Mergers and innovation: Evidence from the Hard Disk Drive market*

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

September 19, 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

This case study is a relatively rare ex-post evaluation of the impact on innovation of two mergers in the world-wide hard disk drive (HDD) industry. We take a holistic view of innovation, employing four different measures: R&D expenditure and patent activity as indicators of innovative inputs, and the number of new products marketed, and their unit costs for end users as indicators of innovative output. This allows us to distinguish the impact of the mergers on the magnitudes of the parties' innovative efforts from the productivity of those efforts. We employ a differences-in-differences methodology, placing considerable emphasis on the robustness of results to alternative control groups. Our main finding is that the mergers had a beneficial effect. They were associated with significantly increased R&D and patent activity by the two acquiring firms (Seagate and Western Digital), and increased innovative productivity for Seagate and less robustly for the Western Digital. These results appear to vindicate the European Commission's decision to allow both mergers.

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

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 find 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

Our paper draws on various literatures. First, there is the enormous general literature on the relationship between competition and innovation, usually traced back to the seminal works of Schumpeter (1934, 1942) and Arrow (1962). This is sufficiently well known not to bear repetition here, and there are many excellent reviews, including Gilbert (2006, 2010).[1] A theme running through some of this literature is that the relationship may be characterised by an inverse U-shape: a way of reconciling Arrow and Schumpeter - increases in competition initially raise the pressure to innovate but after some point, further increases reduce the incentive, unless property rights are protected. Shapiro (2011, p.401) summarises succinctly: "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" (resonating with Arrowian arguments), but "Schumpeter was also quite correct: the prospect of obtaining market power is a necessary reward to innovation". On the implications for how mergers might impact

[1] Some of the most important contributions include Gilbert and Newbery (1982); Levin, Klevorick, Nelson, Winter, Gilbert, and Griliches (1987); Reinganum (1989); Blundell, Griffith, and Van Reenen (1995); Aghion, Harris, and Vickers (1997); Schmidt (1997); Aghion, Dewatripont, and Rey (1999); Boone (2000); Hall and Ziedonis (2001); Gompers, Lerner, and Scharfstein (2005); Aghion, Bloom, Blundell, Griffith, and Howitt (2005); and Griffith, Harrison, and Van Reenen (2006).

on innovation, sweeping generalisations are not justified especially if we recognise that the motives for particular mergers may be very different, e.g. in some cases to dampen competition by securing coordinated effects, but in others to sharpen competition through efficiency savings.

Moving from the general to the specific, a cluster of theoretical models have appeared in very recent years, with the balance suggesting a negative impact of mergers on innovation. Focusing on product innovation, Federico, Langus, and Valletti (2017) and Federico, Langus, and Valletti (2018) identify two effects of the merger: 'price coordination' which tends to favour innovation, and the internalization of the "innovation externality" which depresses innovation. In numerical simulations, they find that the latter is stronger, and thus a merger leads to lower innovation incentives, absent cost efficiencies and spillovers. However, Denicolò and Polo (2018) show how this prediction can be overturned if the positive effect of duplication avoidance is particularly pronounced when there are asymmetries in the R&D intensities of the parties. But, then again, Haucap and Stiebale (2016) present a model which shows that, with a high degree of firm heterogeneity, the merger reduces innovation of both the merged entity and its non-merging competitors in an R&D intensive industry. Turning to process innovation, Motta and Tarantino (2017) employ a model with simultaneous price and cost-reducing investment choices and also find that, absent efficiency gains, the merger lowers total investments. Finally, in the synthesis of their own and others' work, Bourreau, Jullien, Lefouili, et al. (2018) suggest that the overall impact of a merger on innovation may be either positive or negative.

In the empirical literature, Danzon, Epstein, and Nicholson (2007), Ornaghi (2009), and Haucap and Stiebale (2016) all find robustly significant negative impacts of mergers on innovation in the pharmaceutical sector. Szucs (2014) finds that target firms substantially decrease their R&D post merger, and that 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 Haucap and Stiebale (2016).[2]

Finally, in a recent paper most closely related to ours, Igami and Uetake (2017), focus on the HDD market, estimating a dynamic oligopoly model, in which merger decisions are endogenous. By employing hypothetical merger policies as counterfactuals, they show that the optimal policy should block mergers if there are 6 or fewer players in the HDD market.

On the question of measurement, throughout most of the literature, 'innovation' is typically represented by R&D intensity or patents, classic examples include Griliches (1979),

[2]In another study, not directly on mergers, but still relevant, Genakos, Valletti, and Verboven (2018) find that, in the mobile phone industry there is evidence of a larger R&D investments per operator but not at the aggregate industry level in concentrated markets in OECD countries, 2002-14.

Griliches (1990), and Scherer (1983). There have been some advocates of composite measures, e.g. Hagedoorn and Cloodt (2003), construct such a composite but find that "the statistical overlap between these indicators is that strong that future research might also consider using any of these indicators to measure the innovative performance".[3]

On the other hand, there are some strong reasons for caution in interpreting data on both patents and R&D. It is self-evident that, given inevitable technical uncertainty, much R&D will fail to generate any innovation. But maybe less obviously the same is true for patents. Moreover, it has long been recognised that patents are sometimes used to protect an incumbent's market power (see among others Cohen, Nelson, and Walsh (2000), Cohen, Goto, Nagata, Nelson, and Walsh (2002), Gilbert and Newbery (1982)), and likewise, some R&D is therefore essentially defensive devoted to findings ways of denying innovations to others.[4] It is also true that many innovations are not the result of formal R&D. This was first established many years ago by Jewkes's seminal book on twentieth century innovations (Jewkes, Sawers, and Stillerman (1958)) in which he reports that "more than one-half of the cases can be ranked as individual invention in the sense that much of the pioneering work was carried through by men who were working on their own behalf without the backing of research institutions". More recently, focusing on a group of low- and medium technology industries in Spain, Santamaría, Nieto, and Barge-Gil (2009) also show how many activities that lead to innovation are not R&D-based.

