The effect of market consolidation on innovation in the HDD industry

Anna Rita Bennato
Loughborough University

Stephen Davies
Centre for Competition Policy
School of Economics
University of East Anglia

Franco Mariuzzo
Centre for Competition Policy
School of Economics
University of East Anglia

Peter Ormosi
Centre for Competition Policy
Norwich Business School
University of East Anglia

CCP Working Paper 18-4

Abstract

There is an extensive literature on the relationship between market competition and innovation. We contribute to this literature by taking an exhaustive, industry-focused look at how innovation evolved following the 5-to-3 consolidation of the world-wide hard disk drive (HDD) industry. Instead of using a single measure of innovation, the HDD industry offers a rare opportunity to look at possible changes in R&D expenditure, patent activity, the number of new products, and the unit cost of new products at the same time. This is important for two reasons: (1) it allows us a more informed evaluation of how innovation changed as a result of increasing market concentration, and (2) it enables us to test the relative performance of data on R&D spending and patent activity in measuring innovation. For the former, we find no evidence that the consolidation had a negative impact on innovation. For the latter, we provide evidence that R&D spending is a good predictor of unit costs and the number of new products, but patent activity offers little extra explanatory power.

Contact Details:
Peter Ormosi p.ormosi@uea.ac.uk

ISSN 1745-9648

The effect of market consolidation on innovation in the HDD industry*

Anna Rita Bennato[†], Stephen Davies[‡], Franco Mariuzzo[§], and Peter Ormosi[¶]

May 14, 2018

*This paper is partially based on a report prepared by the authors for the European Commission, titled "A feasibility study on the microeconomic impact of enforcement of competition policies on innovation". The views and opinions expressed in this report are strictly those of the authors.

[†]Loughborough University, Epinal Way, Loughborough LE11 3TU, United Kingdom, email: a.bennato@lboro.ac.uk

[‡]School of Economics and Centre for Competition Policy, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom, email: f.mariuzzo@uea.ac.uk

[§]School of Economics and Centre for Competition Policy, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom, email: f.mariuzzo@uea.ac.uk

[¶]Corresponding author, Norwich Business School and Centre for Competition Policy, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom, email: p.ormosi@uea.ac.uk

Abstract

There is an extensive literature on the relationship between market competition and innovation. We contribute to this literature by taking an exhaustive, industry-focused look at how innovation evolved following the 5-to-3 consolidation of the world-wide hard disk drive (HDD) industry. Instead of using a single measure of innovation, the HDD industry offers a rare opportunity to look at possible changes in R&D expenditure, patent activity, the number of new products, and the unit cost of new products at the same time. This is important for two reasons: (1) it allows us a more informed evaluation of how innovation changed as a result of increasing market concentration, and (2) it enables us to test the relative performance of data on R&D spending and patent activity in measuring innovation. For the former, we find no evidence that the consolidation had a negative impact on innovation. For the latter, we provide evidence that R&D spending is a good predictor of unit costs and the number of new products, but patent activity offers little extra explanatory power.

Keywords: mergers, innovation, R&D, patents, evaluation
JEL Classification codes: O30, L10, L40

1 Introduction

Following seminal contributions from two of the giants of 20th century economics, Schumpeter and Arrow, the relationship between competition and innovation has long been hotly debated. There is now considerable amount of literature on measuring how competition affects innovation. This includes a number of studies on the effect of market consolidation on innovation. Remarkably however, only a few of these looked at specific markets, and most have provided aggregate and sometimes rough evidence summarising the average effect in large samples of markets.

In this paper we take a detailed look at how the consolidation of the hard disk drive (HDD) market affected innovation in HDD. Our market specific focus allows us to fully identify the innovation effect of the changing level of competition. We assembled a rich set of data to approximate Schumpeter's innovation trichotomy and measure innovation in its entirety, as opposed to looking only at its component parts in isolation. First we examine how market consolidation affected HDD manufacturers' willingness to invest in R&D. Next we look at the effect of consolidation on the patenting activity of these businesses directly, and through varying R&D investments. Finally we examine how market consolidation, and the level of R&D spending and patent activity drive simple product characteristics. Implicit in this approach is that it brings us closer to Schumpeter's hypotheses about invention and innovation, and their respective and mutual relationship with technological change.

R&D spending and patent activity are widely accepted measures of innovation used in the literature. But are they equally good approximators of innovation and technological improvement? Through our holistic approach we are able to make important contributions to the innovation research literature in general, most importantly by offering evidence on whether R&D expenditure or patent measures are more likely to correlate with measures of innovation such as the number of new products, and the unit cost of new products.

The paper also contributes to a large body of literature evaluating the impact of mergers. Instead of looking at the price effect of mergers we turn our focus to innovation, something that has been left largely untouched in retrospective studies of specific mergers. The findings of this case study prove to be interesting in their own right – shedding some new light on these important mergers. But far more importantly the paper establishes that industry specific and innovation focused ex-post evaluations are viable for policy purposes, while underlining some of the conceptual and methodological challenges. The ex-post evaluation of the innovation impact of mergers has probably never been more timely, when there appears to be a paradigm shift in the European Commission on how merger-related innovations are treated.

To headline our key results, we find no evidence that the 2011/12 consolidation of the HDD market reduced the level of innovation. This is valuable evidence given widespread claims that the European Commission's theory of merger harm is sympathetic to the argument that market consolidation always reduces incentives to innovate. For one of the HDD manufacturers, Seagate, we found that the 2011/12 events had a positive effect on the company's R&D spending, its patent activity, number of new products marketed. We also provide evidence that - at least in this specific market - R&D spending data is a better predictor of the number of new products and of unit cost than patent data, which could have some implications on how future studies are conducted. These findings are robust to a large number of empirical models, research designs, and model specifications.

The paper is structured as follows. We commence with a brief survey of literature, followed by an introduction of the HDD market. Section 3 discusses our study design for analysing each of our four datasets (R&D, patents, number of new products, and unit costs) with a particular focus on finding an adequate Control group. Section 4 delivers the headline results, followed by a detailed discussion of these findings. Throughout this study we have conducted a large number of econometric tests and sensitivity checks. A large number of these are reported in our online appendix, and in Ormosi, Bennato, Davies, and Mariuzzo (2017).

1.1 Literature review

The literature on innovation is immense. Empirically, the relationship between innovation and competition is one of the most (if not the most) researched question in industrial organisation. Here we do not offer a full coverage of this literature – there are many other excellent existing surveys - instead we focus on the developments that help us set up and motivate our own work.

On the theoretical side, growing out of Joseph Schumpeter's two works (1934, 1942) is the assertion that large firms are better placed to invest in innovation, and competition might not be the best platform to boost innovation. Since then a wide range of papers tried to challenge or find support for Schumpeter's proposition. While a unified consensus is difficult to draw, these works are invaluable for identifying the different conditions that influence the relationship between market structure and innovation. Crucial is the level of intellectual property protection – whether the innovator enjoys exclusivity on its innovation (see for example Levin, Klevorick, Nelson, Winter, Gilbert, and Griliches (1987), Hall and Ziedonis (2001)), the role of technical uncertainty (Reinganum (1989)), the level of competition in innovation, i.e. for the market (Gilbert and Newbery (1982)), information asymmetries between owners and managers (e.g. Schmidt (1997), and Aghion, Dewatripont, and Rey (1999)), or firm characteristics (Boone (2000), Gompers, Lerner, and Scharfstein (2005)). Many studies provide exhaustive reviews, such as Gilbert (2006).

Empirically the picture is equally rich. The earliest works looked at the relationship between firm size and R&D intensity (see Gilbert (2006) for a retrospective overview). Blundell, Griffith, and Van Reenen (1995) reveal a complex relationship between competition and innovation: at the firm level, dominant firms tend to innovate more, while at the industry level, concentration dampens innovation; to the extent that growing dominance increases concentration, and hence the level of aggregate innovation will tend to fall. Griffith, Harrison, and Van Reenen (2006) show that the effect of increasing competition on innovation is, within an industry, larger the closer to the global technological frontier. A theme running through some of the literature is that the relationship between competition and innovation may be characterised by an inverse U-shape. Especially in the early days of Structure-Conduct-Performance, this was seen as the way to reconcile Arrow and Schumpeter – following Arrow, increases in competition increase the pressure to innovate, but after some point, increasing competition may begin to reduce the incentive, unless property rights are protected. In more recent years, this was formalised by Aghion, Harris, and Vickers (1997) and Aghion, Bloom, Blundell, Griffith, and Howitt (2005). Our own reading of this literature is slightly differently nuanced, and is influenced by one of the most impressive of the existing surveys: Shapiro (2010, p.401) argues that "a firm with a vested interest in the status quo has a smaller incentive than a new entrant to develop or introduce new technology that disrupts the status quo". This is in line with standard Arrowian arguments. However, Shapiro gives on to add: "Schumpeter was also quite correct: the prospect of obtaining market power is a necessary reward to innovation". He concludes that "There is no conflict whatsoever between these two fundamental insights". This conclusion is perhaps the best balanced summary of the literature surveyed.

The literature evaluating the effect of market consolidation (mergers) on innovation is relatively small but growing. Danzon, Epstein, and Nicholson (2007), Ornaghi (2009), and Stiebale and Haucap (2013) all provide estimates for the pharmaceutical industry. All three studies find a negative effect of mergers on innovation. Their findings are robust to different measures of innovation (such as R&D expenditure, patents, or citation-weighted patents). As such, the negative effect of market power on innovation appears to dominate potential positive effects arising from cost savings. Stiebale and Haucap (2013) further estimate the effects of innovation activity on rival firms active in the same markets. While being smaller in magnitude than the reductions in the merging parties' R&D activity, they find a significant negative effect on rival firms as well. Finally, Szücs (2014) finds that target firms substantially decrease their R&D post merger, while the R&D intensity of acquirers drops due to a sharp increase in sales. On the other hand, other studies find increases in R&D activity after mergers, including Bertrand (2009) and Stiebale and Haucap (2013).

The above studies only provide estimates of average, cross-industry effects rather than identifying the circumstances where the negative/positive effects are particularly pronounced. Such averages offer valuable contributions to academic and policy work but they are unable to pin down the very cases where consolidation positively or negatively affected innovation. To remedy this, we look at a specific industry in the wake of consolidation from 5 to 3 firms. Closer to our paper is Igami and Uetake (2017), who take a closer look at the HDD market and estimate a dynamic oligopoly model using HDD shipment and unit price data, in which merger decisions, along with innovation and entry-exit strategies are endogenous. By employing a set of hypothetical merger policies as counterfactual they show that the optimal merger policy should block mergers where there are 6 or fewer players left in the HDD market. Instead of a structural model we use a reduced-form method, with a meticulously selected set of Control groups to estimate the impact of consolidation. This way we avoid restrictive assumptions about the nature of competition in the HDD market - which we believe is very strongly influenced by external factors, for example competition from neighbouring technologies, such as flash memory based storage. Moreover, we do not limit our analysis to a specific measure of innovation but look at four different factors separately and at their interaction. These are R&D expenditure, patent activity, number of new products, and unit cost of new products. With the latter two measures, we are also able to test the relative power of R&D and patent data in predicting product improvements. Throughout the paper we offer a number of methodological contributions. For example on patents, unlike in many previous studies, we do not arbitrarily choose a single measure of patent activity, instead we synthesise all available measures and we look at the distribution of all results. Finally, we offer insight into how to acquire and use simple product characteristics data to measure the innovation impact of mergers.

2 The Hard Disk Drive and Solid State Drive markets

We look at two mergers (Seagate/Samsung, and Western Digital/Hitachi) in the Hard Disk Drive market. First we briefly introduce the characteristics of the storage market, including Hard Disk Drives. Then we give account of the relevant merger control decisions.

