How long do equity owners hang on to their stocks?

Bernt Arne Ødegaard*

March 2017

Abstract

We characterize the equity holding periods for all equity owners in a stock market over a 15 year period. The median holding period is 0.75 years. The hazard function for equity ownership is characterized by negative aging, an equity owner is less and less likely to realize a position as time passes. There are clear differences between owner types, where private individuals have the longest holding periods, and financial owners are the least patient. Wealthier households have shorter holding periods. Using turnover to estimate holding periods severely over-estimates actual holding periods.

Keywords: Equity Holding Period; Duration; Failure Time; Survival
JEL Codes: D14; G11; G12

Highlights

- The holding periods for all equity owners in a stock market are described and characterized.
- The typical holding period is 0.75 years.
- Holding periods depend on owner type and firm industry.
- Wealth and other measures of the scale of investment influences the holding period.
- Using turnover to estimate holding periods severely over-estimates actual holding periods.

*Professor, University of Stavanger (UiS). Email: bernt.a.odegaard@uis.no. Mailing Address: EOJ House, UiS Business School, NO-4036 Stavanger, Norway.

The *holding period* for an equity owner is of interest in various theoretical context in finance. But there is a lack of empirical work that *describes* equity holding periods in a generic manner, divorced from specific empirical applications. That is the contribution of this paper. The basis for our contribution is access to the complete records of ownership of Norwegian stocks for a 15 year period. We characterize the holding periods for all equity owners in this stock market. The characterization uses the statistical framework of survival analysis.[1] This generic statistical description will be useful for researchers working on theories where the equity holding period enters.

To motivate the need for such a general analysis, let us mention some of the settings in which holding periods appear in finance. We start with the corporate governance literature, where the lack of *large, patient owners* is bemoaned by many commentators on the sorry state of the equity market.[2] The theoretical governance literature on large owners is actually primarily one about the need for *concentrated* ownership, one need owners large enough to have a *voice*, and with holdings large enough that they will have incentives to *use* their voice, not necessarily about the length of the ownership. Ownership length can however be brought in more indirectly. For example, the threat of exit (The Wall Street Walk, Admati and Pfleiderer (2009)) is one potential governance mechanism interacting with concentrated ownership. Another way to link holding period with governance is the literature linking the asset side of a corporation (the length of the firms' investments) with the horizon of its investors.[3] The idea is that one needs long term owners to guarantee the availability of capital for continued investment in long term projects.

Another literature where holding period matters is *asset pricing*. For example, in a recent theoretical study, Brennan and Zhang (2016) argues that one needs to model the investment horizon of individuals as stochastic. The potential importance of investment horizon is empirically confirmed by Kamara, Korajczyk, Lou, and Sadka (2016), which shows that asset price risk depends on investment horizon, i.e. factor exposures varies with investment horizon. Another example is Chakrabarty, Moulton, and Trzcinka (2016), which documents that there is a performance difference between short-term and long-term institutional trades.

The asset pricing literature segues into the microstructure literature through the classical analysis of Amihud and Mendelson (1986), which posits an (inverse) link between holding period and liquidity. Their model argues that investors who buy illiquid stocks, stocks with high spreads, have longer investment horizons, leading them to amortize the higher transaction costs (spreads) over more future periods. There is a very large literature[4] which uses this model as a motivation to introduce a liquidity measure as a priced risk factor in the crossection of assets. One such liquidity measure is the *turnover* of trading in the stock (Datar, Naik, and Radcliffe, 1998).

Actually, the intuition behind the Amihud and Mendelson (1986) model, that one need long holding periods to amortize transaction costs, links back to the governance literature. This link is due to the transaction costs of building up a large equity position. When building up a large position a trader may need to have a long horizon.

[1] This type of analysis is also termed duration, or failure rate analysis.

[2] See for example "Speed, short-termism and other corporate myths", *The Economist* Dec 5, 2015.

[3] See for example Bebchuk and Stole (1993) and the survey by Stein (2003).

[4] See e.g. Amihud, Mendelson, and Pedersen (2005) for a survey.

A final example of a literature where holding periods have been brought in is the *household finance* literature, specifically the literature on households' financial mistakes.[5] One of the issues discussed in this literature is whether (some) investors trade "too much" (Odean, 1999; Barber and Odean, 2000). Well, if they trade too much, that corresponds to their holding periods being "too short."

In all of these literatures, the theory involves holding periods, either directly or indirectly. A general characterization of holding periods should thus be a useful input in many settings, an input which we aim to provide.

The analyses performed in this paper are several. We first characterize holding periods in general, showing that the median holding period for a random equity owner is 0.75 years. This number is based on the estimated survival curve. Such curves are typically characterized in terms of their hazard rates, the likelihood of "failure" (termination) as time passes. We show that hazard rates for equity ownership are not constant, instead they are characterized by negative aging, an equity owner is less and less likely to realize a position as time passes.

Once we have found the holding periods, the next obvious question concerns their potential determinants. Some of the possible determinants we investigate are motivated by the literatures discussed above. We start by asking whether the type of investor matters. We find that the most impatient investors are financial, with a median holding period of a half year. More patient are private individuals, with a median holding period of 0.83 years. We also ask whether there are industry differences in holding periods. Is it the case that some industries have more patient owners? This is motivated by the "maturity matching" literature, that the holding period of owners may match that of the firm's assets. We do find some industry differences, where the owners with longest holding periods tend to own utilities and financial firms, while IT firms have the least patient owners.

We perform two separate investigations of whether the investment *scale* influences duration. First, motivated by the governance literature, we ask whether the very largest owners are different. We find that they do have longer holding periods. Second, motivated by the literature on household investment mistakes, we look at the relationship between holding period and investment scale, which we show affect holding periods. The directions of these relationships are interesting. Looking at the investment in one stock in isolation, we find that the bigger the initial investment, the longer the holding period. This is expected, and consistent with for example a transaction cost story like Amihud and Mendelson (1986). What is more unexpected is what we find when we control for the portfolio wealth of that investor when buying into a new stock. We find that the bigger the wealth (both including and excluding the new asset), the shorter the duration of the new position. The implication is that holding periods for households are decreasing in wealth.