Finally, a recent study of patent statistics particularly relevant to our own case study, Igami and Subrahmanyam (2015), warns "researchers to use caution when comparing patents of different types of firms and across years."

The lessons we draw from this review are as follows. First, there are arguments both ways on whether, in general, mergers encourage or discourage innovation. This is not to deny the value of aggregate empirical studies which attempt to draw estimates of average cross-industry effects. Some might argue that the balance of theory and empirical evidence points to a generally negative effect.[5] Here, however, we are more interested in the detail of a specific case study: what happened in the wake of consolidations from 5 to 3 firms in the HDD market? We prefer not to bring strong priors on whether or not these acquisitions

[3]See also Janger, Schubert, Andries, Rammer, and Hoskens (2017) for analysis of a composite measure advocated by the European Commission.

[4]A more recent paper, Blind, Cremers, and Mueller (2009) show the importance of strategic patenting also in improving a firm's reputation, giving it greater bargaining power in negotiations with other firms. Moreover, strategic patent can be used to create internal incentives for their R&D employees, and to measure their performance.

[5]This might apply to the European Commission which, in its most recent guidelines (http://ec.europa.eu/competition/publications/cpn/), warns against a merger if it combines two important innovators, or eliminates a firm with promising pipeline products, "because the transaction can eliminate an important competitive force and thus lead to a significant impediment of effective competition".

stimulated innovation.

Second, because it is likely that predictions will be very sensitive to the assumptions on which any theoretical model is based, we will not employ a structural model of firm behaviour (as used for example in Igami and Uetake (2017) as discussed above.) Instead we prefer a reduced-form Difference-in-Differences methodology. In this way we avoid introducing restrictive assumptions about the nature of competition in the HDD market. However, because this methodology is essentially atheoretical in its priors, it is essential that there is meticulous selection of a set of control groups to identify the impact of the consolidations. Third, we will not limit our analysis to any one specific measure of innovation, but examine four different measures separately and their interaction: R&D expenditure, patent activity, the number of new products taken to market, and the unit cost to users of those new products. Contrary to much of the previous literature, we do not view either R&D or patent counts as direct measures of innovation outputs, rather they are inputs. As such we investigate not just the impact of the mergers on the magnitudes of R&D and patent activity, but also on their 'productivity' in generating marketed innovations (in the form of new products and their cost to users.) Fourth, because we focus on these various methodological and measurement issues, we believe that the paper provides a practical blueprint for future, much needed, policy evaluations in other innovative industries.[6]

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.

[6]For example, in the competition policy literature, there are scores of merger evaluations in terms of the impact on price but scarcely any on innovation Davies and Ormosi (2012).

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 product 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.[7]

3 Model and data

Our primary objective is to assess the impact of the two mergers on the innovative output of the post-merger entities. We allow for three not mutually exclusive possibilities. The mergers might lead to (1) changes in the levels of innovation inputs, (2) changes in the productivity of those inputs, and/or (3) other, firm specific, reasons, unrelated to the inputs.

3.1 A two-stage framework

The above possibilities are captured by the following simple framework, in which we envisage an innovation production function that relates, for each firm j and time period t, the innovation output (y_{jt}) to research effort or innovation inputs (pat_{jt} and rd_{jt}).

Each of the innovation inputs are associated with a set of observables, x_{kjt}, and unobservable firm characteristics μ_{kj}, along with time shocks, μ_{kt}, with $k = \{1, 2\}$. The relationship between the two innovative inputs and these variables can be expressed as linear panel regressions

$$rd_{jt} = x_{1jt}\phi_1 + \mu_{1j} + \mu_{1t} + \varepsilon_{1jt} \tag{1a}$$

$$pat_{jt} = x_{2jt}\phi_2 + \mu_{2j} + \mu_{2t} + \varepsilon_{2jt}, \tag{1b}$$

where the idiosyncratic error term in each equation, ε_{kjt}, is expected to be uncorrelated with the set of independent variables, after controlling for firm unobserved effects, μ_{kj}. That is, we assume sequential exogeneity.

The two sources of innovative inputs contribute to innovative output y_{jt}, together with a set of control variables, x_{jt}, and unobservables (the μs). Innovative output is also modelled as linear panel regression:

$$y_{jt} = \underbrace{\lambda_1 rd_{jt} + \lambda_2 pat_{jt}}_{\text{innovative inputs}} + x_{jt}\phi + \mu_j + \mu_t + \varepsilon_{jt}, \tag{2}$$

[7] For more on the regulatory background, see Ormosi et al. (2017).

where μ_j captures unobserved firm level differences in innovation productivity, and μ_t controls for time shocks. The assumption of sequential exogeneity is maintained, but this time it implies also that the idiosyncratic error term of innovative output, ε_{jt}, is uncorrelated both with the idiosyncratic error terms of R&D, ε_{1jt}, and with that of patents, ε_{2jt}. For initial presentational simplicity, the above equations are specified as linear without lags and without merger effects.

