The effect of mergers on equilibrium prices: evidence from Israeli provident funds

Roma Poberejsky*

October 2016

Abstract

I study the effect of a medium-size merger between two Israeli financial intermediaries, Meitav and Dash, in 2012. Using a nested logit structure of demand, I model the behavior of agents in the provident fund market, estimating weak substitution patterns between firms and fund specializations. I provide ex ante theoretical predictions regarding post-merger price increase on the scale of 2%-4% for the merging firms, with practically no effect for competitors. Incorporating moderate efficiency gains, the model predicts price drops of about 10%-12%, which are consistent with results from ex post difference-in-differences estimation.

1 Introduction

In July 2012 'DS Apex Holdings' (hereinafter "Dash") had filed a merger request with the Israeli Antitrust Authority (IAA), detailing its intention to acquire one of its competitors (mainly in the exchange-traded note and mutual fund markets) - 'Meitav Investment House' (hereinafter "Meitav"). Less than four months later the merger was approved with remedies, and finalized by March 2013. The newly formed Meitav-Dash, totaling some 120 billion ILS in assets under management, focused its activity in the markets for exchange-traded notes, mutual funds, and provident funds, with market shares of 30%, 18%, and 11%, respectively.

The purpose of this work is to measure the effect of that horizontal merger on market outcomes through merger simulation in the provident fund market. Although Meitav-Dash became only the third largest provident fund manager and no concerns were raised with the regulator regarding that particular market, an analysis of the industry holds special interest: First, the market received

*Department of Economics, Hebrew University. Thesis submitted for the MA degree. Contact: roma.po@mail.huji.ac.il. I greatly appreciate the advice and guidance of Alon Eizenberg, and the comments of David Genesove. This work has benefitted from discussions with Eugene Kandel and Yishay Yafeh, as well as the useful comments of seminar participants at the Hebrew Univeristy and Israel Antitrust Authority.

extensive public and regulatory attention in recent history, culminating with the 2005 report of the 'Bachar Committee',[1] whose main concern was with the highly concentrated ownership structure of Israeli capital markets, and its call for structural reforms. Until that point, provident funds were predominately owned by the two largest banks - Leumi and Hapoalim - covering some 95% of the market. Competition concerns were raised by regulators, pointing out the nearly non-existent cross-bank substitution patterns. The committee's central recommendation - the separation of banks and provident funds - went into effect in 2005, ushering a new competitive setting as insurance firms and investment houses entered the market.

Second, there is evidence to suggest that the proposed merger introduces substantial efficiency gains, in both variable and fixed costs, respectively: The industry is characterized by economies of scale, where unit costs fall as funds' total assets under management increase; In addition, it was reported that some 200-300 of the combined 900 employees of the two firms were laid off following the merger.[2] The distinction between efficiency types is important since savings in variable costs, but not in fixed costs, may almost immediately benefit not only the merging firms, but also pass through to consumers. This paper provides an opportunity to study the extent to which the above cost synergies are present.

Finally, within the scope of this multi-market merger, a study of either the provident fund or mutual fund markets in particular takes advantage of a setting for ex post analysis to complement the ex ante simulation. That approach is, of course, impossible for ETNs, as remedies were enforced in the form of divestiture.

Using an extensive data set from the Ministry of Finance, I estimate demand in the market for provident funds by imposing a simple structure of utility-maximizing price-taking consumer behavior, together with oligopolistic Nash-Bertrand competition between firms. The derived price elasticities and marginal costs provide ex ante predictions of post-merger equilibrium price effects resulting from the merger. Simulations of this sort allow us to separate the merger effect from the overall changes that occur simultaneously. Moreover, this setting helps shed some light on the decision-making and competition policy concerns of antitrust authorities, as well as demonstrate how these tools can be used to address them.

Motivation for this research question indeed lies with competition policy first and foremost. Antitrust regulators are faced with difficult high-impact and high-profile decisions to make in very short time when called to evaluate a proposed merger. On the one hand, economic theory provides little reason to suspect that firms take into account consumer surplus or total welfare when considering a merger, rather than just own profits. It is therefore reasonable that certain mergers need be blocked, as firms will exercise market power and set non-competitive prices. On the other hand, efficiency gains from the concentration process could outweigh the increase in market power and must be

[1] Available at http://ozar.mof.gov.il/bachar/asp/home_en.asp

[2] Source (Hebrew): http://www.globes.co.il/news/article.aspx?did=1000734996

considered, especially if economies of scale are present in the industry, as in the case at hand. Mergers of that sort do benefit society and should ideally be cleared by the regulator. Overall, the total welfare effect of a merger is often ambiguous and regulators are expected to decide according to unknown counterfactual outcomes. The framework this paper presents allows for just that. By imposing several simplifying assumptions, variations of the logit model generate estimates for own- and cross-price elasticities of demand, marginal costs and markups of firms, and predictions for post-merger equilibrium prices.

Finally, this growing body of economic literature on merger simulation, effects, and evaluation, together with the comprehensive set of methodological tools it provides, have seen practically no implementation in Israeli market research. This paper provides a setting to utilize contemporary models and theories of industrial organization using extensive Israeli market data.

Approaches to merger analysis can be generally classified into either event studies (also referred to as 'reduced-form') or structural modeling. The former focus on isolating the effect of the merger on the merging firms, usually within a difference-in-differences framework, by exploiting some exogenous variation. While its results are intuitive and arguably more precise than those of structural modeling for any given merger, finding truly exogenous shocks to the ownership structure of firms has proven to be very difficult. Even when such natural experiments are identified, as in Hastings (2004), generating out-of-sample predictions remains challenging. Structural models, on the other hand, rely on identifying the system of aggregate supply and demand, and have the flexibility to account for changes in firm behavior. This methodology primarily allows researchers to conduct ex ante merger simulation, as well as model calibrations, and not just ex post evaluations. Structural models, however, can be susceptible to their underlying assumptions and this work is no exception in that respect.

One of the more commonly used methods to measure the effect of mergers first requires estimating a system of aggregate demand through assumptions on consumer behavior. The importance of imposing a preference structure over differentiated products was stressed in works as early as Spence (1976) and Dixit and Stiglitz (1977). However, early discrete choice structures such as McFadden (1974) and Trajtenberg (1989) suffered from inherent drawbacks of their overly-simplifying underlying assumptions for the logit model, that were found to be in contrast with what we expect and observe in differentiated products markets. Namely, the independence of irrelevant alternatives (IIA) property, when considered in a logit framework, states that if a good is eliminated from the choice set, individuals who would have chosen it will redistribute themselves across the remaining goods according to their respective market shares. Formally, the structure assumes errors to be independent and identically distributed across both individuals and goods. This specification generates unrealistic substitution patterns that are driven by market shares alone, and not by product characteristics. In the case at hand, it implies that if a firm discontinues one of its stock-specializing funds, consumers would mostly substitute towards funds with the greatest market shares (non-specialized, or 'general' funds), regardless of their specialization or ownership.

Later work has identified and pointed out these weaknesses and developed new tools to overcome them. Berry (1994) demonstrated how in both the aggregate nested logit and random coefficients logit models the error term enters utility in an additive linear form, allowing endogenous prices to be easily instrumented and demand to be consistently estimated while providing richer and more reasonable patterns of substitution. Berry et al. (1995) applied the model to data from the American market for cars. These influential papers illustrated how the model can be estimated using only market-level aggregate data while accounting for endogeneity of prices, and how it generates more realistic cross-price elasticities of demand (i.e. higher for products of similar characteristics).

The first attempts to implement merger simulation as a tool for competition policy soon came about. Hausman et al. (1994) model a particular stratified form of demand, though not at an individual level, while Berry and Pakes (1993), on the other hand, took the first steps in implementing the rich new models previously mentioned to the study of mergers. In more recent years, numerous merger studies have been conducted, spanning a variety of industries and often employing an event study approach. Some notable examples explored the weakened competition following a player exiting the market (Hastings (2004)), the effects of a large scale high-profile merger like Maytag-Whirlpool (Ashenfelter et al. (2013)), and a cross section of five different markets featuring marginal mergers (Ashenfelter and Hosken (2010)). While most of the literature naturally covered low-tech and manufacturing industries - as these constitute the great majority of economic activity - few, like Allen et al. (2014), studied price effects following a merger of financial service providers (commercial banks in the mortgage market).

The growth of empirical merger studies through simulation, witnessed in recent years, can likely be attributed to advances in numerical computation power. These improvements propagated the field forward, allowing models with no analytic closed form solution to be estimated, accommodating for richer models and more reasonable substitution patterns. Notably, Björnerstedt and Verboven (2015) is an example of a study that explores both an ex ante simulation and an ex post evaluation of a large merger. The authors estimated a two-level nested logit and a random coefficients logit models of demand in the Swedish market for painkillers, to simulate and analyze the 2009 merger of GlaxoSmithKline and AstraZeneca. I closely follow a similar approach using only a constant expenditure specification of individual utility in a nested logit model, and complement the analysis with an ex post merger evaluation.

I estimate weak patterns of substitution, suggesting nests play an important role in consumer behavior, and provide additional evidence to support this notion in the form of descriptive statistics and OLS regressions showing that savers do not substitute away when competing funds are introduced. Simulation results imply moderate price effects for the merging parties, with little effect for competitors, in equilibrium. The introduction of efficiency gains yields predicted price drops that are consistent with observed trends in fees and difference-in-differences regressions.