Finally, we consider the practice of using turnover to estimate holding periods (Atkins and Dyl, 1997). We find that turnover based estimates severely over-estimates the length of holding periods. Where the actual average holding period is 0.76, the corresponding estimate using turnover is 1.87. While the turnover based estimates may get the level wrong, they could still get the *ranking* right. We also look at that, and show that the correlation between the ranking using actual holding periods and turnover based estimates is less than a half.

[5] See Barber and Odean (2013) and Guiso and Sodini (2013) for surveys of this literature.

Before we delve into the analysis of our paper, let us mention some related work. Kyrolainen and Perttunen (2006) looks at Finnish investors, and Dias and Ferreira (2005) looks at Portugal. Our dataset is more comprehensive than those used in these two investigations. The Norwegian data is for all owners in the market for a much longer time period. These other papers are also more concerned with particular issues related to market microstructure. Finally, only Dias and Ferreira (2005) actually uses failure rate analysis.

The structure of the paper is as follows: We first, in section 1, present the market and data sources, and provide some descriptive information about the sample. Section 2 characterizes the holding period using survivorship methods. Sections 3 and 4 looks at owner type and industry differences. Section 5 looks at large owners. Section 6 investigates investment scale for households' investments. Finally, section 7 investigates turnover based estimates of holding periods, before we offer a brief conclusion.

1 Market and data

The firms in the sample are listed on the Oslo Stock Exchange (OSE), which is a moderately sized exchange by international standards. In 1997, the 217 listed firms had an aggregate market capitalization which ranked the OSE twelfth among the 21 European stock exchanges for which comparable data are available. The number of companies on the exchange has increased from 172 in 1989 to 292 in 2007.[6]

In the period of our analysis all listed Norwegian companies tracked their equity owners through a electronic centralized registry.[7] All owners at the OSE needed to maintain an account with the registry.[8] Trades at the exchange are reported to the registry, which takes care of transfers of ownership. The registry then provides companies with lists of their owners (stockholders). The primary data used in this paper is monthly snapshots of these ownership data. At each date we observe the number of stocks held by every owner. Each owner has a unique identifier which allows us to follow the owners' holdings over time. For each owner the data include a sector code that allows us to distinguish between such types as financial owners (e.g. mutual funds), industrial (nonfinancial corporate) owners, individual owners (households), state owners and foreign owners. Note that the ownership identifiers are anonymized, they do not allow us to link these equity ownership data with characteristics of the individual owners beyond the sector code. In addition to the ownership data we use stock market prices and trading volumes from the Oslo Stock Exchange Data Service (OBI).

To show that this is a large sample, in figure 1 we count the number of investors observed each month. The number of Norwegian investors with *some* direct equity holdings varies between two and four hundred thousand, and is increasing until about 2002, after which it flattens out. This flattening is probably linked to the bursting of the dotcom bubble at that time. To give some further perspectives on the sample, we also split the data based on portfolio wealth, by providing separate

[6] For some information about the structure of the Norwegian stock market we refer to Bøhren and Ødegaard (2000, 2001), and Næs, Skjeltorp, and Ødegaard (2009).

[7] Verdipapirsentralen (VPS) –the Norwegian Central Securities Registry.

[8] A household will typically have just one VPS account, but there is no regulation outlawing different members of a household, such as a husband and a wife, having separate VPS accounts

counts of the number of portfolios with values above respectively fifty thousand, hundred thousand and one million NOK.[9] The number of investors with sizeable portfolios is increasing in the more recent period, after a short fall around 2003. Note that this figure is at the level of the total portfolio of each investor. When we investigate duration, we look at this at the security level. One investor is typically invested in several securities.

Figure 1 The number of equity investors in the sample

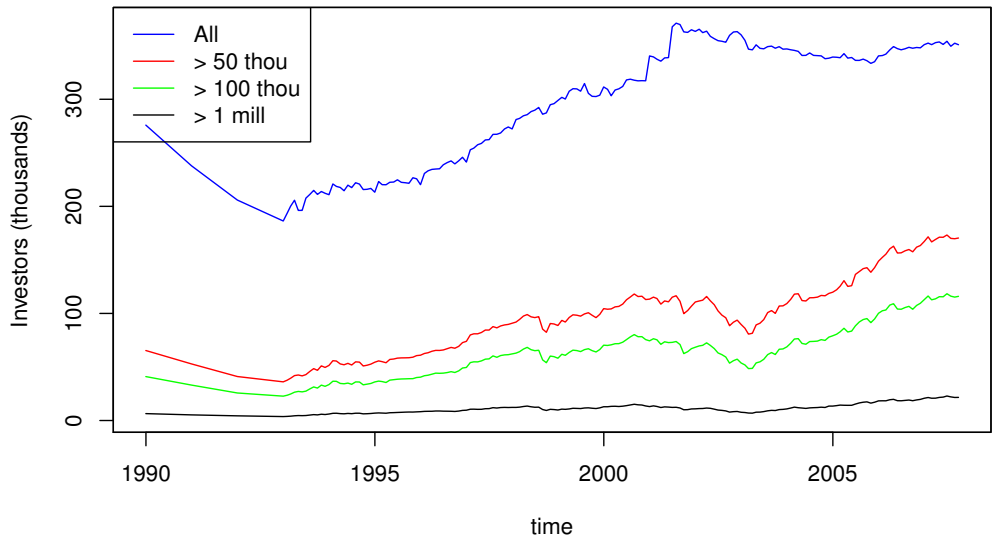

The figure shows the number of investors in the data. The top line (blue) is the total number of investors with nonempty accounts. The lower lines condition on wealth: The red line shows the number of individual accounts with wealth above NOK 50,000 NOK, the green the number of accounts with wealth above 100,000 NOK, and the black line the number of accounts with wealth above 1 million NOK. The wealth boundaries are in constant 2007 kroner. The wealth boundaries correspond to roughly 9, 18 and 180 thousand USD.