This two stage structure reflects a change in emphasis compared to much of the previous literature discussed above, which typically employs R&D and/or patents as self-standing measures of innovative performance (or sometimes R&D is used to represent inputs into patenting). Instead, we prefer more direct measures of innovation, and explore how R&D intensity and patents impact on that output.

Our preference for this approach recognizes the traditional critiques: (i) R&D does not always lead to fruitful outcomes; (ii) not all patents ultimately convert into innovation brought to market; and (iii) patents may often be used as a strategic defensive device to close down foreclose or hinder rival innovation.

3.2 Capturing the impact of the mergers: the econometric model

Our econometric strategy is a difference-in-differences approach applied to Eqs.(1a), (1b) and (2). This requires a carefully constructed Control group. We use data on innovative output (the user unit cost of HDDs and the number of newly marketed products), and input (R&D intensity, and patents) for each calendar quarter t, from Q1 2007 to Q4 2016. Our study period spans over two equal periods pre, and–post merger.[8] Of these $T = 40$ 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 $j \in \{1, \ldots J_0, \ldots, J_0 + J_1\}$.

Denote by x_{jt} a $(K \times 1)$ vector of time-varying firm characteristics. We only include lagged values of these characteristics $(x_{jt-\{1,\cdots,\tau\}})$ 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.

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

To estimate the impact of the merger on our measures of innovation, we would need to model what would have happened to these measures in the absence of the mergers. For this, the choice of the Control group is key to the correct identification of the merger effect. For unbiased estimates the Control has to be sufficiently similar to the Treatment group, and independent of the treatment event. If the Control is not similar enough (i.e. affected by different demand or supply side shocks) then our estimates would also include confounding effects. 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.

If the Control is not independent then the effect of the mergers is spilled over to the Control group. Because there is some substitutability across storage technologies, innovation decisions in one product might trigger a response in the other. This would make 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.

Whereas the independence assumption is difficult to formally verify (and one often relies on economic intuition), the similarity assumption is conventionally tested by looking at the Control and Treatment trends pre-merger. Deviation from parallel trends would imply the presence of a confounding factor that affects either the Control or the Treatment but not both. In our estimated models we choose Control groups that did not violate the assumptions required for unbiased difference-in-differences estimates (for example that they did not violate the parallel trends assumption). This meant we used different Controls in the innovation input and output functions. These will be introduced under the discussion of our output and input estimations.

In our notation $D_j = 0$ 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 and Toshiba, i.e. we are excluding Samsung and Hitachi. As we are studying how innovation (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 in the relevant products post-merger.

Another thing that needs clarifying is 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.[9] 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.

Incorporating the above information into Eq.(2) we get the following innovation production function:

$$y_{jt} = \beta D_j I_t + x_{jt-\{1,\cdots,\tau\}}\phi + rd_{jt-\{1,\cdots,\tau\}}(\lambda_1 + \gamma_1 D_j I_t) + pat_{jt-\{1,\cdots,\tau\}}(\lambda_2 + \gamma_2 D_j I_t) + \mu_j + \mu_t + \varepsilon_{jt} \tag{3}$$

Where λ_1 and λ_2 are the non-merger specific effect of the innovation input on output, γ_1 and γ_2 are the effect of the mergers on the productivity of innovation input, and β gives us a residual firm-specific effect of the merger on innovation.

The two innovation inputs are defined in full as:

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

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

We estimate Eq.(3), and Eqs.(4a)-(4b) separately, using OLS. This decision reflects our main assumption, that is R&D and patent activity are both inputs in the innovation production function.

In what follows we explain for each R&D, patents, and product characteristics, the data used, and introduce a control group for estimating the impact of the mergers on the innovation production function. This is followed by our estimation of the impact of the mergers on innovation input. For this we have estimated a large number of different models, for various measures of innovation, for various Control groups, and different estimation methods. These are provided in the subsequent robustness section.

3.3 Data

3.3.1 R&D intensity data

One of the input measures in our innovation production function is R&D intensity (the ratio of R&D expenditure to total revenue). For all firms in our sample we have complete quarterly data coverage for the period of observation (Q1 2007 - Q4 2016). All the data used for the R&D analysis is from firms' balance sheets, as downloaded from S&P's Capital IQ database.

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

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.[10] Seagate's R&D intensity suffered a slump in Q1-Q2 2012, which was an accounting effect: adding Samsung HDD to Seagate's books meant adding more in total revenue than in R&D expenditure. 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: R&D intensity for Seagate, WD, and Toshiba

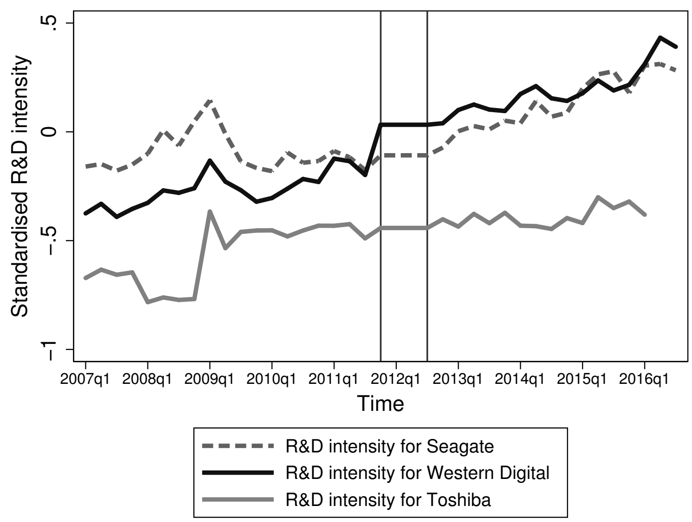

Figure 1 also draws light to two methodological issues. First, when evaluating how R&D intensity changes after a merger, one must not ignore an important artefact of this type of data, that is, following a merger, elements of the financial statement of the acquired company are added to the corresponding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

[10]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).

merging the two financial statements.[11]

Second, 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.