2.1 The storage market

There are two main storage technologies, Hard Disk Drives (HDD), and Flash-based (NAND) storage. An HDD is a device that uses one or more rotating disks with magnetic surfaces (media) to store and allow access to data, whereas Flash storage uses integrated circuit assemblies to store data, which records, stores and retrieves digital data without any moving parts. Solid state drives (SSD) and USB Flash drives (Flash Memory based data storage device with integrated USB interface) are Flash memory based storage. SSDs are built on semiconductor memory arranged as a disk instead of magnetic or optical storage support. Because no mechanical components are involved, SSDs are fast in comparison to rotating media (HDD), providing access to data in microseconds, instead of the several milliseconds requested by HHDs.

The main benefits of SSDs compared to HDDs include increased speed, smaller size, lower power consumption, increased resistance to shock, and reduced noise and heat generation. A major disadvantage of SSDs is their price, although SSD capacity size has been rapidly increasing and unit prices have been dropping. HDDs have been primarily used for archiving, and SSDs are mainly employed in portable devices (laptops, smartphones, tablets). Despite their commercial success, HDDs have always had mechanical limitations, suggesting that their growth would come to an end and would be replaced by a different technology. By their nature, mechanical devices cannot improve as quickly as solid state technologies can. In 20 years (1988-2008) CPU performance increased by 16,800 times, whereas in the same period HDD's performance increased by 11 times.

HDD sales have been dropping since 2011 and SSDs have shown a strong increase in the same period. Part of the reason for HDD's loss is the decline in the sales of desktop PCs – traditionally the main users of HDDs. Nevertheless, even today, HDDs are still the dominant in the market for data storage. SSDs are slowly gaining pace but this is dwarfed by the fact that a large amount of increase in storage demand is for data archives and cloud storage, which rely, to a large extent, on HDDs. Storage used for example in mobile devices, using flash based technologies, is only a tiny fraction of all storage capacity, despite its wide dissemination.

The HDD market has witnessed continuous consolidation since the late 1980's. Before the Seagate/Samsung and the Western Digital (WD)/Hitachi GST (HGST) mergers, there had been five players in the market: Seagate, WD, Toshiba, HGST, and Samsung. Following the two mergers, the market shares of Seagate, Western Digital, and Toshiba have been close to a 40-40-20 split. The SSD market is more fragmented, unsurprisingly, as it is a less mature technology. The major players in SSD are Samsung, Toshiba, SandDisk, Micron, SKHynix, and Intel.

2.2 Regulatory approval

The Seagate/Samsung merger was unconditionally approved in every jurisdiction, with the exception of China (MOFCOM), where approval was subjected to a set of behavioural remedies. The main argument for the unconditional approval outside of China was that Samsung had not exerted effective competitive constraint in the HDD market, and therefore its elimination from the HDD market was not expected to affect the level of competition. The European Commission and the US authorities approved the WD/Hitachi merger subject to the divestiture of the 3.5" desktop HDD manufacturing lines to Toshiba. MOFCOM, again, took a different stance and imposed a set of behavioural remedies. In general the MOFCOM restrictions were more crippling on the WD/HGST merger.[1]

3 Econometric model and data

Our evaluation consists of three inter-dependent stages. First we look at the investment part of innovation, R&D expenditure. This is followed by a discussion of how the mergers affected the invention stage of innovation, as measured by the number of patents. Finally, we test the impact of the 2012 events on HDD product and technology features.

We observe a measure of R&D, patents, and product characteristics for each calendar quarter t, starting with Q1 2007 and finishing with Q4 2016. Out of $T = 40$ total time periods, there are $T_0 - 1$ time periods measured prior to the mergers that take place in period T_0, implying that $t \in \{1, \ldots, T_0 - 1, T_0, T_0 + 1, \ldots, T\}$.

There are J_0 firms in the Control group in the sample and J_1 in the Treatment group. Therefore indexing each firm by j, we have $j \in \{1, \ldots J_0, \ldots, J_0 + J_1\}$.

We measure R&D through its intensity (ratio of R&D expenditure to total revenue), rd_{jt}. We denote patent activity and product characteristics for firm j at time period t as pat_{jt} and $char_{jt}$ respectively. We provide more detailed explanation to these variables below.

Indicate by μ_j firm dummies, and by μ_t time dummies. Denote by x_{jt} a $(K \times 1)$ vector of the following time-varying firm characteristics:

Firm size: There are numerous studies linking various firm characteristics, such as firm size, to innovation (e.g. Shefer, 2005). We measure various dimensions of firm size (total revenue, total assets, gross profit, number of employees, and net income.)[2]

[1]More on the regulatory background in Ormosi et al. (2017).

[2]Gross profit is the difference between total revenue and the cost of revenue. In our regressions we include total revenue and gross profit, which together determine the cost of revenue. Net income includes various earnings on the firms' operations. Total debt refers to various interest bearing obligations. Total operating expenses reflects expenses not directly associated with the production of goods or services. These firm characteristics are closely correlated with each other (larger businesses will have high values, etc). We

Pre-sample R&D activity: Blundell, Griffith, and van Reenen (1995) and Blundell, Griffith, and Windmeijer (2002) both use pre-sample R&D activity as an exogenous control. We aggregate and take firm-level means of the R&D expenditure data preceding Q1 2007, and use it as additional firm specific control in the matching process.

Number of segments: In our data R&D expenditure is reported for the entire company that may have numerous diversified portfolios. This is not a problem for Seagate and WD (only active in HDD at the time) but it is a potential issue for other firms. For example R&D expenditure for Samsung incorporates all R&D spending by Samsung, which includes Samsung's products other than storage. To be able to gauge how much of the given company's total production is related to storage technologies, we used S&P's Capital IQ database for the number of segments the given business is active in. This is a time-constant figure, which means we only include it in finding a matching control and not in the DiD estimations (which control for firm-fixed effects).

We controlled for other firm-level time-variant characteristics. Cost of goods sold represents cost of revenue incurred on all raw materials, work in process, manufacturing expenses and other costs directly attributable to production of finished goods and operating revenues. Gross profit is the difference between total revenue and the cost of revenue. In our regressions we include total revenue and gross profit, which together determine the cost of revenue. Net income includes various earnings on the firms' operations. Total debt refers to various interest bearing obligations. Total operating expenses reflects expenses not directly associated with the production of goods or services. These firm characteristics are closely correlated with each other (larger businesses will have high values, etc.). To handle this we standardise these variables by using their ratio to total revenue rather than their absolute values.

In the headlined model we do not include contemporaneous effect of firm characteristics, which explains lags in the variables $x_{jt-\{1,\cdots,4\}}$. We include lags to avoid issues of simultaneity, but also because we do not believe that any of these variables would have a contemporaneous effect. We normalise each element of x (with the exception of total revenue) by using their ratio to total revenue. We denote by D_j an indicator variable to capture whether firm j was involved in one of the two mergers, and by I_t whether period t was before merger notification (Q2 2011) or after the closure of the approval (Q1 2012), ε_{jt} are idiosyncratic shocks with zero mean.

In this design $D_j = 0$ if $j = \{1, \cdots, J_0\}$=$\{$*Control group*$\}$, and $D_j = 1$ if $j \in \{J_0 + 1, \cdots, J_0 + J_1\}$=$\{$*Seagate*, *Western Digital*, *Toshiba*$\}$. It is important to point out that the Treatment group only contains the two acquiring firms, i.e. we are excluding Samsung and

discuss three firm characteristics in more detail. To handle this we standardise these variables by using their ratio to total revenue rather than their absolute values.

Hitachi. As we are studying how R&D intensity (which is firm, rather than market specific) changed for Seagate and Western Digital, we are uninterested in how innovation develops in Samsung and Hitachi, who no longer have operations and R&D expenditure in the relevant products post-merger.

We use the following model to estimate the following recursive triangular system of 3 equation. Our choice of a recursive system with time lags implies the assumption that there is no reverse causality between R&D, patents, and product characteristics. Later we examine what happens if we relax on this assumption:

$$rd_{jt} = \beta_{10} + \beta_{11} D_{1j} I_t + x_{1jt-\{1,\cdots,4\}} \phi_1 + \mu_{1j} + \mu_{1t} + \varepsilon_{1jt} \tag{1a}$$

$$pat_{jt} = \beta_{20} + \beta_{21} D_{2j} I_t + x_{2jt-\{1,\cdots,4\}} \phi_2 + rd_{jt-\{1,\cdots,4\}} \lambda_2 + \mu_{2j} + \mu_{2t} + \varepsilon_{2jt} \tag{1b}$$

$$char_{jt} = \beta_{30} + \beta_{31} D_{3j} I_t + x_{3jt-\{1,\cdots,4\}} \phi_3 + rd_{jt-\{1,\cdots,4\}} \lambda_3 + pat_{jt-\{1,\cdots,2\}} \gamma_3 + \mu_{3j} + \mu_{3t} + \varepsilon_{3jt} \tag{1c}$$

In this model β_{n1} is the treatment effect.

For each of the n error terms we assume that $E(\varepsilon_{jt}|\mathbf{x}_{jt}) = 0$, and $E(\varepsilon_{jt}\varepsilon'_{jt}|\mathbf{x}_{jt}) = \mathbf{\Sigma}$. For simplicity, as a first step, we also assume that $\mathbf{\Sigma}$ is diagonal, which would mean that estimating the system equation-by-equation using OLS would give consistent, and asymptotically efficient estimates. If the error terms are correlated (and $\mathbf{\Sigma}$ is not diagonal), OLS estimators will be inconsistent. We look at the possibility where R&D spending and patent citations are endogenous in Equation (1c). For this reason we also estimate Equation (1c), in which case patents and R&D will appear as endogenous variables the following way:

$$\begin{aligned} rd_{jt} &= \beta_{10} + \beta_{11} D_{1j} I_t + x_{1jt-\{1,\cdots,4\}} \phi_1 + \mu_{1j} + \mu_{1t} + \varepsilon_{1jt} \\ pat_{jt} &= \beta_{20} + \beta_{21} D_{2j} I_t + + x_{2jt-\{1,\cdots,4\}} \phi_2 + rd_{jt-\{1,\cdots,4\}} \lambda_2 + \mu_{2j} + \mu_{2t} + \varepsilon_{2jt} \end{aligned} \tag{2}$$

In the results section we report estimates from both the sequential and the simultaneous estimations.

At the heart of our econometric strategy is finding an adequate Control group. In what follows we explain for each R&D, patents, and product characteristics, the data and Control groups used, then we provide the regression results together. We have estimated a very large number of different models, for various measures of innovation, for various Control groups, different model specifications, and different estimation methods. In the main results section we only report those that were the best fit, but in the subsequent Section and in the Appendix we show that our results are robust to our exact model and estimation choice.

3.1 R&D intensity data

To measure the first stage of innovation we look at R&D intensity (the ratio of R&D expenditure to total revenue), which is frequently used as a proxy for innovation, although it has been subjected to some criticism, mainly because not all R&D spending leads to innovation (Gilbert, 2006), or that increasing R&D spending (over-time) may simply reflect diminishing returns (e.g. Strumsky et al. 2010). Moreover, some innovation is achieved through investment in own R&D, and some are purchased from other innovators. On the other hand, there is evidence that in the absence of changes in firm size, employment or exogenous factors, firm-level R&D expenditure follows a random walk with a small error variance – i.e. R&D expenditure in the short run (5-10 years) is roughly constant or increasing slightly.[3] This constant or linearized trend could facilitate the identification of the effect of changes caused by the mergers.

For all firms in our sample we have complete quarterly data coverage for the period of observation Q1 2007 to Q4 2016, which spans over two equal periods pre, and–post merger[4]). All the data used for the R&D analysis is from firms' balance sheets, as downloaded from S&P's Capital IQ database. Below, we provide a brief introduction of the main variables used for our analysis.