2 Characterizing holding periods

2.1 Survival analysis

The econometric framework suited for analyzing questions about the length of time an investor chooses to keep his or her stake in a company, and what economic factors affect this decision, is survival (or duration/failure rate) analysis. In survival analysis, one models the decision to terminate a relationship. In our setting, termination is the decision to liquidate an equity holding in a company.[10]

Survival analysis is the preferable method for analyzing holding period decisions because it is designed to alleviate the problem of censoring. In our setting, the censoring problem stems from

[9] At the 2007 exchange rate of 5.57, these levels correspond to roughly 9, 18 and 180 thousand USD.

[10] The methods underlying survival analysis originated in medicine, and much of the terminology comes from these applications. In economics, duration models are used on e.g. labor market data to analyze determinants of the time spent unemployed, in which case the pertinent termination is movement between employment and unemployment (see Lancaster (1979) and Nickell (1979) for examples and Kiefer (1988) and van den Berg (2001) for surveys.)

the fact that we only observe investors for a limited period of time. Figure 2 illustrates the problem. Of the investors illustrated in the figure, it is only the holding period of investor A which will be measured correctly. The holding period of investor B will be *right censored*; all we see is that the investor was present at the last date and we do not know the final termination date. For investor C we correctly observe the terminal date, but we do not observe when the relationship is initiated, which is termed *left censoring*. Duration analysis involves the estimation of the probability distribution of the termination decision, taking the censoring problem into account.

Figure 2 Illustrating the censoring problem

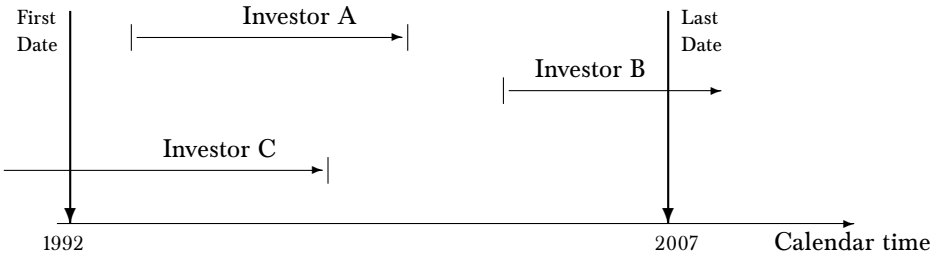

The figure illustrates some conceptual problems in our estimation of holding periods using monthly observations. In calendar time our sample starts in 1992 and ends in 2007. We illustrate the holding periods of 3 example investors, A, B and C. For investor A the holding period is contained within 1992–2007, and therefore estimated correctly. For investor B we correctly observe the initial date but as the investor keeps his stake until after the last date, all we know is that we observe the stake on the last date. The holding period of this owners is underestimated due to right censoring. For owner C we correctly observe the terminal date, but we do not observe the first date, only that this owner was present in the first date of the sample, in 1992. Hence the holding period is underestimated due to left censoring.

This probability distribution of the termination decision is characterized in a number of ways. The most important is the *survival function*: the probability of surviving beyond a given date. Another is the *hazard function*: the probability of termination, conditional on having survived so far. When we want to ask if a covariate affects duration, this is done by measuring that covariate's contribution to the survival function.[11]

2.2 Estimated hazard and survival functions

We apply duration analysis to the holding periods of individual investors using monthly data for all investors at the OSE over the period 1992-2007.[12] To reduce noise, investors with less than hundred shares are removed from the sample. Thus, we count as initiation the first time an investor is observed holding 100 or more shares, and termination when he or she reduces the stake to less than 100 shares.[13] This leaves about 3.2 million observations of investor-stock durations. An investor

[11] It may be helpful to show the more formal definitions of these concepts (See Kalbfleisch and Prentice (2002)). Let T be the failure time. The survivor function $F(t)$ is the proportion of the original sample still alive at time t, $F(t) = P(T > t)$. The hazard function $\lambda(t)$ is the instantaneous risk of death, or more specifically, the instantaneous rate of change in the log of the number of survivors per unit time. In the case of a continuous distribution of T, the hazard is specified as $\lambda(t) = \lim_{h\to 0^+} \frac{P(t\leq T<t+h|T\geq t)}{h}$. We will alternative characterize the hazard by the cumulative hazard: $\Lambda(t) = \int_0^t \lambda(s)ds$.

[12] In survival analysis terms, our data set is a an example of spell data, where there is interval censoring since we only observe once every month, and there are some (identified) spells which may be left or right censored.

[13] At the Oslo Stock Exchange, the typical minimal trading lot is 100 shares. Looking only at complete sellouts of stakes is of course a simple definition of termination. One could think of alternatives, such as a stake decrease by a given percentage.

can have several durations, in the same stock, or in other stocks. In Table 1 we show the median holding period estimated for all owners, which is the holding period of the *typical investor*. We find that the typical investor holds a position for 0.75 years.

Table 1 Median holding periods, all owners

	Median	n
All	0.753	3,203,356

The table reports the median survival time for the sample (where the survival curve crosses 0.5). The estimation is done with R's `survival` package.

In Figure 3 we show the estimated survival and hazard functions for the complete sample of investors. From the survival function, shown in the left panel of the figure, we can read off the median holding period of 0.75 as the point where the survival function crosses the 0.5 line. Other interesting properties of holding periods are, however, better illustrated by the hazard function shown in the right panel of the figure. If the probabilities of liquidating an equity position, conditional on the length of time the ownership has lasted, are time independent, the hazard function will be flat. This is clearly not the case for our sample. Instead, we see a systematic time variation. The conditional likelihood of exit starts high, and then decreases steadily over time. In the failure time literature, this behavior of the hazard rate is called negative aging.

