of a limit buy order is symmetric.²⁴

There are multiple possible implementations of our equilibrium. For instance, there is a trivial implementation, where all traders desiring to change their holdings continuously submit market orders or limit orders at the current equilibrium price. While this has the advantage of simplicity, this implementation is at odds with important stylized facts; for example it leads to an empty limit order book. To narrow down the set of possible implementations while giving rise to realistic dynamics, we restrict our attention to the case where market participants alter their trading strategies only when their information process jumps. This is quite natural in our framework, and it follows if, between two jumps of their information process, traders' cognitive

²⁴Note also that a limit order to sell and a limit order to buy placed at the same price cannot both remain together in the limit order book: as soon as offers cross each other they are executed, in accordance with price and time priority rules.

resources are devoted to the complex task of assessing the liquidity status of their institution, which leaves little opportunity to alter trading plans.

4.1 Implementing the equilibrium with limit orders and algorithms

First we assume that, when their information process jumps, in addition to market and limit orders, traders can place trading algorithms. The latter are computer programs feeding orders to the market over time, in response to pre-specified future changes in market variables.²⁵ Note that both limit orders and algorithms enable trades to happen without direct human intervention while the trader is engaged in information collection and processing.²⁶ Lastly, in keeping with our limited-cognition assumption, we do not allow algorithms to condition their orders on changes in the liquidity status of the institution occurring between jumps of its information process.

It is straightforward to implement the holding plan of high-valuation traders. Before T_f , as soon as their information process reveals that their institution has recovered from the shock, they place a limit order to buy at the current price, which is immediately executed.²⁷ At that point in time, since they now have their optimal holdings, they cancel any limit order they would have previously placed in the book.

The case of low-valuation traders is more intricate. Indeed, one must bear in mind that the equilibrium price given in Proposition 1 is strictly increasing over $(0, T_f)$. This implies that any limit order to buy submitted at time t is either immediately executed (if the limit price is greater than p_t) or never executed (if the limit price is lower than p_t). Consequently, if a trader only places (limit or market) orders when her information process jumps, she cannot implement increasing holding plans. In contrast, when the equilibrium holding plan is decreasing, it can be implemented by placing at time t an immediately executed market order to sell, as well as limit orders to *sell* at price $p_u > p_t$ which will be executed later. Now, Proposition 3 states that equilibrium holdings are decreasing if and only if condition (21) holds. This leads to the following Proposition:

²⁵Because we have no aggregate uncertainty, market-level variables (price, volume, quote, etc.) are in our model deterministic functions of time. Thus, conditioning on time is enough to make the algorithm also depend on the state of the market. In a more general model with aggregate uncertainty, one would need to explicitly allow algorithms to depend on market-level variables at time u .

²⁶This is in line with the view of Harris (2003), who argues that “Limit orders represent absent traders [enabling them] to participate in the markets while they attend to business elsewhere.”

²⁷Equivalently, they could place a market order to buy. Assuming they place a limit order to buy at p_t pins down the market clearing price.

Proposition 7. *The equilibrium characterized in Proposition 1 can be implemented by traders placing market and limit orders (only) when their information process jumps if and only if (21) holds.*

If (21) does not hold, asset holding plans are hump shaped. As argued above, in order to implement the increasing branch of the hump, traders cannot use limit buy orders: instead, they must rely on algorithms.

For concreteness, consider the example illustrated in Figure 5. Interpret t_1 and t_2 as the time of two consecutive jumps of the information process of a trader, both revealing low valuation. At t_2 the trader's asset holdings undergo a discrete downward jump, from q_{t_1, t_2} to q_{t_2, t_2} . Then they continuously increase, reach a flat, and finally decrease again. To implement this holding plan, the trader places a market sell order at time t_2 , which implements the initial downward jump in her optimal holdings. She also programs her algorithm to place a sequence of market buy orders after t_2 , which implement the increasing part of her holding plan. Finally, to implement the smoothly decreasing part of her holding plan, she uses a schedule of limit sell orders.

As shown on Figure 5, the limit sell orders placed at time t_2 start executing before those placed at time t_1 . Since the equilibrium price process is increasing, this means that these orders are placed at lower prices than the previous ones. The figure also shows that the holding plan set at time t_2 declines less steeply than its t_1 counterpart. This reflects that the quantity offered in the book at corresponding prices is lower for the plan set at t_2 than for the plan set at t_1 . Thus, to implement the new holding plan, at t_2 the trader cancels some of the orders placed at t_1 , replacing them with more aggressive orders. Finally note that, at t_2 , the trader also modifies the trading algorithm generating the purchases necessary to implement the increasing part of her holding plan. This can be interpreted in terms of human intervention resetting the parameters of the trading algorithm. This discussion is summarized in our next proposition.

Proposition 8. *Consider a trader observing low valuation at time $t < T_f$. If (21) does not hold such trader can implement the optimal holding plan arising in Proposition 1 by placing market and limit sell orders at t and programming her trading algorithm to trigger market buy orders at times $u > t$.*

Returning to Figure 5, now interpret t_1 and t_2 as the times at which the information processes of two traders jump, each time revealing low valuation. The figure illustrates that the late trader starts selling before the early trader implying that the limit sell orders of the late trader are

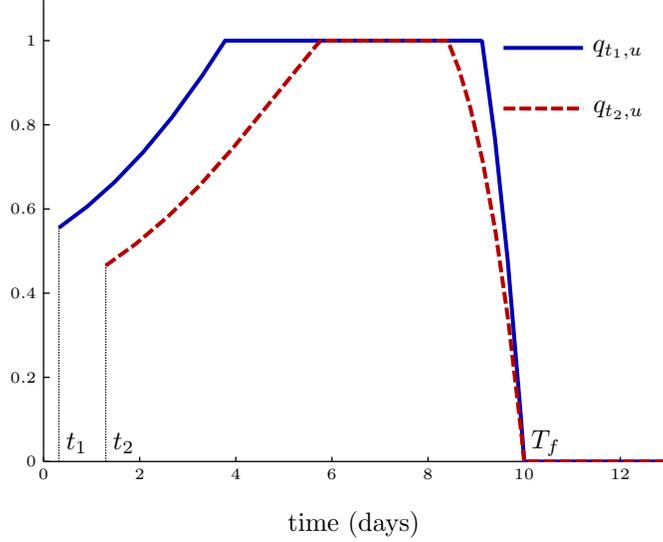


Figure 5: The holding plans of an early low-valuation trader and of a late low-valuation trader, when $\sigma = 0.3$

placed at lower prices than the limit sell orders of the early trader, i.e., there is undercutting.²⁸ To see why this is optimal, compare the time- u expected valuations of the time- t_1 trader and the time- t_2 trader, assuming that for both there has been no new information event by time u . For the trader who observed low valuation at time t_1 the probability that her valuation is high now is: $1 - e^{-\gamma(u-t_1)}$. For the other trader it is $1 - e^{-\gamma(u-t_2)}$, which is lower since $t_2 > t_1$. Hence the expected valuation of the time- t_1 trader is greater than that of the time- t_2 trader. This is why the time- t_1 trader sells later.²⁹

Taking stock of the above results, we now describe the overall market dynamics prevailing when (21) does not hold and $Q_u > 1$ for some u . There are four successive phases in the market, as illustrated in Figure 6.

- There exists a time $T_1 < T_f$ such that, from time 0 to T_1 , low valuation traders place limit sell orders at lower and lower prices, i.e., there is undercutting. These limit orders accumulate in the book, without immediate execution. Correspondingly, the best ask decreases and the depth on the ask side of the book increases. During this period, low valuation traders also place market sell orders, which are executed against limit buy orders placed by high-valuation traders and algorithms.

²⁸This is in line with the empirical results of Biais, Hillion, and Spatt (1995), Griffiths, Smith, Turnbull, and White (2000) and Ellul, Holden, Jain, and Jennings (2007).

²⁹This goes along the same lines as the intuition why low valuation traders placing trading plans at t will initially sell and then buy from agents placing their trading plans later.

- Denoting by $T_\psi \in (T_1, T_f)$ the time at which Q_u achieves its maximum, during $[T_1, T_\psi]$ the best ask quote remains constant and the depth on the ask side of the book declines. Indeed, during this period, low valuation traders stop undercutting the best quote, while high valuation traders cancel their limit orders. During this second phase of the market, there are still no executions at the best ask, and trades are initiated by low valuation traders placing market sell orders, executed against the limit buy orders placed by high valuation traders.
- The third phase is between T_ψ and T_f . During this period, high valuation traders hit the limit sell orders outstanding in the book with market orders to buy (or marketable limit orders). Correspondingly, the depth on the ask side continues to decline, and the best ask price increases.³⁰
- Finally, after time T_f , the market has recovered from the shock, and the price remains constant at $1/r$.

The algorithmic trading strategies generated by our model are in line with stylized facts and empirical findings. That algorithms progressively build up an increasing position via successive buy orders can be interpreted in terms of order splitting. That they buy progressively as the price deterministically trends upwards can be interpreted in terms of short-term momentum trading. That after the liquidity shock traders build up long positions which they will eventually unwind is in line with the findings of Brogaard (2010) that algorithms do not withdraw after large price drops and benefit from price reversals. Our theoretical results are also in line with the empirical findings of Hendershott and Riordan (2010) that trading algorithms provide liquidity when it is scarce and rewarded. Indeed, the strategies followed by our algo-traders (who buy initially while simultaneously placing limit orders to sell to be executed later) is a form of market-making, similar to that arising in Grossman and Miller (1988). Note however that, while in Grossman and Miller some market participants are exogenously assumed to be liquidity providers and other liquidity consumers, in our model all participants are identical

³⁰Note that there is no bid-ask bounce in this market. Until time T_ψ all trades are executed at the limit price of buy orders stemming from traders who just realized their valuation was high. After T_ψ all trades are executed against the best limit order to sell outstanding in the book. Note also that the tick size is zero, so that prices are not restricted to a discrete grid. Thus there is a sequence of executions at smoothly increasingly high bid prices, followed by a sequence of executions at smoothly increasing ask prices. Thus in equilibrium there is positive serial correlation in successive order types, in line with the empirical findings of Biais, Hillion, and Spatt (1995), Griffiths, Smith, Turnbull, and White (2000) and Ellul, Holden, Jain, and Jennings (2007).