The rest of the paper is structured as follows: section 2 presents the data;

section 3 discusses the institutional background for the industry and merger; section 4 details the two-level nested logit model, introducing the framework used to study consumer and firm behavior, and addresses the model's specifications; section 5 presents the main set of results from demand estimation and merger simulation; and section 6 provides a discussion of the findings and concludes.

2 Data

The data set is a panel of monthly fund-level observations, covering all existing provident funds between 1999 and 2014, reported to the Ministry of Finance by the funds and are publicly available.[3] Observed fund characteristics include type (central or personal), main and sub-specializations (e.g. general, bonds, stocks; and sovereign bonds, foreign stocks, etc.), deposits, withdrawals, and total assets under management in ILS, fees, returns and past performance (3 and 5 year average returns and their standard deviation), as well as categorical variables for fund managers and controlling firms. Additional macro-level data, such as central bank interest rate and consumer price index, were obtained from the Bank of Israel. Table 1 provides summary statistics.

Market shares, crucial for the analysis, are commonly derived from either quantities sold (volume) or revenues generated (value) at every period, i.e. flow variables. In the case at hand, this straightforward approach would miss the fact that savers continue to receive financial services for their accumulated capital as well. Every month, individuals indeed choose a provident fund to deposit their savings in, but also where all their previous deposits — the accrued funds — are managed. Often, though not necessarily, this would be the same fund. I argue that specifying market shares as portions of total assets under management better reflects the economies of scale present in the industry, and their derived efficiencies and market power - fund managers with a larger asset balance, rather than periodical flows, face lower transaction costs and can consequently offer lower fees to savers.

Table 2 provides a breakdown of 2011 average firm and specialization market shares, weighted by assets under management. The dominance of general funds is immediately evident, with nearly 85% of all savings allocated. Far behind are funds specializing in bonds (4.54%), nominal (2.3%) and indexed (2.25%) securities, and stocks with a mere 1.08% share. Roughly 5% of market capitalization is in funds with an emphasis on designated bonds, but this figure is mostly driven by a minority of players, and is a relic of the time when provident funds could enjoy the state-subsidized returns of pension plans. The two largest players, Psagot and Clal, hold almost half of the market, with respective shares of 25.72% and 20.38%, followed by Harel (11.41%), Dash (8.64%) and Delek (7.6%). Meitav only constituted 3.34% at the time. The merged Meitav-Dash, according to these numbers, is predicted to become the third largest player,

[3] Accessible at http://gemelnet.mof.gov.il.

with a market share of 11.98%. The table also illustrates the similarity in specialization composition of different players. It can be shown that, as a share of firm total assets under management, the figures are remarkably similar across competitors.

Advanced study funds, which are reported in the same category by the Ministry of Finance, were omitted from the data (39830 observations). Their initial purpose was indeed short term savings for advanced studies, although nowadays they serve as a popular channel for general short term saving and constitute a separate market, reaching roughly 158 billion ILS by the end of 2014, compared to 211 billion ILS for provident funds (Figure 1). I also discard 13723 observations of funds that are not available to the general public, i.e. sector- or firm-specific funds, since individuals generally cannot substitute towards them. These consistently comprise some 20% of provident fund market capitalization, as shown in Figure 2. The final sample consists of 73066 observations, spanning over 1999-2014 and covering 781 funds by 18 firms.

Fees were unfortunately reported as annual, rather than monthly, aggregates by the Ministry of Finance. Discussions with scholars in the field indicate that this move was intentional: an attempt to deter savers from making long-term decisions based on short-term performance and discounts. Practically, this move creates difficulties on two separate levels: it introduces measurement error that biases downwards the estimated merger effects, and it makes ex post evaluation challenging.

If fees vary across months, the use of annual averages leads to measurement error bias in demand estimation, where prices are correlated with the error term. This time-aggregation measurement error is added to the standard measurement error caused by aggregating over individuals observing different prices (e.g., list versus transaction prices in Berry et al. (1995)). Both sources of measurement error introduce attenuation bias in estimates, making demand appear less elastic to price. It could possibly be the case, as deposit rates are much higher in December, probably due to savers wishing to maximize their tax benefits for the ending fiscal year.[4] If provident funds raise fees accordingly to meet this rising seasonal demand, my results could be biased downwards, providing a lower bound for the effect of the merger. Instrumental variables are often a simple way to correct for the bias, which I indeed implement, though as pointed out in Berry (1994), since fees are already treated as endogenous, measurement error in prices may not be a serious problem.

The second challenge is ex post analysis. For this part, it is standard practice to compare the predicted changes in equilibrium fees to those observed in the following periods (Early 2013, in this case). This exercise often allows for an evaluation of the model and calibration of its parameters. Instead of comparing pre- and post-merger fees in a somewhat narrow window around the merger to get a reasonably clean estimate of the merger effect, say October 2012 and January 2013, I will only be able to compare the averages of 2012 and 2013. The

[4] Israeli authorities offer subsidies for long term savings in the form of tax benefits, incentivizing greater deposits to pension, provident and advanced study funds.

low frequency of fee data, unfortunately, provides a rather challenging setting for any such calibrations.

3 Industry Background

3.1 Provident fund market

In the Israeli savings market, provident funds, like pension funds, are a type of long term savings plan for retirement. While the accumulated savings in either may only be withdrawn upon reaching retirement age (currently set at 62 for women and 67 for men), the two differ in the required complementary insurance. Savers in a pension fund must also purchase disability and life insurance, while savers in provident funds are not mandated to. Long term savings markets have seen several regulatory changes and reforms in recent years, including one where savings accumulated in provident funds until 2008 can be immediately withdrawn on retirement ('central funds'), while any savings after that point may only be received as monthly payments ('personal funds'). Another major policy change that went into effect in 2008 dictates that all employees must hold a pension plan, effectively removing the outside option for that market.

Institutional investors, providing these services, accumulate large amounts of capital, which is subsequently invested - mostly in listed and private companies' stocks and bonds, as well as in sovereign bonds. As of December 2014, Israeli provident funds' assets under management comprised a large portion of all institutional investor types, totaling roughly 211 billion ILS by the end of 2014 (Figure 1). Meanwhile, other product categories, like life insurance[5] and pension funds amounted to about 221 billion and 191 billion[6] ILS, respectively. Consequences of the 2008 policy change in pension funds are also illustrated in the same figure, where many new plans were opened and pension funds' assets began to grow more rapidly.

Unlike choosing a pension plan in recent years, individuals may still choose the outside good when it comes to provident funds. The process of selecting and purchasing these financial products, and the timing of their actual consumption are what distinguishes this market from most others. Consider an individual buying a car: even without any assumption or specification of her utility function or choice process, it has to be some proactive decision of selecting a certain car over all other options, including not buying a car at all. In other words, one must opt-in to deviate from the outside option. In the savings market, however, it is common practice for employers to offer their employees some default service provider which, in turn, offers some general savings plan as the default, together with a few additional specialized options like stock- or bond-oriented plans. Some even provide a richer choice set, including higher resolution

[5] 'Life insurance' is a misleading title: these are defined-benefit retirement plans, where the funds pre-commit to the rate at which savings will be converted to an annuity upon retirement, thus insuring savers against rising life expectancy. 'Longevity insurance' is probably a better name.

[6] Excluding "old" pensions plans - those which could only be joined until 1994.

specialization options, e.g. sovereign bonds or foreign stocks. If an employee wishes to deviate from the default, additional forms must be filled. Figure 3 illustrate the overwhelming majority that stays with the general plan, suggests that this inertia is not negligible.

Switching costs of that kind are not specific to provident funds nor to Israel and have proven difficult to differentiate from variation in individual preferences. Farrell and Klemperer (2007) survey the literature on switching costs and discuss difficulties in identifying and separating these costs from unobserved preference heterogeneity. In closely related fields, Madrian and Shea (2001) study enrollment in 401(k)-type savings plans and find that a substantial fraction of participants retain the default contribution rate and portfolio allocation, while Handel (2013) estimates a dynamic model of consumer inertia in health insurance markets and studies the impact of policies that nudge consumers toward better decisions by reducing inertia. Though dynamic approaches allow for richer patterns of behavior and paths of convergence to market equilibrium, static merger simulation is a significantly simpler model to estimate.

Horizontal mergers of financial intermediaries are not uncommon, and the industry is often characterized by substantial economies of scale, suggesting efficiency gains from the Meitav-Dash merger are to be expected. In a recent study of a medium-size merger in the Chilean pension industry, Agostini et al. (2014) provide evidence of economies of scale for personnel, marketing and administration expenses, and that marginal costs are decreasing at least in some range; while Jayaraman et al. (2002) study mutual fund mergers, estimating ambiguous performance effects and a significant reduction in expenses. Amel et al. (2004) review the evidence on consolidation and efficiency in the financial sector, showing that, for asset management companies, the ratio of operating expenses to fund assets, a proxy for the managerial and administrative efficiency of a fund, declines steadily as assets grow. The possibility of exploiting scale economies exists since many operating costs of asset management are constant, explaining why firms with more assets under management have higher margins.

