



NOVEMBER 2017

**LIQUIDITY RISK AND INVESTOR
BEHAVIOUR: ISSUES, DATA AND
MODELS**

**Scientific
Advisory Board
Review**

AUTORITÉ
DES MARCHÉS FINANCIERS



This publication was coordinated by the AMF's Analysis, Strategy and Risk Division and was written by the authors named. Although all sources used are considered reliable, their completeness and accuracy cannot be guaranteed. Unless otherwise indicated, figures and other information are only valid as at the date of publication of each of these analyses and are subject to change over time. All copying, distribution or reproduction of these analyses, in full or in part, requires express prior written agreement from the AMF.

Warning — The "Scientific Advisory Board Review" is a publication presenting analyses and research carried out by the Scientific Advisory Board's members or by researchers from the financial academic or professional sector. Each analysis reflects the personal views of its author(s) and does not express the AMF's position; nor does it necessarily reflect the position of the institution for which the authors work.

Serge Darolles
Université Paris-Dauphine

Summary

An investment fund's liquidity is often equated with the liquidity of the financial instruments that it holds. However, fund management can in some ways resemble a liquidity transformation activity. A change in the structure of a fund's client base affects the potential mismatch between the liquidity of its assets and liabilities. An asset/liability approach for liquidity management is therefore critical and requires a client behaviour model. This paper addresses the various challenges facing management companies: information systems, behavioural models, risk models and regulation. It presents, for a significant number of funds, the results of a statistical study of client behaviour and the impact of this behaviour on a fund's exposure to liquidity risk.

1. THE ISSUES

An investment fund's liquidity risk is often assessed solely by analysing the financial instruments that it holds. The structure of a fund's assets does in fact directly expose its performance to a deterioration in the liquidity of the financial instruments. This approach implicitly assumes that the size of the fund remains constant, and that the deterioration in liquidity does not prompt clients to exit. However, in practice, the size of a fund changes in response to investors' reactions, subject to any liquidity constraints they face. These factors, relating to the composition of a fund's liabilities, therefore have a direct impact on the level of liquidity risk, in particular when the fund invests in illiquid assets while also offering its clients favourable liquidity. Fund management thus resembles a liquidity transformation activity, and any deterioration in the liquidity of the portfolio securities can increase the mismatch between the liquidity of the fund's assets and liabilities. Liquidity risk is a direct consequence of this mismatch, and a combined analysis of the fund's assets and liabilities is needed to quantify it. This suggests, therefore, a form of asset/liability management of liquidity that is possible only if fund managers incorporate information about their clients' behaviour into their analysis.

As such, and given their risk departments' current practices, management companies are facing a variety of challenges. We believe there are four major issues. The first is their information system. Historical data for financial instruments held by the fund are stored in this system. They are used to determine the fund's exposure to various types of risks — market, credit, liquidity, and so on — through statistical calculations based on historical series of transaction prices, spreads, volumes and ratings. In contrast, information about the fund's clients is generally not as easily accessible, particularly when the aim is to work at the most granular level, i.e. the client. In many cases, only total assets under management are stored in the information system as they are used directly to calculate leverage and diversification constraints. Risk teams consequently do not have the tools to systematically monitor changes in the structure of the liabilities. However, these changes may have different causes and different implications for liquidity. Consider the example of a fund whose total size is stable between two dates. This stability might suggest that the fund's exposure to liquidity risk is stable. But it could also be the result of redemptions initiated by long-term investors and of subscriptions by clients investing for very short periods. It is clear, in that case, that the quality of the liabilities has deteriorated and the fund's exposure to liquidity risk has increased. The only way to see this is to track individual investor behaviour and to then infer the changes in the structure of the fund's liabilities. The challenge for management companies is therefore to upgrade their information systems to include liability tracking. This requires a better understanding of their clients. Order marking systems can be used, for example, to identify which client sent a given buy or sell order, then to store subscription/redemption records by client/specific client profile. This is a substantial investment for the management company and presents a number of technical challenges. But this information also simplifies how distributors are treated and paid, and adds to the value of the portfolio of existing clients. In this respect, it could provide a second justification for the investment required.