Figure 3 Estimated hazard and survival functions

Panel A: Survival function

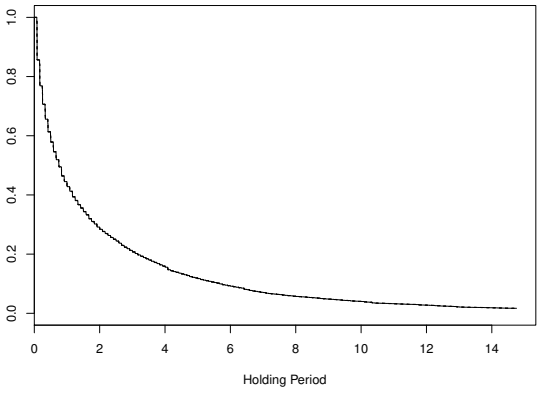

Panel B: Hazard function

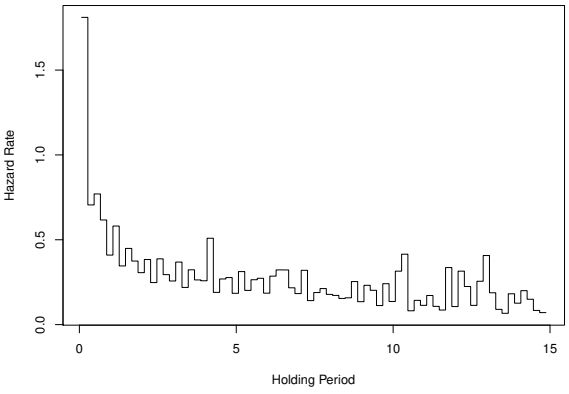

Estimated survival and hazard functions using all investor-company holding periods at the OSE in the period. The figure on the left is the estimated survival function. The figure on the right is the estimated hazard function. Analysis time in years. The analysis is based on 3,305,752 observations. The estimates are corrected for right censoring. The estimation is done with R's `survival` package. The hazard function is estimated with R's `pehaz` estimator.

3 Does holding period differ with owner type?

We now ask whether the holding periods differ, depending on investor type. To answer this we estimate separate survivor functions for the five different owner types in our data: individual, financial, foreign, nonfinancial and state (public) owners. The differences are most clearly illustrated by the estimates of the cumulative hazard rates, shown in Figure 4. The curves show that the individual owners are at the bottom, and the financial owners at the top, a pattern which is consistent for the whole sample period.

Figure 4 Cumulative hazard curves, by owner type

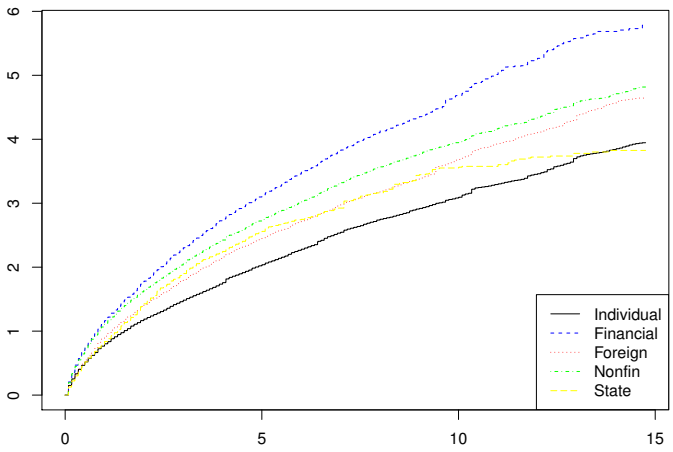

Estimated cumulative hazard (Λ) curves, grouped by owner type. The different owner types are private (individual), financials, industrial (nonfinancial corporate), state and foreign owners. Estimated using R's `survfit` function.

To quantify the differences, Panel A in Table 2 shows the estimates of median holding period for the five owner types. Confirming the figure, the most patient investors are private individuals, who hold their positions for 0.83 years, while the typical corporate investor, be it financial or nonfinancial, holds a position for only half a year. A different method to quantify differences between owner types is to estimate contributions of dummy variables for the different owner type to the survival curve. The result of such a regression is shown in Panel B of Table 2. The differences in coefficient estimates confirms that individual owners are the most patient, while financial owners are the least patient, and the other owner types are in between.

Table 2 Differences in holding periods with owner type

Panel A: Median Holding Periods

	Median	n
Individual	0.833	2,498,745
Financial	0.499	89,911
Foreign	0.748	277,823
Nonfinancial	0.501	328,528
State	0.830	8,349

Panel B: Contributions to survivorship function

Individual	0.541***
	(0.001)
Financial	−0.056***
	(0.005)
Foreign	0.316***
	(0.003)
Nonfinancial	0.071***
	(0.003)
State	0.337***
	(0.016)
Observations	3,203,356
Note:	*p<0.1; **p<0.05; ***p<0.01

The table in Panel A shows estimated median holding periods for the subsample of that owner type. For each sample we estimate a separate survivorship curve. The owner types are private (individual), financials, industrial (nonfinancial corporate), state and foreign owners. The table reports estimates of median holding period for these separate estimations. The column on the left reports medians, the column on the right the number of observations in each sample.The table in Panel B shows estimated contributions to the survival function for the different owner types. The owner types are private (individual), financials, industrial (nonfinancial corporate), state and foreign owners. We run a single survivorship regression, using all the data, with dummies for each of the five types. (The estimation is done without a constant term.) The estimation is done with R's `survival` package. The regression is done with `survreg` specifying a Weibull distribution.

4 Are there industry differences?

Motivated by the possibility of industry differences in holding period driven by "duration matching," we group the firms at the OSE into ten different industry categories.[14] As we did for the owner types, we compare the ten industries by plotting the cumulative hazard rates, as shown in figure 5.