Some initial impression regarding the expected results can be drawn at this stage. Figure 4 illustrates how assets correspond to the number of competing funds, where the estimated slope is not significantly different than zero. In an unintuitive manner, it suggests that an opening of a new fund by a competitor does not correspond to a decrease in assets under management. In other words, savers tend not to substitute towards alternatives when these are introduced, and funds likely cannibalize on own assets or attract new savers to the market, rather than those of competitors. Though not causal, this interpretation is consistent with the low elasticity of demand I hypothesize and indeed measure.

Employer-level data on offered funds were not available to me for the purposes of this work and thus I am unable to learn about deviation patterns from default managers. In most cases, however, the default plan offered is 'general' (i.e. a balanced mix of assets, including stocks, corporate and sovereign bonds, and cash), in which case the data suggest that the overwhelming majority does not deviate from it.

An important feature in the process of fund shopping and fee negotiation is

that many firms have trade unions which often negotiate with service providers on behalf of employees and likely manage to get better rates. In ignoring this delegated aspect of fee negotiation, I effectively assume that unions are acting as agents for savers and that no agency problems exist. In a similar manner, large enough employers (or whole sectors) may offer their employees the option of firm- and sector-specific provident funds - unavailable to the general public and usually offer much lower rates. Funds of this type constitute around 20% of the market (Figure 2) and were omitted from the analysis. Delegated negotiation with unions also creates incentives for fund managers to compete over markets, rather than within them: by offering lucrative initial terms, acquiring employees of whole firms as clients could promise a pool of captive savers that will not easily switch to a competitor.

A second characteristic that sets this market apart from many others, be it goods or services, is the timing of product consumption. Individuals, in fact, consume the service only many years after the purchase, sometimes up to half a century later. Together with results indicating that savers rarely deviate from the default, this is consistent with studies in the field of behavioral economics, showing that individuals are generally myopic and poorly discount future values.

Finally, prices (fees) are denoted as a percentage of assets under management, and beginning in 2013, also as a share of deposits. While this fact does not complicate the modeling, since revenues can still be computed regardless of how many factors are separately priced, it does raise concerns of the estimated change in price and its comparison to post-merger observed fees. Since fees were reported at an annual resolution, I can, at best, compare the average fees of 2012, where only the stock of assets was charged, to the 2013 averages, which consist of both stock and flow fees. I compute a synthetic price for these purposes by dividing total revenues by assets.

3.2 The merger

On July 22, 2012 Dash and Meitav expressed their plans to merge and filed a request with the IAA, requesting the horizontal merger to be cleared.[7] The core business of the two firms was providing private, business and institutional clients with an array of financial services, such as brokerage, investment banking, and portfolio and fund management. Meitav and Dash were major competitors in the market for exchange-traded notes[8] (ETNs), but also significant players in the mutual and provident funds markets, with combined market shares of 43%, 18%, and 11%, respectively, while assets under management totaled 120

[7] Application available at http://www.antitrust.gov.il/files/11579/8969.pdf (Hebrew).

[8] Exchange-traded notes are a rare sight in the US capital market, compared to exchange-traded funds (ETFs), but are prevalent in Israel. ETNs are similar to ETFs in that they both track an underlying asset, offer lower fees than actively managed mutual funds, and are traded on the Tel Aviv Stock Exchange. However, an ETN is an unsecured debt note issued by an institution and is not a direct claim on the asset. On the one hand, this both minimizes tracking errors and provides tax benefits to investors, since the ETN does not buy and sell assets within the fund. On the other hand, investors are exposed to credit risk of the ETN underwriter.

billion ILS at the time of the merger. While no concerns were raised with the Antitrust Authority regarding the provident fund market, Meitav and Dash were still significant players and measuring the merger effects holds interest as it allows us to measure synergy gains in the new equilibrium.

The merger was approved on November 14th, 2012, conditional on divesting the exchange-traded note activity of either one of the merging firms to a third party. Selling Meitav's ETN business (13% of the market) to Psagot was approved in March 2013 and the merger was finalized on March 19th, 2013, with a significantly lower yet still substantial market share of 30% in the ETN market.

In order to simulate the merger, I must specify a certain point in time for the merger, whereas the process took nearly eight months from initial application to completion. Following the literature, I specify the merger date to be November 2012, when the request was cleared with conditions by the IAA. Even though merging parties are strictly prohibited by law from coordinating behavior during the review period, there still remain incentives to cooperate before consummation and there is evidence in the literature that firms exhibit changes in pricing behavior following the announcement of a planned merger. Kim and Singal (1993), for example, find that many airlines increased fare prices before their mergers were completed. The estimated demand parameters are nevertheless robust to the chosen date.

Having the merger cleared on the condition of selling part of the ETN business suggests that there were indeed some competition concerns with the regulator, which having Meitav divest its ETN activity ameliorated. From the IAA 2012 Annual Report,[9] it is evident that the concern was solely with the ETN market:

> In November 2012 the Director General conditioned the merger between Dash and Meitav, two of the four substantial players in market for exchange traded notes, on the condition that the traded notes operations of one of the companies must be divested to a third party. This prevented a situation in which two dominant firms would jointly hold 80% of the relevant market (the merged entity 43% and the second firm 34%). Both companies are investment banks that are engaged in the management of pension funds, mutual funds, brokerage and exchange traded notes. With respect to the former fields, no reasonable concern for substantial harm to competition was found, while such concern was present in the market for exchange traded notes.

On the other hand, evidence to support the notion that the merger was expected to negatively affect competition and raise prices is provided in Figure 5, depicting the downward trend in fees across practically all firms in the months and years leading to the merger. Prices dropped from a monthly level of 0.096% in late 2009, to about 0.082% in mid-2012 — a decrease of 14.6%. In an environment with constantly decreasing margins, the merger could have been an

[9] Available at http://www.antitrust.gov.il/files/32807/OECD Annual Report 2012.pdf.

attempt to gain market power and stop the decline, while reducing personnel and operation costs. It was later recorded that fees continued their descent further down to roughly 0.063% by the end of 2014 — an additional decrease of some 23%.

Finally, the Israeli provident fund market would be considered only moderately concentrated, according to US Department of Justice guidelines (a Herfindahl-Hirschman Index in the range of 1000 to 1800), and steadily becoming more competitive over time (Figure 6). Even though, by definition, the market's post-merger HHI did increase, it does not indicate a substantial gain in market power, according to the same guidelines (an increase of under 100 points).

4 Empirical strategy

The nested logit structural form I impose on consumer demand follows Björnerstedt and Verboven (2015). Using their notation and proposed functional form for demand, I posit a constant expenditure specification, rather than the standard unit demand. This difference, however, affects the interpretation of my results, as market shares would not be in terms of market volume, but rather market value. Though the constant expenditure assumption is less common in the literature and is not a convincing functional form for many products, the market for provident funds is one of the easiest settings to argue for it: deposits are defined as shares of gross income and it is not uncommon for savers to deposit into one fund, while keeping previously accumulated capital in another.

An alternative demand structure, random coefficients logit (RCL), was introduced by Berry (1994) and allows for flexible substitution patterns, unconstrained by predetermined nesting segmentation. The main drawback is the lack of a closed form for market shares, complicating computation. Nevo (2000) carefully discusses methods of estimating random-coefficients logit models.

4.1 Savers

Let us consider an economy with T national-level markets, representing different time periods, I consumers and $J+1$ differentiated products. I denote $i = 1 \dots I$ and $j = 0 \dots J$ respectively, where $j = 0$ is the outside good. I omit the market subscript t for a clearer notation. Let the utility of consumer i from good j follow:

$$u_{ij} = x_j\beta + \xi_j + \alpha f(y_i, p_j) + \psi_{ij} \tag{1}$$

Where x_j is a vector of observed product characteristics, ξ_j is the product-specific fixed effect, capturing unobserved characteristics, $f(y_i, p_j)$ is a function of individual income y_i and product price p_j. Finally, ψ_{ij} is a heteroskedastic random error term that follows an extreme value distribution. Stemming from Roy's Identity, and conditional on buying product j, demand of consumer i for

product j takes the form:

$$d_{ij} = -\frac{\partial f(y_i, p_j)}{\partial p_j} / \frac{\partial f(y_i, p_j)}{\partial y_i} \tag{2}$$

It is common practice to specify $f(y_i, p_j)$ such that demand d_{ij} would equal one, meaning that if individual i chose to purchase product j, she would only consume one unit of it. This perfectly inelastic conditional demand form is a great fit for many markets, such as cars or computers. In the savings market, however, there are no discrete units of product. Moreover, in practically all cases, savings are defined as a portion of a worker's wage.[10] I therefore desire a specific form for $f(y_i, p_j)$ that would satisfy this condition. If y_i/p_j is the quantity purchased (individual income over price), I require $d_{ij} = \gamma \frac{y_i}{p_j}$, where γ is constant and denotes the share of income allocated to provident funds. As in Björnerstedt and Verboven (2015), I define

$$f(y_i, p_j) = \gamma^{-1} \ln y_i - \ln p_j$$

such that

$$d_{ij} = -\frac{-1}{p_j} / \frac{\gamma^{-1}}{y_i} = \gamma \frac{y_i}{p_j}$$

and utility now takes the explicit form of

$$u_{ij} = x_j \beta + \xi_j + \alpha(\gamma^{-1} \ln y_i - \ln p_j) + \psi_{ij}$$

To simplify the notation and distinguish between individual-specific and product-specific characteristics, utility can be rewritten as

$$u_{ij} = K_i + \delta_j + \psi_{ij} \tag{3}$$

$$where\ K_i = \alpha \gamma^{-1} \ln y_i \tag{4}$$

$$and\ \delta_j = x_j \beta - \alpha \ln p_j + \xi_j \tag{5}$$

The two-level nested logit model enables substitution patterns to manifest through the assumed distribution of the random utility term ψ_{ij}, allowing for partial relaxation of its independence assumption across alternatives. I partition products into $g = 0 \dots G$ exhaustive and mutually exclusive groups, within which the error terms are correlated. For convenience, group 0 contains only the outside option and has a normalized mean utility, $\delta_0 = 0$, such that $u_{i0} = K_i + \psi_{i0}$. Each group g is subsequently partitioned into $h_g = 1 \dots H_g$ mutually exclusive subgroups, where error terms show additional correlation.