The second issue is to understand client behaviour. Creating a subscription/redemption database merely gives a historical view of client behaviour. Effective liquidity risk management at a given date requires a forward-looking view to be able to anticipate future subscriptions/redemptions. Without a very clear idea of likely redemption scenarios, fund managers have to invest some of their funds in immediately available liquid instruments to help cover potential redemptions without having to sell illiquid assets at a loss¹. A better understanding of client behaviour allows fund managers to more effectively manage their fund's exposure to liquidity risk and to invest a larger percentage of the portfolio in illiquid financial instruments. Statistical models of varying degrees of sophistication need to be developed to anticipate client behaviour. The first basic approach is to assume stable client behaviour over time. Historical data are used to estimate the marginal distribution of subscriptions/redemptions, then to calculate the probability of extreme events, such as a 50% decrease in assets under management. Much as is done for volatility, it is also possible to include a time component in these simple models. Historical data can be used to determine whether, for example, significant outflows on one date are on average followed by significant outflows on subsequent dates. Time series models can be used to statistically connect observations made on different dates and to capture these timing effects. Finally, a last approach is to

¹ See Darolles, Roussellet (2017) for an analysis of how investment funds manage their cash.

identify the economic factors associated with subscriptions/redemptions, then include these exogenous explanatory factors in the behavioural model. In addition to measuring liquidity risk, developing these models could also establish consistency between the tracking of a fund's liabilities and the calculation of commercial performance indicators. Clients typically do not all have the same value for a management company, and those that tend to remain invested for longer are more attractive than the others. A better understanding of client behaviour therefore also provides effective tools for managing and spurring business development.

The third issue for management companies concerns the connection between an understanding of client behaviour and management of liquidity risk. In absolute terms, risk models should account for both the characteristics of the portfolio, i.e. market liquidity risk, and the way this portfolio is funded by clients, i.e. funding liquidity risk. A fund can invest in illiquid assets without suffering from overexposure to liquidity if its funding relies on inflows of investment from clients with a long-term horizon. Conversely, a fund invested in highly liquid assets is exposed to a liquidity risk if clients have a short-term investment goal. It is therefore important to take full account of both dimensions and to find a way to tie the final risk calculation to the structure of the liabilities. Consider the simple example of the calculation of a fund's value at risk. It is easy to account for the market liquidity of the securities held by the fund and to penalise the most illiquid securities. But, in doing so, we obtain the same value at risk for two funds that hold identical portfolios but have very different client structures. However, a fund whose clients invest in the short term or are very sensitive to the fund's performance should be penalised, because a shock to the market value of the fund implies outflows on subsequent dates, and hence future asset sales at potentially unfavourable prices². This example shows that only an asset/liability approach helps control the investment fund's liquidity risk. Measuring the assets' liquidity risk by accounting for the structure of the liabilities thus helps establish mechanisms to manage liquidity in order to align the percentage of illiquid investments and the structure of the fund's liabilities.

The fourth is open-ended fund regulations. It is possible to reduce the liquidity gap between a fund's assets and liabilities by asking the fund manager to impose liquidity constraints on clients. Various mechanisms can be used to extend a client's average investment horizon. For example, funds can require a minimum holding period or prohibit clients from redeeming more than a certain percentage of their position on a given date. Reducing the liquidity offered protects the client from too wide a gap between asset and liability liquidity and allows the fund manager to continue to invest in long-term assets. These constraints should be imposed only if they are necessary, i.e. when the calculations made for a given fund show that there is in fact a significant gap between liquidity on the asset side and on the liability side. They therefore require a better understanding of the liabilities and their structure at a given point in time, as well as of the liability dynamics, including contagion effects that could trigger massive outflows in response to market shocks. To that end, access to information about liabilities is quite crucial, and it is important that this information be tracked to enable the statistical models to estimate client behaviour.