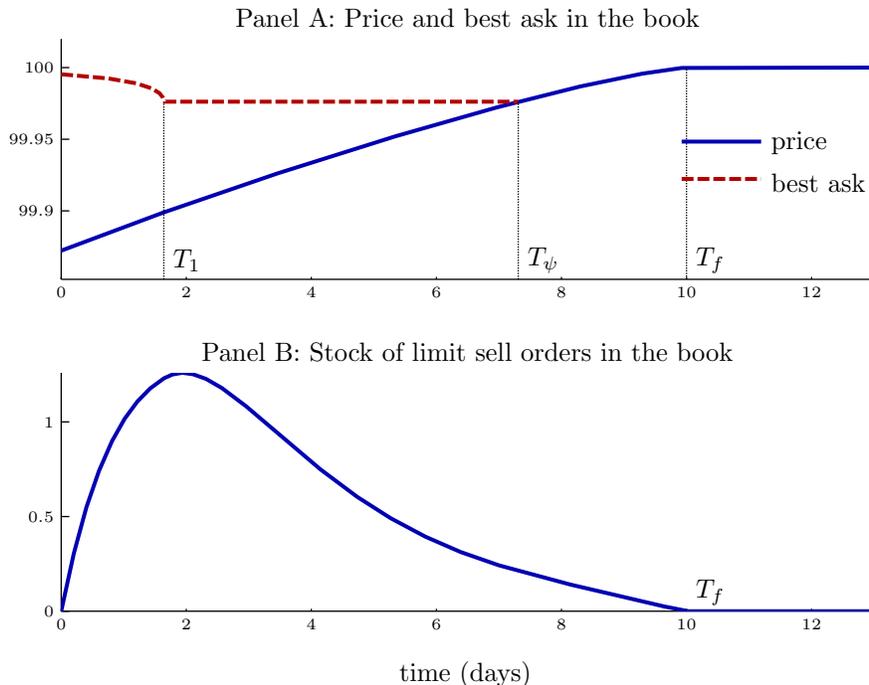


Figure 6: Price dynamics (Panel A) and limit order book activity (Panel B), when $\sigma = 0.3$. Supplementary appendix VIII.18 provides the formulas leading to Panel B.

ex-ante, yet they play different roles in the market because of differences in the realizations of their information and valuation processes.

4.2 Equilibrium without trading algorithms

What happens if trading algorithms are not available and traders can only place limit and market orders at the time of information events? Suppose for now that the equilibrium price is increasing (which will turn out to be the case in equilibrium). As mentioned in Section 4.1 if traders place orders only when their information process jumps, this precludes them from implementing increasing holding plans. Instead, they can implement only decreasing holding plans. Therefore, if condition (21) does not hold, the equilibrium will differ from that arising in Proposition 1. Instead, the equilibrium is as in the next proposition.

Proposition 9. *If $s > \sigma/(1 + \sigma)$ and traders can place only limit and market orders when their information process jumps, there exists an equilibrium in which the price path is strictly increasing over $(0, T_f)$, and equal to $1/r$ over $[T_f, \infty)$. High-valuation traders, and low-valuation traders after T_f , follow the same asset holding plan as in Proposition 1. For low-valuation traders before T_f , the optimal holding plan $q_{t,u}$ is continuous in (t, u) and strictly less than 1.*

Moreover, there exists $T_\phi \in (0, T_f)$ and a strictly decreasing function $\phi : (0, T_\phi] \mapsto \mathbb{R}_+$, such that:

- If $t \in (0, T_\phi]$, then $q_{t,u}$ is constant for $u \in [t, \phi_t]$, and strictly decreasing for $u \in (\phi_t, T_f)$.
- If $t \in (T_\phi, T_f]$, then $q_{t,u}$ is strictly decreasing for $u \in (t, T_f)$.

While the proposition describes the equilibrium in abstract terms, in the the appendix we provide closed-form analytical solutions for all relevant equilibrium objects. Also, one may wonder whether the equilibrium in Proposition 9 is unique. We provide a partial answer to this question in the supplementary appendix to this paper. We show that the equilibrium of Proposition 9 is unique in the class of Markov equilibria, i.e., equilibria where traders find it optimal to choose holding plans that depend only on the information-event time (the current aggregate state) and their current utility status (their current idiosyncratic state).³¹

The intuition for Proposition 9 is as follows. When $s > \sigma/(1+\sigma)$, time- $t < T_\phi$ low-valuation traders would choose hump-shaped holding plans if their holdings were not constrained to be decreasing (the solid curve in Figure 7). Faced with the constraint of choosing a decreasing holding plan, they “iron” the increasing part of the hump-shaped plans (the dashed curve in Figure 7).

To implement the holding plans of this proposition, time- t low-valuation traders place market sell orders, as well as schedules of limit sell orders, which begin executing at time ϕ_t . Proposition 9 also implies that, similarly to the case where traders can use algorithms, there is undercutting in equilibrium: since ϕ_t is strictly decreasing, successive traders place limit sell orders at lower and lower prices.³²

Note that, both in Proposition 1 and Proposition 9, the time at which the price fully recovers from the liquidity shock is T_f , the time at which the residual supply S_u (defined in (12)) reaches 0. This is because S_u is a function of the total quantity of the asset brought to the market and of the unit demand of high-valuation traders, both of which are unaffected by whether or not low-valuation traders can use algorithms. The next proposition offers further insights into the comparison of the price paths arising in Proposition 9 and Proposition 1.

³¹Importantly, our proof makes no *a priori* monotonicity restriction on the price path. Instead, we consider general and possibly non-monotonic price paths. We then show, via elementary optimality and market clearing considerations, that the preference dynamics and the focus on Markov equilibria imply that the price path is continuous, strictly increasing until time T_f , and flat after time T_f .

³²Note that in Proposition 9 we obtain undercutting for all (s, σ) such that $\sigma/(1+\sigma) < s$. In Proposition 1, by contrast, we obtain undercutting for the smaller set of parameters such that the maximum of Q_u is greater than 1.

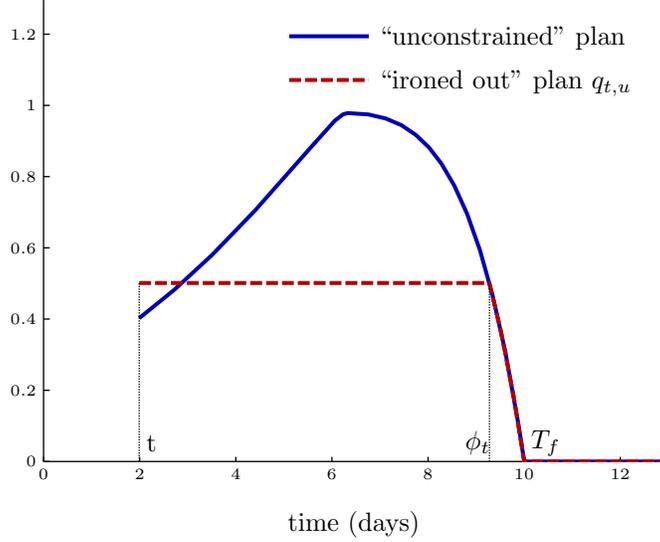


Figure 7: Holding plans of algorithmic versus limit order traders, when $\sigma = 0.3$, given the equilibrium price of Proposition 9.

Proposition 10. *When $s > \sigma/(1 + \sigma)$:*

- *The price arising in Proposition 1 is strictly lower than its counterpart in Proposition 9 around time zero for ρ close to 0.*
- *The price arising in Proposition 1 is strictly higher than its counterpart in Proposition 9 between T_ϕ and ϕ_0 ;*
- *The price is identical in Proposition 1 and in Proposition 9 after time ϕ_0 .*

This Proposition shows that, shortly after the liquidity shock, the price can be lower when institutions use trading algorithms.³³ The intuition is as follows. When they cannot use algorithms, traders know they will not be able to buy back before their next information event. Hence they initially sell less, which reduces the selling pressure on the price. Consequently the price can be higher than when traders are able to use algorithms. An outside observer may conclude that banning the use of algorithms could be socially beneficial, since it would alleviate the initial price pressure created by the liquidity shock. However, inferring aggregate welfare implications from price effects is misleading. Indeed, Proposition 2 implies that the equilibrium arising with algorithms Pareto dominates that arising when traders can place only limit and market orders when their information process jumps.

³³While we are able to establish this result only analytically for small ρ , our numerical calculations suggests that it can hold for larger value of ρ . In particular, it does hold for the value $\rho = 250$ that we have chosen for the numerical calculations presented in this paper as well as in our supplementary appendix.