I decompose the error term, ψ_{ij}, into an i.i.d shock at the consumer-product level (ε_{ij}) together with nest- and subnest-specific components, that reflect the structure of demand, and give it the following explicit form:

$$\psi_{ij} = \zeta_{ig} + (1 - \sigma_2)\zeta_{ih_g} + (1 - \sigma_2)(1 - \sigma_1)\varepsilon_{ij}. \tag{6}$$

[10] Exceptions to this would probably be the self-employed, which usually put aside a fixed amount, if anything at all, rather than a portion of their somewhat volatile monthly income.

Where ζ_{ig} and ζ_{ih_g} capture unobserved (to the econometricians) fixed effects, common to all products in nest g and subnest h_g, respectively.

The nesting parameters, σ_1 and σ_2, capture random taste patterns and reflect the nested form of demand. The two parameters measure the importance of the structure that defines products within subnests and nests, respectively, as being distinct from products outside. As σ_2 (σ_1) approaches one, the products in the same nest (subnest) can be considered stronger substitutes than products across nests (subnests). If σ_2 (σ_1) is zero, preferences for the products within the nest (subnest) are not correlated in any way distinct from the rest of the products. In that case, where consumers are just as likely to substitute towards a product within the nest (subnest) as to an outside product, the model simplifies to either a one-level nested logit (if only one σ is zero) or the standard conditional logit model (if both). For consistency with random utility maximization, the model requires $0 \leq \sigma_2 \leq \sigma_1 < 1$, such that products within a subgroup are the closest substitutes, followed by products within a group and only then by products outside the group.

I explore several nesting structures, including single-nest specifications, defined at either the firm, main or secondary fund specialization levels. The preferred model considers groups at the firm level, and subgroups at the main specialization level. Motivation mainly comes from the fact that firms actively address savers wishing to switch to a competitor, exerting marketing efforts, as in Hastings et al. (2013), and offering discounts or other products as substitutes, which is not likely the case when savers move from one specialization to another within a firm.

Finally, the conditional probability of individual i choosing product j over any other alternative l, including the outside option, is rearranged such that we get an expression of choice probabilities. The probability that the selected product z_i by consumer i is j is given by:

$$P(z_i = j|x; \delta, \sigma) = P(u_{ij} > u_{il} \quad \forall l \neq j) =$$

$$P(K_i + \delta_j + \psi_{ij} > K_i + \delta_l + \psi_{il} \quad \forall l \neq j) = P(\psi_{il} - \psi_{ij} < \delta_j - \delta_l \quad \forall l \neq j)$$

Market shares s_j of all products are given by the above conditional probability:

$$s_j(\delta, \sigma) = \begin{cases} P(z_i = j|x; \delta, \sigma), & j \neq 0 \\ 1 - \sum_{j=1}^{J} s_j, & j = 0 \end{cases}$$

The nested logit probability of individual i choosing product $j \in h_g$ can be written as the product of standard logit probabilities: (a) the probability that an alternative within group g is chosen; (b) the probability that an alternative within subgroup g_h is chosen, given that an alternative in g is chosen; and (c) the probability that product j is chosen, given that an alternative in g_h is chosen. Conditional on observables and model parameters, this is denoted as

$$s_j(\delta, \sigma) = P(g_i = g) * P(h_{g_i} = h|g_i = g) * P(y_i = j|g_i = g, h_{g_i} = h),$$

and in the form of logit probabilities:

$$s_j(\delta,\sigma) = \frac{e^{I_g}}{e^I} * \frac{e^{I_{gh}/(1-\sigma_2)}}{e^{I_g/(1-\sigma_2)}} * \frac{e^{\delta_j/(1-\sigma_1)}}{e^{I_{gh}/(1-\sigma_1)}}, \tag{7}$$

where

$$I_{gh} = (1-\sigma_1)\ln\sum_{k=1}^{J_{gh}} e^{\frac{\delta_k}{1-\sigma_1}}, \quad I_g = (1-\sigma_2)\ln\sum_{h=1}^{H_g} e^{\frac{I_{gh}}{1-\sigma_2}}, \quad I = \ln(1+\sum_{g=1}^{G} e^{I_g}).$$

We can now derive aggregate demand, q_j: the probability that a consumer buys product j multiplied by the quantity purchased d_{ij}, summed over all consumers:

$$q_j = \sum_{i=1}^{I} s_j(\delta,\sigma)d_{ij} = \sum_{i=1}^{I} s_j(\delta,\sigma)\frac{\gamma y_i}{p_j} = s_j(\delta,\sigma)\frac{\gamma}{p_j}\sum_{i=1}^{I} y_i, \tag{8}$$

where a useful and often easily available measure for the total income of all consumers in the economy, $\sum_{i=1}^{I} y_i$, is the gross domestic product. I use GDP for an estimate of the total market size, including the outside good, in terms of value, rather than volume.

Next, I follow Berry (1994) and invert $s_j = s_j(\delta,\sigma)$ to solve for $\delta_j = \delta_j(s,\sigma)$:

$$\delta_j = \ln(s_j/s_0) - \sigma_1\ln(s_{j|gh}) - \sigma_2\ln(s_{h|g}), \tag{9}$$

where

$$s_{j|gh} = \frac{s_j}{\sum_{l=1}^{J_{gh}} s_l}$$

and

$$s_{h|g} = \frac{\sum_{l=1}^{J_{gh}} s_l}{\sum_{m=1}^{H_g}\sum_{l=1}^{J_{gh}} s_l}$$

are the respective probabilities of choosing product j given that an alternative from subgroup h of group g is chosen, and choosing subgroup h of group g given that group g is chosen. Substituting the aggregate demand equation (8) into the above expression, we can rewrite the unobserved choice probabilities in terms of observables:

$$s_j = \frac{p_j q_j}{\gamma\sum_{i=1}^{I} y_i}, \qquad s_{j|gh} = \frac{p_j q_j}{\sum_{l=1}^{J_{gh}} p_l q_l}, \qquad s_{h|g} = \frac{\sum_{l=1}^{J_{gh}} p_j q_j}{\sum_{m=1}^{H_g}\sum_{l=1}^{J_{gh}} p_l q_l},$$

Inserting these expressions together with the specification (5) for the mean utility back into the inverted choice probability (9), we derive the empirical estimation equation

$$\begin{gathered}\ln\frac{q_j p_j}{\gamma\sum_{i=1}^{I} y_i - \sum_{j=1}^{J} q_j p_j} = \\ x_j\beta - \alpha\ln p_j + \sigma_1\ln\frac{p_j q_j}{\sum_{l=1}^{J_{gh}} p_l q_l} + \sigma_2\ln\frac{\sum_{l=1}^{J_{gh}} p_j q_j}{\sum_{m=1}^{H_g}\sum_{l=1}^{J_{gh}} p_l q_l} + \xi_j,\end{gathered} \tag{10}$$

all elements of which are observed.

This specification provides some basic intuition and theoretical predictions for the sign and magnitude of marginal effects, from which we can derive a term for own- and cross-price elasticities. First, (8) gives us

$$\frac{\partial q_j}{\partial p_j} = \gamma \sum_{i=1}^{I} y_i \frac{\partial \left(\frac{s_j(\delta, \sigma)}{p_j} \right)}{\partial p_j} = \gamma \sum_{i=1}^{I} y_i \left(\frac{\partial s_j}{\partial p_j} \frac{1}{p_j} - \frac{s_j}{p_j^2} \right);$$

next, using (7) we can show that

$$\begin{gathered} \frac{\partial s_j}{\partial p_j} = \frac{\partial s_j}{\partial \delta_j} \frac{\partial \delta_j}{\partial p_j} = \\ -\frac{\alpha}{p_j} s_j \left(\frac{1}{1-\sigma_1} - \left(\frac{1}{1-\sigma_1} - \frac{1}{1-\sigma_2} \right) s_{j|gh} - \frac{\sigma_2}{1-\sigma_2} s_{j|g} - s_j \right). \end{gathered} \tag{11}$$

Using the two expressions above, we derive aggregate own-price elasticity to be

$$\frac{\partial q_j}{\partial p_j} \frac{p_j}{q_j} = -\alpha \left(\frac{1}{1-\sigma_1} - \left(\frac{1}{1-\sigma_1} - \frac{1}{1-\sigma_2} \right) s_{j|gh} - \frac{\sigma_2}{1-\sigma_2} s_{j|g} - s_j \right) - 1. \tag{12}$$

Cross-price elasticities are obtained in a similar fashion. Generally, in the two-level nested logit elasticity of demand for product j with respect to the price of product k takes the form of

$$\begin{gathered} \frac{\partial q_j}{\partial p_k} \frac{p_k}{q_j} = \\ -\alpha \left(\frac{1}{1-\sigma_1} D_1 - \left(\frac{1}{1-\sigma_1} - \frac{1}{1-\sigma_2} \right) s_{j|gh} D_2 - \frac{\sigma_2}{1-\sigma_2} s_{j|g} D_3 - s_j \right) - D_1 \end{gathered} \tag{13}$$

with D_1, D_2, and D_3 are binary indicators where $D_1 = 1$ if j and k are the same product, $D_2 = 1$ if j and k are in the same subgroup, and $D_3 = 1$ if j and k are in same group.