This paper seeks to present an overview of the research project launched in 2015 at Université Paris-Dauphine, in collaboration with participants from the asset management world, on modelling investment funds' liability liquidity risk. The first step in the thought process was to build a historical database to archive the past subscription/redemption flows of clients of various types of investment funds. A number of management companies participated and gave us access to these data at the individual client level, which is unusual in the world of academic research. This allowed us, in particular, to measure the benefit of working with disaggregated data, at the client rather than fund level. Understanding the links between the decisions made by clients based on type of client builds on the work already underway on buy/sell order marking. This database is described in Section 2 of this paper. Once it was ready, we used it to study how clients behave, with the aim of anticipating their future investment decisions based on observations of their past and present behaviour. The dynamic investor behaviour model is introduced in Section 3 of this paper. The model's estimate for different funds helped identify different client effects depending on the type of fund. This allowed us to quantify the differences observed among the various major families of investors, and thus to project the amount of potential outflows based on the composition of the liabilities. The results obtained therefore made it possible to optimise the length

² Of course, this reasoning works only when it is expensive to reduce the size of the portfolio and is therefore not valid for portfolios invested in highly liquid assets.

of investments on the asset side or to anticipate marketing actions. Lastly in Section 4, we address the various approaches currently being explored.

2. BUILDING THE DATABASE

Management companies currently collect and store very limited data about the behaviour of their clients. But, owing to the digital transformation and regulatory developments, they will quickly have access to increasingly large volumes of data about their clients and how they manage their investments. While these data are very useful for the funds' business development, they are also expected to help the management teams better control liquidity risk. In this section, we describe the various steps in the thought process that led to the creation of the database of the historical behaviour of the clients of their funds.

2.1 Using the information on the liability side of the funds

It is always useful for a management company to identify exactly who the end investor is, that is, who makes the subscription and redemption decisions. A better understanding of liabilities allows for more targeted marketing actions and more effective communication in times of crisis. But, while this need is not new, the existing solutions for developing a better understanding of liabilities only partially meet expectations. The current fund distribution model includes a number of intermediaries between the management company and the end investor, which makes the quality of the order tracking highly unreliable.

Nevertheless, there is currently a great deal of interest in these matters. The first reason is regulatory. The Autorité des Marchés Financiers (AMF) reiterated in February 2017 that "understanding and analysing the fund's liabilities is essential to identifying the risks faced by the asset management company"³. Future European regulations (MiFID II, PRIIPs) will also affect the distribution channel and the producer/distributor relationship, and thus increase the understanding of liabilities. MiFID II, in particular, could be an opportunity to establish a system for distributors to report to producers on the breakdown by investor risk criteria.

The second reason is strategic. The emergence of FinTechs and their direct distribution model could disrupt the traditional channels and offer strong competition to existing players. Having direct access to the end client means they can request information not available to traditional management companies. This information can then be used to better target the investment offering based on the clients' characteristics. Once aggregated, it also provides a better view of the structure of a fund's liabilities and can be used to better anticipate how they will change under stressed market conditions. The inability to conduct this type of analysis puts fund managers at a distinct disadvantage, as their distribution method breaks their direct link to their clients.

The AMF guide⁴ published in February 2017 provides a number of ways to effectively track liabilities. The approaches are mostly qualitative, through a better understanding of this link between manager and investor. However, a statistical analysis of historical subscriptions and redemptions for a fund, based on client characteristics, can be used to better assess redemption risks. A quantitative approach can also be a precondition for a more qualitative, targeted and case-by-case approach based on the weight of different clients and the size of the funds. Developing a quantitative approach nonetheless requires access to resources, data and expertise that a single management company may not have within its teams. It is therefore essential to work together to expand the scope of the statistical study in terms of funds, types of management and historical depth through data-sharing and a substantial collection, anonymisation and standardisation effort.

³ Autorité des marchés financiers (2017), *The use of stress tests as part of risk management: Guide for asset management companies*, February 2017, AMF publications.

⁴ Autorité des marchés financiers (2017), *The use of stress tests as part of risk management: Guide for asset management companies*, February 2017, AMF publications.

2.2 How the consortium benefits academic research

Building a historical database of subscriptions/redemptions, using management company data, was the first step in the collaborative project that formed the basis of this research⁵. Researchers had previously not had access to a database of clients' investment choices. This information is highly strategic for management companies, and it is understandable that they would rather not take the risk of sharing it. The only information that was public was the historical change in the size of investment funds. While this information may be most useful when studying the links between a fund's performance and its size, all liquidity analyses require a more precise picture of funds' liabilities, and of their structure in particular. But this structure changes over time based on clients' entries and exits. A full analysis cannot be conducted using only aggregated data.