3.3.2 Patent activity data

We extracted patent data for each technology (HDD, Flash), and, only subsequently, grouped the data by firms.[12] 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 for the period Q1 2007 - Q4 2016. 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 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.

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, our primary measure, used in the headline results. This approach allows us to remain agnostic about what the best measure of patent activity is. In this spirit we also 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.

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

[12] 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.

3.3.3 Product innovation data

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 affect 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 look at two of the simplest ways of measuring product innovation: the number of new products marketed, and the unit price for HDD users ($ 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.[13] 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.[14] We have access to the following product 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

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

[14]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.

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.[15] 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 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.

3.3.4 Firm characteristics

We use the following firm characteristics as control variables in our estimations.

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.)[16]

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

[16] 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

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.

3.4 Choice of SSD as Control

We use SSD drives as Control group in estimating the impact of the mergers on HDD innovation output. 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 for our purposes. In mature industries product differentiation is typically no longer driven by innovation. 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

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 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.

example in areal density).[17] 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 bring out 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."[18]

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 an appropriate choice as Control.

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

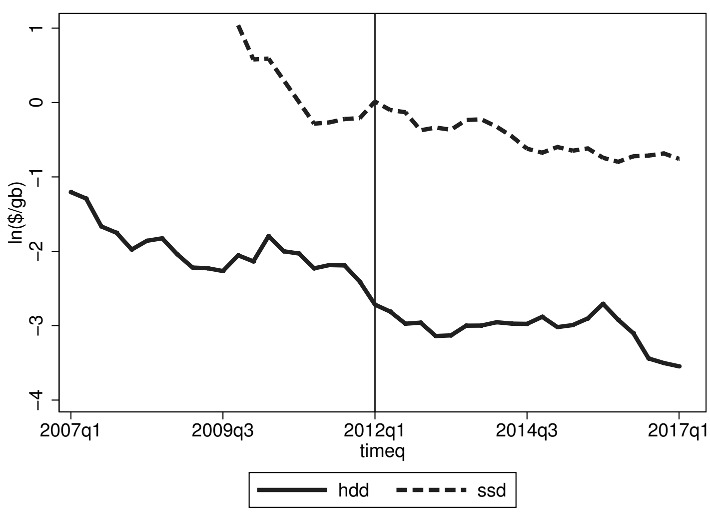

Regarding the independence of SSDs from the HDD mergers, we rely on the same argument as above. From industry references it clearly appears that there is competition between the two products,[19] which makes it more likely that the two are strategic complements in

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

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

[19]http://www.pcworld.com/article/3184464/storage/intel-optane-memory-has-a-mission-make-hard-

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 Results

4.1 Impact on innovation output (user cost and number of products

In this section we present the results of estimating the model introduced in Eq.(3). Table 1 summarises these results for our main variables of interest. Table 1 has three panels, one for each Treatment firm (Seagate, WD, Toshiba). The columns ln(cost) and ln(number) indicate our two measures of innovation output, the unit price of storage, and the number of new HDD models respectively.[20]

We do not have a priori knowledge on how quickly the merger effect trickled down to a change in innovation output. One can think of this time period as being composed of two parts: (1) the duration between the merger and a change (if any) in innovation input; and (2) the lag between a change in innovation input and a change in innovation output. We use a distributed lag model (i.e. in estimating output, we control for lagged input), which addresses the possibility of (2). In Table 1 we report a sum of the lagged effects (and the significance of this sum). We have no a priori information on (1). Since this lag might vary from industry to industry, we turned to the data for more information. We ran several experiments, for 3 different 'treatment times' $W \in \{$Q1 2012, Q1 2013, and Q1 2014$\}$. Through our handling of the treatment time we can acquire information on the duration of this delay.

We focus on three main sets of coefficients, as explained in Section 3. The productivity of R&D and the productivity of patents (denoted as λ_1 and λ_2 respectively in Eq.(3)), the effect of the mergers on R&D and patent productivity (denoted as γ_1 and γ_2 in Eq.(3)), and finally, the residual merger effect (denoted as β in Eq.(3)). Table 1 also shows the models where the pre-treatment parallel trends were rejected, and we focus on the results where the parallel trends assumption was not violated). In our estimation we standardised all non-binary dependent and independent variables - the coefficients can be interpreted accordingly.

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/

[20]Note that ln(cost) is an inverse indicator of innovation, a lower unit cost implies higher innovation.