Figure 5 Cumulative hazard curves by firm type

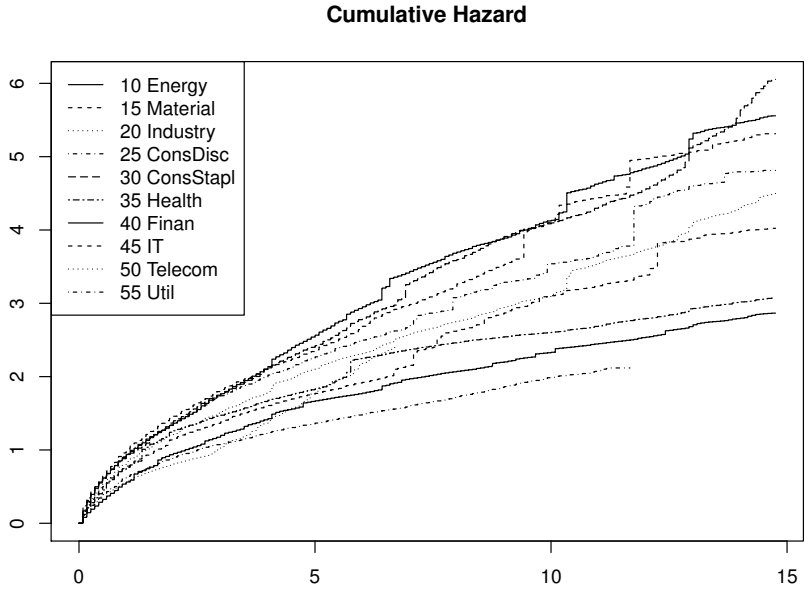

Estimated cumulative hazard (Λ) curves, grouped by firm category. Firms are categorized using the GICS standard. Estimated using R's `survfit` function.

The figure shows that there are clear industry differences, but the picture is a bit crowded. The relative differences may be better gauged by the estimated medians reported in Panel A of Table 3 and the contributions to the survivorship function reported in Panel B of the same table. The more patient owners are in utilities, telecoms and financials, while the least patient owners are in IT firms.

[14] We use the GICS categorization of industries. The GICS categorization was introduced by MSCI and S&P, and is in common usage in the asset management industry. The GICS standard was updated in 2016, adding a separate group for real-estate firms. In our sample this group is not present. Instead, in the period we consider here, pure real estate firms will be in the financials group.

Table 3 Holding periods by firm category (GICS)

Panel A: Median holding periods

	Median	n
10 Energy	0.666	833,533
15 Material	0.833	175,674
20 Industry	0.756	570,180
25 ConsDisc	0.666	175,229
30 ConsStapl	0.666	231,164
35 Health	0.833	118,587
40 Finan	1.255	434,304
45 IT	0.586	614,805
50 Telecom	1.422	88,913
55 Util	1.252	21,987

Panel B: Contributions to survivorship function

10 Energy	0.251***
	(0.002)
15 Material	0.654***
	(0.004)
20 Industry	0.495***
	(0.002)
25 ConsDisc	0.358***
	(0.003)
30 ConsStapl	0.255***
	(0.003)
35 Health	0.663***
	(0.004)
40 Finan	0.971***
	(0.002)
45 IT	0.237***
	(0.002)
50 Telecom	0.858***
	(0.005)
55 Util	1.043***
	(0.013)
Observations	3,196,576
Note:	*p<0.1; **p<0.05; ***p<0.01

The table in Panel A shows estimated median holding periods for the subsample of firms in a given industry. The industry grouping uses the 10 GICS categories in the original MSCI/S&P GICS categorization. For each sample we estimate a separate survivorship curve. The table reports estimates of median holding period for these separate estimations. The column on the left reports medians, the column on the right the number of observations in each sample. The table in Panel B report coefficient estimates for a survivorship regression using as covariates ten dummy variables for the GICS industries. The regression is estimated without a constant term. The estimation is done with R's `survival` package. The regression is done with `survreg` specifying a Weibull distribution.

5 The largest owners

In this section we look at the largest owners. Are the owners with the largest investment in a corporation likely to have longer holding periods? This is primarily motivated by the governance literature. Is it necessary to have a longer horizon when building up a larger stake? A simple investigation of this is to pick a subsample of only the largest owners in a corporation. How does this subsample compare to the sample of all owners? To this end we pick all owners which has at one time been the largest, and redo some of the analyses we did for all owners for this sample. Table 4 shows estimates of median holding periods corresponding to the estimates in Tables 1 and 2. Not surprisingly, the estimated holding periods are on average longer. The typical holding period for this sample is 1.75 years, one year longer than for the typical owner. In some sense this seems a low figure, many of these investors are dominant owners with representation on the board of a corporation.

Table 4 Median holding periods. Largest owners, different owner groups.

	Median	n
All	1.753	1,935
Individual	2.504	102
Financial	1.082	338
Foreign	1.501	810
Nonfinancial	2.003	604
State	4.255	81

For each sample we estimate a separate survivorship curve. The table reports estimates of median holding period for these separate estimations. The column on the left reports medians, the column on the right the number of observations in each sample. Only using owners which at some point have been the largest in a corporation. The estimation is done with R's `survival` package.

The table also shows medians for the different owner types. For all of them the length of holding periods are much longer for the large owners. Crossectionally, between types, the main difference between these numbers and the earlier ones is that the state holding period is relatively higher here. That is driven by the Norwegian State's large stakes in some of the larger Norwegian corporations.[15] To look in more detail at the differences between owner types, the best illustration is the cumulative hazard curves, shown in figure 6. One observation that is particularly obvious in the hazard curve is that the main source of the difference between state and the individual owners is the behavior at the long end.

[15] Examples are Statoil, DNB and Telenor, companies where the state retains a large portion after privatization. Some of the patient state holdings are also driven by stakes by the Government Pension Fund Norway (GPFN), better knows as "Folketrygdfondet", which is a large, passive institutional investor in the Norwegian equity market.

Figure 6 Cumulative hazard curves. Largest owners, different owner types.