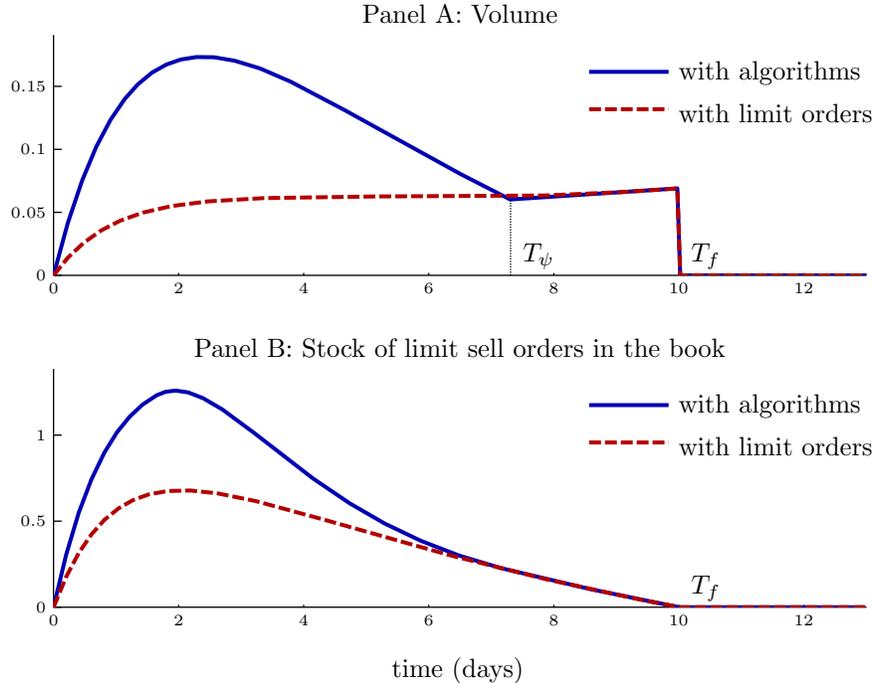


Figure 8: Trading volume (upper panel) and volume of limit orders outstanding in the book (lower panel), when $\sigma = 0.3$

Figure 8 plots the trading volume (upper panel) and the volume of limit orders outstanding in the book (lower panel) in Propositions 1 and 9. Trading volume is higher when institutions are able to use algorithms. This reflects the additional trading volume generated by round trip trades. Qualitatively, the stock of limit orders in the book has the same shape in Propositions 1 and 9. The book is filled progressively as limit orders to sell are placed by low-valuation traders. Then, cancellations and executions lead to a decrease in the stock of orders in the book. During the early phase, the amount of limit orders outstanding is higher in Proposition 1 than in Proposition 9. Indeed, with algorithms low valuation traders buy after their information event, this induces them to place more limit order to sell, in order to unwind this position towards the end of the price recovery process.

5 Conclusion

This paper studies the reaction of traders and markets to liquidity shocks under cognition limits. We model the aggregate liquidity shock as a transient decline in the valuation of the asset by all participants. While institutions recover from the shock at random times, traders with limited cognition observe the status of their institution only when their own information process

jumps. We interpret this delay as the time it takes traders to collect and process information about positions, counterparties, risk exposure and compliance. We characterize equilibrium prices and asset holding plans under two alternative market structures: when traders can use algorithms to automatically change their holding plan while investigating the liquidity status of their institution, and when they can place only market and limit orders when their information process jumps.

Our analysis suggests that, to an outside observer, algorithms may seem to be price destabilizing. We show however that they play a socially useful role by facilitating market-making: the equilibrium prevailing when traders can use algorithms always Pareto dominates that prevailing if traders can only place limit orders. This reflects the absence of externalities, as the asset holding constraints imposed by cognition limits do not depend on other traders' actions.

Note however that in our analysis there is no adverse selection. In a more general model with changes in common values and where some “slow” traders could not use algorithms, while “fast traders” could, information asymmetries could arise.³⁴ In this context, algorithmic trading could inflict negative externalities on slow traders, and equilibrium might no longer be optimal. Biais, Foucault, and Moinas (2010a) analyze these issues in a one period model. Because of their static setup, however, they are unable to consider rich dynamic order placement policies such as those arising in the present model. An interesting, but challenging, avenue of further research would be to extend the present model to the asymmetric information case. In doing so one could take stock of the insights of Goettler, Parlour, and Rajan (2009) and Pagnotta (2010) who study dynamic order placement under asymmetric information.

³⁴Consistent with the view that algorithmic traders have superior information Hendershott and Riordan (2010) and Brogaard (2010) find empirically that algorithmic trades have greater permanent price impact than slow trades and lead price discovery.

A Proofs

This appendix provides the proofs omitted in the main body of the paper. To keep the exposition concise, some intermediate results and calculations are gathered in our supplementary appendix, Biais, Hombert, and Weill (2010b).

A.1 Proof of Lemma 1

Let us begin by deriving a convenient expression for the intertemporal cost of buying and selling assets. For this we let $\tau_0 \equiv 0 < \tau_1 < \tau_2 < \dots$ denote the sequence of information events. For accounting purposes, we can always assume that, at her n -th information event, the trader sells all of her assets, q_{τ_{n-1}, τ_n} , and purchases a new initial holding q_{τ_n, τ_n} . Thus, the expected inter-temporal cost of the holding plan can be written:

$$\mathbb{E}_0 \left[-p_{\tau_1} e^{-r\tau_1} q_{0, \tau_1} + \sum_{n=1}^{\infty} \left\{ e^{-r\tau_n} p_{\tau_n} q_{\tau_n, \tau_n} + \int_{\tau_n}^{\tau_{n+1}} p_u dq_{\tau_n, u} e^{-ru} - e^{-r\tau_{n+1}} p_{\tau_{n+1}} q_{\tau_n, \tau_{n+1}} \right\} \right].$$

Given that p_u is continuous and piecewise continuously differentiable, and that $u \mapsto q_{\tau_n, u}$ has bounded variations, we can integrate by part (see Theorem 6.2.2 in Carter and Van Brunt, 2000), keeping in mind that $d/du(e^{-ru} p_u) = -e^{-ru} \xi_u$. This leads to:

$$\begin{aligned} & \mathbb{E}_0 \left[-p_{\tau_1} e^{-r\tau_1} q_{0, \tau_1} + \sum_{n=1}^{\infty} \int_{\tau_n}^{\tau_{n+1}} e^{-ru} \xi_u q_{\tau_n, u} du \right] = -p_0 s + \mathbb{E}_0 \left[\sum_{n=0}^{\infty} \int_{\tau_n}^{\tau_{n+1}} e^{-ru} \xi_u q_{\tau_n, u} du \right] \\ & = \text{constant} + \mathbb{E}_0 \left[\int_0^{\infty} e^{-ru} \xi_u q_{\tau_u, u} du \right]. \end{aligned}$$

In the above, the first equality follows by adding and subtracting $q_{0, u} = s$, and by noting that $q_{0, u}$ is constant; the second equality follows by using our “ τ_u ” notation for the last contact time before u . With the above result in mind, we find that, up to a constant, we can rewrite the intertemporal payoff net of cost as:

$$\mathbb{E}_0 \left[\int_0^{\infty} e^{-ru} \left(v(\theta_u, q_{\tau_u, u}) - \xi_u q_{\tau_u, u} \right) du \right]. \quad (\text{A.1})$$

Now for any $t \leq u$, the probability that $\tau_u \leq t$ is the probability that $N_u - N_t = 0$, which is equal to $e^{-\rho(u-t)}$. Thus, the distribution of τ_u has an atom of mass $e^{-\rho u}$ at $t = 0$, and then the density $\rho e^{-\rho(u-t)}$ for $t \in (0, u]$. Hence, after applying Baye’s rule, (A.1) rewrites as:

$$\int_0^{\infty} e^{-ru} \left\{ e^{-\rho u} \mathbb{E} \left[v(\theta_u, q_{0, u}) - \xi_u q_{0, u} \mid \tau_u = 0 \right] + \int_0^u \rho e^{-\rho(u-t)} \mathbb{E} \left[v(\theta_u, q_{t, u}) - \xi_u q_{t, u} \mid \tau_u = t \right] dt \right\} du. \quad (\text{A.2})$$

To simplify this expression, we rely on the following lemma, proved in Section VIII.1, page 44 in the supplementary appendix.

Lemma A.1. $\mathbb{E} [v(\theta_u, q_{t, u}) - \xi_u q_{t, u} \mid \tau_u = t] = \mathbb{E} [v(\theta_u, q_{t, u}) - \xi_u q_{t, u}]$ for all $t \geq 0$.

This lemma is clearly true for $t = 0$ since $q_{0, u} = s$ for all u , and since the information event process is independent from the type process. In Section VIII.1 in the supplementary appendix, we show that it also holds for $t > 0$. Intuitively, this is because of two facts. First, as noted above, the information event process is independent from the type process. Second, $\{\tau_u = t\} = \{N_t - N_{t-} = 1 \text{ and } N_u - N_t = 0\}$ only depends on increments of the information process at and after t , which are independent from the trader’s information one instant before t , and hence independent from the predictable process $q_{t, u}$. Relying on Lemma A.1, equation

(A.2) becomes:

$$\begin{aligned} & \int_0^\infty e^{-(r+\rho)u} \mathbb{E}[v(\theta_u, s) - \xi_u s] du + \int_0^\infty e^{-ru} \int_0^u \rho e^{-\rho(u-t)} \mathbb{E}[v(\theta_u, q_{t,u}) - \xi_u q_{t,u}] dt du \\ &= \int_0^\infty e^{-(r+\rho)u} \mathbb{E}[v(\theta_u, s) - \xi_u s] du + \mathbb{E}\left[\int_0^\infty e^{-rt} \int_t^\infty e^{-(r+\rho)(u-t)} \mathbb{E}_t[v(\theta_u, q_{t,u}) - \xi_u q_{t,u}] du \rho dt\right], \end{aligned}$$

where $\mathbb{E}_t[\cdot]$ refers to the expectation conditional upon \mathcal{F}_t and the second line follows from changing the order of integration and applying the law of iterated expectations.

A.2 Preliminary results about Q_u

We begin with preliminary results that we use repeatedly in this appendix. For the first preliminary result, let $\Psi(Q) \equiv \inf\{\psi \geq 0 : (1 - \mu_{h\psi})^{1/\sigma} Q \leq 1\}$, and $\psi_u \equiv \Psi(Q_u)$. We have:

Lemma A.2 (Preliminary results about Q_u). *Equation (17) has a unique solution, Q_u . Moreover, $0 \leq Q_u < (1 - \mu_{hu})^{-1/\sigma}$ and Q_u is continuously differentiable with*

$$Q'_u = \frac{e^{\rho u}(s - \mu_{hu}) - e^{\rho u}(1 - \mu_{hu})^{1+1/\sigma} Q_u}{\int_{\psi_u}^u e^{\rho t} (1 - \mu_{ht})^{1+1/\sigma} dt}. \quad (\text{A.3})$$

In the proof of this Lemma, in Section VIII.2 page 44 of the supplementary appendix, we first show existence, uniqueness, and the inequality $Q_u < (1 - \mu_{hu})^{1/\sigma}$ using the monotonicity and continuity of equation (17) with respect to Q . We then show that the solution is continuously differentiable using the Implicit Function Theorem.