Figure 1 plots R&D intensity for Seagate, Western Digital and Toshiba between 2007 and 2016. The two vertical lines show the start and the closure of the merger approval process. Figure 1 reveals a few interesting patterns. R&D intensity for WD and Seagate is parallel until Q4 2009, then WD starts its ascending trail. This seems to correspond to industry news of WD's dedication to increasing innovation.[5] It appears that the acquisition of HGST was not the cause but a part of WD's path of increasing innovation. Seagate's R&D intensity suffered a slump in Q1-Q2 2012, which was an accounting effect, Total Revenue increased more than R&D expenditure as a result of adding Samsung HDD to Seagate's books. Post-merger Seagate's R&D intensity follows an increasing trend. Finally, Toshiba had a leap in 2009, much sharper than Seagate and WD, possibly the result of Toshiba's acquisition of Fujitsu. This is followed by a fairly constant level of R&D intensity both before and after 2012.

Figure 1 also draws light to a couple of methodological issues. When evaluating how R&D intensity changes after a merger, one must not ignore an important artefact of this type of data, that is, following a merger, elements of the financial statement of the acquired company

[3]Hall et al. (1986), Coad and Rao (2009).

[4]WD acquired Sandisk to boost its SSD/Flash portfolio in 2016, which is another reason why we excluded post-2017 data.

[5]In February 2011 WD opened a new HDD R&D centre in Singapore, and in December 2011 it set up its first overseas SSD R&D centre in Taiwan (focusing on R&D enterprise applications).

Figure 1: R&D intensity for Seagate, WD, and Toshiba

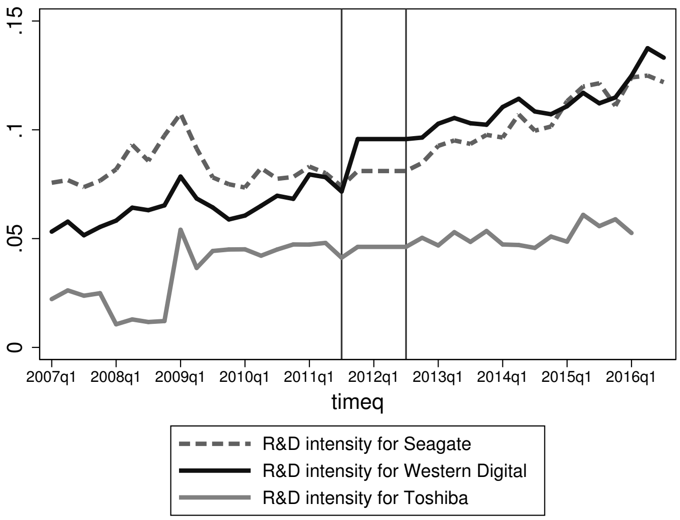

are added to the corresponding elements of the financial statement of the acquiring company. This means that for simple arithmetic reasons R&D expenditure and total revenue will be higher in the post-merger period even if the merger does not increase the R&D intensity of the relevant businesses. For this reason we ignore the period of the treatment (the merger approval period) when estimating the impact of treatment, to take out the hikes caused by merging the two financial statements. [6]

A further methodological point is worth raising when using R&D data. It is very difficult (if possible at all) to acquire data specifically for the relevant segments or products of the analysed firms. Therefore such data might be more fitting in cases where the relevant firms are less diverse, where R&D expenditure figures in financial statements can be safely attributed to the relevant product. In our case, Seagate and Western Digital fit this bill and so do many of our Control firms (e.g. Sandisk, Kingston, Micron, Hynix) but Toshiba is active in many different areas, and storage only constitutes around a quarter of its total operating revenue and R&D expenditure.

3.1.1 A synthetic Control group

Finding an adequate Control group is not a trivial exercise. There are no rivals in the HDD market that are not affected by the mergers, and the HDD market is worldwide, which means that local markets cannot be used. Instead we explore product differentiation (HDD,

[6] We remove Q4 2011 and Q1 2012 from our analysis, and we also disregard the growth in R&D intensity in this period.

SSD, Flash drives) to find a Control group. To start with, we use a synthetic control group, as described in Abadie and Gardeazabal (2003), and Abadie, Diamond, and Hainmueller (2010). The idea behind this method is to generate a weighted sample of firms that are most similar to the Treatment firm based on a set of observable characteristics. We use all IT firms (as above) as a pool for potential Controls, and calculate the weights based on total revenue, gross profit, total assets, net income, total debt, expenses, pre-sample R&D expenditure, and the proportion of relevant segments to find the synthetic control. Our comparisons across these Control groups (as presented in Section A.1 in the Appendix) suggest that the synthetic control group is closest to the characteristics of a perfectly designed Control group.

The main conditions for an adequate Control is similarity to the Treatment group and independence of the treatment. In terms of similarity, we assume that these technologies were exposed to the same demand and supply side shocks, except for the effect of the merger. We control for the cost of revenue (revenue minus profit) which should pick up some of the supply side shocks. On the demand side we assume that the main determinants of demand, income and substitutability change in parallel for buyers of high-tech products, driven by the same underlying economic conditions. On the supply side, we have less intuition, but we control for changes in costs, which should pick up at least some of the shocks.

Regarding independence, because there is some substitutability across storage technologies, innovation decisions in one product might trigger a response in the other. This would make these a biased counterfactual. The sign of the bias would depend on whether innovation in other firms is a strategic substitute or complement (i.e. if Seagate innovates more, will WD follow suit). Our intuition is that these are strategic substitutes, therefore the estimate is biased towards zero but have the correct sign.

3.2 Patent activity data

We extracted patent data for each technology (HDD, Flash), and, only subsequently, grouped the data by firms.[7] This approach enables an analysis at firm level and thus grants the matching of patent data with firm R&D expenditure and other firm characteristics. In the analysis that follows we have 53,107 observations of HDD-relevant patents owned by almost 16,000 firms. Given our interest in the effect of the merger we confine the time period to four years before the merger and four years after the merger. Our database refers to patent families, including patent applications taken in multiple countries to protect the invention, which is relatively common for inventors or applications. The effective date of each patent

[7]Relevant data on patents have been collected and cleared by an Italian start-up, BigFlo, which works in collaboration with the University of Bergamo in Italy. They gathered full information on patents related to HDDs, SSDs, and Flash drives.

application refers to the quarter when a first application is registered in a country. The date of subsequent applications for the same patent are also relevant as they can inform us about changes in patent ownership.

Unlike R&D spending, there is no unique way to measure patent activity, and, as such, various measures have been proposed and employed. A non-comprehensive list includes: patent counts, patents weighted by citations, patent intensity (the ratio between patent count and revenues), and stock of patents net of patent depreciation. Unlike previous works, we do not arbitrarily choose one or two measures. Instead we remain agnostic about what the best measure of patent activity is. In this spirit we take a novel approach, and estimate the effect of the mergers on every possible patent measure and then synthesise all the estimated effects in a single estimate.

We use factor analysis of number of patents, number of citations, patent literature, number of inventors, patent claims, number of applicants and number of countries, and find that variation across these factors mainly reflects variation in one underlying factor, which we then use as a factor of patent activity. This gives us three variables to begin with: patent count, patent citation, and patent factor. For each of these variables we: (i) generate stocks; (ii) smooth out shocks by employing a moving average over 4 (quarterly) lags and 4 (quarterly) leads; (iii) normalize the three variables by total revenues, as to obtain measures of patent intensity. Furthermore, with no insight on whether the causal effect of the merger on patent activity should be measured in levels, in logs (proportions), or in growth, we transform the $4 \times 3 = 12$ variables in each of these three possibilities. This exercise gives us 36 different measures of patent activities.

Then, separately for Seagate, Toshiba and Western Digital, we estimate the causal effect of the merger on each of the counts of patent activities. In order to make results comparable we standardise all continuous variables, including the control variables total debts, total assets, total revenue and total R&D intensity (all up to four lags). The procedure gives, for each of the three companies, $36 \times 3 = 108$ standardised causal estimates of patent activities. Finally, we combine these 108 causal estimates using a meta-analysis approach and obtain the average effects and the distribution of estimates.

3.2.1 Flash patents as Control group

Finding a Control group that is independent of the mergers in terms of patent activity but is sufficiently similar to the HDD market is not trivial, mainly because all HDD manufacturers were involved in the treatment (the mergers and the related events) and therefore we have to rely on different technologies, and patenting activity is sensitive to the underlying technology.

We have 40,655 Flash Memory related patent applications in our sample, which we used as Control group. Our data (Figure 6 in the Appendix) suggests that Flash is sufficiently similar to the Treatment for the purposes of measuring patent activity. We run various tests of similarity (explained in the Appendix), and in our final results exclude the ones where it was rejected. Regarding independence, there is a strong possibility that this Control group is not independent, and if there is a bias, its sign depends on whether innovation in the Control is strategic complement or substitute to innovation in the Treatment group. Our intuition, and implied assumption is that they are strategic complements (i.e. a rise in innovation in HDD is accompanied by a rise in innovation in Flash storage) or even more, their relationship is sequential. This means that even if there is bias, the bias only affects the magnitude and not the sign of our estimates.

3.3 Product innovation data

Whilst R&D expenditure identifies the breath and intensity of innovation, patents and product characteristics capture the direction of the innovation. Having information on the evolution of product characteristics offers an insight into technological diffusion and an altogether more accurate measure of innovation. Moreover, it allows us to test how R&D spending and patent activity affects these characteristics - i.e. which of the two measures is a better approximation of innovation in the HDD market. Product characteristics are much less studied in the economics literature on innovation, probably due to the difficulty of accessing this type of data in many industries. Here we are only looking at two of the simplest ways of measuring product innovation: the number of new products marketed, and the unit price ($ price of a Gb of storage).

We collected information on 1931 HDDs and on 1353 SSDs that were sold on Amazon between 2001 and 2016.[8] Using retail data has a disadvantage that we only capture consumer sales of HDDs and ignore the enterprise applications of HDD. On the other hand, innovations in HDD are likely to have uniform effect across all applications: enterprise, desktop, mobile and consumer electronics. For this reason we expect that our selective data on desktop and mobile applications is representative of the whole industry in terms of technological innovations. For 98 HDDs and 54 SSDs we could not identify a brand from the scraped data and these were removed from the sample. We removed brands with fewer than 10 products, and we also removed hybrid drives as they represent a combination of the two technologies. The sample consists of 33 SSD and 5 HDD brands.[9] We have access to the following product

[8]The sample accounts for the mergers that happened before 2012, for example Fujitsu is recorded as Toshiba as a result of their 2009 merger.

[9]This is as expected, industrial organisation literature, such as Jovanovic and MacDonald (1994), or Klepper and Simons (2000) have shown that as industries and technologies mature, markets tend to become more concentrated.

characteristics for HDDs and SSDs:

Date first marketed on Amazon: There is some grouping in the way firms market new HDDs and SSDs. For example, 17 different Intel SSDs appeared on Amazon on 27 March 2016. However more than 2/3 of all drives in our sample were marketed on unique days, and most groupings happened in 2s and 3s (i.e. two or three products in the same day).

Form factor: The form factor refers to the physical size of the drive. Both HDDs and SSDs come in the following form factors: 5.25-inch, 3.5-inch, 2.5-inch or 1.8-inch. In our sample we only have the latter three. The remedy in the WD/HGST merger was the divestiture of the 3.5-inch form factor HDD manufacturing to Toshiba. WD retained the 2.5-inch manufacturing lines.

Storage capacity: Ideally, one would have looked at areal density. However using retail data we had limited access to technological details and could only measure formatted capacity (expressed gigabytes). Capacity alone does not give an unambiguous picture of innovation because newer products do not necessarily mean larger capacity. Moreover, the fact that there is a larger capacity storage does not mean that demand for smaller capacities disappears. Therefore firms continuously market smaller and larger capacity drives at the same time.