Table 1: Effect of mergers on innovation output

	ln(cost)			ln(numbers)		
Treatment time	Q1 2012	Q1 2013	Q1 2014	Q1 2012	Q1 2013	Q1 2014
Seagate						
R&D productivity (λ_1)	0.282	0.236	-0.013	**0.737*** **	**0.589***	0.441
p-val	(0.556)	(0.74)	(0.988)	(0.006)	(0.053)	(0.237)
Merger effect on R&D productivity (γ_1)	**-1.110***	**-2.434** **	**-3.552*** **	**1.459*** **	**1.307** **	**1.201***
p-val	(0.073)	(0.046)	(0.000)	(0.000)	(0.04)	(0.075)
Patent productivity (λ_2)	0.189	0.061	0.091	0.023	0.029	0.076
p-val	(0.306)	(0.628)	(0.268)	(0.608)	(0.768)	(0.372)
Merger effect on patent productivity (γ_2)	**-0.566** **	**-0.739***	**-1.058*** **	0.124	-0.093	0.051
p-val	(0.039)	(0.089)	(0.002)	(0.139)	(0.603)	(0.878)
Residual effect (β)	0.756	1.062	**1.071** **	-0.233	-0.107	0.272
p-val	(0.260)	(0.124)	(0.026)	(0.233)	(0.795)	(0.240)
observations	171	164	157	173	166	159
parallel trend rejected?	Y	N	N	Y	Y	N
parallel test (p-val)	(0.009)	(0.733)	(0.150)	(0.000)	(0.000)	(0.873)
Western Digital						
R&D productivity (λ_1)	0.657	0.454	-0.269	**0.436** **	0.146	0.295
p-val	(0.160)	(0.563)	(0.779)	(0.011)	(0.595)	(0.412)
Merger effect on R&D productivity (γ_1)	0.047	**2.437***	0.725	**-2.104*** **	**-4.67*** **	**4.132** **
p-val	(0.932)	(0.070)	(0.804)	(0.000)	(0.000)	(0.014)
Patent productivity (λ_2)	**0.267***	**0.165***	0.083	-0.029	-0.071	-0.011
p-val	(0.072)	(0.089)	(0.42)	(0.406)	(0.345)	(0.833)
Merger effect on patent productivity (γ_2)	-0.084	**0.799*** **	0.789	**0.462*** **	**-0.739*** **	**3.106*** **
p-val	(0.65)	(0.008)	(0.355)	(0.000)	(0.009)	(0.000)
Residual effect (β)	**1.534** **	-0.810	0.170	**-0.589*** **	**1.243*** **	**-2.585*** **
p-val	(0.043)	(0.263)	(0.921)	(0.008)	(0.007)	(0.004)
observations	153	146	139	155	148	141
parallel trend rejected?	Y	N	N	Y	Y	N
parallel test (p-val)	(0.002)	(0.275)	(0.602)	(0.019)	(0.000)	(0.244)
Toshiba						
R&D productivity (λ_1)	0.55	0.474	0.056	**0.393** **	0.337	0.279
p-val	(0.178)	(0.567)	(0.954)	(0.026)	(0.154)	(0.323)
Merger effect on R&D productivity (γ_1)	**14.057*** **	**15.482** **	**69.961*** **	**-2.801*** **	**-13.334*** **	**14.058** **
p-val	(0.000)	(0.016)	(0.000)	(0.001)	(0.000)	(0.012)
Patent productivity (λ_2)	0.169	0.190	0.135	0.076	-0.01	-0.064
p-val	(0.373)	(0.063)	(0.215)	(0.521)	(0.914)	(0.387)
Merger effect on patent productivity (γ_2)	**0.745*** **	0.402	**5.875*** **	-0.112	-0.131	**-3.778*** **
p-val	(0.007)	(0.207)	(0.000)	(0.251)	(0.492)	(0.000)
Residual effect (β)	**7.839*** **	**8.553** **	**32.69*** **	**-1.568*** **	**-7.355*** **	**11.74*** **
p-val	(0.000)	(0.012)	(0.000)	(0.009)	(0.000)	(0.001)
observations	149	142	135	151	144	137
parallel trend rejected?	Y	Y	Y	Y	Y	Y
parallel test (p-val)	(0.001)	(0.046)	(0.001)	(0.000)	(0.077)	(0.000)

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

Regarding the firm-specific impact of the mergers, innovation productivity of R&D spending increased for Seagate. This is true for both measures of innovation output, and these results are robust to our choice of treatment time. Moreover, the mergers increased Seagate's productivity of patent activity but only regarding one of our output measures (unit cost for customers).

For WD there is evidence that the merger increased R&D and patent productivity but only for the number of new products. For unit cost, the merger seems to have lowered innovation productivity. For Toshiba there is some evidence of deteriorating R&D productivity (increasing unit costs and fewer new products) but this is an unrealistically large effect, which might be explained by the issues regarding the measurement of R&D for Toshiba, as we will show later.

Table 1 also offers some more general findings: R&D spending (productivity of R&D) contributes to increased innovation output, when output is measured by the number of new products, and that patent activity does not have an impact on either measure of innovation output.

To conclude these results, the mergers seemed to have unambiguously positive effect on Seagate's, and some positive effect on WD's innovation activity. Our evidence on Toshiba appears statistically questionable. We discuss these findings following our robustness section.

4.2 Impact on innovation input

To estimate the impact of the merger on R&D spending and patent activity, we need to find a Control group for each innovation input. We used SSD (the closest possible substitute) as a Control in our estimation on the effect of the mergers on innovation output. Having such a narrow scope for potential Control groups was important because innovation output is likely to be product specific (i.e. we believe it is unlikely that, for example, a different computer component such as a RAM, which is based on similar technology as SSDs (flash memory), is similar enough in the technical process of innovation. For innovation input we are able to expand our search for an ideal Control. We still use SSD/Flash memory as one of the Controls in our estimations, but for R&D and patents we have experimented with other possible Control groups. This is especially the case for R&D, where access to data allowed us to design an artificially created synthetic control group in order to produce unbiased causal inferences.