For the second preliminary result, consider equation (17) after removing the min operator in the integral:

$$\int_0^u \rho e^{-\rho(u-t)} (1 - \mu_{ht})^{1+1/\sigma} \bar{Q}_u dt = S_u \iff \bar{Q}_u \equiv \frac{\int_0^u e^{\rho t} (s - \mu_{ht}) dt}{\int_0^u e^{\rho t} (1 - \mu_{ht})^{1+1/\sigma} dt}. \quad (\text{A.4})$$

Now, whenever $\bar{Q}_u \leq 1$, it is clear that \bar{Q}_u also solves equation (17). Given that the solution of (17) is unique it follows that $Q_u = \bar{Q}_u$. Conversely if $Q_u = \bar{Q}_u$, subtracting (17) from (A.4) shows that:

$$\int_0^u \rho e^{-\rho(u-t)} (1 - \mu_{ht}) \left((1 - \mu_{ht})^{1/\sigma} \bar{Q}_u - \min\{(1 - \mu_{ht})^{1/\sigma} \bar{Q}_u, 1\} \right) dt = 0.$$

Since the integrand is positive, this can only be true if $(1 - \mu_{ht})^{1/\sigma} \bar{Q}_u = \min\{(1 - \mu_{ht})^{1/\sigma} \bar{Q}_u, 1\}$ for almost all t . Letting $t \rightarrow 0$ delivers $\bar{Q}_u \leq 1$. Taken together, we obtain:

Lemma A.3 (A useful equivalence). $\bar{Q}_u \leq 1$ if and only if $Q_u = \bar{Q}_u$.

The next Lemma, proved in Section VIII.3 page 46 of the supplementary appendix, provides basic properties of \bar{Q}_u :

Lemma A.4 (Preliminary results about \bar{Q}_u). *The function \bar{Q}_u is continuous, satisfies $\bar{Q}_{0+} = s$ and $\bar{Q}_{T_f} = 0$. It is strictly decreasing over $(0, T_f]$ if $s \leq \sigma/(1 + \sigma)$ and hump-shaped otherwise.*

A.3 Proof of Proposition 1

The asset holding plans are determined according to Lemmas 4-6. The asset holding plan of high-valuation traders are uniquely determined for $u \in (0, T_f]$, and are indeterminate for $u > T_f$. The holding plan of low-valuation traders are uniquely determined, as Lemma A.2 shows that equation (17) has a unique solution. Now, turning to the price, the definition of Q_u implies that the price solves $rp_u = 1 - \delta(1 - \mu_{hu})Q_u^\sigma + \dot{p}_u$ for $u < T_f$. For $u \geq T_f$, the fact that high-valuation traders are indifferent between any asset holdings in $[0, 1]$ implies that

$rp_u = 1 + \dot{p}_u$. But the price is bounded and positive, so it follows that $p_u = 1/r$. Since the price is continuous at T_f , this provides a unique candidate equilibrium price path. Clearly this candidate is C^1 over $(0, T_f)$ and (T_f, ∞) . To show that it is continuously differentiable at T_f note that, given $Q_{T_f} = 0$ and $p_{T_f} = 1/r$, ODE (18) implies that $\dot{p}_{T_f^-} = 0$. Obviously, since the price is constant after T_f , $\dot{p}_{T_f^+} = 0$ as well. We conclude that \dot{p}_u is continuous at $u = T_f$ as well.

Next, we show that the candidate equilibrium thus constructed is indeed an equilibrium. For this recall from Lemma A.2 that $0 < Q_u < (1 - \mu_{hu})^{-1/\sigma}$, which immediately implies that

$$0 < 1 - rp_u + \dot{p}_u < 1$$

for $u < T_f$. It follows that high-valuation traders find it optimal to hold one unit. Now one can directly verify that, for $u < T_f$, the problem of low-valuation traders is solved by $q_{t,u} = \min\{(1 - \mu_{ht})^{-1/\sigma} Q_u, 1\}$. For $u \geq T_f$, $1 - rp_u + \dot{p}_u = 0$ and so the problem of high-valuation traders is solved by any $q_{t,u} \in [0, 1]$, while the the problem of low-valuation traders is clearly solved by $q_{t,u} = 0$. The asset market clears at all dates by construction.

We already know that the price is equal to $1/r$ for $t > T_f$ so the last thing to show is that it is strictly increasing for $t < T_f$. Letting $\Delta_u \equiv (1 - \mu_{hu})^{1/\sigma} Q_u$ for $u \leq T_f$, and $\Delta_u = 0$ for $u \geq T_f$. In Section VIII.4, page 47 in the supplementary appendix, we show that:

Lemma A.5. *The function Δ_u is strictly decreasing over $(0, T_f]$.*

Now, in terms of Δ_u , the price writes:

$$p_u = \int_u^\infty e^{-r(y-u)} (1 - \delta\Delta_y^\sigma) dy = \int_0^\infty e^{-rz} (1 - \delta\Delta_{z+u}^\sigma) dz,$$

after the change of variable $y - u = z$. Since Δ_u is strictly decreasing over $u \in (0, T_f)$, and constant over $[T_f, \infty)$, it clearly follows from the above formula that p_u is strictly increasing over $u \in (0, T_f)$.

A.4 Proof of Proposition 2

Let us begin with a preliminary remark. By definition, any feasible allocation \tilde{q} satisfies the market-clearing condition $\mathbb{E}_0[\tilde{q}_{\tau_u, u}] = s$. Taken together with the fact that $\xi_u = rp_u - \dot{p}_u$ is deterministic, this implies:

$$\mathbb{E} \left[\int_0^\infty e^{-ru} q_{\tau_u, u} \xi_u du \right] \tag{A.5}$$

$$= \int_0^\infty e^{-ru} \mathbb{E}[q_{\tau_u, u}] \xi_u du = \int_0^\infty e^{-ru} s \xi_u du = p_0 s. \tag{A.6}$$

With this in mind, consider the equilibrium asset holding plan of Proposition 1, q , and suppose it does not solve the planning problem. Then there is a feasible asset holding plan q' that achieves a strictly higher value of the objective (20). But, by (A.6), this asset holding plan has the same cost as the equilibrium asset holding plan. Subtracting the inter-temporal cost (A.6) from the inter-temporal utility (A.1), we obtain that $V(q') > V(q)$, which contradicts individual optimality.

A.5 Proof of Proposition 3

Taken together, Lemma A.3 and Lemma A.4 immediately imply that

Lemma A.6. *The function Q_u satisfies $Q_{0^+} = s$, $Q_{T_f} = 0$. If $s \leq \sigma/(1 + \sigma)$, then it is strictly decreasing over $(0, T_f]$. If $s > \sigma/(1 + \sigma)$ and $\bar{Q}_u \leq 1$ for all $u \in (0, T_f]$, Q_u is hump-shaped over $(0, T_f]$.*

The only case that is not covered by the Lemma is when $s > \sigma/(1 + \sigma)$ and $\bar{Q}_u > 1$ for some $u \in (0, T_f]$. In this case, note that for u small and u close to T_f , we have that $\bar{Q}_u < 1$. Given that \bar{Q}_u is hump-shaped, it follows that the equation $\bar{Q}_u = 1$ has two solutions, $0 < T_1 < T_2 < T_f$. For $u \in (0, T_1]$ (resp. $u \in [T_2, T_f]$),

$\bar{Q}_u \leq 1$ and is increasing (resp. decreasing), and thus Lemma A.3 implies that $Q_u = \bar{Q}_u$ and increasing (resp. decreasing) as well.

In particular, since we know from Lemma A.2 that Q_u is continuously differentiable, we have that $Q'_{T_1} > 0$ and $Q'_{T_2} < 0$. Thus, Q_u changes sign at least once in (T_1, T_2) . To conclude, in Section VIII.5 we establish:

Lemma A.7. *The derivative Q'_u changes sign only once in (T_1, T_2) .*

A.6 Proof of Proposition 4

In Section VIII.6, page 50 in the supplementary appendix, we prove the following asymptotic results:

Lemma A.8. *As ρ goes to infinity:*

$$T_f(\rho) \downarrow T_s \tag{A.7}$$

$$Q_u(\rho) = \frac{s - \mu_{hu}}{(1 - \mu_{hu})^{1+1/\sigma}} - \frac{1}{\rho} \frac{\gamma}{(1 - \mu_{hu})^{1/\sigma}} \left[\left(1 + \frac{1}{\sigma}\right) \frac{s - \mu_{hu}}{1 - \mu_{hu}} - 1 \right] + o\left(\frac{1}{\rho}\right), \quad \forall u \in [0, T_s] \tag{A.8}$$

$$T_\psi(\rho) = \arg \max_{u \in [0, T_f(\rho)]} Q_u(\rho) \longrightarrow \arg \max_{u \in [0, T_s]} \frac{s - \mu_{hu}}{(1 - \mu_{hu})^{1+1/\sigma}}. \tag{A.9}$$

With this in mind we can study the asymptotic behavior of price, allocation, and volume.

A.6.1 Asymptotic price

Proposition 1 shows that the price at all times $t \in [T_f, \infty)$ is equal to $1/r$. But we know from Lemma A.8 that $T_f(\rho) \rightarrow T_s$ as $\rho \rightarrow \infty$. Clearly, this implies that, for all $t \in (T_s, \infty)$, as $\rho \rightarrow \infty$ the price converges towards $1/r$, its unlimited cognition counterpart.