Price: We recorded the prices of all products in the sample as they were collected in May 2017. For example for an HDD that was first marketed in 2010, we had the price as it appeared in 2017. One could argue that this way for older products we record the final price (i.e. the price in 2017), which might not be the same as the introduction price (e.g. price in 2010). However, the pace of introducing new HDDs is very fast. When a HDD manufacturer comes out with a new product it risks cannibalising into the sales of old products. Despite this, HDDs are introduced at a fast rate. On average, the same manufacturer introduced a new product of exactly the *same* capacity every 6 months (5 months when only looking at the three Treatment firms), and the same manufacturer introduced a new product of *any* capacity every month (less than 10 days when looking across the three Treatment firms). If manufacturers dropped the prices of their older products, they would have cannibalised into the sales of their newly introduced products. In situations like this (where the same firm offers products that are substitutes), firms are unlikely to engage in price competition between their own products.[10] For this reason we believe that the price of older products still available on Amazon gives a good approximation of their original price. Moreover, even if there is a price drop, the technological depreciation of HDDs is so fast that demand for older products very rapidly disappears. Therefore the price reduction – if exists – must

[10]See for example Douglas and Pavcnik (2001).

quickly take place, i.e. even for the relatively new products the prices already reflect the final (lowest) retail price.

From the above, we derive our two variables used for measuring technological progress, the number of new products, and the unit price of new products ($/Gb). Both variables are recorded by firm j in period t.

3.3.1 SSDs as Control group

For an adequate Control group one would need to use a product, for which new models are marketed at a similar pace to HDDs, but are different enough so they were not affected by the mergers.

SSD is a less mature technology than HDDs, and therefore it is possible that the pace of innovation for SSDs is different from HDDs. The question is how much this matters. In mature industries product differentiation is not driven by innovation any more. However HDDs are different. In the HDD market competition is still driven by differences in technology (unlike in typical mature industries where technology tends to be static), and therefore there is still intensive technological progress in HDDs (for example in areal density).[11] For both our variables of interest (number of new products and unit price) expansion is still ongoing both in HDD and SSD.

Take the number of new products. Firms come out with new products as a response to demand conditions. There is a significant overlap between the two technologies in the demand for storage. SSDs have been converging to HDDs both in price and capacity and have exerted increasing competitive pressure on HDDs. Even at the time of the merger, the merging parties argued that "SSDs will become "mainstream" in the coming years, replacing HDDs in many applications." [12] This would put SSDs in a favourable position to be used as Control.

Regarding the unit price of capacity, technological differences have a central role. HDDs are mechanical devices and as such, their development is limited at some point. However, so far the pace of increase in areal density (HDD) has been fast. Figure 2 compares how the unit cost of disk capacity evolved in HDDs and SSDs. Visually, the two lines follow a similar trend, with the exception of 2009, where there are only a few observations for SSD. We will formally confirm this parallel trend later. This would suggest that – at least for this particular characteristic – SSD is not an outlandish choice as Control.

Regarding the independence of SSDs from the HDD mergers, we rely on the same argu-

[11] https://www.tomcoughlin.com/Techpapers/HDD%20Market%20Down%20to%20Three%20Suppliers,%20042011.pdf

[12] Para. 231, European Commission, Seagate/HDD Business of Samsung COMP/M.6214, Decision October 19, 2011

Figure 2: Unit cost of storage [ln($/Gb)

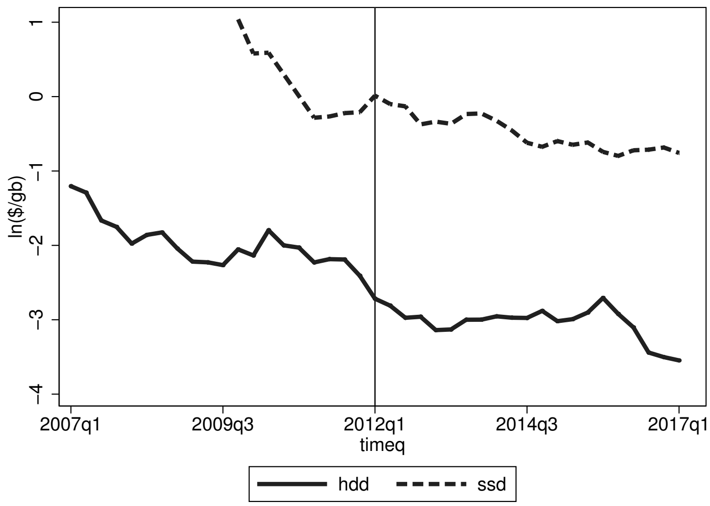

ment as above. From industry references it clearly appears that there is competition between the two products,[13] which makes it more likely that the two are strategic complements in innovation (e.g. if the pace of innovation increases in one technology it also increases in the other). This would mean that even if SSDs are a biased Control, the bias would only affect the magnitude and not the sign of the estimated effects.

4 Estimating the impact of consolidation on innovation

In this section we present the results of estimating the sequential system introduced in Equations (1a)-(1c) and the corresponding simultaneous system. In this section we only present the best performing models and Control groups. We provide further robustness checks in the subsequent section and in the Appendix. Table 1 has two main blocks. The left one contains the sequential, the right one reports the simultaneous estimates. The left-hand block reports the best fitting one of all the models we estimated. On the right-hand side we provide our estimates where we instrument the R&D and patent variables. The table has three sections, one for each Treatment firm (Seagate, WD, Toshiba).

We report the difference-in-difference coefficients (DD) which is the treatment effect of the 2012 events on each of the Treatment firms. RD_int shows the effect of lagged

[13]http://www.pcworld.com/article/3184464/storage/intel-optane-memory-has-a-mission-make-hard-drives-faster-than-ssds.html and http://www.financialexpress.com/industry/technology/data-storage-solid-state-drives-can-now-compete-with-hard-disk-drives/648502/

R&D intensity in general, and *RD_int* x *treat* is the effect of lagged R&D intensity for the Treatment firm following the 2012 events. Finally *patcit* is the effect of the lagged patent factor (derived using a factor analysis as explained above) in general, and *patcit* x *treat* is the effect of the number of lagged patent citations for the Treatment firm following the 2012 events. We also report test results on the hypothesis that the sum of these lagged effects is zero. The treatment effect, *DD* should be interpreted as the residual effect of the mergers (all effects that are not related to R&D intensity or patent citations).

Table 1: The effect of the 2012 consolidation - headline results

	Sequential				**Simultaneous**	
	RnD_int	Pat_cit	ln(cost)	ln(number)	ln(cost)	ln(number)
Seagate						
DD	**0.0316*** **	**0.084** **	1.062	0.127	**1.038***	**-0.742***
(p-val)	(0.005)	(0.026)	(0.124)	(0.688)	(0.094)	(0.091)
RD_int		**0.269*** **	0.236	**0.547** **	**-0.262*** **	**0.250*** **
(p-val)		(0.000)	(0.74)	(0.042)	(0.006)	(0.000)
RD_int x treat		**-1.44*** **	**-2.434** **	**1.146***	**-2.698** **	**2.716*** **
(p-val)		(0.000)	(0.046)	(0.062)	(0.024)	(0.001)
patcit			0.061	0.034	0.179	-0.108
(p-val)			(0.628)	(0.622)	(0.274)	(0.365)
patcit x treat			**-0.739***	-0.049	**-0.776** **	0.147
(p-val)			(0.089)	(0.862)	(0.024)	(0.553)
Observations	78	558	164	163	157	159
Western Digital						
DD	**0.0479*** **	**-0.583*** **	-0.810	**-2.585*** **	0.479	-0.408
(p-val)	(0.006)	(0.000)	(0.263)	(0.004)	(0.649)	(0.584)
RD_int		**0.262*** **	0.454	0.295	**-0.286*** **	**0.267*** **
(p-val)		(0.000)	(0.563)	(0.412)	(0.001)	(0.000)
RD_int x treat		-0.147	**2.437***	**4.132** **	0.582	-1.208
(p-val)		(0.250)	(0.070)	(0.014)	(0.719)	(0.293)
patcit			**0.165***	-0.011	0.103	-0.0339
(p-val)			(0.089)	(0.833)	(0.495)	(0.755)
patcit x treat			**0.799*** **	**3.106*** **	-0.384	0.565
(p-val)			(0.008)	(0.000)	(0.524)	(0.187)
Observations	78	558	146	141	140	142
Toshiba						
DD	**0.0166*** **	**2.739*** **	**14.73*** **	**-10.45*** **	0.912	-1.081
(p-val)	(0.005)	(0.000)	(0.001)	(0.000)	(0.727)	(0.548)
RD_int		**0.246*** **	0.338	**0.538**	**-0.317*** **	**0.269*** **
(p-val)		(0.000)	(0.625)	(0.042)	(0.002)	(0.000)
RD_int x treat		6.112***	**28.414*** **	**-20.369*** **	0.686	-2.191
(p-val)		(0.000)	(0.001)	(0.000)	(0.897)	(0.547)
patcit			0.153	0.104	0.210	0.00865
(p-val)			(0.256)	(0.383)	(0.205)	(0.940)
patcit x treat			**1.114** **	**-0.619** **	**-0.883***	-0.0633
(p-val)			(0.011)	(0.024)	(0.091)	(0.860)
Observations	74	591	146	148	137	139

pvals in parentheses

The first column is the treatment effect on R&D intensity, using the synthetic control group, which performed best in our tests. The results suggest that all three firms increased their R&D intensity as a result of the 2012 events.[14] However, when testing for parallel trends (an assumption required for unbiased DiD estimates) we find that the assumption is violated for both Western Digital and Toshiba. This implies that we are only able to attribute the increase in R&D expenditure to the mergers in the case of Seagate. For WD and Toshiba the increase was likely to have been triggered by something before the 2011/12 mergers.

The second column shows the impact of the mergers (and previous R&D spending) on patent activity, where the Control group is Flash patents. It appears consistent across our various models that R&D spending increases patent activity in general. Regarding the impact of the mergers, we found evidence that the merger increased patent activity in Seagate but reduced their ability to convert R&D spending into patents. The merger seems to have reduced WD's and improved Toshiba's patent activity. We show later that the positive impact on Seagate's patent activity remains robust across all model specifications.

For the unit cost of HDD storage we have evidence that R$D spending contributes to lower costs in general but patents do not have similar general effect. Regarding the firm-specific effects, there is evidence that Seagate improved its ability to convert increased R&D spending into lower unit-costs. We found no robust effect for the other two firms.

For the number of new products R&D but not patents have a positive effect in general. Of the firm specific effects the only thing that appears robust across our two models is that the mergers improved Seagate's ability to convert R&D spending into more new products.

4.1 Discussion of results

In general the results provide evidence that R&D spending has a positive impact on patent activity, and that R&D spending - but not patent activity - boosts the number of new products and reduces the unit cost of these new products. One interpretation of this finding would be in line with Griliches (1998), i.e. once controlling for R&D expenditure, the residual effect of patents disappears because R&D already contains the information that one can get from controlling for patent activity. However, even when we take out R&D expenditure from our model, patent activity still does not explain much of the variation in the number of new products or the unit cost of these products. This is an important contribution to the existing literature for several reasons. Firstly, it offers evidence that R&D spending might be a better approximation of product characteristics (at least for the purposes of our two

[14] The Toshiba R&D results should be treated with more caution for two reasons: (1) as explained above, the matching for Toshiba did not work as well as for the other two firms; and (2) storage production is only a small segment of Toshiba, whereas our R&D data is firm-level, therefore these effects could be picking up changes in other segments (i.e. a general, not HDD specific drop in Toshiba's R&D spending).

measures, unit costs, and the number of new products). This could be useful information for future research using either R&D or patent measures to approximate innovation.