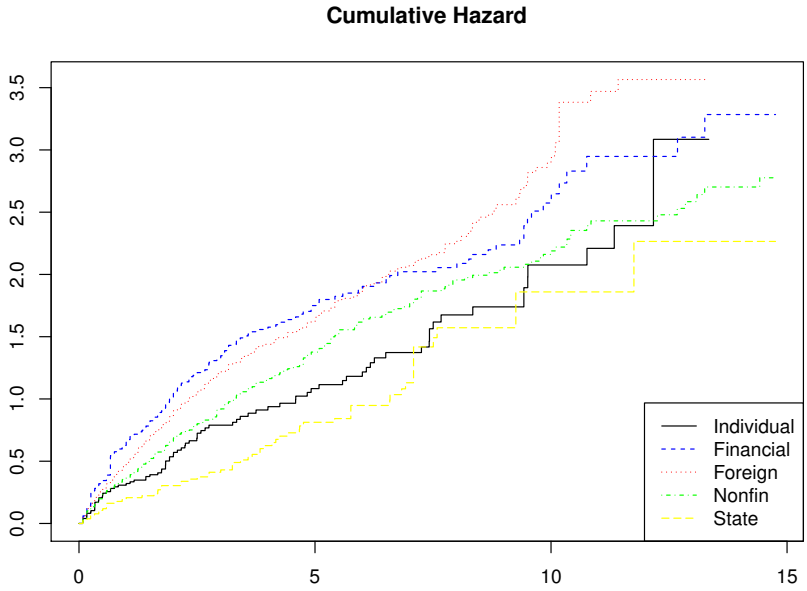

Estimated cumulative hazard (Λ) curves, grouped by owner type. The different owner types are private (individual), financials, industrial (nonfinancial corporate), state and foreign owners. Only using owners which at some point have been the largest in a corporation. Estimated using R's `survfit` function.

6 Investment scale and holding periods for households

In this section we relate holding period to investment scale, looking separately at the holding periods by private investors (households). This is linked to the household finance literature where one investigates investment decisions of households. Are there systematic "mistakes" made by households? Do these mistakes vary e.g. by the financial sophistication of a household (Calvet, Campbell, and Sodini, 2009), or by the importance of the decision in the household's portfolio? If one wants to argue that households are making mistakes, one need a model specification in which one can argue that particular decisions amounts to suboptimal behaviour. In one popular branch of this literature, which looks at diversification decisions, it is clear that an undiversified portfolio does not use what is closest to a free lunch in finance, the opportunity to reduce risk at no cost in expected returns.

For holding period, on the other hand, there is no such clear specification of what would be suboptimal behaviour. What is closest is the literature arguing that some investors "trade too much" (Odean, 1999; Barber and Odean, 2000). Their argument is primarily about the scale of the transaction costs paid by "churning" the portfolio. For our purposes, we will not try to argue about what would be optimal behavior. We merely investigate whether scale affects holding periods.

We consider different measures of wealth. The first is the initial investment in a stock when a position is established, without looking at that investor's total investment in the stock market. The second looks at that household's total investment in the stock market at the time of the entry into a position. The third subtracts the initial investment in the stock analyzed, so that the wealth is the amount invested in *other* stocks at the OSE. We use the logarithm of the amounts to reduce scaling issues in the estimation.

One possible source of noise relates to the usual issue in household finance, that for small stakes, an investment decision may be unimportant relative to the total portfolio of an investor.[16] To potentially remove some of the less informed investment decisions, we look at the sample of investments where the investment amount is above 100 000 NOK.[17]

For the case of the initial investment, holding period is increasing with the initial amount invested. But for total portfolio wealth, both including and excluding the new investment, the sign is negative. The implication is that wealthier individuals hold stocks for shorter periods. This is an interesting observation which may provide some food for theorizing.

[16] For the same sample, Ødegaard (2017) shows that diversification is increasing in the wealth of an investor.

[17] At the 2007 exchange rate this cotoff corresponds to USD 12,500. We provide corresponding estimates for *all* owners, including those below 100,000 NOK, in an Internet Appendix. In the regression where we use total wealth as a covariate, the restriction is on total wealth, which explains that that sample is larger.

Table 5 Explaining survivor curve by investment scale for family owners

	(1)	(2)	(3)	(4)
ln(Initial Investment)	0.539***			0.059***
	(0.003)			(0.004)
ln(Portfolio Wealth) (including new asset)		−0.092***		
		(0.001)		
ln(Portfolio Wealth) (excluding new asset)			−0.121***	−0.129***
			(0.001)	(0.002)
Constant	−6.177***	1.488***	1.630***	1.020***
	(0.040)	(0.014)	(0.020)	(0.042)
Observations	475,913	1,393,497	345,469	345,469
Note:				*p<0.1; **p<0.05; ***p<0.01

The table shows estimated contributions to the survival function. We only use data for private (individual) owners with investment (wealth) above 100,000 NOK. Each column represents the result of a survivorship regression explaining survivorship with various measures of investment amounts. Wealth is measured three ways: The first is the initial investment in a stock (the initial investment is the number of stocks held at the end of the first month times the stock price). The second is the total portfolio wealth of the investor. Portfolio wealth measured by summing the number of shares multiplied by the stock price at the end of the first month of an investment. The third is the total portfolio wealth less the initial investment. We take logs of all wealth measures to avoid scaling issues in the estimation. The estimation is done with R's `survival` package. The regression is done with `survreg` specifying a Weibull distribution.

7 Holding period and turnover

Stock turnover is often used as a proxy for holding period. The argument is simple: Turnover measures how much of the firm's equity is traded during a given period. If for example a stock's annual turnover is 50%, that means that 50% of the firm's outstanding equity was traded during one year. We would then say that the expected holding period for an investor in this stock was 2 years, because it would take 2 years for all of the firm's equity (100%) to be traded. In this section we look at to what extent this type of estimate of holding period from stock turnover is related to actual holding periods.

7.1 Turnover data

Turnover is calculated on an annual basis. Figure 7 describes the turnovers at the OSE in the period in question, by a time series of annual crossectional medians. In the period, the median turnover varies between 50% and 90%.