Integrating the ODEs of Proposition 1 shows that, at all times $t \in [0, T_s]$, the price is equal to:

$$p_t = \int_t^{T_f(\rho)} e^{-r(u-t)} (1 - \delta(1 - \mu_{hu}) Q_u^\sigma(\rho)) du + \frac{e^{-r(T_f(\rho)-t)}}{r}.$$

But Lemma A.8 shows that $(1 - \mu_{hu}) Q_u^\sigma(\rho)$ converges point-wise towards $[(s - \mu_{hu})/(1 - \mu_{hu})]^\sigma$. Moreover, we know that $(1 - \mu_{hu}) Q_u^\sigma \in [0, 1]$. It thus follows from dominated convergence that p_t converges to its unlimited cognition counterpart.

A.6.2 Asymptotic distribution of asset holdings

Let us start with some $u \leq T_s$. With unlimited cognition, traders whose valuation is low at time u hold $(s - \mu_{hu})/(1 - \mu_{hu})$, and traders whose valuation is high hold one share. With limited cognition, the time- u cross-sectional distribution of asset holdings across low-valuation traders is, by our usual law of large numbers argument, the time- u distribution of asset holding generated by the holding plan of a representative trader, conditional on $\theta_u = \ell$. In Section VIII.7, page 52 in the supplementary appendix, we establish:

Lemma A.9. *For all $\varepsilon > 0$, as $\rho \rightarrow \infty$:*

$$\text{Proba} \left(\left| q_{\tau_u, u} - \frac{s - \mu_{hu}}{1 - \mu_{hu}} \right| > \varepsilon \mid \theta_u = \ell \right) \rightarrow 0 \tag{A.10}$$

$$\text{Proba} \left(\left| q_{\tau_u, u} - 1 \right| > \varepsilon \mid \theta_u = h \right) \rightarrow 0. \tag{A.11}$$

A.6.3 Asymptotic volume

Basic formulas. As above, let T_ψ denote the arg max of the function Q_u . For any time $u < T_\psi$ and some time interval $[u, u + du]$, the only traders who *sell* are those who have an information event during this time interval, and who find out that they have a *low* valuation. Thus, trading volume during $[u, u + du]$ can be computed as the volume of assets sold by these traders, as follows. Just before their information event, low-valuation traders hold on average:

$$\mathbb{E}[q_{\tau_u, u} | \theta_u = \ell] = e^{-\rho u} s + \int_0^u \rho e^{-\rho(u-t)} q_{\ell, t, u} dt, \quad (\text{A.12})$$

since a fraction $e^{-\rho u}$ of them have had no information event and still hold s unit of the asset and a density $\rho e^{-\rho(u-t)}$ of them had their last information event at time t and hold $q_{\ell, t, u}$. Instantaneous trading volume is then:

$$V_u = \rho(1 - \mu_{hu}) \left(\mathbb{E}[q_{\tau_u, u} | \theta_u = \ell] - q_{\ell, u, u} \right), \quad (\text{A.13})$$

where $\rho(1 - \mu_{hu})$ is the measure of low-valuation investors having an information event, the term in large parentheses is the average size of low-valuation traders' sell orders, and $q_{\ell, u, u}$ is their asset holding right after the information event.

For any time $u \in (T_\psi, T_f)$ and some time interval $[u, u + du]$, the only traders who *buy* are those who have an information event during this time interval, and who find out that they have a *high* valuation. Trading volume during $[u, u + du]$ can be computed as the volume of assets purchased by these traders:

$$V_u = \rho \mu_{hu} \left(1 - \mathbb{E}[q_{\tau_u, u} | \theta_u = h] \right), \quad (\text{A.14})$$

where the time- u average asset holdings of high-valuation investors is found by plugging (A.12) into the market clearing condition:

$$\mu_{hu} \mathbb{E}[q_{\tau_u, u} | \theta_u = h] + (1 - \mu_{hu}) \mathbb{E}[q_{\tau_u, u} | \theta_u = \ell] = s.$$

For $u > T_f$, the trading volume is not zero since high-valuation traders continue to buy from the low valuation traders having an information event:

$$V_u = \rho(1 - \mu_{hu}) \mathbb{E}[q_{\tau_u, u} | \theta_u = \ell]. \quad (\text{A.15})$$

Taking the $\rho \rightarrow \infty$ limit. We first note that $q_{\ell, u, u}(\rho) = \min\{(1 - \mu_{hu})^{1/\sigma} Q_u(\rho), 1\} = (1 - \mu_{hu})^{1/\sigma} Q_u(\rho)$. Next, we need to calculate an approximation for:

$$\mathbb{E}[q_{\tau_u, u} | \theta_u = \ell] = s e^{-\rho u} + \int_0^u \min\{(1 - \mu_{ht})^{1/\sigma} Q_u(\rho), 1\} \rho e^{-\rho(u-t)} dt.$$

For this we follow the same calculations leading to equation (VIII.12) in the proof of Lemma A.8, page 51 in the supplementary appendix, but with $f(t, \rho) = \min\{(1 - \mu_{ht})^{1/\sigma} Q_u(\rho), 1\}$. This gives:

$$\begin{aligned} \mathbb{E}[q_{\tau_u, u} | \theta_u = \ell] &= s e^{-\rho u} + \int_0^u \rho e^{-\rho(u-t)} f(t, \rho) dt = f(u, \rho) - \frac{1}{\rho} f_t(u, \rho) + o\left(\frac{1}{\rho}\right) \\ &= (1 - \mu_{hu})^{1/\sigma} Q_u(\rho) + \frac{1}{\rho} \frac{\gamma}{\sigma} \frac{s - \mu_{hu}}{1 - \mu_{hu}} + o\left(\frac{1}{\rho}\right) = \frac{s - \mu_{hu}}{1 - \mu_{hu}} + \frac{\gamma}{\rho} \frac{1 - s}{1 - \mu_{hu}} + o\left(\frac{1}{\rho}\right), \end{aligned}$$

where the second line follows from plugging equation (A.8) into the first line. Substituting this expression into equations (A.13) and (A.14), we find after some straightforward manipulation that, when ρ goes to infinity,

$V_u \rightarrow \gamma(s - \mu_{hu})/\sigma$ for $u < T_\psi(\infty)$, $V_u \rightarrow \gamma(1 - s)$ for $u \in (\lim T_\psi(\infty), T_s)$, and $V_u \rightarrow 0$ for $u > T_s$.

The trading volume in the Walrasian equilibrium is equal to the measure of low-valuation institutions who become high-valuation institutions: $\gamma(1 - \mu_{hu})$, times the amount of asset they buy at that time: $1 - (s - \mu_{hu})/(1 - \mu_{hu})$. Thus the trading volume is $\gamma(1 - s)$. To conclude the proof, note that after taking derivatives of $Q_u(\infty)$ with respect to u , it follows that

$$Q'_{T_\psi(\infty)}(\infty) = 0 \Leftrightarrow \frac{s - \mu_{hT_\psi(\infty)}}{\sigma} = 1 - s$$

which implies in turn that $\gamma(s - \mu_{hu})/\sigma > \gamma(1 - s)$ for $u < T_\psi(\infty)$.

A.7 Proof of Proposition 5

A.7.1 First point: when (21) holds

Given the ODEs satisfied by the price path in both markets, it suffices to show that, for all $u < T_s$,

$$(1 - \mu_{hu})Q_u^\sigma > \left(\frac{s - \mu_{hu}}{1 - \mu_{hu}}\right)^\sigma.$$

Besides, when $s \leq \sigma/(\sigma + 1)$, it follows from Lemma A.3 and Lemma A.4 that:

$$Q_u = \bar{Q}_u = \frac{\int_0^u e^{\rho t}(s - \mu_{ht}) dt}{\int_0^u e^{\rho t}(1 - \mu_{ht})^{1+1/\sigma} dt}.$$

Plugging the above and rearranging, we are left with showing that:

$$F_u = (s - \mu_{hu}) \int_0^u e^{\rho t}(1 - \mu_{ht})^{1+1/\sigma} dt - (1 - \mu_{hu})^{1+1/\sigma} \int_0^u e^{\rho t}(s - \mu_{ht}) dt < 0.$$

But we know from the proof of Lemma A.4, equation (VIII.5), page 46 in the supplementary appendix, that F_u has the same sign as Q'_u , which we know is negative at all $u > 0$ since $s \leq \sigma/(\sigma + 1)$.

A.7.2 Second point: when s is close to 1 and σ is close to 0

The price at time 0 is equal to:

$$p_0 = \int_0^{+\infty} e^{-ru} \xi_u du,$$

where we make the dependence of s and σ explicit. With perfect cognition, $\xi_u = 1 - \delta((s - \mu_{hu})/(s - \mu_{hu}))^\sigma = 1 - \delta(1 - (1 - s)e^{\gamma u})^\sigma$ for $u < T_s$, and $\xi_u = 1$ for $u > T_s$. Therefore $p_0 = 1/s - \delta I(s)$, where:

$$I(s) \equiv \int_0^{T_s} e^{-ru} (1 - (1 - s)e^{\gamma u})^\sigma du,$$

where we make explicit the dependence on s . Similarly, with limited cognition, the price at time 0 is equal to $p_0 = 1/s - \delta J(s)$, where:

$$J(s) \equiv \int_0^{T_f} e^{-ru} (1 - \mu_{hu})Q_u^\sigma du.$$

We begin with a Lemma proved in Section VIII.8.1, page 53 in the supplementary appendix:

Lemma A.10. *When s goes to 1, both $I(s)$ and $J(s)$ go to $1/r$.*

Therefore, p_0 goes to $(1 - \delta)/r$ both with perfect cognition and with limited cognition. Besides, with perfect cognition:

$$p_0(s_0) = (1 - \delta)/r + \delta \int_{s_0}^1 I'(s) ds,$$

and with limited cognition:

$$p_0(s_0) = (1 - \delta)/r + \delta \int_{s_0}^1 J'(s) ds.$$

The next two lemmas compare $I'(s)$ and $J'(s)$ for s in the neighborhood of 1 when σ is not too large. The first Lemma is proved in Section VIII.8.2, page 54 in the supplementary appendix:

Lemma A.11. *When s goes to 1:*

$$\begin{aligned} I'(s) &\sim \text{constant} && \text{if } r > \gamma, \\ I'(s) &\sim \Gamma_1(\sigma) \log((1 - s)^{-1}) && \text{if } r = \gamma, \\ I'(s) &\sim \Gamma_2(\sigma)(1 - s)^{-1+r/\gamma} && \text{if } r < \gamma, \end{aligned}$$

where the constant terms $\Gamma_1(\sigma)$ and $\Gamma_2(\sigma)$ go to 0 when $\sigma \rightarrow 0$.