The first step in building the database focused on gathering data about the behaviour of a fund's clients. While records of subscriptions/redemptions could easily be obtained from the partner management companies, the challenge was to classify each individual client by type. Institutional clients and retail clients respond differently to shocks to the value of their investments. To complete this classification, we worked closely with each of the partner management companies. This work could only have been done by a consortium that brought together academia and industry. Additionally, large volumes of data are needed to be able to model inflow/outflow probabilities. Collaborating with a number of management companies that were prepared to give us their records allowed us to build this large shared database and make it available to a team of academic researchers.

Once the database was created, the second step was to define a common client nomenclature. We came to realise that none of the management companies used the same client typology. We therefore had to create this single typology, along with rules we could use to assign clients to one of the classifications within this new nomenclature.

2.3 The main characteristics of the database

Access to data from the partner management companies allowed us to select a group of 15 funds, with different management styles, from a very broad universe of funds. The main purpose of this selection was to work with funds with variable exposure to asset liquidity risk. Money-market funds, for example, invest in liquid assets and therefore have very little exposure to a deterioration in liquidity. Large-cap equity funds also have little exposure to market liquidity risk, and this exposure increases for small- and mid-cap equity funds. At the other end of the scale, bond funds are much more sensitive to a deterioration in liquidity as they invest in illiquid assets. On the liabilities side, all the funds selected in theory offer daily liquidity. But an analysis of the behaviour of each fund's client base allowed us to measure how much of this liquidity is actually taken.

Total assets under management for the 15 funds selected as a whole exceeded EUR 10 billion, with some of the funds that are quite large in size due to the inclusion of several money-market funds.

⁵ The project named "MGPF", for *Modélisation de la Gestion du Passif des Fonds* (fund liability management modelling), received the Pôle Finance Innovation label in 2013. It then obtained funding from the French government and from the Ile de France region through an "FUI", for *Fond Unique Interministériel* (single interministerial fund), in June 2014, allowing the Consortium to begin work in earnest in November 2014. See Darolles, Le Fol, Lu, Sun (2017) for a detailed description of this project.

Table 1: Descriptive statistics

Management Company	Fund	Category	Number of sub-segments	Inception	AUM
Firm 1			16		
	Fund 1	Euro equity large cap	3	02/10/1998	329 723 439
	Fund 2	Euro equity mid/small ca	2	06/09/1991	376 326 122
	Fund 3	Euro Fixed Income	2	03/07/1992	255 141 000
	Fund 4	Euro Fixed Income	3	24/02/1982	375 685 999
	Fund 5	Euro Fixed Income	3	05/02/1990	935 044 376
	Fund 6	Euro Money Market	2	31/12/1985	450 074 000
	Fund 7	Euro Money Market	1	07/07/1995	1 570 300 000
Firm 2			19		
	Fund 8	Euro equity large cap	5	20/11/2001	280 424 002
	Fund 9	Euro equity mid/small ca	5	11/05/1994	333 368 999
	Fund 10	Euro Fixed Income	5	25/10/2000	354 900 000
	Fund 11	Euro Money Market	4	08/03/2006	3 415 839 000
Firm 3			8		
	Fund 12	Euro Money Market	2	04/01/2013	1 319 876 994
	Fund 13	Euro equity large cap	3	09/01/2001	295 271 161
	Fund 14	Euro equity mid/small ca	2	14/02/1997	34 287 000
	Fund 15	Euro Fixed Income	1	30/11/2001	388 074 000
Total			43		10 714 336 092

For example, the largest fund in our sample manages more than EUR 3.4 billion, while the smallest is in the Euro Equity Mid/Small Cap category with EUR 34 million under management. At transaction level, for the period under review we gathered and analysed more than 930,000 buy and sell transactions (578,000 buy and 357,000 sell transactions). The following table shows a very wide range of levels of activity for the funds, with a very high number of transactions per day for some funds and a very low number for others. An analysis of this table also shows that daily activity is therefore very high for money-market funds but much lower for bond funds. Additionally, we can see from Table 2 that levels of activity can also vary from one management company to the next.