Figure 7 Summarizing turnover at the OSE

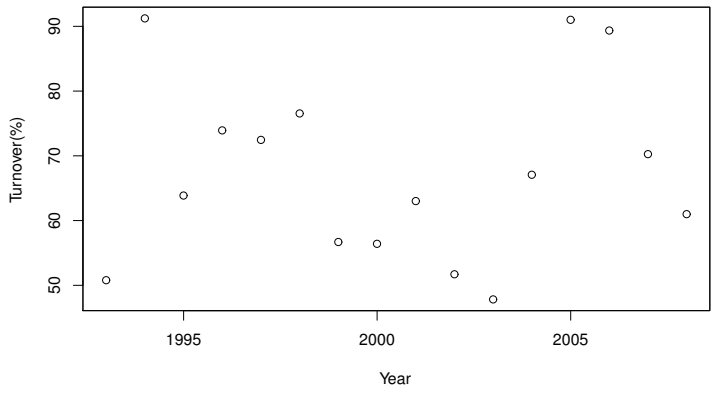

The figure plots annual averages of percentage turnover for OSE firms. The annual turnover is the sum of daily turnovers, calculated as the number of shares traded in a day divided by number of shares outstanding on that day. The turnover is calculated across all stocks at the OSE that traded more than 100 times during the year. See Ødegaard (2016) for more details on liquidity at the OSE, and detailed definitions.

7.2 Point Estimates

Let us start by looking at the magnitude of estimated holding period using the two methods. We first calculate the holding period median using survival analysis. This is done by picking all owners for a given security, and estimate a separate survival curve for that security. Of interest is the median holding period, which we calculate for each stock. This median is then compared to the implicit holding period estimated from turnover, which is estimated as

$$\text{Holding Period} = \frac{1}{\text{Annual Turnover}}$$

We perform these two estimations on a stock by stock basis for all stocks on the OSE in the period. We only include stocks where we can estimate holding period using both methods, which leaves us slightly less than 400 stocks in total.

Figure 8 shows the distribution of the stock by stock estimates of holding period using the two methods. The estimates using turnover are clearly much higher than the median holding periods. This is confirmed by the descriptive statistics in Panel A of Table 6, where we find that the estimate using turnover is on average 1.2 years higher than the actual median holding period. The crossectional average holding period estimated from the survivorship curve is 1.2, while the corresponding estimate from turnover is 2.4. The turnover estimates thus over-estimate actual holding periods by slightly more than a year.

Figure 8 Distributions of estimated holding periods

Panel A: HP estimate using survival methods

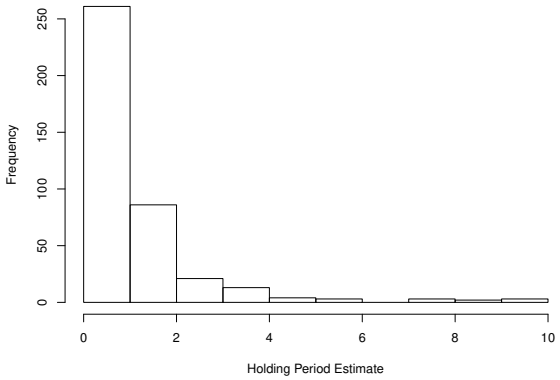

Panel B: HP estimate from turnover

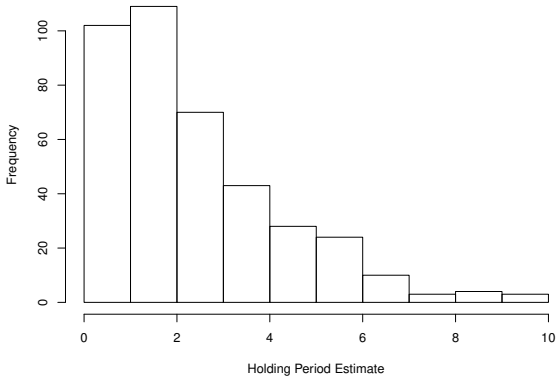

The figures shows the distribution of estimated holding periods, using two different methods for estimating holding period. The first, in panel A, uses the median holding period from a survival curve estimation. The second, in panel B, uses the holding period estimated from turnover, calculated as Holding Period $= 1/$Annual Turnover. Estimation is done for all stocks on the OSE where both estimators can be calculated. The figures shows the distribution of the stock by stock point estimates.

7.3 What about the crossection?

While the holding period estimates using turnover do not get levels right, the method may still be reasonable in terms of identifying stocks with long term vs short term owners (As done in asset pricing studies using turnover as a liquidity measure, which is argued to test Amihud and Mendelson (1986) implications). To investigate that we look at the sample of estimated holding period using the two alternative methods. Do the turnover based estimates "get the ranking right" relative to the more correct method? The answer is no, as the simple correlations in panel B of Table 6 shows. The estimated correlation is between 0.37 and 0.52.

Table 6 Comparing actual holding period estimates and turnover based holding period estimates

Panel A: Summarizing estimated holding periods

Statistic	Estimate (Survival)	Estimate (Turnover)	Difference
N	396	396	396
Mean	1.212	2.402	1.190
St. Dev.	1.353	1.846	1.705
Min	0.162	0.253	−4.672
Median	0.756	1.866	0.813
Max	9.671	9.170	8.512

Panel B: Crossectional correlations of alternative holding period estimates

	Correlation
Standard	0.467
Rank(Kendall)	0.366
Rank(Spearman)	0.516

Panel A: Summary statistics for stock by stock estimates of average holding periods calculated with two methods. The first, in column 2, is the median holding period from a survival curve estimation. The second, in column 3, is the holding period estimated from turnover, calculated as Holding Period $= 1/$Annual Turnover. Column 4 describes the difference between the two measures.
Panel B: The table lists correlations between two different methods for estimating holding period. The first is the median holding period from a survival curve estimation. The second is the holding period estimated from turnover, calculated as Holding Period $=$ $1/$Annual Turnover. These calculations are calculated for each stock in the crossection, for the whole period 1992–2007. The table lists the correlations between these estimates, using three methods for calculating the correlation: 1) Standard correlation, 2) Kendall Rank Correlation and 3) Spearman Rank Correlation.

8 Conclusion

We have, using a dataset that covers 15 years of the holdings of all owners at the Oslo Stock Exchange, characterized the distribution of equity holding periods.

We estimate the generic holding period, and find that the typical investor has a holding period of 0.75 years. The hazard function is falling with time, implying that the distribution has a tail of very long period owners. Holding periods varies with such variables as investor type and the industry of a given firm. Owners with the largest stakes in a company have longer holding periods.