In this Lemma and all what follows $f(s) \sim g(s)$ means that $f(s)/g(s) \rightarrow 1$ when $s \rightarrow 1$. The second Lemma is proved in Section VIII.8.3, page 55 in the supplementary appendix:

Lemma A.12. *Assume $\gamma + \gamma/\sigma - \rho > 0$. There exists a function $\tilde{J}'(s) \leq J'(s)$ such that, when s goes to 1:*

$$\begin{aligned} \tilde{J}'(s) &\rightarrow +\infty && \text{if } r > \gamma, \\ \tilde{J}'(s) &\sim \Gamma_3(\sigma) \log((1 - s)^{-1}) && \text{if } r = \gamma, \\ \tilde{J}'(s) &\sim \Gamma_4(\sigma)(1 - s)^{-1+r/\gamma} && \text{if } r < \gamma, \end{aligned}$$

where the constant terms $\Gamma_3(\sigma)$ and $\Gamma_4(\sigma)$ go to strictly positive limits when $\sigma \rightarrow 0$.

Lemmas A.11 and A.12 imply that, if σ is close to 0, then $J'(s) > I'(s)$ for s in the left-neighborhood of 1. The second point of the proposition then follows.

A.8 Proofs of Proposition 9

A.8.1 The candidate equilibrium

Before formulating our guess, we need the following Lemma, proved in in Section VIII.9, page 59 in the supplementary appendix:

Lemma A.13. *Suppose $s > \sigma/(1 + \sigma)$ and let $T_\phi > 0$ be the unique solution of:*

$$\frac{\mu_h T_\phi}{\sigma} = \left(1 + \frac{1}{\sigma}\right) s - 1. \tag{A.16}$$

Then $T_\phi < T_s$ and, for all $t \in [0, T_\phi)$ there exists a unique $\phi_t \in (T_\phi, T_f)$ such that

$$\int_t^{\phi_t} e^{\rho u} \left[(1 - \mu_{hu})^{1+1/\sigma} (s - \mu_{ht}) - (1 - \mu_{ht})^{1+1/\sigma} (s - \mu_{hu}) \right] du = 0. \tag{A.17}$$

In addition $t \mapsto \phi_t$ is continuously differentiable, strictly decreasing over $[0, T_\phi]$, and $\lim_{t \rightarrow T_\phi} \phi_t = T_\phi$.

For concision, we directly state below the analytical formula defining the candidate equilibrium (a heuristic constructive proof is available from the authors upon request).

The price path. The price path is continuous and solves the following ODEs. When $t \in (0, T_\phi)$:

$$rp_t - \dot{p}_t = \xi_t = 1 - \delta \left(\frac{s - \mu_{ht}}{1 - \mu_{ht}} \right)^\sigma + \delta \frac{d}{dt} \left[\left(\frac{s - \mu_{ht}}{(1 - \mu_{ht})^{1+1/\sigma}} \right)^\sigma \right] \int_t^{\phi_t} e^{-(r+\rho)(u-t)} (1 - \mu_{hu}) du, \quad (\text{A.18})$$

When $t \in (T_\phi, \phi_0)$:

$$rp_t - \dot{p}_t = \xi_t = 1 - \delta \frac{1 - \mu_{ht}}{1 - \mu_{h\phi_t^{-1}}} \left(\frac{s - \mu_{h\phi_t^{-1}}}{1 - \mu_{h\phi_t^{-1}}} \right)^\sigma. \quad (\text{A.19})$$

When $t \in (\phi_0, T_f)$:

$$rp_t - \dot{p}_t = \xi_t = 1 - \delta(1 - \mu_{ht})\overline{Q}_t^\sigma. \quad (\text{A.20})$$

where \overline{Q}_t is defined in equation (A.4). Lastly, when $t \geq T_f$, $p_t = 1/r$. Clearly, together with the continuity conditions at the boundaries of each intervals (T_ϕ , ϕ_0 , and T_f), the above ODEs uniquely define the price path.

High-valuation traders' holding plans. The asset holding plans of high-valuation traders are as follows. For $t \in [0, T_f]$, the time- t high-valuation trader holds $q_{t,u} = 1$ for all $u \geq t$. For $t \geq T_f$, the time- t high-valuation trader holds $q_{t,u} = s/\mu_{ht}$ for all $u \geq t$. Since these asset holding plan are constant, we guess that they are implemented by submitting market orders at time t .

Low-valuation traders' holding plans. The asset holding plans of low-valuation traders are as follows. For $t \in (0, T_\phi]$, the time- t low-valuation trader holds:

$$q_{t,u} = \frac{s - \mu_{ht}}{1 - \mu_{ht}} \quad \text{for all } u \in [t, \phi_t] \quad (\text{A.21})$$

$$= (1 - \mu_{ht})^{1/\sigma} Q_u \quad \text{for all } u \in [\phi_t, T_f], \quad \text{where } Q_u \equiv \left(\frac{1 - \xi_u}{\delta(1 - \mu_{hu})} \right)^{1/\sigma} \quad (\text{A.22})$$

$$= 0 \quad \text{for all } u \geq T_f, \quad (\text{A.23})$$

and where ξ_u is defined in equations (A.18) and (A.19). For $t \in [T_\phi, T_f]$

$$q_{t,u} = (1 - \mu_{ht})^{1/\sigma} Q_u \quad \text{for all } u \in [t, T_f] \quad (\text{A.24})$$

$$= 0 \quad \text{for all } u \geq T_f. \quad (\text{A.25})$$

Lastly, for $t \geq T_f$ and $u \geq t$, $q_{t,u} = 0$. We guess that, in order to implement her holding plan at the information event time t , the trader submits market orders and a schedule of limit sell orders.

Note that, after plugging in the definition of ξ_{ϕ_t} given in equation (A.19), one sees that the above-defined asset holding plan is continuous at $u = \phi_t$. Since $Q_{T_f} = \overline{Q}_{T_f} = 0$, we also have continuity at $u = T_f$. Lastly in Lemma A.16 below, we will show that when $t \in [0, T_\phi]$ ($t \in [T_\phi, T_f]$) and $u \in [\phi_t, T_f]$ ($u \in [t, T_f]$), the holding plan $u \mapsto q_{t,u}$ is strictly decreasing.

The verification proof is organized as follows. Section A.8.2 provides preliminary properties of equilibrium objects. Section A.8.3 proves the optimality of the candidate asset holding plans. Section A.8.4 concludes by showing that the market clears at all dates.

A.8.2 Preliminary results

We begin by noting that, by direct inspection of our guess, the candidate equilibrium coincides with the equilibrium of Proposition 1 at all times $t > \phi_0$. Next, a result we use repeatedly is:

Lemma A.14. *For $t \in [0, T_\phi]$,*

$$\frac{d}{dt} \left[\frac{s - \mu_{ht}}{(1 - \mu_{ht})^{1+1/\sigma}} \right] \geq 0, \quad \text{with an equality only if } t = T_\phi.$$

The proof is in Section VIII.10, page 61 in the supplementary appendix. Next, we prove in Section VIII.11, page 61 in the supplementary appendix:

Lemma A.15. *For $t \in (0, T_f)$, $rp_t - \dot{p}_t = \xi_t \in (0, 1)$.*

This Lemma ensures that Q_u , in equation (A.22), is well defined, and is also helpful to establish the optimality of trading strategies. To show that asset holding plans are decreasing when $u \geq \phi_t$, we will need the following Lemma, proved in Section VIII.12, page 62 in the supplementary appendix:

Lemma A.16. *$u \mapsto Q_u$ is continuous over $(0, T_f)$, strictly increasing over $(0, T_\phi)$, and strictly decreasing over (T_ϕ, T_f) .*

Equipped with the above Lemma, we show in Section VIII.13, page 63 in the supplementary appendix, another crucial property of the price path:

Lemma A.17. *The candidate equilibrium price, p_t , is continuously differentiable and strictly increasing over $[0, T_f]$.*

A.8.3 Asset holding plans are optimal

Execution times of limit orders. The first step is to verify that the candidate asset holding plans can be implemented using market and limit orders only, as explained in Section A.8.1. To that end, we begin by deriving the execution times associated with alternative limit orders to buy or sell.

Limit orders to buy or sell submitted at $t \geq T_f$. Consider the case of limit sell orders (the case of limit buy order is symmetric). In the candidate equilibrium, the price is constant for all $t \geq T_f$. Clearly, this means that a limit order to sell at the ask price $a > p_{T_f}$ is never executed. By the price priority rule, a limit order to sell at price $a < p_{T_f}$ is executed immediately. We can always ignore such limit orders because they are clearly dominated by a market order to sell, which is also executed immediately, but has a strictly higher execution price, p_{T_f} . Lastly, we argue that a limit order to sell at price $a = p_{T_f}$ submitted at time $t \geq T_f$ is executed immediately, at time t . First we know from the price priority rule that this order is executed at or after time t . But note that, at all times $u > t \geq T_f$, the limit order book must be empty at price p_{T_f} : otherwise, by the time priority rule, the earliest submitted limit order at price p_{T_f} would be executed at time u , which would contradict the fact that asset holding plans are constant after T_f . Thus, by any time $u > t$, all limit orders to sell at price p_{T_f} , and in particular those submitted at time t , have been executed. Since this is true for all $u > t$, this implies that a limit order to sell at price p_{T_f} is executed immediately at time t , and is thus equivalent to a market order.