Identification comes from the assumption that the matrix of product characteristics is exogenous. Based on this assumption, I follow Berry et al. (1995) and use as instruments the products' own characteristics, and counts and sums of the other products' characteristics, overall and of the same firm, to account for endogenous prices and shares.[11] I also implement the suggestions of Verboven (1996) for the nested logit model, using counts and sums by groups and subgroups as additional instruments. Finally, I include potential cost-shifters, in the form of exogenous macroeconomic variables (such as the central bank and risk-free interest rates, GDP, and inflation), their squares and interactions, yielded comparable results.

[11] These include capital flows, mean and cumulative returns, volatility (standard deviation of past return), sharpe ratios, current and excess returns, and main and secondary specializations.

4.2 Fund managers

Next, a structure of firm behavior is imposed to get an estimate of marginal costs. Further following the notation of Berry et al. (1995), Nevo (2000), and Björnerstedt and Verboven (2015), let F_f denote the set of products produced by firm f. I assume firms maximize profits in a Nash-Bertrand oligopolistic competition setting. The total profits Π of each firm f, as a function of price vector p, are given by the sum of profits for each product $k \in F_f$ the firm produces:

$$\Pi_f(p) = \sum_{k \in F_f} (p_k - c_k) q_k(p) \tag{14}$$

Where p_k is the price of product k, c_k is the marginal cost in its production, and q_k is the demand for product k. The profit-maximizing price vector p must satisfy the standard first-order condition:

$$q_j(p) + \sum_{k \in F_f} (p_k - c_k) \frac{\partial q_k(p)}{\partial p_j} = 0 \tag{15}$$

To rewrite this system of J first-order conditions, let us define the $J \times J$ ownership matrix θ^F, where element $\theta^F(j,k)$ equals to one if products j and k are produced by the same firm, and zero otherwise. Let $q(p)$ be the $J \times 1$ demand vector, $\Delta(p) \equiv \partial q(p)/\partial p'$ is the $J \times J$ Jacobian matrix of first derivatives, and c is the $J \times 1$ marginal cost vector. The $\odot$ notation denotes element-by-element matrix multiplication.

$$q(p) + (\theta^F \odot \Delta(p))(p - c) = 0$$

can be inverted to get

$$p = c - (\theta^F \odot \Delta(p))^{-1} q(p). \tag{16}$$

Now, equation (16) can first be used to uncover pre-merger marginal cost vector c:

$$c^{pre} = p^{pre} + (\theta^{F,pre} \odot \Delta(p^{pre}))^{-1} q(p^{pre}), \tag{17}$$

and then to predict post-merger equilibrium prices by:

$$p^{post} = c^{pre} - (\theta^{F,post} \odot \Delta(p^{post}))^{-1} q(p^{post}). \tag{18}$$

This framework also allows for efficiency gains, as a result of a merger, by replacing the marginal costs vector c^{pre} with c^{post}. That flexibility is useful in the Meitav-Dash merger, as efficiency gains are likely present. Similarly, for purposes of ex post evaluation, the model parameters can be calibrated to best fit the observed outcomes and to account for external institutional details.

4.3 Specifications

To successfully take the model to the data, I must address and specify several underlying assumptions, starting with the potential market. In the most common case in the literature - the unit demand model - this would represent the total number of potential consumers or savers, I. In a constant expenditure model, on the other hand, the potential market would refer to the total potential budget available for allocation to savings in provident funds. Though I opt for the constant expenditure approach, I argue that markets for financial services may bring the two specifications closer together. For the unit demand specification, instead of assuming the potential number of individuals in the economy, and consequently computing the average quantities consumed, one might consider each unit of currency as an end consumer. Instead of decision-making individuals, consider decision-making shekels. This standardization in terms of ILS may allow for a more direct definition, with less assumptions imposed: the researcher no longer needs to assume both the number of potential consumers in the economy and that individuals choose one service provider, at most, for the unit demand specification. Instead, only the amount of capital savers could potentially allocate to provident funds (as in the constant expenditure specification) need be assumed.

Nests, the pivotal feature of the demand model, are assumed to be firms and main specialization, with groups defined at the firm level ('brand'), and subgroups at the specialization level. This specification seems highly plausible and is consistent with how likely are individuals to substitute: firms approach savers wishing to leave for a competitor, lowering their sensitivity to price through marketing efforts, which does not happen when savers move between specializations. I estimated models with the reverse nesting order (where savers substitute more to a different firm than to another specialization), one-level nesting and combinations with secondary specialization. Some were found to be inconsistent with random utility maximization, with nesting parameters greater than one, or $\sigma_1 < \sigma_2$, others produced very similar simulation results (likely due to the importance of firm nests, in a form of brand loyalty). The importance of firm-level nests is stressed in an OLS regression of assets under management on the number of funds by competitors: the introduction of new products (funds) by competitors is shown to be uncorrelated with assets under management, suggesting weak patterns of substitution between players (Figure 4). Fund size does not, contrary to what we would expect, negatively correspond with the opening of new funds by competitors. In other words, it implies that savers rarely substitute and that when new funds are introduced, they likely cannibalize on own assets or attract new savers to the market, rather than those of competitors.

By the end of 2014, total value of assets under management of provident funds was on the order of 220 billion ILS, while institutional investors as a whole managed roughly 780 billion ILS. For convenience, I assume the potential market size to be in terms of GDP, which is about 1 trillion ILS, as of 2014. I test the robustness of my results to this assumption by setting the potential market value to be different multiples of GDP (0.5, 1, 1.5, and 2) and receive

similar results.

Unlike in classic product markets, defining units of consumption in service markets often requires special attention. A useful way to think of this is to define products $1 \dots J$ at the fund-month level as if funds were single-product firms. Each shekel's allocation must be determined every month anew. The observed heterogeneity in the provident fund market comes through variation in fund characteristics, each offering a single product. I account for time-invariant unobserved product characteristics using fund-level fixed effects.

5 Results

5.1 Demand estimation

I explore the constant expenditure specification of the two-level nested logit described above, several one-level nesting models to test how susceptible my results are to the imposed demand structure,[12] and also compare them to the un-nested conditional logit estimates. Estimated demand parameters are reported in Table 3, most with the expected sign. Exceptions were negative effects for 5-year mean returns (though positive for the 3-year) and excess returns, which we would expect to positively affect demand, and a positive effect for 3-year performance volatility (standard deviation of returns), that goes against intuition that savers are risk averse.

High values of the nesting parameters are estimated across the models, with 0.91 and 0.85 for σ_1 and σ_2, respectively, in the preferred two-level nesting specification (firm-specialization). A high value of σ_1 indicates that consumer preferences are strongly correlated across products within a subgroup, and when σ_2 is positive, consumer preferences show additional correlation across products of the same group. If $\sigma_1 = \sigma_2$, preferences do not show correlation across subgroups within a group and the model is reduced to a one-level nested logit, with groups as the nests. If $\sigma_2 = 0$, the model is again reduced to single-level nesting, but with subgroups as the nests. Finally, if $\sigma_1 = \sigma_2 = 0$, consumer preferences do not show any correlation across subgroups nor groups and the model is reduced to a non-nested conditional logit.

Results from the different models lie in relatively close proximity with each other, in terms of estimated demand parameters, all suggesting high correlation and rather weak substitution patterns across groups and subgroups. A possible interpretation of this consistency is again the dominance of non-specialized funds together with the pro-activity required to switch manager. Alternatively, this could be the successful product of marketing efforts by the managing firms, creating brand loyalty, coupled with the relative ignorance or myopia of the general public in this field.

[12] One-level nested logit models include main specialization, secondary specialization and firm nests. Two-level specifications are [main specialization, secondary specialization], [main specialization, firm] and [firm, main specialization] as the group and subgroup nests.