We also investigate to what degree investment scale influences holding periods for households. We find that a larger initial stake leads to longer holding periods. What is more surprising is that the *portfolio wealth* of an investor has an opposite effect. Wealthier households (with larger wealth in other OSE stocks) have shorter holding periods.

We finally investigate to what extent the traditional estimate of holding periods from turnover does a reasonable job. We find it does not. Using the turnover based estimate leads to an over-estimate of holding periods by about a year. Using turnover also does not do a good job in a crossectional sorting of firms by holding period. The correlation in the ranking using turnover and the ranking one gets from the actual holding periods is less than a half.

We motivated this paper by saying that generic properties of holding periods may give useful inputs in various settings. Let us close by pointing out some interesting avenues for further research based on our findings. First, the fact that holding period seem to vary with industry. Is this consistent with "maturity matching"? Second, the fact that wealthier individuals seem to have shorter investment horizons is intriguing. Finally, that we see that turnover based measures do a bad job in linking the crossection of stocks to their actual holding periods may lead to a re-assessment of the empirical asset pricing work where turnover is used as a liquidity proxy justified by the Amihud and Mendelson (1986) argument that liquidity is related to holding periods.

References

Anat R Admati and Paul Pfleiderer. The "Wall Street Walk" and shareholder activism: Exit as a form of voice. *Review of Financial Studies*, (1):3–41, 2009.

Yakov Amihud and Yakov Mendelson. Asset pricing and the bid/ask spread. *Journal of Financial Economics*, 17:223–249, 1986.

Yakov Amihud, Haim Mendelson, and Lasse Heje Pedersen. Liquidity and asset prices. *Foundations and Trends in Finance*, 1(4):269–363, 2005.

Allen B Atkins and Edward A Dyl. Transactions costs and holding periods for common stocks. *Journal of Finance*, 52(1):309–325, March 1997.

Brad M Barber and Terrance Odean. The behavior of individual investors. In Constantinides, Harris, and Stulz (2013), pages 1533–1570.

Brad M Barber and Terrence Odean. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance*, 55(2): 773–806, April 2000.

L A Bebchuk and L A Stole. Do short-term objectives lead to under- or overinvestment in long-term projects? *Journal of Finance*, 48:719–729, 1993.

Øyvind Bøhren and Bernt Arne Ødegaard. The ownership structure of Norwegian firms: Characteristics of an outlier. Research Report Nr 13/2000, Norwegian School of Management, September 2000.

Øyvind Bøhren and Bernt Arne Ødegaard. Patterns of corporate ownership: Insights from a unique data set. *Nordic Journal of Political Economy*, pages 57–88, 2001.

Michael J Brennan and Yuzhao Zhang. Capital asset pricing with a stochastic horizon. Working Paper, SSRN, 2016.

Laurent E Calvet, John Y Campbell, and Paolo Sodini. Measuring the financial sophistication of households. *American Economic Review*, 99(2):393–398, May 2009.

Bidisha Chakrabarty, Pamela C Moulton, and Charles Trzcinka. The performance of short-term institutional trades. *Journal of Financial and Quantitative Analysis*, 2016. Forthcoming.

George M Constantinides, Milton Harris, and René M Stulz, editors. *Handbook of the Economics of Finance volume 2B. Asset Pricing*, volume 21 of *Handbooks in Economics*. Elsevier, Amsterdam, 2013.

Vinay T Datar, Narayan Y Naik, and Robert Radcliffe. Liquidity and stock returns: An alternative test. *Journal of Financial Markets*, 1:203–219, 1998.

Jorge D Dias and Miguel A Ferreira. Timing and holding periods for commmon stocks: A duration-based analysis. Working Paper, ECB and ISCTE Business School-Lisbon, February 2005.

Luigi Guiso and Paolo Sodini. Household finance: An emerging field. In Constantinides et al. (2013), pages 1397–1532.

John D Kalbfleisch and Ross L Prentice. *The Statistical Analysis of Failure Time Data*. Wiley, second edition, 2002.

Avraham Kamara, Robert A Korajczyk, Xiaoxia Lou, and Ronnie Sadka. Horizon pricing. *Journal of Financial and Quantitative Analysis*, 51:1769–1793, 2016.

Nicholas M Kiefer. Economic duration data and hazard functions. *Journal of Economic Literature*, XXVI:646–679, June 1988.

Petri J Kyrolainen and Jukka Perttunen. Do Individual Investors Care About Transaction Costs? Bid-Ask Spreads and Holding Periods for Common Stock. *SSRN eLibrary*, December 2006.

Tony Lancaster. Econometric models for the duration of unemplyment. *Econometrica*, 47(4):939–56, July 1979.

Randi Næs, Johannes Skjeltorp, and Bernt Arne Ødegaard. What factors affect the Oslo Stock Exchange? Working Paper, Norges Bank (Central Bank of Norway), December 2009.

Stephen J Nickell. Estimating the probability of leaving unemployment. *Econometrica*, 47(5):1249–66, September 1979.

Terrence Odean. Do investors trade too much? *American Economic Review*, 89:1279–1278, 1999.

Bernt Arne Ødegaard. The liquidity of the Oslo Stock Exchange – a source book 1980-2015. Manuscript, University of Stavanger, August 2016.

Bernt Arne Ødegaard. Is diversification increasing in wealth? norwegian evidence. Working Paper, University of Stavanger, March 2017.

Jeremy C Stein. Agency, information, and corporate investment. In George M Constantinides, Milton Harris, and René M Stulz, editors, *Handbook of the Economics of Finance*, volume 21 of *Handbooks in Economics*, chapter 2, pages 111–166. Elsevier, Amsterdam, 2003.

G J van den Berg. Duration models: Specification, identification and multiple durations. In James J Heckman and Edward Leamer, editors, *Handbook of Econometrics*, volume 5 of *Handbooks in Economics*, chapter 55, pages 3381–3453. Elsevier, 2001.