Table 2: Level of activity

Fund	Period	Days	Total number of			Daily number of		
			Subscriptions	Redemptions	Transactions	Subscriptions	Redemptions	Transactions
			522 015	283 034	805 049			
Fund 1	2013-2014	497	174 903	22 134	197 037	351.91	44.53	396.44
Fund 2	2013-2014	497	144 992	20 880	165 872	297.73	42.01	339.74
Fund 3	2013-2014	497	18 942	6 436	25 378	38.11	12.95	51.06
Fund 4	2013-2014	497	3 709	5 983	9 692	7.46	12.04	19.50
Fund 5	2013-2014	497	5 671	7 323	12 994	11.41	14.73	26.14
Fund 6	2013-2014	497	36 779	54 044	90 823	74.00	108.74	182.74
Fund 7	2013-2014	497	137 019	166 234	303 253	275.69	334.47	610.17
			52 773	68 955	121 728			
Fund 8	2010-2014	1252	7 005	6 354	13 359	5.60	10.67	16.27
Fund 9	2010-2014	1252	5 663	4 037	9 700	4.52	7.75	12.27
Fund 10	2010-2014	1252	1 399	6 952	8 351	1.12	6.67	7.79
Fund 11	2010-2014	1252	38 706	51 612	90 318	30.92	72.14	103.05
			45 768	62 601	108 369			
Fund 12	2013-2014	493	210	312	522	0.43	0.63	1.06
Fund 13	2010-2014	1249	1 468	1 400	2 868	1.18	1.12	2.30
Fund 14	2010-2014	1115	1 877	3 023	4 900	1.68	2.71	4.39
Fund 15	2010-2014	1233	564	564	1 128	0.46	0.46	0.91
			620 556	414 590	1 035 146			

The second step involved putting clients into similar groups. For each of the funds, we had access to several pieces of information: all of the buy/sell transactions made by the fund's investors, a client identifier⁶, the number of units in question, the corresponding price and the transaction date. Each management company classifies all of its clients into different groups, but this classification varies from one management company to the next. We had to create a new common classification for all the companies and apply it to all investors trading in the selected funds. Based on this work, we defined 16 types of investors, ranging from institutionals, private banks and independent wealth management advisors to retail clients. This classification enabled us to observe, on an aggregate basis, the behaviour of a large number of investors and hence to draw a number of conclusions about their average behaviour. For example, do the institutional and retail clients of small- and large-cap equity funds have a short-term investment horizon? It is clear that the only way to answer this question is to work with disaggregated data. Average holding periods cannot be calculated merely by observing investment flows at fund level. Another benefit of the disaggregated approach is the ability to calculate the amount of subscriptions and redemptions separately. A fund's assets under management may remain stable even as a significant percentage of institutional clients is replaced by retail clients. This change in the structure of a fund's liabilities can only be observed from individual flows and is impossible to identify using aggregated flows. Lastly, the statistical treatment of disaggregated data also makes it possible to track the history of a given client within a single management company⁷. We can therefore track any switches between asset classes or types of funds by a given client or group of clients. For example, does a deterioration in market liquidity prompt certain investors to reduce their exposure to this risk and invest in money-market funds? Here as well, we can only answer this type of question through the use of disaggregated data.

Access to disaggregated data also allowed us to address questions about contagion effects. Consider the example of two types of clients — institutional investors and retail clients — invested in a given fund. We could speculate that the best-informed clients will react quickly to a deterioration in the financial environment and thus reduce their risk more quickly, mainly by scaling back their investments in the riskiest assets. Is this response a sign that retail clients will exit in the future? And what do we expect the amount involved in these future exits to be? These questions can also be asked about the investment universe as a whole or just for a given fund. Ultimately, the aim is to use the buy/sell time series for different investments based on type. We can also assess the risks

⁶ For confidentiality reasons, an identifier was assigned to each client within the management company.

⁷ The identifiers used in client anonymisation are defined within each management company. It is therefore impossible to track any client transfers between two different management companies.

associated with the openness of investment funds, and those that exist when the same investors are present in different funds. This can create channels of contagion in the funds' liabilities, which can only be measured with a disaggregated approach.

However, and we are already seeing this with the very limited universe of 15 funds used in this initial study, the size of the samples to be processed quickly becomes enormous, and in our case reached nearly one million transactions. Extending the study to all funds would bring us to several billion transactions. We need to develop models to use these data effectively. They should be simple ones to start with, to make them more accessible to as many people as possible. These models are presented in the next section.