Limit orders to sell submitted at $t < T_f$. Lemma A.17 showed that the price path is strictly increasing over $[0, T_f]$. Therefore, the price–priority rule implies that: a limit order to sell at price $a > p_{T_f}$ is never executed; a limit order to sell at price $a \in [p_t, p_{T_f})$ is executed when the price reaches a , i.e. at the time u such that $p_u = a$; a limit order to sell at price $a = p_{T_f}$ is executed at or after time T_f . But, as noted in the previous paragraph, the limit order book is empty at all times $u > T_f$, and the same reasoning as above shows that this limit order must be executed at or before T_f . Taken together, these two remarks imply that a limit order to sell at price $a = p_{T_f}$, submitted at time $t < T_f$, is executed exactly at time T_f .

As above, we can ignore limit orders to sell at price $a < p_t$. By the price priority rule, such limit orders are executed immediately at price $a < p_t$, and therefore are clearly dominated by a market order to sell at price p_t .

Limit orders to buy submitted at $t < T_f$. By the price–priority rule: a limit order to buy at price $b \leq p_t$ is never executed; limit order to buy at price $b > p_t$ is executed immediately at price b , and is thus clearly dominated by market order to buy at price p_t .

A re-statement of the trader’s problem. Based on the above, we show that the trader’s problem reduces to maximizing her utility from one information event to the next:

Lemma A.18. *In the candidate equilibrium, an asset holding plan $q_{t,u}$ solves the trader’s problem if and only if it maximizes*

$$\int_t^\infty e^{-(r+\rho)(u-t)} \left\{ \mathbb{E}_t [v(\theta_u, q_{t,u})] - \xi_u q_{t,u} \right\} du \quad (\text{A.26})$$

for almost all $(t, \omega) \in \mathbb{R} \times \Omega$, and subject to the constraint that $u \mapsto q_{t,u}$ is decreasing for $u \in [t, t \vee T_f]$, and constant for $u \geq t \vee T_f$.

For the “if” part, note that the execution times of Section A.8.3 imply that, upon an information event at time t :

- a trader cannot submit a limit order to *buy* executed at time $u > t$;
- if $t < T_f$, a trader can submit a limit order to *sell* at any time $u \in (t, T_f]$;
- a trader cannot submit a limit order to *sell* at any time $u > T_f$.

Therefore, in the candidate equilibrium, a trader’s asset holding plan $u \mapsto q_{t,u}$ can never increase, can decrease in an arbitrary fashion over $[t, t \vee T_f]$, and must stay constant after $t \vee T_f$. This means that the value of the time- t asset holding plan, $u \mapsto q_{t,u}$, is less than the maximum value of the maximization problem in Lemma A.18. Therefore, an asset holding plan solves the trader’s problem if it solves the program of Lemma A.18 for almost all $(t, \omega) \in \mathbb{R} \times \Omega$.

To prove the “only if” part, we proceed by contrapositive. Suppose that $q_{t,u}$ does not maximize (A.26) for some positive measure set of $\mathbb{R} \times \Omega$. Then, consider the following change of holding plan: for all times and events in that set, switch to a plan that achieves a higher value in the objective (A.26), and keep your original holding plan the same otherwise. To see that the change of holding plan is feasible, note that the execution times of Section A.8.3 imply that, if the change of holding plan requires modifying orders submitted at some information event t_1 , all of these modifications can be undone at any subsequent information event $t_2 > t_1$. Indeed a limit sell order at price $p_u \in (p_0, p_{T_f}]$ is executed at time u regardless of its submission time $t_1 < u$. Thus, if such an order is modified at t_1 and is still unfilled at time t_2 , then $t_2 < u$ and the modification can be undone. Other kind of orders (limit sell orders at different prices, limit buy orders, or market orders) are either immediately executed at information event t_1 , or never executed, so any modification at t_1 can be undone with market orders at t_2 . Taken together, this shows that the change of holding plan is feasible and, clearly, it increases the value of the trader’s objective.

Optimality of high-valuation traders’ asset holding plans. By construction, the candidate asset holding plan is constant for all u given t , so it can be implemented by only submitting market orders at each information event. To show that the candidate asset holding plan is optimal, it suffices to show that it maximizes a high-valuation trader’s utility flow at each time u . To see that this is indeed the case, recall from Lemma A.15 that $\xi_u \in (0, 1)$ for $u \in (0, T_f)$, and that by construction we have that $\xi_u = 1$ for $u > T_f$. Therefore, the flow utility

$$v(h, q) - \xi_u q = \min\{q, 1\} - \xi_u q,$$

is maximized by $q = 1$ if $u < T_f$, and by any $q \in [0, 1]$ if $u \geq T_f$. In particular, $q_{t,u} = 1$ and, as long as $t \geq T_f > T_s$, $q_{t,u} = s/\mu_{ht} \leq 1$ maximizes the flow utility for $u \geq T_f$. Thus, the candidate asset holding plan is optimal for high-valuation traders.

Optimality of low-valuation traders’ asset holding plans. After time T_f , the price is constant and equal to $1/r$, so the opportunity cost is $\xi_u = 1$. Given that $v_q(\ell, 0) = v(h, 0) = 1$ and $v_q(\ell, q) < 1$ for $q > 0$, this immediately implies that $q = 0$ is the asset holding maximizing the flow payoff, $\mathbb{E}[v_q(\theta_u, q) | \theta_t = \ell] - \xi_u q$, for any $t < u$. Clearly, given the execution time of Section A.8.3, a low-valuation trader can always implement this zero asset holding for all $u \geq T_f$. Indeed, if $t < T_f$, such trader simply needs to submit a limit order to sell all her remaining assets at price p_{T_f} . If $t > T_f$, such trader simply needs to submit a market order to sell all her assets at time t . We conclude that we can restrict attention to holding plans such that $q_{t,u} = 0$ for all

$u \geq T_f$. Plugging this restriction into the objective (A.26), we find that the low-valuation trader's problem at time $t < T_f$ reduces to choosing $\tilde{q}_{t,u}$ in order to maximize

$$\int_t^{T_f} e^{-(r+\rho)(u-t)} M(u, \tilde{q}_{t,u}) du \quad (\text{A.27})$$

$$\text{where } M(u, q) \equiv \mathbb{E}[v(\theta_u, q) | \theta_t = \ell] - \xi_u q = (1 - \xi_u) q - \delta \frac{1 - \mu_{hu}}{1 - \mu_{ht}} \frac{q^{1+\sigma}}{1 + \sigma}, \quad (\text{A.28})$$

subject to the constraint that $u \mapsto \tilde{q}_{t,u}$ is decreasing over $[t, T_f]$.

We now verify that our candidate asset holding plan solves this optimization problem. First, we note that, by Lemma A.16, $q_{t,u}$ is, as required, decreasing over $[t, T_f]$. Second, we show below that two sufficient conditions for optimality are

$$\begin{aligned} (i) \quad & \int_t^{T_f} e^{-(r+\rho)(u-t)} \frac{\partial M}{\partial q}(u, q_{t,u}) q_{t,u} du = 0, \\ (ii) \quad & \int_t^{T_f} e^{-(r+\rho)(u-t)} \frac{\partial M}{\partial q}(u, q_{t,u}) \tilde{q}_{t,u} du \leq 0 \quad \text{for any decreasing function } u \mapsto \tilde{q}_{t,u}, \end{aligned}$$

where the partial derivative is well defined because, since $q_{t,u}$ is decreasing, $q_{t,u} \leq q_{t,t} = (s - \mu_{ht}) / (1 - \mu_{ht}) < 1$. To see why condition (i) and (ii) are sufficient, consider another decreasing holding plan $\tilde{q}_{t,u}$. We have

$$\int_t^{T_f} e^{-(r+\rho)(u-t)} [M(u, q_{t,u}) - M(u, \tilde{q}_{t,u})] du \geq \int_t^{T_f} e^{-(r+\rho)(u-t)} \frac{\partial M}{\partial q}(u, q_{t,u}) (q_{t,u} - \tilde{q}_{t,u}) du \geq 0,$$

where the first inequality follows from the concavity of $q \mapsto M(u, q)$, and the second inequality follows from (i) and (ii). The proof that condition (i) and (ii) hold follows from algebraic manipulations that we gather in Section VIII.14.

A.8.4 The market clears at all times

For all $u \leq T_f$, high-valuation traders who had at least one information event hold $q_{t,u} = 1$. Plugging this into equation (8) and rearranging, this leads to the market-clearing condition:

$$\int_0^u \rho e^{-\rho(u-t)} (1 - \mu_{ht}) \mathbb{E}[q_{t,u} | \theta_t = \ell] dt = \int_0^u \rho e^{-\rho(u-t)} (s - \mu_{ht}) dt. \quad (\text{A.29})$$

Market clearing at $u \in (0, T_\phi)$. Then, for all $t \leq u$, it follows from equation (A.21) that low-valuation traders hold $q_{t,u} = (s - \mu_{ht}) / (1 - \mu_{ht})$ and so clearly the market-clearing condition (A.29) holds.

Market clearing at $u \in (T_\phi, \phi_0)$. Then, it follows from equation (A.21) and (A.22) that low-valuation traders hold $q_{t,u} = (s - \mu_{ht}) / (1 - \mu_{ht})$ if $t \leq \phi_u^{-1}$, and $(1 - \mu_{ht})^{1/\sigma} Q_u$ if $t \in [\phi_u^{-1}, u]$. Thus, the left-hand side of (A.29) writes:

$$\int_0^{\phi_u^{-1}} \rho e^{-\rho(u-t)} (1 - \mu_{ht}) \frac{s - \mu_{ht}}{1 - \mu_{ht}} dt + \int_{\phi_u^{-1}}^u \rho e^{-\rho(u-t)} (1 - \mu_{ht})^{1+1/\sigma} Q_u dt. \quad (\text{A.30})$$

Plugging ξ_u , as defined in equation (A.19), into the definition of Q_u , as defined in equation (A.22), we obtain:

$$Q_u = \frac{s - \mu_{h\phi_u^{-1}}}{(1 - \mu_{h\phi_u^{-1}})^{1+1/\sigma}}$$

Next, using the implicit equation (A.17) defining ϕ_u^{-1} , we obtain that:

$$Q_u = \frac{\int_{\phi_u^{-1}}^u \rho e^{-\rho(u-t)} (s - \mu_{ht}) dt}{\int_{\phi_u^{-1}}^u \rho e^{-\rho(u-t)} (1 - \mu_{ht})^{1+1/\sigma} dt}.$$

Plugging in (A.30) shows that the market-clearing condition (A.29) holds.