My estimates comply with the underlying theoretical nested model of aggregated individual demand. The coefficient of log price on utility, α, is estimated with the expected sign in all specifications, and is around 0.06 for the preferred model, in absolute terms. The subgroup and group nesting parameters (σ_1 and σ_2, respectively) satisfy the requirements for the model to be consistent with random utility theory: $0 \leq \sigma_2 \leq \sigma_1 \leq 1$. Strict inequalities imply that consumers do perceive funds of the same subgroup as the closest substitutes, funds of a different subgroup but within a group as weaker substitutes, and products of a different group as the weakest substitutes.[13]

While I explore both single- and two-level nesting specification for robustness, I shall focus on and discuss the preferred model: two-level nested logit, with groups at the firm level and subgroups at the main specialization level. Implied own- and cross-price elasticities are derived according to equation (13) and reported in Table 4. I estimate the elasticity of demand with respect to own price to be, on average, 1.538 in absolute terms, with a range of 1.1-1.62. Consistent with economic theory and previous intuition, the cross price elasticities are higher for funds within the same subgroup (0.078) than for funds of a different subgroup but of the same group (0.019) and especially than for funds of a different group (0.003). Results are robust to the inclusion of additional instrumental variables, and are unlikely to be driven by the chosen specification. These estimates also correspond with previous work in closely related fields: Hastings and Tejeda-Ashton (2008) study how financial literacy impacts savers' choice in Mexicos privatized social security system and estimate own-price demand elasticities in the range of -1.45 to -1.86; Agostini et al. (2014) analyze a merger between two medium-sized pension fund managers in Chile and estimated own-price elasticities in the range of -1.68 to -2.36, and cross-price elasticities between 0.08 and 0.77, using a proportionally calibrated almost ideal demand system (PCAIDS); and Hastings et al. (2013) return to the Mexican privatized pension market and estimate own-price elasticities between -0.75 and -1.93, highlighting how price sensitivity of savers can be lowered, leading to inelastic demand and high equilibrium fees.

Though the preferred model was consistent with random utility theory, alternative specifications showed some discrepancies. Certain unreported models, especially those incorporating secondary fund specialization, produced inelastic estimates of demand with respect to own price. These findings are seemingly indicative of internal inconsistencies within the model and require some attention. Inelastic estimates of demand are a cause for some concern since they imply that firms do not behave optimally and do not exert their market power to the fullest. However, such behavior could be explained by compliance with regulation directives on the one hand, or by fear of retribution by savers on the other. Both are plausible and highlight some weaknesses of the static Nash-Bertrand competition assumption on the supply side. Undoubtedly, firms do not always set prices by solving for the static equation (14) at each period sepa-

[13] Under all two-level specifications, the difference between σ_1 and σ_2 was found to be statistically significant, with a p-value under 0.001.

rately. Rather, they are likely to have a dynamic profit function in mind, taking into account adaptive expectations and simultaneous responses of other players in the market, including regulators. Indeed, hands-on intervention on the side of central government is not unprecedented in the provident fund market, with the 2005 structural reform following the Bachar Committee in recent history.

5.2 Merger simulation

Using the results detailed in the previous section, the theoretical framework allows me to generate out-of-sample predictions regarding market equilibrium, following changes such as firm entry, divestiture and merger. I exploit this setting to estimate the ex ante predicted price effects from the merger of Meitav and Dash, using only data available up to that time: with no closed-form analytical solution, I simulate the new price equilibrium using fixed-point iteration on equation (17). Predicted merger price effects are derived from equation (18) and illustrated in Figure 7. I estimate the increase in market power of the merging parties allows Meitav-Dash to raise prices by 1.6% for funds previously managed by Dash and by 4.2% for funds previously managed by Meitav, in the new equilibrium. Competitors are predicted to have no significant response to the merger in their pricing.[14]

I test the robustness of my results to different assumptions and choices made. In addition to assumptions made in the process of demand estimation, the simulation itself introduces two more: by not altering the vector of marginal costs ($c^{post} = c^{pre}$), I assumed there are no efficiency gains due to the merger; and by not allowing firms to take competitors' profits into account, I assumed there is no cooperation in the market.

In a set of additional results, I explore my model's sensitivity to different values of the above properties. While partial coordination between competitors in a relatively concentrated market is plausible at the very least, it is hard to provide quantifiable evidence in its support. Efficiency gains, on the other hand, are explicitly present, judging by the previously mentioned layoffs that ensued. The precise extent of those, however, is not publicly available, and perhaps not even known to the merging parties themselves. Additional intuition to support their existence is provided by a pooled difference-in-differences regression of log price on dummy indicators for the merging parties (MD_j), post-merger period, and an interaction term, using a two-year window around the merger:

$$\ln price_{jt} = \alpha + \beta_1 MD_j + \beta_2 PostMerger_t + \beta_3 MD_j * PostMerger_t + \epsilon_{jt}.$$

Estimates are reported in Table 5, suggesting that while prices overall were about 13% lower in the years and months following the merger, fees of funds by Meitav-Dash had dropped by an additional 14.6%, as expected if efficiency gains were present.

[14] As an execise, I also simulate the most anti-competitive merger possible in the market: Clal and Psagot. Results suggest price increases of 14.4% and 6.2% for funds managed by Clal and Psagot, respectively, with minor price increases by competitors.

With the lack of external data, I resort to studying select values and calibrate the model to correspond with post-merger data. I run my simulation for several possible values, deviating from the baseline at 0 (no efficiency gains) up to a 25 percent reduction in marginal costs. For simplicity, I assume the same percentage for both Meitav and Dash, though it does not have to be symmetric. This exercise allows me to estimate cost synergy by calibrating the model such that its predictions align with observed outcomes.

Efficiency gains are found to be transferred almost entirely to savers through reduced fees, while I measure no additional effect for competitors. In the preferred specification of a 15% reduction in marginal costs (Figure 8), I measure predicted price decreases of around 10%-12%, for funds belonging to Meitav-Dash, with insignificant responses from competitors. These predictions are comparable to the observed 14.6% drop in fees. It could be argued that in the ongoing market dynamics, the increased efficiency following consolidation generated downward pressure on fees, or at the very least contributed to it. These synergy estimates, resulting from economies of scale, are the subsequent pass-through rate, are a main contribution of this paper to the literature.

6 Discussion and conclusion

This work set out to study the effects of a consummated merger in the Israeli provident funds market between Meitav and Dash in 2012 through merger simulation. Following the exercise of Björnerstedt and Verboven (2015), I estimate a two-level nested logit structure of aggregate consumer demand for differentiated products using comprehensive fund-level data gathered from the Ministry of Finance. I impose their suggested constant expenditure specification, instead of the more common unit demand, arguing for its merit in this setting. I estimate weak substitution patterns between groups (defined at the firm, or 'brand' level) and subgroups (defined at the specialization level) that are consistent with industry norms of default specialization and delegated union bargaining. My results also correspond to existing literature on switching costs and inertia, and are suggested by patterns observed in the data.

Using parameters from demand estimation, I simulate the merger and measure moderate price effects for the merging parties (an increase on the scale of 2% - 4%) with minor effects for their competitors. I support the findings and argue for the existence of efficiency gains with basic difference-in-differences estimates, which suggest that while prices after the merger decreased on average throughout the market, the merging parties lowered fees by a greater extent than their competitors.

My findings provide evidence to suggest that Israeli provident funds operate in a setting with weak substitution patterns. The implicit preference of the general public for non-specialized funds, is illustrated in Figure 3 and captured by the low cross-firm and cross-specialization price elasticities in table 3. This behavior is not intuitive: we would expect greater variation in consumer preferences, that corresponds with the age distribution of population. Young savers

with a longer investment horizon should favor greater risk taking (e.g. selecting stock-specialized funds) than savers on the verge of retirement, who would likely emphasize sovereign bonds. The observed trends are consistent with the existence of inertia and non-negligible switching costs, and could possibly be better modeled using a random coefficients logit approach. That, however, is beyond the scope of this work, which is a first attempt to study the industry using structural analysis, and utilizes a more restricted model.

Competition policy was one of the main motivations for this paper. In that perspective, the observed consumer behavior creates incentives for fund managers to compete over markets, rather than within them: through delegated negotiations with unions and lucrative initial terms, acquiring employees of whole firms as clients would often guarantee a pool of captive savers that will not easily switch to a competitor.[15] Mergers, on the other hand, seem to generate efficiency gains and correspond with lower prices in markets such as this one, where economies of scale are present. It could therefore be argued that, to maximize total welfare, competition authorities should focus on lowering barriers to entry and switching costs, rather than blocking or limiting efficient mergers.

[15] Switching costs and rates between funds in Israel were not thoroughly studied are not the focus of this paper. However, the implied mechanism is consistent with recent literature, such as Hastings et al. (2013) and Illanes (2016), that suggests these markets exhibit significant price dispersion and very low switching rates.

References

Agostini, C., Saavedra, E., and Willington, M. (2014). Economies of Scale and Merger Efficiencies in the Chilean Pension Funds Market. *Journal of Competition Law and Economics*, 10(1):137–159.

Allen, J., Clark, R., and Houde, J.-F. (2014). The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry. *American Economic Review*, 104(10):3365–96.

Amel, D., Barnes, C., Panetta, F., and Salleo, C. (2004). Consolidation and efficiency in the financial sector: A review of the international evidence. *Journal of Banking & Finance*, 28(10):2493–2519.

Ashenfelter, O. and Hosken, D. (2010). The effect of mergers on consumer prices: Evidence from five mergers on the enforcement margin. *Journal of Law and Economics*, 53(3):417–466.

Ashenfelter, O. C., Hosken, D. S., and Weinberg, M. C. (2013). The Price Effects of a Large Merger of Manufacturers: A Case Study of Maytag-Whirlpool. *American Economic Journal: Economic Policy*, 5(1):239–61.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890.

Berry, S. and Pakes, A. (1993). Some applications and limitations of recent advances in empirical industrial organization: Merger analysis. *The American Economic Review*, pages 247–252.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262.