Regarding firm specific effects, most importantly, we found no evidence that the merger would have led to a fall in innovation activity for any of the firms. Seagate seems to be the only firm where the mergers triggered an increase in R&D expenditure, which we show positively effects patent activity and product characteristics (number of new products, and unit costs). We do not have conclusive evidence on what exactly triggered this increase, but we can offer a number of alternative interpretations. The 2012 events and the start of the consummation of the merger with Samsung triggered an increase in innovation activity. There were innovation synergies between Seagate and Samsung, which were corroborated by the merger. The two firms had cross-licensing agreements even before the merger. With the merger, the shared pool of IP was conducive to increased R&D spending. Another explanation is that there was increasing competitive pressure from SSD. It is also possible that Seagate spent on R&D more intensively and experienced improved product characteristics than WD because their merger was less restricted by regulatory approval and therefore the consummation of the merger advanced further than for WD.[15] Seagate was able to access Samsung's stock of intellectual property (patents). Seagate became the assignee on more than 20% of Samsung's HDD-related patents with the mergers (in 2012). Many of the patents that Samsung kept were not strictly on HDDs, but on complementary products that use HDDs. It was therefore safe to expect that the Samsung/Seagate merger had the potential to affect Seagate's innovation activities, if not least, through the synergies resulting from shared access to some key HDD patents.

For WD we do not find strong and consistent evidence of an effect of the consolidation. This does not mean that WD's R&D spending did not change after 2012. It means that the 2012 change in WD's R&D intensity in comparison to our Control group was not significantly different. For patents, we found an increase and no effect on product characteristics. One possible explanation is that for WD, the acquisition of HGST was part of a longer term trend of increasing R&D. This could have been a response to intensifying competition from SSD. The 2012 events did not trigger the increase, it started earlier than that. For WD it is also possible that the MOFCOM decisions particularly hindered the consummation of the WD/HGST merger until October 2015. Remedies were much stricter and they fundamentally required that WD duplicate their R&D, production, marketing, and sales operations. This

[15] Although there were remedies in place to ensure that the brands were kept separately and that the acquired brand does not suffer as a result of the merger, property rights (including intellectual property) were transferred with the conditional approval of the merger (which is evidenced by the fact that revenues were received by the acquiring firms post-merger).

was crippling for WD's efficiency.[16] We did not find any evidence of patent transfers from HGST to Toshiba at the time of the merger, despite the requirement that relevant HGST IP rights should have been transferred to Toshiba upon their purchase of the divested 3.5-in HDD operations. HGST as a brand existed until Q4 2015, and the cut-off point of patent data is 2 years (data that is less than 2 years old may not have been included in the relevant patent registers). It is therefore possible that licensing rights were given to Toshiba, but HGST remained the assignee. Other events might also affected WD's R&D spending. For example the divestiture of the 3.5in operations to Toshiba had to include all 3.5in related IP rights. This might have negatively affected how innovation, and indeed R&D spending evolved post-2012 for WD. Our own interpretation is closest to the first one. Figure 1 shows that WD's R&D expenditure had been on a steady increasing path since Q4 2009. The 2012 did not further affect this increase. This however does not mean that a combination of other things are not confounding this interpretation.

Finally, we found no evidence of a change in Toshiba's R&D spending and patent activity, and product characteristics after the mergers. However, for Toshiba we could not establish an unbiased Control group (violation of parallel trend assumption), therefore the R&D estimates are potentially biased. Moreover, R&D figures include Toshiba's other segments (around 25% of Toshiba's revenue comes from storage related operations). For this reason it would be far-fetched to go into a detailed discussion of the causes of finding a potential drop in R&D intensity.

5 Robustness checks

To evaluate the robustness of our results, we looked at a number of potential Control groups. Details of these groups and how we constructed them are given in Ormosi et al. (2017).

5.1 R&D expenditure

Selected storage firms as Control: Our first Control group consists of a sample of SSD, and USB Flash drive firms. This is not an exhaustive list of all storage producers, but these are the largest firms in these markets (making them most similar to the Treatment firms), and the ones where R&D expenditure data was available.[17]

[16] For WD and the number of employees jumped from less than 60,000 to over 100,000 after the merger. For Seagate, the pre and post-merger figures are very similar (around 55,000). At the same time there was only 8 per cent difference between the two firms in post-merger capacity shipped .

[17] The Control group includes the following firms: Transcend Information Inc., Intel Corporation, Sandisk Corporation, Kingston Technology, Micron Technology Inc., Imation Corp., Verbatim (Mitsubishi Kagaku Media), SK Hynix Inc., Sony Corporation, Lite-On Technology Corp., Powerchip, Barun Electronics, I-

Large sample IT firms as control: To circumvent the problem of a possible bias, as an alternative Control group, we looked at a more extended sample of firms. We selected all firms classified under 'Information technology' on S&P's Capital IQ database. Being in different product markets we expected these to be more likely to be independent and thus unaffected by the HDD consolidation. We had over 200,000 such IT firms in this sample. We eliminated very small businesses ($1 million total revenue) and businesses where balanced data was not available. This left us with a sample of 1701 firms, plus the 5 Treatment firms (Western Digital, Hitachi, Seagate, Samsung, and Toshiba). We distinguished between 4 potential Control groups here. First of all, we only included firms that were most similar to the Treatment firms in their primary industry (SIC codes 357x). Second, using a larger group, we included firms with SIC codes 35xx. Our third Control includes firms with SIC codes 3xxx, and finally our fourth Control includes all 1701 IT firms. As we show in the Appendix, this latter sample performs best, therefore in the analysis below we only use that Control.

A weighted sample of IT firms: This is a reduced version on the previous Control group, containing only the most similar firms (based on Propensity Score Matching with replacement). Matching is conducted based on total revenue, pre-sample R&D expenditure,[18] revenue growth, total assets, gross profit, net income, and number of segments[19]. Figure X in the Appendix shows the firms included in the weighted Control group, when matched against WD, Seagate, and Toshiba. We used equal weights for the firms with the nearest 30 propensity scores. We tried different matching and weighting methods but they provided worse fits.

Table 2 shows the estimated effect of the mergers on the annual R&D intensity of WD, Seagate, and Toshiba with all four Control groups. The coefficients seem robust to our choice in the exact specification of the Control group and under different model specifications. Although there is evidence of an increase for all firms, we show later we can only attribute this to the merger in the case of Seagate.

5.2 Patents

We looked at two more potential Control groups to check the robustness of our results.[20]

O Data Device Inc., Quanta Storage Inc., Ritek Corp., Panram Int., Power Quotient Int., Silicon Power Computer & Communications, Trek 2000 Int. Ltd.

[18]Similar to Blundell, Griffith, and Van Reenen (1995) and Blundell, Griffith, and Windmeijer (2002), we aggregate and take firm-level means of the R&D expenditure data preceding Q1 2007 to control for some of the unobserved firm heterogeneity.

[19]To be able to gauge how much of the given company's total production is related to storage technologies, we used S&P's Capital IQ database for the number of segments the given business is active in.

[20]Detailed explanation of these Control groups is given in Ormosi et al. (2017).

Table 2: Effect of the mergers on R&D intensity of WD, Seagate, and Toshiba

Control	Other storage firms	IT firms	Weighted IT firms	Synthetic
Seagate				
DiD	0.176***	0.212***	0.321***	**0.0316*** **
std.err.	(0.007)	(0.000)	(0.002)	(0.000)
obs.	517	23267	641	78
Western Digital				
DiD	0.469***	0.346***	1.190***	**0.0479*** **
std.err.	(0.000)	(0.000)	(0.000)	(0.000)
obs.	440	24053	623	78
Toshiba				
DiD	0.0413	0.0193	0.0708***	**0.0166*** **
std.err.	(0.541)	(0.394)	(0.002)	(0.002)
obs.	517	23267	607	74

Other HDD patents: This Control group consists of the top 10 firms in terms of the number of HDD-related patents held in our data.[21] These patents are not innovations of the HDD units but innovations on something complementary to HDD.[22] It is important to emphasise that this is not to be confused with complementary patents. Complementary patents are relatively common in specific technological areas, like the semiconductor industry, to protect the innovation proposed in the patent applications. Such types of patents are introduced simultaneously with essential patents, and the use of the created patent pools allows their independent application via licensing contracts. We are not looking at complementary patents but patents on complementary products.

Top storage firms' patents: This group includes patents of the top storage firms that we also used as Control in the R&D section above.

Table 3 shows the mean and the 95% confidence intervals for the distribution of our estimates.[23] The results for Seagate are robust to the choice of the Control group. The mergers resulted in an increase in patent activity, and the rate at which R&D spending is transferred to patents also improved for Seagate. For WD the results are not robust to the Control choice, in some models patent activity improved (other HDD), in some of them it lessened (other Flash). One possible explanation is that WD's patent activity has increased, but not as fast as the increase in Flash patents. Finally, for Toshiba there is evidence of the merger increasing patent activity, but the effect of R&D cannot be interpreted due to the problems with Toshiba's R&D figures, as discussed above.

[21] The list of Control firms includes: Canon, Funai, Hon Hai Precision Industry, IBM, Inventec, Lenovo, LG Electronics, Panasonic, Ricoh, and Sony.

[22] For example, Sony has a large number of HDD related patents. Many of this are related to game consoles such as Playstation or PSP, which use HDD's for data storage.

[23] Once we acquired all DiD estimates, we selected the ones which satisfy the parallel trend assumption, and then examined the robustness of results only across this selected group of Controls. Finally, we standardise the estimates, and synthesise them (using a meta study approach) into one estimate.

Table 3: The effect of the mergers on patent activity

Control		Seagate	Western Digital	Toshiba
Other HDD	DD	**0.407**	**0.230**	0.495
		[0.342;0.473]	[0.130;0.329]	[-0.555;1.544]
	R&D	**0.190**	**0.197**	**0.150**
		[0.140;0.240]	[0.147;0.247]	[0.100;0.200]
	R&D x Treat	-0.189	0.348	1.604
		[-0.419;0.042]	[-0.130;0.827]	[-0.519;3.727]
Other Flash	DD	**0.084**	**-0.583**	**2.739**
		[0.002;0.167]	[-0.700;-0.467]	[1.516;3.961]
	R&D	**0.269**	**0.262**	**0.246**
		[0.215;0.323]	[0.207;0.316]	[0.192;0.300]
	R&D x Treat	**-1.440**	-0.147	**6.112**
		[-1.660;-1.221]	[-0.572;0.278]	[3.558;8.667]
Top Storage	DD	**0.465**	0.034	**1.100**
		[0.414;0.515]	[-0.042;0.11]	[0.425;1.775]
	R&D	-0.009	-0.005	0.013
		[-0.041;0.022]	[-0.035;0.025]	[-0.020;0.045]
	R&D x Treat	**-0.349**	0.213	**2.670**
		[-0.505;-0.193]	[-0.124;0.55]	[1.320;4.021]

5.3 Product characteristics

When looking at product characteristics we had no a priori information on the lag between mergers and the change in product features. This lag might vary from industry to industry, so we turned to the data for more information. We ran several experiments, for 5 different 'treatment times' $W \in \{$Jul 2012, Jan 2013, Jul 2013, Jan 2014, and Jul 2014$\}$. This could be informative of the lag of the effect of mergers on new products and product unit costs. Moreover, pre-event parallel trend was more likely violated in 2009 where we had fewer observations, by shifting the window we are more likely to find comparably periods with parallel trends.

Another thing that needed clarifying was how R&D spending and patent citations affected the number of new products and unit costs. Previous literature typically use lags up to 4-6 periods for R&D and patents when looking at their impact on company valuation.[24] We turn to data to find which number of distributed lags offers the best fitting model. This turns out to be the one with up to 5 lags on R&D spending, and up to 3 lags on patent citations, which is what we use in our reported estimates.