4.2.1 R&D intensity

Regarding R&D intensity, there are no direct rivals in the HDD market unaffected by the mergers, and the HDD market is worldwide, which means that local markets cannot be used as Control. Instead we explore product differentiation (HDD, 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 business segments to find the synthetic control. We looked at a number of other potential Control groups: a group of other storage firms (SSD and Flash Drive), and an unweighted and a weighted group of other IT firms. Our comparisons across these Control groups (as presented in Section A.1 in the Appendix) suggest that the synthetic control group performs best.

4.2.2 Patent activity

Measuring HDD patent activity In the innovation production function as an explanatory variable we used a composite measure of patents. But when looking at the impact of the mergers on patents one might wonder if other measures give us a different result. For this reason, instead of relying on a factor variable, we set out to look at all possible patent measures. We started with three variables: (1) patent count, (2) patent citation, and (3) patent factor. For each of these variables we: (i) use a simple count, (ii) generate stocks; (iii) smooth out shocks by employing a moving average over 4 (quarterly) lags and 4 (quarterly) leads; (iv) 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 these $4 \times 3 = 12$ 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 standardised causal estimates of patent activity. We combine these 36 causal estimates using a meta-analysis approach and obtain the average effects and the distribution of estimates.

Choosing a Control group All HDD manufacturers were involved in the treatment (the mergers and the related events), therefore we have to rely on different technologies to compose our Control group, which might be a problem, if patent activity is sensitive to the underlying technology. We headline our results where we used a different storage technology (SSD/Flash) as Control group (later we examine the robustness of these results by presenting estimates with different control groups). We have 40,655 Flash Memory related patent applications in our sample. Our data (Figure 6 in the Appendix) suggests that Flash might be sufficiently similar to the Treatment for the purposes of measuring patent activity.

Of course there is a strong possibility that this Control group is not independent of the mergers, 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.

4.2.3 Results

Table 2 shows the effect of the two mergers on innovation input (R&D and patent activity). The row *DD* shows the effect of the merger on R&D intensity and patent activity.

Table 2: Effect of the mergers on innovation input

	Seagate		**Western Digital**		**Toshiba**	
	R&D	Patents	R&D	Patents	R&D	Patents
DD	**0.0316*****	**0.549*****	**0.0479*****	**0.517*****	**0.0166*****	**-0.383*****
(p-val)	(0.005)	(0.000)	(0.006)	(0.000)	(0.005)	(0.000)
observations	78	558	78	558	74	591

p-vals in parentheses
$*p < 0.10$, $**p < 0.05$, $***p < 0.01$

It appears that the merger positively contributed to R&D intensity at all three firms. Below we will show that we rejected the pre-merger parallel trend assumptions for Toshiba. This does not mean that R&D spending did not increase post-merger for Toshiba, only that we do not have evidence that this increase was caused by the merger. Regarding patent activity, the general effect of R&D intensity on patent activity is positive and around the same rate across all three models. We discuss these results in more detail following our robustness checks.

5 Robustness checks: alternative Control groups

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

5.1 Robustness of R&D expenditure results

Our main choice of Control above was a synthetic group but we look at another three potential groups in order to demonstrate the robustness of our findings.

Selected storage firms as Control: This 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.[21]

Large sample IT firms as control: To circumvent the problem of a possible bias (no independence), 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. At the same time, these IT firms were potentially affected by similar supply and demand shocks (which we will test later). 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,[22]

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

[22] 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.

revenue growth, total assets, gross profit, net income, and number of segments[23]. 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 3 shows the estimated effect of the mergers on R&D intensity for the three alternative Control groups, along with the previously reported synthetic control. The coefficients seem robust to our choice in the exact specification of the Control group. Although there is evidence of an increase for all firms, the violation of the parallel trends assumption for Toshiba means that we can only attribute this to the mergers in the case of Seagate and WD (although the evidence seems more sensitive to the choice of Control group in the case of WD).[24]

Table 3: Effect of the mergers on R&D intensity for alternative Control groups

Control	Other storage firms	IT firms	Weighted IT firms	Synthetic
Seagate				
DiD	0.176***	0.212***	0.321***	0.0316***
p-val	(0.007)	(0.000)	(0.002)	(0.000)
obs.	517	23267	641	78
parallel trends rejected?	N	Y	N	
parallel test p-val	(0.493)	(0.001)	(0.288)	
Western Digital				
DiD	0.469***	0.346***	1.190***	0.0479***
p-val	(0.000)	(0.000)	(0.000)	(0.000)
obs.	440	24053	623	78
parallel trends rejected?	Y	Y	N	
parallel test p-val	(0.013)	(0.000)	(0.118)	
Toshiba				
DiD	0.0413	0.0193	0.0708***	0.0166***
p-val	(0.541)	(0.394)	(0.002)	(0.002)
obs.	517	23267	607	74
parallel trends rejected?	Y	Y	Y	
parallel test p-val	(0.000)	(0.000)	(0.000)	

5.2 Patents

In our preferred results we used Flash based storage as Control, but here we look at another two potential Control groups.[25]

[23]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.

[24]Details of testing the assumptions are given in the Appendix A.