Björnerstedt, J. and Verboven, F. (2015). Does Merger Simulation Work? Evidence from the Swedish Analgesics Market. *American Economic Journal: Applied Economics*, forthcoming.

Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American Economic Review*, pages 297–308.

Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. *Handbook of industrial organization*, 3:1967–2072.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *The American Economic Review*, 103(7):2643–2682.

Hastings, J. S. (2004). Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California. *American Economic Review*, 94(1):317–328.

Hastings, J. S., Hortasu, A., and Syverson, C. (2013). Sales Force and Competition in Financial Product Markets: The Case Of Mexico's Social Security Privatization. Working Paper 18881, National Bureau of Economic Research.

Hastings, J. S. and Tejeda-Ashton, L. (2008). Financial literacy, information, and demand elasticity: survey and experimental evidence from Mexico. Working Paper 14538, National Bureau of Economic Research.

Hausman, J., Leonard, G., and Zona, J. D. (1994). Competitive analysis with differenciated products. *Annales d'Economie et de Statistique*, pages 159–180.

Illanes, G. (2016). Switching costs in pension plan choice. MIT job market paper, 2016.

Jayaraman, N., Khorana, A., and Nelling, E. (2002). An Analysis of the Determinants and Shareholder Wealth Effects of Mutual Fund Mergers. *The Journal of Finance*, 57(3):1521–1551.

Kim, E. H. and Singal, V. (1993). Mergers and Market Power: Evidence from the Airline Industry. *The American Economic Review*, 83(3):549–569.

Madrian, B. C. and Shea, D. F. (2001). The Power of Suggestion: Inertia in 401 (k) Participation and Savings Behavior. *The Quarterly Journal of Economics*, 116(4):1149–1187.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, pages 105–142.

Nevo, A. (2000). A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand. *Journal of Economics & Management Strategy*, 9(4):513–548.

Spence, M. (1976). Product selection, fixed costs, and monopolistic competition. *The Review of Economic Studies*, pages 217–235.

Trajtenberg, M. (1989). The welfare analysis of product innovations, with an application to computed tomography scanners. *The journal of political economy*, pages 444–479.

Verboven, F. (1996). International price discrimination in the European car market. *Rand journal of economics*, 27(2):240–268.

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max
fee (stock)	0.94	0.42	0.01	6.57
fee (flow)	1.35	0.85	0.02	3.68
return	0.54	2.12	-22.22	85.31
alpha (excess return)	0.09	1.63	-7.44	11.88
return (3 years mean)	6.45	4.60	-16.71	48.36
return (5 years mean)	6.74	3.40	-8.10	28.86
return (3 years std.)	1.77	1.19	0.00	14.38
return (5 years std.)	1.85	1.10	0.07	12.20
liquidity ratio	88.78	7.77	25.80	100.00
Sharpe ratio	-0.07	0.64	-2.08	2.55
deposits	1.36	7.18	0.00	376.08
AUM	339.41	1385.46	0.00	28981.25
revenue	281.23	1076.96	0.00	19896.46
price (revenue / AUM)	0.08	0.03	0.00	0.53
GDP	760.34	174.37	481.35	1088.50
inflation	0.17	0.46	-0.86	1.54
BOI interest rate	3.97	2.80	0.25	13.50
risk-free interest rate	4.17	2.59	0.18	12.25

Note: 73,066 monthly observations of 781 provident funds by 18 firms. Fees, return, alpha, inflation, risk-free rate, and BOI interest rate are in annual percentage points. Liquidity ratio is defined as liquid assets over short-term liabilities and used as a proxy for a fund's short-term ability to pay its retired savers their benefits. Sharpe ratio is the industry standard measure for risk adjusted returns GDP is in billion current ILS. Deposits and AUM are in million current ILS. Revenue is in thousand current ILS. Price is in monthly percentage points.

Table 2: Provident fund asset-weighted market shares, 2011, by firm and specialization

	Bonds	General	Designated	Indexed	Stocks	Nominal	Other	Total
Psagot	0.91	23.33	0.00	0.50	0.26	0.67	0.05	25.72
Clal	0.58	16.65	1.44	0.98	0.10	0.60	0.02	20.38
Harel	0.01	10.90	0.00	0.05	0.06	0.39	0.01	11.41
Dash	0.61	7.60	0.06	0.17	0.05	0.10	0.05	8.64
Menora	0.02	3.01	2.98	0.00	0.01	0.07	0.00	6.09
Altshuler-Shaham	0.64	3.41	0.00	0.00	0.05	0.01	0.00	4.11
Halman-Aldubi	0.05	4.35	0.00	0.18	0.04	0.15	0.00	4.77
Yelin-Lapidot	0.33	1.65	0.00	0.00	0.15	0.00	0.00	2.13
Delek (excellence/phoenix)	0.07	6.45	0.48	0.27	0.11	0.22	0.01	7.60
Eliyahu (Migdal)	0.05	1.03	0.00	0.03	0.01	0.03	0.00	1.15
Meitav	0.42	2.76	0.00	0.07	0.06	0.02	0.00	3.34
Other	0.84	3.61	0.00	0.00	0.18	0.04	0.00	4.67
Total	4.54	84.73	4.96	2.25	1.08	2.30	0.13	100.00

Note: Weighted by assets under management, percent of market total. Other specializations include foreign currency, Jewish/Islamic law and unclassified. Other firms include IBI, Infinity, Analyst and Ayalon.

Table 3: Nested logit results, constant expenditure

	(1)	(2)	(3)	(4)	(5)
α	-0.956***	-0.074***	-0.055***	-0.074***	-0.058***
	(0.068)	(0.011)	(0.013)	(0.011)	(0.014)
σ (single nest)		0.971***	0.894***		
		(0.006)	(0.007)		
σ_1 (subgroup)				0.987***	0.905***
				(0.007)	(0.007)
σ_2 (group)				0.909***	0.853***
				(0.010)	(0.011)
deposits	0.011***	0.005***	0.006***	0.005***	0.006***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
return (3y mean)	0.670	1.100*	1.319*	1.164*	1.351*
	(1.422)	(0.471)	(0.534)	(0.469)	(0.534)
return (5y mean)	-3.998**	-0.175	-1.945***	-0.243	-1.918***
	(1.476)	(0.495)	(0.561)	(0.493)	(0.561)
return (3y s.d.)	0.268***	0.194***	0.149***	0.190***	0.148***
	(0.013)	(0.007)	(0.007)	(0.006)	(0.007)
return (5y s.d.)	-0.032	-0.117***	0.023	-0.117***	0.017
	(0.020)	(0.022)	(0.025)	(0.022)	(0.025)
excess return	-0.008	-0.030***	-0.049***	-0.032***	-0.049***
	(0.006)	(0.003)	(0.004)	(0.003)	(0.004)
sharpe ratio	-0.105	0.077*	0.036	0.084*	0.043
	(0.059)	(0.039)	(0.044)	(0.039)	(0.044)
liquidity ratio	-0.017***	0.002*	0.003**	0.002**	0.003**
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Upper nest	N/A	Main spec.	Firm	Main spec.	Firm
Lower nest	N/A	N/A	N/A	Firm	Main spec.
R^2	0.02	0.15	0.73	0.15	0.72

b coefficients; se in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: All specifications include observed covariates that enter individual utility (past fund performance, volatility, excess return, fund size and capital flows) in a fund-level fixed-effects regression. Endogenous variables such as price and shares are instrumented according to Berry et al. (1995) and Verboven (1996).

Table 4: Pre-merger implied price elasticities (October 2012), CES

	Mean	Std. Err.	Min	Max
Own-price elasticity	-1.538	0.095	-1.616	-1.100
Cross: same subgroup	0.078	0.095	0.000	0.515
Cross: same group	0.019	0.041	0.000	0.312
Cross: different group	0.000	0.000	0.000	0.003

Note: Estimated elasticities from the preferred nesting specification of nests and subnests at the firm and specialization levels, respectively.

Table 5: Difference-in-differences

MD	-0.030
	(0.019)
PostMerger	-0.130***
	(0.011)
MD * PostMerger	-0.146***
	(0.018)
Constant	-7.277***
	(0.008)
Observations	6725

b coefficients; se in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Results from a difference-in-differences regression of log price on merging firm (MD) and post-merger dummies, and an interaction term, using a two-year window around the merger.

Figure 1: Total AUM

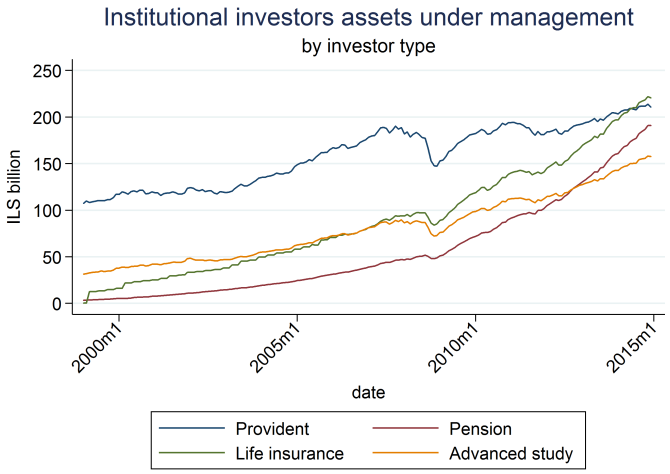

Note: Total assets under management, 1999-2014. Life insurance, provident and pension funds are long-term saving plans for retirement. Advanced study funds are aimed at a shorter horizon and can be withdrawn, without losing tax benefits, after 6 years.