[24]Pakes (1981), Pakes and Griliches (1980, 1984a), Wang and Hagedoorn (2014)

Tables 4 shows five key parameters for our 5 treatment times. The sum of the lagged effects of R&D spending and patent activity (in this case we measured patent activity through a factor variable). Both of these are interacted with the treatment, which then give us the sum of the lagged effects of R&D spending and patent factor on the treatment unit following the mergers (sum of R&D lags $\times$ treatment, sum of patent lags $\times$ treatment). The table also shows the models where the pre-treatment parallel trends were rejected (we focus on the results where the parallel trends assumption was not violated).

Table 4: Effect on unit cost and number of new models

	ln(cost)					ln(numbers)				
Treatment time	Q1 2012	Q3 2012	Q1 2013	Q3 2013	Q1 2014	Q1 2012	Q3 2012	Q1 2013	Q3 2013	Q1 2014
Seagate										
DD	0.756	0.417	1.062	**1.090****	**1.071****	-0.233	0.409	-0.107	0.127	0.272
std.err.	(0.260)	(0.596)	(0.124)	(0.046)	(0.026)	(0.233)	(0.381)	(0.795)	(0.688)	(0.240)
Sum of R&D lags	0.282	0.072	0.236	0.304	-0.013	**0.737*****	**0.848****	**0.589***	**0.547****	0.441
pval	(0.556)	(0.914)	(0.74)	(0.682)	(0.988)	(0.006)	(0.021)	(0.053)	(0.042)	(0.237)
Sum of R&D lags x treatment	**-1.110***	-1.192	**-2.434****	**-3.740*****	**-3.552*****	**1.459*****	0.726	**1.307****	**1.146***	**1.201***
pval	(0.073)	(0.284)	(0.046)	(0.000)	(0.000)	(0.000)	(0.291)	(0.04)	(0.062)	(0.075)
Sum of patent lags	0.189	0.041	0.061	0.133	0.091	0.023	0.118	0.029	0.034	0.076
pval	(0.306)	(0.813)	(0.628)	(0.142)	(0.268)	(0.608)	(0.262)	(0.768)	(0.622)	(0.372)
Sum of patent lags x treatment	**-0.566****	-0.312	**-0.739***	**-1.431*****	**-1.058*****	0.124	-0.091	-0.093	-0.049	0.051
pval	(0.039)	(0.311)	(0.089)	(0.003)	(0.002)	(0.139)	(0.438)	(0.603)	(0.862)	(0.878)
observations	171	168	164	161	157	173	170	166	163	159
parallel trend rejected?	Y	Y	N	Y	N	Y	Y	Y	N	N
Western Digital										
DD	**1.534****	0.862	-0.810	-0.565	0.170	**-0.589*****	-0.128	**1.243*****	**-2.286*****	**-2.585*****
std.err.	(0.043)	(0.166)	(0.263)	(0.663)	(0.921)	(0.008)	(0.745)	(0.007)	(0.004)	(0.004)
Sum of R&D lags	0.657	0.348	0.454	0.169	-0.269	**0.436****	**0.475***	0.146	0.24	0.295
pval	(0.160)	(0.614)	(0.563)	(0.852)	(0.779)	(0.011)	(0.079)	(0.595)	(0.260)	(0.412)
Sum of R&D lags x treatment	0.047	0.129	**2.437***	2.075	0.725	**-2.104*****	**-2.419*****	**-4.67*****	**2.917***	**4.132****
pval	(0.932)	(0.895)	(0.070)	(0.45)	(0.804)	(0.000)	(0.008)	(0.000)	(0.079)	(0.014)
Sum of patent lags	**0.267***	0.172	**0.165***	0.160	0.083	-0.029	0.024	-0.071	-0.026	-0.011
pval	(0.072)	(0.235)	(0.089)	(0.105)	(0.42)	(0.406)	(0.795)	(0.345)	(0.469)	(0.833)
Sum of patent lags x treatment	-0.084	0.017	**0.799*****	0.839	0.789	**0.462*****	**0.422*****	**-0.739*****	**2.172*****	**3.106*****
pval	(0.65)	(0.923)	(0.008)	(0.52)	(0.355)	(0.000)	(0.003)	(0.009)	(0.004)	(0.000)
observations	153	150	146	143	139	155	152	148	145	141
parallel trend rejected?	Y	Y	N	N	N	Y	Y	Y	Y	N
Toshiba										
DD	**7.839*****	**14.73*****	**8.553****	-1.651	**32.69*****	**-1.568*****	**-10.45*****	**-7.355*****	**5.947*****	**11.74*****
std.err.	(0.000)	(0.001)	(0.012)	(0.593)	(0.000)	(0.009)	(0.000)	(0.000)	(0.001)	(0.001)
Sum of R&D lags	0.55	0.338	0.474	0.370	0.056	**0.393****	**0.538**	0.337	0.302	0.279
pval	(0.178)	(0.625)	(0.567)	(0.677)	(0.954)	(0.026)	(0.042)	(0.154)	(0.15)	(0.323)
Sum of R&D lags x treatment	**14.057*****	**28.414*****	**15.482****	2.775	**69.961*****	**-2.801*****	**-20.369*****	**-13.334*****	**4.219***	**14.058****
pval	(0.000)	(0.001)	(0.016)	(0.616)	(0.000)	(0.001)	(0.000)	(0.000)	(0.069)	(0.012)
Sum of patent lags	0.169	0.153	0.190	0.246	0.135	0.076	0.104	-0.01	-0.043	-0.064
pval	(0.373)	(0.256)	(0.063)	(0.041)	(0.215)	(0.521)	(0.383)	(0.914)	(0.209)	(0.387)
Sum of patent lags x treatment	**0.745*****	**1.114****	0.402	**3.376*****	**5.875*****	-0.112	**-0.619****	-0.131	**-3.526*****	**-3.778*****
pval	(0.007)	(0.011)	(0.207)	(0.000)	(0.000)	(0.251)	(0.024)	(0.492)	(0.000)	(0.000)
observations	149	146	142	139	135	151	148	144	141	137
parallel trend rejected?	Y	N	Y	Y	Y	Y	N	Y	Y	Y

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4 shows that in general R&D spending increases the number of new products but found no impact on unit costs. We found no similar evidence for patent activity. Seagate improved both in terms of unit costs and in the number of new products after the merger. These results are robust to our choice of treatment time (although its significant changes) even where the parallel trend assumption is rejected. For WD the evidence that unit costs are increasing appears in all but the last model. The effect on the number of new WD drives is ambiguous. For Toshiba there is some evidence of dropping unit costs and fewer new products. The latter seems robust to our choice of treatment time.

6 Conclusion

This paper offered a rare opportunity to examine three levels of innovation: R&D spending, patent activity, and the characteristics of new products. We used this for two main objectives. On the one hand this unique dataset provided a novel evaluation of the relationship between competition and innovation, and offered evidence that increasing concentration (and a reduction in the number of competitors) did not lessen innovation in the HDD market. Our interpretation is that this is due to the strong competitive pressure exerted on HDD manufacturers from the SSD market. On the other hand, the breadth of the data allowed us to estimate the relative performance of R&D spending and patent activity data in predicting changes in innovation. We found that - at least in the HDD market - R&D expenditure is a good proxy for innovation, but patent activity offers little explanatory power. These findings are robust to a large number of different model specifications, control groups, and study designs.

The European Commissions 2017 decision on the Dow and Dupont merger has triggered a lively debate among academics and practitioners. At the heart of the debate is the Commissions new innovation based theory of harm, which lead to the conclusion that the merger would have lessened the merging firms incentive to spend on R&D, which in turn would have led to a reduction in the number of new pesticide products. The theoretical underpinnings of this theory of harm are provided in Federico et al. (2017), which posits that horizontal mergers can be expected to reduce innovation incentives as a result of a standard unilateral effect. We found no evidence to support the general applicability of this claim.

This paper also demonstrated the difficulty of claiming a one-size-fits-all relationship between competition and innovation. The three HDD manufacturers responded differently to the market consolidation. Quantitative studies like this one are useful but a key lesson is that they are often not enough. To identify what is causing the effects estimated in these quantitative studies one would need more information, which could be acquired with case specific qualitative studies (for example interviews) on each firm.

Finally, by showing that it is possible to estimate retrospectively the impact of mergers and acquisitions on innovation, we hope that this paper will be followed by a number of similar papers in other industries, similar to the wide-spread nature of papers on the price impact of consolidation.

7 References

References

Abadie, A., A. Diamond, and J. Hainmueller (2010): "Synthetic control methods for comparative case studies: Estimating the effect of Californias tobacco control program," *Journal of the American Statistical Association*, 105(490), 493–505.

Abadie, A., and J. Gardeazabal (2003): "The economic costs of conflict: A case study of the Basque Country," *American Economic Review*, 93(1), 113–132.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005): "Competition and innovation: An inverted-U relationship," *The Quarterly Journal of Economics*, 120(2), 701–728.

Aghion, P., M. Dewatripont, and P. Rey (1999): "Competition, financial discipline and growth," *The Review of Economic Studies*, 66(4), 825–852.

Aghion, P., C. Harris, and J. Vickers (1997): "Competition and growth with step-by-step innovation: An example," *European Economic Review*, 41(3-5), 771–782.

Bertrand, O. (2009): "Effects of foreign acquisitions on R&D activity: Evidence from firm-level data for France," *Research Policy*, 38(6), 1021–1031.

Blundell, R., R. Griffith, and J. Van Reenen (1995): "Dynamic count data models of technological innovation," *The Economic Journal*, pp. 333–344.

Blundell, R., R. Griffith, and F. Windmeijer (2002): "Individual effects and dynamics in count data models," *Journal of Econometrics*, 108(1), 113–131.

Boone, J. (2000): "Competitive pressure: the effects on investments in product and process innovation," *The RAND Journal of Economics*, pp. 549–569.

Coad, A., and R. Rao (2008): "Innovation and firm growth in high-tech sectors: A quantile regression approach," *Research Policy*, 37(4), 633–648.

Danzon, P. M., A. Epstein, and S. Nicholson (2007): "Mergers and acquisitions in the pharmaceutical and biotech industries," *Managerial and Decision Economics*, 28(4-5), 307–328.

Gilbert, R. (2006): "Looking for Mr. Schumpeter: Where Are We in the Competition–Innovation Debate?," *Innovation Policy and the Economy*, 6, 159–215.

Gilbert, R. J., and D. M. Newbery (1982): "Preemptive patenting and the persistence of monopoly," *The American Economic Review*, pp. 514–526.

Gompers, P., J. Lerner, and D. Scharfstein (2005): "Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999," *The journal of Finance*, 60(2), 577–614.

Griffith, R., R. Harrison, and J. Van Reenen (2006): "How special is the special relationship? Using the impact of US R&D spillovers on UK firms as a test of technology sourcing," *American Economic Review*, 96(5), 1859–1875.

Griliches, Z. (1998): "Patent statistics as economic indicators: a survey," in *R&D and productivity: the econometric evidence*, pp. 287–343. University of Chicago Press.

Griliches, Z., A. Pakes, and B. H. Hall (1986): "The value of patents as indicators of inventive activity," .

Hall, B. H., and R. H. Ziedonis (2001): "The patent paradox revisited: an empirical study of patenting in the US semiconductor industry, 1979-1995," *Rand Journal of Economics*, pp. 101–128.

Igami, M., and K. Uetake (2017): "Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996-2016," .

Lanjouw, J. O., and M. Schankerman (2004): "Patent quality and research productivity: Measuring innovation with multiple indicators," *The Economic Journal*, 114(495), 441–465.

Levin, R. C., A. K. Klevorick, R. R. Nelson, S. G. Winter, R. Gilbert, and Z. Griliches (1987): "Appropriating the returns from industrial research and development," *Brookings Papers on Economic Activity*, 1987(3), 783–831.

ORMOSI, P. L., A. BENNATO, S. DAVIES, AND F. MARIUZZO (2017): "A feasibility study on the microeconomic impact of enforcement of competition policies on innovation," *European Commission*.

ORNAGHI, C. (2009): "Mergers and innovation in big pharma," *International Journal of Industrial Organization*, 27(1), 70–79.