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

Other HDD patents: This group consists of the top 10 firms in terms of the number of HDD-related patents held in our data.[26] These patents are not innovations of the HDD units but innovations on something complementary to HDD.[27] 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 4: The effect of the mergers on patent activity

Control	Seagate	WD	Toshiba
Other HDD	**0.549*****	**0.517*****	**-0.383*****
	(0.000)	(0.000)	(0.000)
Other Flash	**0.375*****	**0.168*****	**-0.491*****
	(0.000)	(0.000)	(0.000)
Top Storage	**0.386*****	**0.309*****	**-0.546*****
	(0.000)	(0.000)	(0.000)

pvals in parentheses
$*p < 0.10$, $**p < 0.05$, $***p < 0.01$

Table 4 shows the mean and the 95% confidence intervals for the distribution of our estimates.[28]

The main findings seem consistent across the three models with alternative Control groups, i.e. the merger increased R&D and patent activity for Seagate and Western Digital, and reduced (patent) or did not affect it (R&D) for the non-merging third firm, Toshiba.

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

[27] 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.

[28] 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.

6 Discussion of results

Regarding firm specific effects, we found no evidence that the merger would have led to a fall in innovation activity for any of the merging firms. For Seagate the mergers triggered an increase in the productivity of R&D intensity. Moreover, the merger also increased Seagate's innovation input (both R&D intensity and patents). We do not have conclusive evidence on what exactly ignited 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.[29] 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 found evidence that the consolidation increased R&D productivity for some factors (number of new products), but it also led to increased unit cost. We also found evidence of a positive effect on innovation input (R&D and patents). One possible explanation for the partially negative effect on productivity is that the MOFCOM decisions particularly hindered the consummation of the WD/HGST merger until October 2015. Remedies were much stricter than for Seagate and they fundamentally required that WD duplicate their R&D, production, marketing, and sales operations. This was crippling for WD's efficiency.[30] Other events might also affected WD's innovation activities. For example the divestiture of the 3.5in operations to Toshiba had to include all 3.5in related IP rights. This might have

[29] 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).

[30] 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 .

negatively affected how innovation, and indeed R&D spending evolved post-2012 for WD.

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.

We found no unbiased and robust evidence of a change in Toshiba's R&D spending and innovation output after the mergers, and found evidence of a drop in patent activity. 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. Finally, the matching process that we applied for the R&D estimations did not work well for Toshiba. All of these could indicate that the effects are potentially picking up changes in other segments (i.e. a general, not HDD specific drop in Toshiba's R&D spending). The drop in Toshiba's patent activity seems to indicate that Toshiba's acquisition of the divested Hitachi assets had a counter-productive effect of patenting activity. It is important to bear in mind that the Toshiba acquisition of HGST was the result of a divestiture condition imposed on the other two merging firms, rather than the result of organic business expansion.

7 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. Our findings are robust to a large number of different model specifications, control groups, and study designs.

The European Commission's 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 alleged new innovation based theory of harm, which led 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. Our evidence contributes to this debate, inasmuch as it offers an example of a merger in a concentrated market (which at the time of the merger was deemed to have high entry barriers) that did not lead to a loss in innovation. This would support the view that a case-by-case approach rather than a theory of harm approach is preferred in merger innovation claims.

As a more general contribution to the literature, we found that innovation input (R&D intensity) positively boosts innovation output. Moreover, we found no evidence that patent activity would have an effect on innovation output. 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 as it would imply that R&D spending might be a good predictor of innovation output, which could be useful information for future research using either R&D or patent measures to approximate innovation.

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, somewhat akin to the way papers proliferated on ex-post estimating the price impact of consolidation.

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.

Arrow, K. J. (1962): "The economic implications of learning by doing," *The Review of Economic Studies*, 29(3), 155–173.

Blind, K., K. Cremers, and E. Mueller (2009): "The influence of strategic patenting on companies patent portfolios," *Research Policy*, 38(2), 428–436.

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.

Bourreau, M., B. Jullien, Y. Lefouili, et al. (2018): "Mergers and Demand-Enhancing Innovation," Discussion paper, Toulouse School of Economics (TSE).

Cohen, W. M., A. Goto, A. Nagata, R. R. Nelson, and J. P. Walsh (2002): "R&D spillovers, patents and the incentives to innovate in Japan and the United States," *Research Policy*, 31(8-9), 1349–1367.

Cohen, W. M., R. R. Nelson, and J. P. Walsh (2000): "Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not)," Discussion paper, National Bureau of Economic Research.

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.

Davies, S. W., and P. L. Ormosi (2012): "A comparative assessment of methodologies used to evaluate competition policy," *Journal of Competition Law and Economics*, 8(4), 769–803.

Denicolò, V., and M. Polo (2018): "Duplicative research, mergers and innovation," *Economics Letters*, 166, 56–59.

Federico, G., G. Langus, and T. Valletti (2017): "A simple model of mergers and innovation," *Economics Letters*, 157, 136–140.

——— (2018): "Horizontal mergers and product innovation," *International Journal of Industrial Organization*, 59, 1–23.

Genakos, C., T. Valletti, and F. Verboven (2018): "Evaluating market consolidation in mobile communications," *Economic Policy*, 33(93), 45–100.

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. (1979): "Issues in Assessing the Contribution of R&D to Productivity," *Bell Journal of Economics*, 10, 925116.

——— (1990): "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28, 4.

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

Hagedoorn, J., and M. Cloodt (2003): "Measuring innovative performance: is there an advantage in using multiple indicators?," *Research Policy*, 32(8), 1365–1379.

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.

Haucap, J., and J. Stiebale (2016): *How mergers affect innovation: Theory and evidence from the pharmaceutical industry*, no. 218. DICE Discussion Paper.