Figure 2: Provident funds' AUM shares, by target population

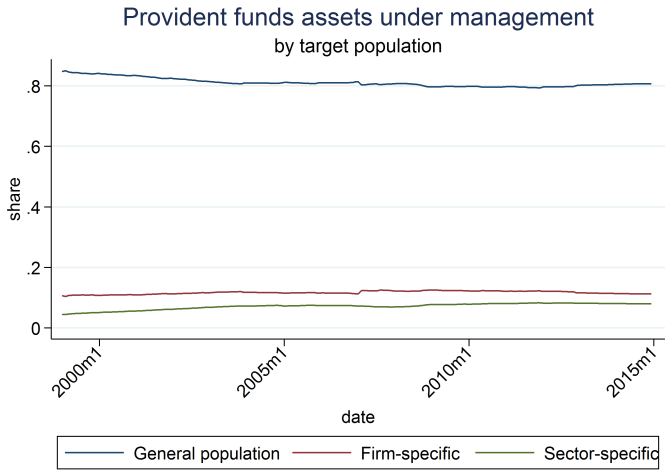

Note: Share of assets under management, by target population, 1999-2014. Sector- or firm-specific funds, that constitute some 20% of the market, are unavailable for the general public and were omitted from the analysis.

Figure 3: Provident funds' AUM shares, by specialization

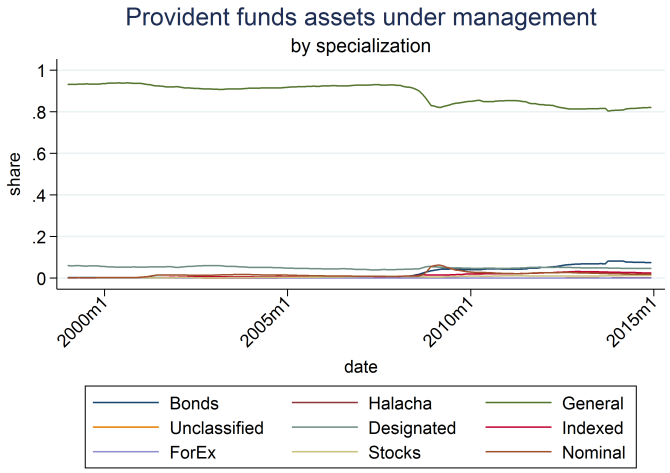

Note: Share of assets under management, by fund specialization, 1999-2014. Sector- or firm-specific funds excluded. General, non-specialized funds, which are the default, comprise some 80% of assets, down from 90%+ levels before the 2008 crisis.

Figure 4: Log AUM on number of funds by competitors

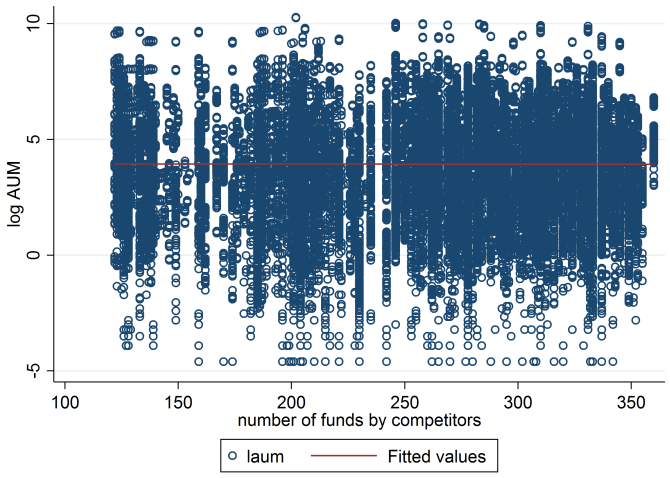

Note: Log assets under management and number of competing provident funds, scatter and linear fit ($logAUM = 3.953 - .0000987 * numfunds$). Sector- or firm-specific funds excluded. Slope of linear fit is not statistically different than zero (p-value of 0.55).

Figure 5: Mean fees, 1999-2014

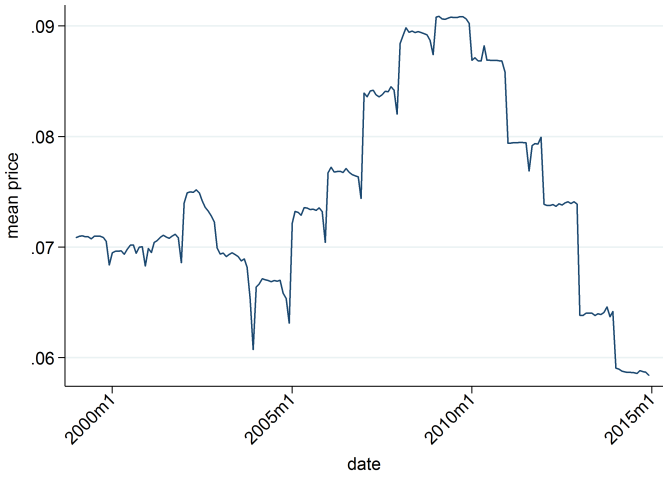

Note: Mean monthly fee, percentage points. A downward trend has begun in 2010, two years before the Meitav-Dash merger. Fee data were reported as annual aggregates; within-year variation is a product of fund opening and closing. The observed pattern is detected in practically all firms and specializations.

Figure 6: Provident fund Herfindahl-Hirschman Index

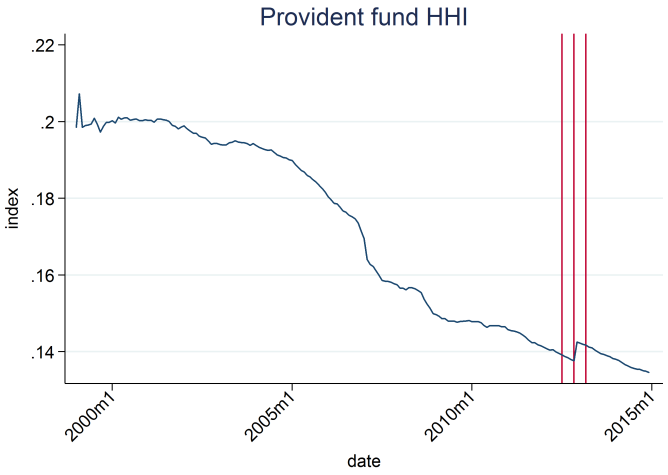

Note: Herfindahl-Hirschman Index, 1999-2014. Values in the range of 0.1-0.18 are considered moderately concentrated by DoJ guidelines. Index is computed as the sum of squared marked shares. Vertical lines are Merger request application, clearance, and completion: July 2012, November 2012, and March 2013, respectively.

Figure 7: Predicted price changes, no efficiency gains

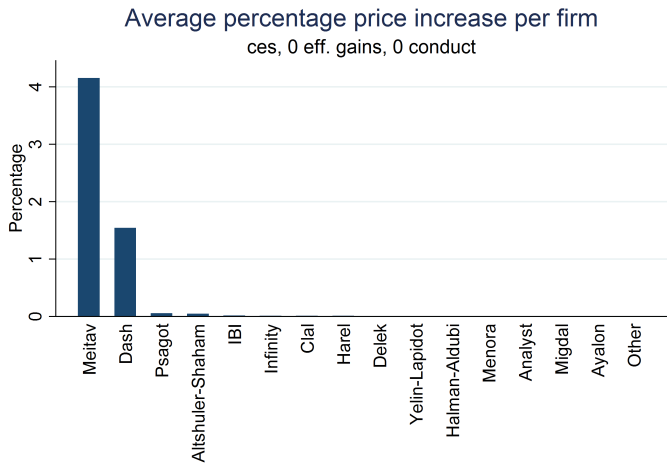

Note: Simulation results of a two-level constant expenditure nested logit model with nests and subnests at the firm and specialization levels, respectively. No efficiency gains or coordination assumed.

Figure 8: Predicted price changes, 15% efficiency gains

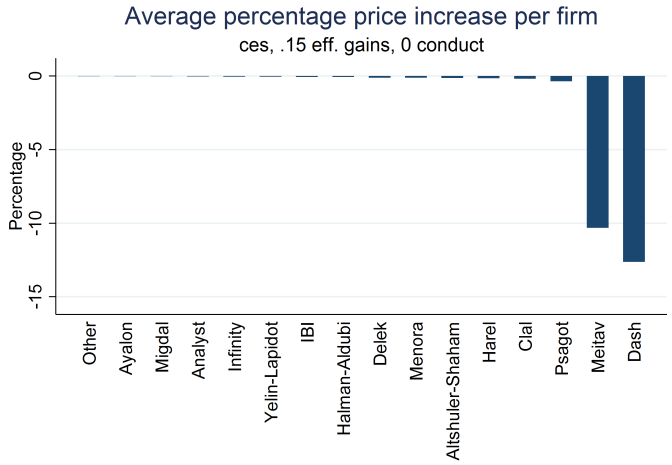

Note: Simulation results of a two-level constant expenditure nested logit with nests and subnests at the firm and specialization levels, respectively. 15% efficiency gains for both merging parties and no coordination assumed.