REINGANUM, J. F. (1989): "The timing of innovation: Research, development, and diffusion," *Handbook of Industrial Organization*, 1, 849–908.

SCHMIDT, K. M. (1997): "Managerial incentives and product market competition," *The Review of Economic Studies*, 64(2), 191–213.

SCHUMPETER, J. A. (1934): "Theorie der Wirtschaftlichen Entwicklung. Leipzig: Dunker & Humblot, The Theory of Economic Development: An Inquiry into Profits, Capital, Interest, and the Business Cycle, translated by Redvers Opie," .

SCHUMPETER, J. A. (1942): *Socialism, capitalism and democracy*. Harper and Brothers.

SHAPIRO, C. (2011): "Competition and innovation: Did Arrow hit the bull's eye?," in *The rate and direction of inventive activity revisited*, pp. 361–404. University of Chicago Press.

STIEBALE, J., AND J. HAUCAP (2013): "How Mergers Affect Innovation: Theory and Evidence," .

STRUMSKY, D., J. LOBO, AND J. A. TAINTER (2010): "Complexity and the productivity of innovation," *Systems Research and Behavioral Science*, 27(5), 496–509.

SZÜCS, F. (2014): "M&A and R&D: Asymmetric effects on acquirers and targets?," *Research Policy*, 43(7), 1264–1273.

WOOLDRIDGE, J. M. (2010): *Econometric analysis of cross section and panel data*. MIT press.

8 Online appendix

A Further tests on the R&D expenditure estimates

A.1 Evaluating the Control groups

Below we plot R&D intensity against these Control groups. Figure 3 commences with WD. The two vertical lines mark the start and end of the merger procedure – this period was excluded from the analysis as explained in the main text. Visually, the synthetic control seems to be most similar to the Treatment group in terms of pre-merger R&D intensity.

Looking at the plotted R&D intensity values, our visual conclusion of the evolution of R&D is that WD's R&D intensity grows faster than the Control's. However, this seems to have started before the mergers, and were therefore less likely to have been caused by them.

Figure 3: R&D intensity plot for WD and four different Control groups

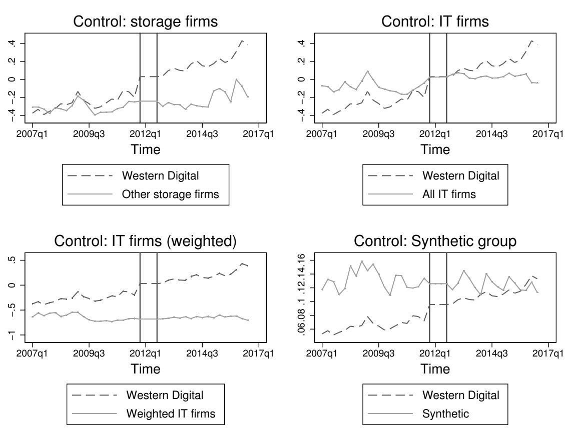

Figure 4 shows the Control groups for Seagate. Again, in terms of pre-merger similarity, the synthetic control performs best. What Figure 4 suggests is that Seagate's R&D intensity moved around the same level as the Treatment group pre-merger. Post-merger there is a higher level of growth for Seagate than for the Control groups. We will test this formally, but in any case, this would suggest that something happened between Q3 2011 and Q2 2012, which lead to Seagate increasing its R&D intensity.

Finally, we look at Toshiba on Figure 5. The Treatment line shows a jump in 2009 when Toshiba acquired Fujitsu's HDD operations. After the merger the Toshiba line seems to go together with the Control line. The figures show that the matching process was not as effective as for WD or for Seagate. We will formally test this later.

Figure 4: R&D intensity plot for Seagate and four different Control groups

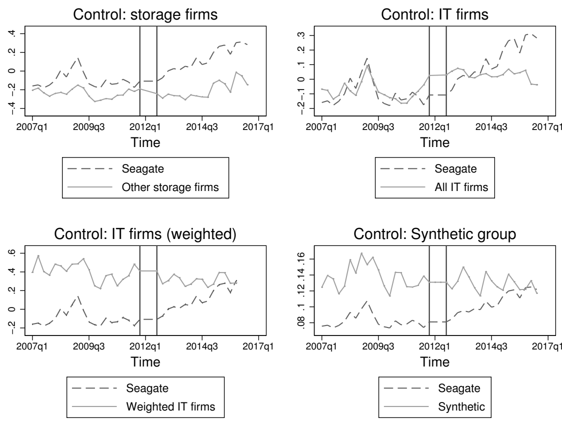

Figure 5: R&D intensity plot for Toshiba and four different Control groups

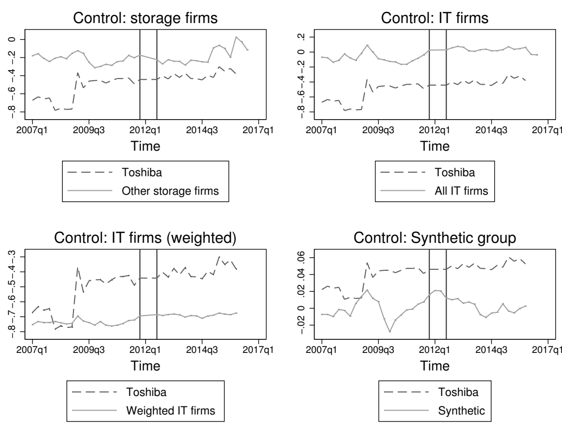

A.2 Evaluating the assumptions

A.2.1 No serial correlation

If there is positive serial correlation in the R&D intensity data, then the standard errors of the above coefficient estimates will be lower than the unbiased standard errors. This would imply that the effect of the mergers might be found significant even when, in an unbiased model, it would not be. Similarly, negative serial correlation in the price data may overestimate the standard error of the merger effect. We used Wooldridge's (2002)Wooldridge (2010) autocorrelation test for panel data.

A.2.2 Independence

A spill-over effect occurs where the effect of the treatment spills into the Control group. This may be problematic in markets with strategic interaction, as typically are those studied in most of the merger literature. This is more difficult – if at all possible – to test formally. If there is a spill-over effect, the sign of the bias will depend on whether the Treatment and Control groups are complements or substitutes in innovation. If it is the former, then the estimates will be downward biased because an increase in innovation in the Treatment group is followed by an increase in innovation in the Control group. Therefore the real effect is likely to be higher than the estimated effect. If they are substitutes, then the bias will be upwards, and therefore it will be more difficult to decide how it would affect the estimates without knowing the magnitude of the bias. It is clear that in the former case, the researcher still gets useful information out of the estimates even if they are biased.

It is possible that there was a spill-over effect into other parts of the storage market (SSD and/or other Flash), which is our first Control group. However, we offer three other Control groups (unweighted and weighted IT firms, and a synthetic control) based on the assumption that it is very unlikely that the Treatment affected non-storage product markets. This is why we chose a sample of IT firms, as the independence assumption is much less likely to be violated for this Control group. The similarity assumption might be more of an issue for this case, which we test by looking at parallel trends.

A.2.3 Parallel trends

For DiD to provide unbiased estimates one would need Treatment and Control to follow parallel trends in the absence of the merger. Obviously, we do not observe the Treatment group without the merger after 2012. For this reason we can only test whether the parallel trend exists before the merger. Figures 3, 4, and 5 above provide a first visual test. For a formal test we look at annual deviations from parallel trends in the pre-merger data. The intuition is that if the vertical distance between the two trendlines significantly changes in any year, it would be a violation of the parallel trend assumption. To run a formal test we look at pre-merger R&D intensity data, and estimate a fixed effects model with yearly dummies, and interactions between the yearly dummies and the treatment. If the pre-merger trends are parallel, then the interaction coefficients (te x 2008-2011) should be jointly non-significant. The results are reported in Table 5. We run the tests for the three Control groups: (1) other storage firms, (2) IT firms, and (3) weighted IT firms. We also apply a less restrictive test. Because our DD model compares before and after means, it would suffice to test if the linear approximation of pre-merger trends are parallel. This is reported as parallel trend in Table 5.

Table 5: Testing the parallel trend assumption

Control	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
		Seagate			Western Digital			Toshiba	
te2008	0.169***	0.112***	0.148	0.144***	0.0433***	0.176**	-0.144***	-0.163***	-0.117***
p-val	(0.000)	(0.000)	(0.150)	(0.000)	(0.002)	(0.041)	(0.000)	(0.000)	(0.000)
te2009	-0.0247	-0.132	-0.177	0.206***	0.0797***	0.359**	0.207***	0.130***	0.178***
p-val	(0.818)	(0.103)	(0.541)	(0.007)	(0.000)	(0.016)	(0.000)	(0.000)	(0.000)
te2010	0.0173	0.0796***	0.186	0.259***	0.168***	0.362**	0.271***	0.244***	0.211***
p-val	(0.787)	(0.000)	(0.142)	(0.001)	(0.000)	(0.025)	(0.000)	(0.000)	(0.000)
te2011	-0.0521	0.0724***	0.248	0.322***	0.232***	0.429***	0.200**	0.202***	0.195***
p-val	(0.441)	(0.000)	(0.209)	(0.001)	(0.000)	(0.001)	(0.020)	(0.000)	(0.000)
F-test of joint significance	17.276	13.852	1.101	9.438	48.236	14.402	22.335	173.219	406.911
p-val	0.000	0.000	0.378	0.000	0.000	0.000	0.000	0.000	0.000
Parallel trend	0.004	0.006***	0.009	0.014**	0.015***	0.013	0.018***	0.021***	0.020***
p-val	(0.493)	(0.001)	(0.288)	(0.013)	(0.000)	(0.118)	(0.000)	(0.000)	(0.000)
Observations	509	27720	734	571	27782	772	558	27769	735

Table 5 shows the estimated coefficients (full regression results are available from the authors). There is no evidence that there was a deviation from parallel trend for Seagate for control groups (1) and (3). For Western Digital the parallel trend assumption is not violated under our less restrictive test for Control group (3). For Toshiba however, as our visual analysis has already suggested, the parallel trends assumption is violated for virtually all 4 pre-merger years.

A.3 Robustness checks

We estimate treatment effects using four different Control groups and find that the estimates are robust to changes in the composition of the Control. We offer three more robustness checks.

A.3.1 Placebo treatment

First we tested whether a placebo Treatment group returns significant treatment effect. We used the total sample of IT firms and re-run the DiD model assuming in each iteration that another firm was the 'Treatment'. With each iteration we generated a new weighted sample (matching the 'Treatment' firm) and then estimated the treatment effects. The idea is that if our treatment effect for Seagate is a fluke then we would find a large number of other firms producing similarly significant treatment effects. On the other hand, if the other firms did not receive the same treatment as Seagate then there would only be a small proportion of firms with statistically significant positive treatment effects.

This resulted in a sample of 1701 'Treatment effects'. Less than 15% of these produced results similar to Seagate's (positive and statistically significant treatment effect). Given the large number of firms (some in very different IT markets) this is a very good piece of evidence for two reasons:

- It shows that at most there are only few confounding effects, i.e. there were unlikely to be any other major shocks in Q1-Q2 2012 that would have affected IT firms that same way the merger affected Seagate.
- More importantly, even where estimates for other firms were also significantly different from zero, they were evenly spread between negative and positive values. Therefore when the pool of IT firms is used as a control, even when there are other firms in the sample that reacted to something in 2012, the sign of these reactions cancelled each other out in their total effect, therefore our choice of using weighted or unweighted IT firms as Control is a good one and should provide unbiased results.