3. MODELLING INVESTOR BEHAVIOUR

This section seeks to describe the different steps in the modelling process. Our intention was, first, to propose simple models, and then to gradually add levels of complexity based on the stylised facts observed. The goal is to develop a standard method and then retain the ability to use more sophisticated internal models. An analogy can be drawn with the calculation of a portfolio's value at risk. We can see the different steps in the modelling process, such as the calculation of a Gaussian value at risk, followed by an ARCH value at risk, and so on. This approach allows us to assess, at each stage, the role played by the new parameters included in the model.

Our aim in this paper is only to model the number of subscriptions and redemptions observed for a given fund, for all clients. It is, of course, possible to apply the model to data from a group of funds — large-cap equity funds, for example — or from the universe of funds as a whole. For example, we discuss only the results obtained for well-identified funds in our sample, based on the investment strategy pursued and the degree of exposure to market liquidity risk. It is also possible to apply this approach to different types of investors, and thus to study the causal effects on the behaviour of these investors. Here we present only the initial results obtained for all investors without distinguishing between types.

Applying this approach to historical subscriptions/redemptions observed for different types of funds and for all types of clients yielded significant results. For example, we observed very different client behaviour based on the liquidity of the funds in which they had invested, even when the funds had identical liquidity conditions.

3.1. The Poisson model

The most simple statistical count model is the Poisson model. It assumes that the number of subscriptions or redemptions observed each day, represented as N_t , is the realisation of a random variable that is independent and identically distributed according to a Poisson distribution with parameter λ . However, one of the stylised facts observed in the transaction data runs counter to the theoretical properties of the Poisson distribution, in particular, that the variance is equal to the mean⁸. We see empirically that the variance in the number of subscriptions or redemptions is far above its mean. Using a simple Poisson distribution would therefore have the disadvantage of incorrectly calibrating the variability of the series of interest and of underestimating, for example, the probability of the liquidity stress scenarios. This must be taken into account when working with models permitting the replication of the empirical characteristics of the series. We therefore use a Poisson model with overdispersion, i.e. the observations are assumed to be drawn from a Poisson distribution with a parameter equal to λF_t , where F_t is the value taken at time t by an unobservable factor drawn from a negative binomial distribution with expectation 1 and parameter⁹ γ . Latent factor F_t makes it possible to create overdispersion and therefore to better calibrate the moments of the empirical distribution. The estimate of the two parameters λ and γ for the different funds in the sample do in fact show that the overdispersion is effective, with the γ parameters statistically different from zero. The estimated values of this parameter, and thus of the levels of

⁸ See, for example, Cox (1983).

⁹ See Johnson et al. (1992).

overdispersion, are higher for redemptions than for subscriptions. There is therefore much more time variability for subscriptions.

This model has the advantage of being very easy to estimate, as well as the disadvantage of being purely static, i.e. the observation distribution on date t does not depend on observations on previous dates. The direct implication is problematic. The best forecast for the number of subscriptions or redemptions on the subsequent date is equal to the mean, constant over time, of the subscriptions or redemptions. However, an analysis of the corresponding time series shows that subscriptions and redemptions can be concentrated during given periods. As with volatility, a period of high redemptions seems to increase the probability of observing significant redemptions to follow. There are therefore persistence phenomena in subscriptions and redemptions that are impossible to capture using a static model. The solution is to include a dynamic component in the model.

3.2. Univariate autoregressive model

The dynamic component of the model is included by adopting the philosophy of the ARCH¹⁰ models, where volatility on a given date depends on the square of returns observed on previous dates. In our framework, we propose replacing the constant Poisson distribution parameter λ with parameter λ_t which varies over time as a function of the previous observation in the time series. The specification used is as follows:

$$\lambda_t = \lambda_0 + \rho N_{t-1},$$

where parameter λ_0 is a constant intensity and additional parameter ρ captures the time persistence in subscriptions or redemptions. If $\rho > 0$, we can clearly see that an increase in the number of transactions on the past date will have a positive impact on intensity λ_t , and thus increase the average number of transactions on the subsequent date. This channel creates both persistence and clusters in the time series of transactions, similar to what we obtain for volatility using ARCH models. It is not difficult to include additional delays in the above specification. In this study, we confine ourselves to including one single delay so as to keep the model as simple and as parsimonious as possible. If $\rho = 0$, we have the simple case described above: the past has no impact on the number of transactions on the current date. It is therefore very easy to determine whether or not there is persistence in the series observed. We only have to estimate parameter ρ and test whether it is statistically different from zero. The estimate of this specification for the series of subscriptions and redemptions in our sample shows that, for the vast majority of funds, persistence parameters ρ are significantly different from zero, and that the highest levels of persistence are observed for redemptions. The risk of seeing orders concentrated over short periods of time is therefore greater for redemptions than for subscriptions.

From a practical standpoint, the advantage of dynamic models is that they provide non-constant forecasts for subscriptions or redemptions. Once parameter ρ becomes significant, the forecast of future redemptions depends on current redemptions, and that of future subscriptions on current subscriptions. This forecast can of course be used by managers of funds that are able to anticipate what the amount of redemptions will be on the subsequent date. They can therefore begin to adjust the size of their portfolio to be able to easily handle their clients' redemption orders.

3.3. Autoregressive model with cross effects

In the previous approaches, the subscription and redemption series are modelled separately. It is interesting to see whether there are cross effects between these two series. To that end, we can expand the proposed approach to include new parameters. We then designate N_t^{in} , the number of subscriptions in a fund on date t , and N_t^{out} , the number of redemptions for the same fund also on date t . We then assume that the intensity of

¹⁰ See Engle (1982) for an introduction to ARCH models.

the Poisson distributions describing subscriptions on date t , λ_t^{in} and that describing redemptions λ_t^{out} satisfy the following two equations:

$$\lambda_t^{in} = \lambda_0^{in} + \rho^{in-in} N_{t-1}^{in} + \rho^{in-out} N_{t-1}^{out},$$

$$\lambda_t^{out} = \lambda_0^{out} + \rho^{out-in} N_{t-1}^{in} + \rho^{out-out} N_{t-1}^{out},$$

where the two parameters ρ^{in-out} and ρ^{out-in} capture the dependencies between subscriptions (or redemptions) on date t and redemptions (or subscriptions) on the previous date. All the other parameters in the model keep their previous interpretations.

It is now possible to discuss the financial interpretation of the four parameters ρ^{in-in} included in the most general version of the model. The first parameter ρ^{in-in} can easily be interpreted in terms of reputation. Past subscriptions on average increase the number of current subscriptions. Parameter $\rho^{out-out}$ captures the effects of panic. Investors, when they see significant outflows, interpret this as a negative signal and tend to exit the fund as well. Cross effects can also be interpreted. Parameter ρ^{in-out} measures the fund manager's ability to stabilise the size of the fund, for example, by initiating marketing actions to offset past outflows with more subscriptions. Finally, the last parameter ρ^{out-in} captures the behaviour of investors that exit the fund after a massive influx of other investors. This can be seen as contrarian behaviour by certain investors that leave the fund when an unusually strong performance attracts an unusually large number of new investors. They anticipate capacity issues and a deterioration in performance due to an increase in the size of the fund.

In terms of liquidity risk management, the presence of positive cross effects is generally beneficial since it tends to stabilise the fund's assets under management. The most critical case corresponds to a value $\rho^{in-out} < 0$, which means fewer subscriptions on average when redemptions increase. This model allows us to separate the negative effects of past outflows into two possibilities. They can increase future outflows or reduce future inflows. An empirical analysis will show which of the two effects is greater.

Table 3: Estimators by type of fund

Category	λ^{in}	λ^{out}	In-In	In-Out	Out-In	Out-Out	g^{in}	g^{out}
Money-Market Fund	8.1***	10.1***	0.32***	0.30***	0.19***	0.46***	35.5***	44.9***
Large-Cap Equity Fund	4.3***	4.8*	0.22***	0.00	0.78***	0.18***	35.9***	68.8***
Small/Mid-Cap Equity Fund	0.5***	1.1***	0.52***	0.06**	0.16***	0.19***	2.5***	4.4***
Bond Fund	0.4***	0.4***	0.10*	0.00	0.00	0.02	0.5***	0.7***

Table 3 presents the estimators¹¹ for this last model for four funds in different categories, and therefore with varying exposure to market liquidity risk. The objective of the exercise is to compare differences in the behaviour of the clients of these funds based on their liquidity, while they all offer daily liquidity.