Market clearing at $u \in (\phi_0, T_f)$. First note that by equation (A.20), we have $Q_u = \bar{Q}_u$, where \bar{Q}_u is defined in equation (A.4). Then, the demand from low-valuation investors is

$$\int_0^u \rho e^{-\rho(u-t)} (1 - \mu_{ht})^{1+1/\sigma} \bar{Q}_u dt = \int_0^u \rho e^{-\rho(u-t)} (s - \mu_{ht}) dt$$

by definition of \bar{Q}_u ,

Market clearing at $u > T_f$. The market clears by construction of the holding plans of high- and low-valuation investors after time T_f .

A.9 Proof of Proposition 10

For this proof the superscript *ATE* (algorithmic trading equilibrium) refers to the equilibrium objects of Proposition 1, while the superscript *LOE* (limit order equilibrium) refers to the equilibrium object of Proposition 9.

First point: $p_u^{ATE} = p_u^{LOE}$ for $u \geq \phi_0$. By construction $p_u^{ATE} = p_u^{LOE} = 1/r$ for $u \geq T_f$, and $\xi_u^{ATE} = \xi_u^{LOE}$ for $u \in (\phi_0, T_f)$. This immediately imply that $p_u^{ATE} = p_u^{LOE}$ for all $u \geq \phi_0$.

Second point: $p_u^{ATE} > p_u^{LOE}$ for $u \in (T_\phi, \phi_0)$. To obtain the result, we prove the following Lemma.

Lemma A.19. *Let $u \in (T_\phi, \phi_0)$. If $Q_u^{ATE} \geq Q_u^{LOE}$, then time- t low-valuation traders' asset holdings satisfy $q_{t,u}^{ATE} \geq q_{t,u}^{LOE}$ for all $t \in (0, u)$ with a strict equality $t \in (0, \phi_u^{-1})$.*

This Lemma is proved in Section VIII.15. Since high-valuation traders hold 1 unit of the asset in both the ATE and the LOE, it cannot be that both markets clear at a date $u \in (T_\phi, \phi_0)$ such that $Q_u^{ATE} \geq Q_u^{LOE}$. Thus, $Q_u^{ATE} < Q_u^{LOE}$ for all $u \in (T_\phi, \phi_0)$. From the definition of Q_u^{ATE} , we have $\xi_u^{ATE} = 1 - (1 - \mu_{hu}) Q_u^{ATE \sigma}$. From (A.23), in the LOE we have $\xi_u^{LOE} = 1 - (1 - \mu_{hu}) Q_u^{LOE \sigma}$. Thus $\xi_u^{ATE} > \xi_u^{LOE}$ for $u \in (T_\phi, \phi_0)$, and $\xi_u^{ATE} = \xi_u^{LOE}$ for $u \geq \phi_0$. This clearly implies that $p_u^{ATE} > p_u^{LOE}$ for $u \in (T_\phi, \phi_0)$.

Third point: $p_0^{ATE} < p_0^{LOE}$ for ρ close to zero. We start by establishing a preliminary result. Since all the equilibrium variables are defined as roots of continuously differentiable functions that are well-defined in $\rho = 0$, with non-zero partial derivatives, we can apply the Implicit Function Theorem to show that:

Lemma A.20. *Every equilibrium object has a well-defined limit when $\rho \rightarrow 0$ and is continuous in ρ over $\rho \in [0, +\infty)$. Moreover, those limits satisfy the same equations as in the case $\rho > 0$ after letting $\rho = 0$.*

The proof is in Section VIII.16, page 68 in the supplementary appendix. Denoting all these limits with a hat, “ $\hat{\cdot}$ ”, we show that

Lemma A.21. *The limiting prices at time zero satisfy $\hat{p}_0^{ATE} < \hat{p}_0^{LOE}$.*

And the third point of Proposition 10 follows by continuity.

References

- Antje Berndt, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz. Measuring default risk premia from default swap rates and edfs. Working paper, Stanford University, 2005.
- Bruno Biais, Piere Hillion, and Chester Spatt. An empirical analysis of the limit order book and the order flow in the paris bourse. *Journal of Finance*, 50:1655–1689, 1995.
- Bruno Biais, Thierry Foucault, and Sophie Moinas. Equilibrium algorithmic trading. Working paper, TSE and HEC, 2010a.
- Bruno Biais, Johan Hombert, and Pierre-Olivier Weill. Supplementary material for “trading and liquidity with limited cognition”. Working Paper, TSE, HEC, UCLA, 2010b.
- Pierre Brémaud. *Point Processes and Queues*. Springer-Verlag, New-York, 1981.
- Jonathan Brogaard. High frequency trading and its impact on market quality. Working paper, Northwestern University, 2010.
- Michael Carter and Bruce Van Brunt. *The Lebesgue-Stieltjes Integral: a Practical Introduction*. Springer-Verlag, New York, 2000.
- Carole Comerton Forde, Terrence Hendershott, Charles M. Jones, Pamela C. Moulton, and Mark S. Seasholes. Time variation in liquidity: The role of market makers inventories and revenues. *Journal of Finance*, 65:295–332, 2010.
- Joshua D. Coval and Erik Stafford. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86:479–512, 2007.
- René Descartes. *Meditations and other philosophical writings*. Penguin Classics, 1641.
- Darrell Duffie, Nicolae Gârleanu, and Lasse H. Pedersen. Over-the-counter markets. *Econometrica*, 73:1815–1847, 2005.
- Darrell Duffie, Nicolae Gârleanu, and Lasse H. Pedersen. Valuation in over-the-counter markets. *Review of Financial Studies*, 20:1865–1900, 2007.
- Andrew Ellul, Graig W. Holden, Pankaj K. Jain, and Robert H. Jennings. Order dynamics: Recent evidence from the nyse. Working paper, Indiana University Bloomington, 2007.

- Thierry Foucault. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets*, 2:99–134, 1999.
- Thierry Foucault, Ohad Kadan, and Eugene Kandel. Limit order book as a market for liquidity. *Review of Financial Studies*, 18:1171–1217, 2005.
- Xavier Gabaix and David Laibson. The bid bias and the equity premium puzzle. *NBER Macroeconomics Annual*, 16:257–312, 2002.
- Nicolae Gârleanu. Portfolio choice and pricing in illiquid markets. *Journal of Economic Theory*, 144:532–564, 2009.
- Ronald L. Goettler, Christine A. Parlour, and Uday Rajan. Equilibrium in a dynamic limit order market. *Journal of Finance*, 60:2149–2192, 2005.
- Ronald L. Goettler, Christine A. Parlour, and Uday Rajan. Informed traders and limit order markets. *Journal of Financial Economics*, 2009. Forthcoming.
- Robin Greenwood. Short and long-term demand curves for stocks: Theory and evidence on the dynamics of arbitrage. *Journal of Financial Economics*, 75:607–650, 2005.
- Mark D. Griffiths, Brian F. Smith, D. Alasdair Turnbull, and Robert W. White. The costs and determinants of order aggressiveness. *Journal of Financial Economics*, 56:65–88, 2000.
- Sanford J. Grossman and Merton H. Miller. Liquidity and market structure. *Journal of Finance*, 43:617–637, 1988.
- Larry Harris. *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press, New York, New York, 2003.
- Terrence Hendershott and Albert J. Menkveld. Price pressures. Working paper, UC Berkeley and University of Amsterdam, 2010.
- Terrence Hendershott and Ryan Riordan. Algorithmic trading and information. Working paper, UC Berkeley, 2010.
- Terry Hendershott and Mark S. Seasholes. Marketmaker inventories and stock prices, working paper. *American Economic Review (PEP)*, 97:210–214, 2007.
- Amir E. Khandani and Andrew W. Lo. What happened to the quants in August 2007? *Forthcoming, Journal of Financial Markets*, 2008.

- Andrei Kirilenko, Pete Kyle, Mehrdad Samadi, and Tugkan Tuzun. The flash crash: The impact of high frequency trading on an electronic market. Working paper, University of Maryland, 2010.
- Ricardo Lagos and Guillaume Rocheteau. Liquidity in asset markets with search frictions. *Econometrica*, 77:403–426, 2009.
- Ricardo Lagos, Guillaume Rocheteau, and Pierre-Olivier Weill. Crises and liquidity in otc markets. Working Paper, NYU, UCI, UCLA, 2007.
- Gregory Mankiw and Ricardo Reis. Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, 4(1295-1328), 2002.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, 1995.
- Basel Committee on Banking Supervision. Findings on the interaction of market and credit risk. *Bank for International Settlement Working Paper*, (16), 2009.
- Emiliano Pagnotta. Information and liquidity trading at optimal frequencies. Working paper, NYU Stern, 2010.
- Christine A. Parlour. Price dynamics in limit order markets. *Review of Financial Studies*, 11: 789–816, 1998.
- Philip Protter. *Stochastic Integration and Differential Equations*. Springer-Verlag, New York, 1990.
- Ricardo Reis. Inattentive producers. *Review of Economic Studies*, 3:793–821, 2006a.
- Ricardo Reis. Inattentive consumers. *Journal of Monetary Economics*, 8:1761–1800, 2006b.
- Ioanid Rosu. A dynamic model of the limit-order book. *Review of Financial Studies*, pages 4601–4641, 2009.
- Yeneng Sun. The exact law of large numbers via fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126:31–69, 2006.
- Dimitri Vayanos. The decentralization of information processing in the presence of interactions. *The Review of Economic Studies*, 70(3):667–695, 2003.

Pierre-Olivier Weill. Leaning against the wind. *Review of Economic Studies*, 74:1329–1354, 2007.