Table 6: Placebo treatment times

treatment time	Q1 2007	Q1 2008	Q1 2009	Q1 2010	Q1 2011
Seagate	-0.116	-0.0675	-0.104	0.0631	0.0788
(p-val)	(0.346)	(0.570)	(0.293)	(0.460)	(0.447)
n	770	738	693	639	570
WD	-0.00144	0.133	0.195	0.232*	0.192***
(p-val)	(0.988)	(0.163)	(0.058)	(0.046)	(0.000)
n	843	803	759	706	637
Toshiba	0.133***	0.174***	0.224***	0.153***	0.0505***
(p-val)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
n	760	739	704	663	601

We also tested for placebo Treatment times. This involves checking what happens if we assume that Treatment (mergers) happened in a different year before the merger. We re-run our regressions for five different pre-merger years and for two different dependent variables (quarterly change, annual change in R&D intensity) using the weighted IT Control. We used pre-Q3-2012 data as we did not want the actual merger effects confound the placebo effects.

Table 6 shows the effect of these placebo Treatment times. The results are good for Seagate and for WD but for Toshiba we estimated significant placebo effects. This would reiterate our previous stance that our identification strategy did not work for Toshiba.

A.3.2 Different matching assumptions

As explained above, in one of the models we matched IT firms with each of the Treatment firms and used a weighted sample of the IT firms that were most similar based on observed characteristics. In the matching exercise we matched with the Treatment firms the 30 most similar (nearest neighbours) and acquired their weights. To see whether our choice of 30 firms affected the results, Table 7 shows the DiD estimates for Seagate and Western Digital, under different matching assumptions. [25] The table shows that the results were not sensitive to the choice of the number of matched firms.

Table 7: Treatment effects under different matching assumptions

	number of nearest neighbours				
	20	25	30	35	40
Seagate	0.475***	0.340**	0.339***	0.327***	0.340***
(p-val)	(0.009)	(0.024)	(0.007)	(0.004)	(0.000)
n	539	697	837	926	1085
WD	0.268*	0.208	0.173*	0.169**	0.165**
(p-val)	(0.082)	(0.120)	(0.053)	(0.041)	(0.036)
n	579	726	886	1006	1105

B Further tests on the patent results

We tested for *serial correlation* in all models. Models and Control groups where serial correlation could not be rejected were filtered out.

Where the Control group was other firms' HDD patents, the *independence* assumption would mean that the merger only affected HDD producers' patent activity, and not the HDD patent activity of producers of other goods as well. In this Control, firms produce goods that are complementary to HDD. There is a viable argument that when HDDs improve through innovation, they will trigger complementary goods also to boost their innovation. If innovation manifests in new technologies, complementary goods will have to innovate to link to these new technologies. For this reason it would seem credible that if the mergers increase innovation in HDDs, it would trigger an increase in innovation in complementary goods – although this may come with a time lag. This would mean that the estimated effect would be biased downwards. As we are not particularly concerned about the magnitude of the effect, rather than its sign, this is sufficient for us to conclude that a positive effect remains positive even after eliminating the bias.

Another Control group is the Treatment firms' Flash related patents. Here there might be some spill-over effects. Increased R&D spending may contribute to an increase in innovation in both HDD and Flash (including SSD). This would result in a downward bias. Finally, the

[25] We omitted Toshiba as we have rejected the reliability of those results above.

Control group with the top NAND Flash patenting firms is probably where spill-over effects are less likely. These include firms and patents that are on a different product market.

To test *parallel trends* we used a different assumption to the R&D section. Patent data is different from R&D spending. It is very often the case that a firm active in patenting one year, files no patent in the following year. This is related to the nature of the discovering process, the path of which is very difficult to predict. This is compared with a Control group of many firms, for which the distribution of patents is smoothened out over time. If one looked at annual deviation from the parallel trend, we would inevitably pick up the deviations caused by the volatility of firm-level patent data. To remedy this, we assume a linear pre-merger trend for both the Treatment and the Control groups and test if these linear trends are parallel. In synthesising the estimates from all our models, we only kept the ones where we could not rejected the parallel trend assumption.

Figure 6 shows the evolution of patenting for HDD, all Flash Memory, SSD, and USB Flash Drives. The latter two categories are sub-sets of Flash Memory, which also contains other technologies based on Flash Memory, for example DRAM. Unsurprisingly it stands out that HDD is a more mature technology than SSD or USB Flash Memory. HDD patenting peaked in 2005 then had a small decline and has stabilised on a relatively steady path (due to the time lag in updating the patent office registers, 2015 and 2016 data are not complete). On the other hand SSD patenting really picked up in 2008, peaked in 2013, and dropped in 2014. Similarly, USB Flash patenting increased until 2013 and dropped in 2014. It appears that SSD and USB Flash alone follow an altogether different innovation trajectory. On the other hand, the sample of All Flash Memory patents might satisfy parallel trend assumptions.

Figure 6: Number of HDD, Flash Memory, SSD, and USB Flash patents per year

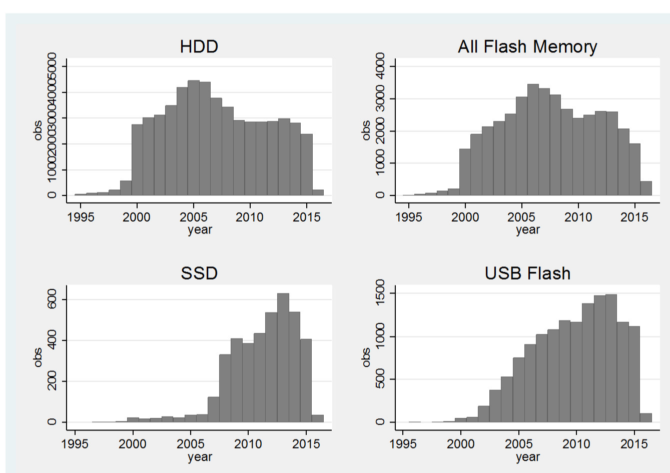

We deliver results for a large number of different model specifications, and Control groups, which alone acts as an extensive robustness check. However, as a final step we construct a patent indicator that brings together the richness of patent data into one variable, we construct a patent indicator by following an approach similar to the multiple-indicator factor model in Lanjouw and Schankerman (2004). We make use of a complete set of variables collating information on patent counts, patent citations (distinguishing citations from attorneys and from the literature), patent inventors (number), patent claims (number), patent applications (number) and application countries (number). However, in contrast to their paper we choose to utilise factor analysis as the methodology in order to reduce the number of patent-related correlated variables. The justification for using this methodology (instead of principal component analysis) is that we have a set of original variables that together contribute in explaining innovation, while all those variables on their own would have limited contribution and be subject to criticism.

Table 8 shows the DD estimates for each Treatment firm, using a factor variable. The results are qualitatively the same as in our headline table.

Table 8: Effect of the mergers using a factor variable

		Other HDD	Other Flash	Top storage
Seagate	Coeff	0.594***	0.443*	0.541***
	Std.err	0.069	0.207	0.107
	Obs	317	319	728
WD	Coeff	0.560***	0.451*	0.547***
	Std.err	0.069	0.207	0.107
	Obs	317	319	728
Toshiba	Coeff	-0.308***	-0.458*	-0.407***
	Std.err	0.069	0.209	0.109
	Obs	317	319	728

C Further tests on product characteristics

C.1 Assumptions required for unbiased DD estimates

Figure 7 shows how the number of newly marketed drives changes for the Treatment firms and for all SSD firms. As previously with the patent data, the data is highly volatile, this time due to the fact that firms often market products in clusters, therefore some calendar quarters might have a high number of new products appearing on Amazon, and some others, none. However, if this volatility is random across the two trends (HDD and SSD) that are otherwise parallel, then the DiD estimator should be unbiased. We will test this formally later.

Figure 7: Per-firm quarterly average number of new products marketed on Amazon

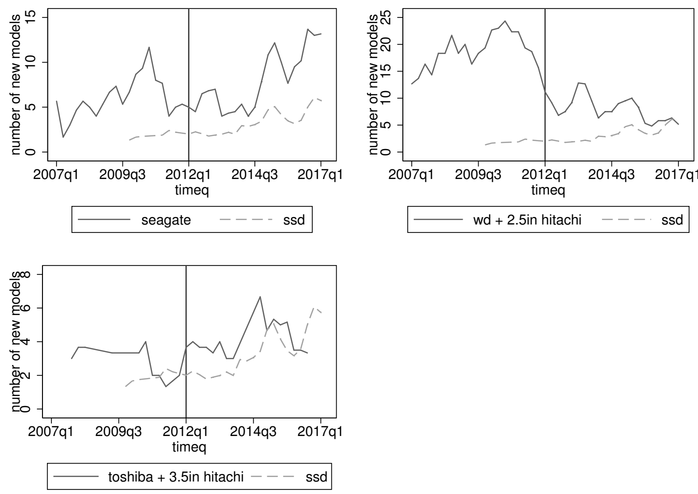

It immediately stands out from Figure 7 that SSD's only appear in our sample from 2009, and especially at the beginning the per-firm average number of SSDs was very low. Figure 7 shows that for Seagate there was an increase in the number of new drives marketed roughly 2 years after the merger. This is important because this could be an indication of the length of lag between R&D spending and its effect on production. A similar (but much less pronounced) increase can be seen for Toshiba. For WD there has been a drop in the number of new HDDs marketed on Amazon.

Figure 8 shows that there has been a continuous decrease in the unit price of SSD capacity. HDD on the other hand displays a mixed picture. Unit capacity price has been steadily falling for Seagate and WD (steeper for Seagate), and fell first then levelled out for Toshiba.

We tested separately whether the Treatment and Control follow a parallel trend before the treatment(s). We found that out of our 5 treatment events, the first two estimates are likely to be biased because pre-treatment trends were not parallel (this is visually confirmed on Figure X and Y). This would allow us to use the other 3 models. However, as shown above, the main story here does not hinge on our DD estimate. Rather, it is about the effect of previous R&D spending on product numbers and unit prices. This is also important regarding the independence assumption required for unbiased DD, because, strictly speaking, in this respect even the choice of our Control group is irrelevant here. To illustrate why, take the example of Seagate. For our R&D spending estimates in Section X we had a better selection of Control groups and there we have shown how the mergers increased R&D spending. Here we show that this increased R&D activity is associated with an increased

Figure 8: Quarterly lowest price of unit capacity - Treatment firms against SSD

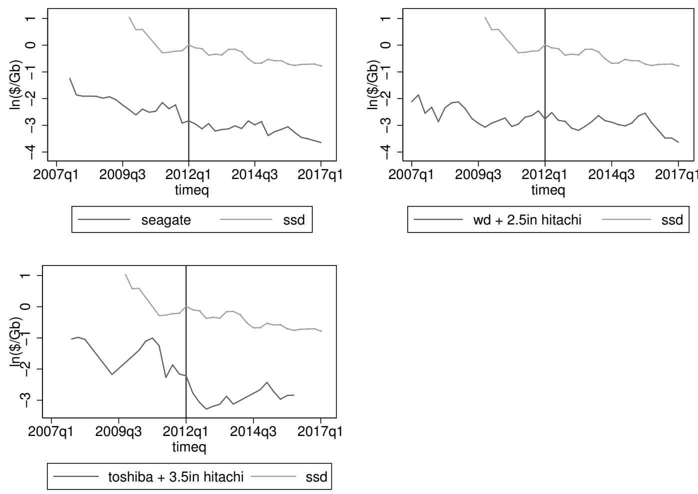

number of new products and lower unit prices.

We tested for serial correlation. In general, using logs of the dependent variable eliminated serial correlation (at least when using Wooldridge's (2002) test for serial correlation in panel data).

We did do some simple robustness checks within the possibilities given by our data. We re-run the above regressions for two slightly different Control groups. The first one only included the 5 largest SSD producers (in terms of number of SSDs marketed). These are firms that are more comparable in size to the Treatment firms. In another experiment we took the Treatment firms' SSD production as Control (Samsung and Toshiba are also active in SSD). The intuition is that if the 2012 HDD mergers affected HDD innovation, it might not have triggered the same response in SSD innovation.