Our analysis of the table begins with the marginal components captured via parameters λ in the table. We can see that the greater the increase in the fund's liquidity risk exposure, the steeper the decline in the value of the parameters. We therefore observe on average far less activity in funds where assets are illiquid, such as bond funds. If we turn now to the risk associated with the effects of panic, it is important to focus on the out-out column showing the estimators of the $\rho^{out-out}$ for the four funds. We see that this parameter is significant for the first three funds, which have the least exposure to market liquidity, whereas it becomes insignificant for the last fund. It therefore seems that clients do take into account in their investor behaviour the challenges fund managers may face in managing the liquidity offered. A bond fund's clients will have less of a reaction to exits by other investors than a liquid fund's clients. We can interpret the very high value obtained for money-market

¹¹ Parameters statistically significant at 10% (or 5%, 1%) are indicated in the table by * (or **, ***).

funds as the effects of seasonality, which explains why clients exit funds at similar times. The first finding of this study is therefore that fund clients do effectively incorporate the liquidity dimension into their investment policy and that the daily liquidity offered by the fund is not used in the same way depending on the type of fund in question.

If we now look at the cross effects between subscriptions and redemptions, we see that these effects are significantly different from zero and positive for the three most liquid funds. They therefore play a stabilising role and, in particular, offset past outflows with a larger volume of new inflows. This stabilising mechanism does not exist, however, for the fund that is most exposed to liquidity risk. Lastly, to conclude, we can comment on the overdispersion parameter values γ . The highest values are observed for the liquid funds, with the peak reached by large-cap equity funds. Additionally, outflows are systematically more dispersed than subscriptions.

4. CONCLUSION

Working on the behaviour of investment fund clients at the most granular level possible allowed us to obtain a number of new results. First, we were able to assess the difficulty of building a single database using information from different management companies. The specific features of the distribution systems make it very difficult to characterise the end client, and thus to create uniform groups of clients in order to study their behaviour. We then developed relatively simple and easily estimable models of investor behaviour. These models clearly use only a fraction of the information contained in the data, but we were nevertheless able to draw a number of conclusions about how investors behave based on the liquidity of the fund in which they have placed their money. These behavioural models form one of the basic building blocks required to analyse the liquidity transformation performed by investment funds. They can be used to evaluate the liquidity mismatch between the portfolio's level of liquidity and the liquidity actually requested by the fund's clients.

We are clearly still in the early stages of the statistical analysis of investor behaviour. The predictive capabilities of the models could certainly be improved by looking, for example, at the correlations between behaviours of different types of investors, e.g. institutional and retail. Working with much larger data samples would also help improve the quality of the results, and would allow us to use much more complex models. We nevertheless believe that, before we do any more work on the modelling side, we should help educate participants in the asset management world and start to explain clearly what they would gain from a better statistical understanding of the behaviour of a fund's clients. To that end, workshops have been held to give the risk teams at the project's partner management companies an opportunity to discuss the uses of the work presented in this article. A prototype of a risk management tool has also been developed. Based on a big data platform, the tool uses its storage and computing power to give fund managers, salespeople and risk teams real-time knowledge of redemption probabilities and to enable them to simulate various liquidity stress scenarios by incorporating shocks to the structure of their liabilities.

References

Autorité des Marchés Financiers (2017), The use of stress tests as part of risk management: Guide for asset management companies, February 2017, AMF publications.

Cox, D. (1983), Some Remarks on Overdispersion, *Biometrika*, 70, 269-274.

Darolles, S., Le Fol, G., Lu, Y., Sun, T. (2017), Self-Exciting Model for Mutual Fund Flows: Investor Behaviour and Liability Risk, Université Paris-Dauphine Working Paper.

Darolles, S., Roussellet, G. (2017), Managing Hedge Fund Liquidity Risks, Université Paris-Dauphine Working Paper.

Engle, R. (1982), Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation, *Econometrica*, 50, 987-1008.

Johnson, N., Kotz, S., Kemp, A. (1992), *Univariate Discrete Distribution*, Second Edition, Wiley, New York.