Igami, M., and J. Subrahmanyam (2015): "Patent Statistics as an Innovation Indicator? Evidence from the Hard Disk Drive Industry," *https://ssrn.com/abstract=2263318*.

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

Janger, J., T. Schubert, P. Andries, C. Rammer, and M. Hoskens (2017): "The EU 2020 innovation indicator: A step forward in measuring innovation outputs and outcomes?," *Research Policy*, 46(1), 30–42.

Jewkes, J., D. Sawers, and R. Stillerman (1958): *The sources of invention*, vol. 11. Macmillan London.

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.

Motta, M., and E. Tarantino (2017): "The effect of horizontal mergers, when firms compete in prices and investments," *Working paper series*, 17.

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.

Santamaría, L., M. J. Nieto, and A. Barge-Gil (2009): "Beyond formal R&D: Taking advantage of other sources of innovation in low-and medium-technology industries," *Research Policy*, 38(3), 507–517.

Scherer, F. M. (1983): "The propensity to patent," *International Journal of Industrial Organization*, 1(1), 107–128.

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.

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 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.

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

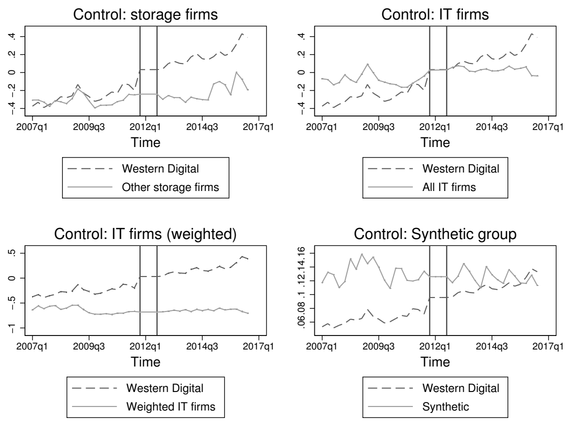

Figure 4: R&D intensity plot for Seagate and four different Control groups

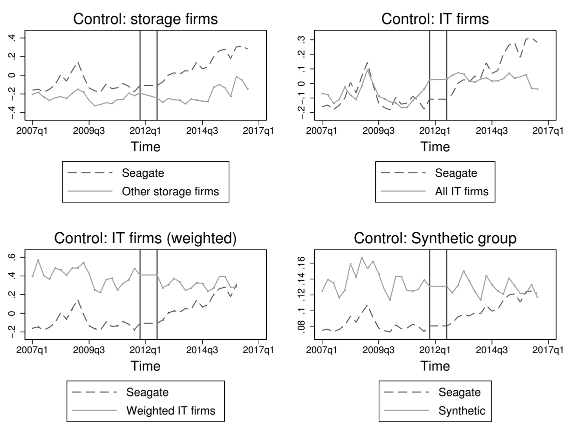

Finally, we look at Toshiba on Figure 5. The Treatment line shows a jump in 2009 when Toshiba acquired Fujitsu's HDD operations. After the merger the Toshiba line seems to go together with the Control line. The figures show that the matching process was not as effective as for WD or for Seagate. We will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

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

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

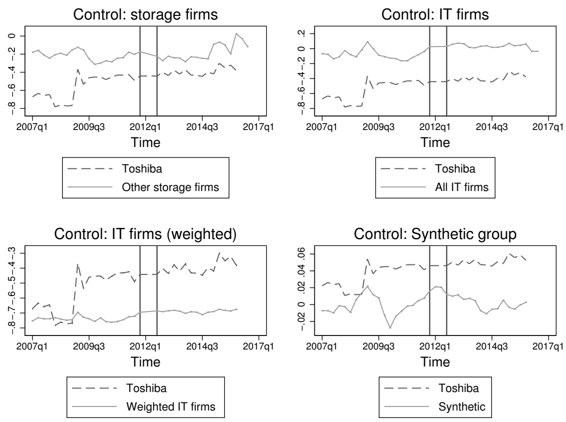

that the effect of the mergers might be found significant even when, in an unbiased model, it would not be. Similarly, negative serial correlation in the price data may overestimate the standard error of the merger effect. We used Wooldridge's (2002)Wooldridge (2010) autocorrelation test for panel data.

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.

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

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

(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. [31] 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

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

A.3.3 Firms used in synthetic control

Table 8: Firms (weights) used in synthetic control

Firm	Weight
Seagate	
Advanced Micro Devices, Inc. (NasdaqCM:	0.419
Alphabet Inc. (NasdaqGS:GOOGL)	0.038
Compal Electronics, Inc. (TSEC:2324)	0.37
Intuit Inc. (NasdaqGS:INTU)	0.165
NVIDIA Corporation (NasdaqGS:NVDA)	0.008
WD	
Alphabet Inc. (NasdaqGS:GOOGL)	0.042
Compal Electronics, Inc. (TSEC:2324)	0.128
Intuit Inc. (NasdaqGS:INTU)	0.124
Inventec Corporation (TSEC:2356)	0.339
NVIDIA Corporation (NasdaqGS:NVDA)	0.368
Toshiba	
Avaya Inc.	0.093
NetApp, Inc. (NasdaqGS:NTAP)	0.046
Symantec Corporation (NasdaqGS:SYMC)	0.106
TCL Multimedia Technology Holdings Limi	0.293
Unisys Corporation (NYSE:UIS)	0.414
salesforce.com, inc. (NYSE:CRM)	0.048

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 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.

We deliver results for 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

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

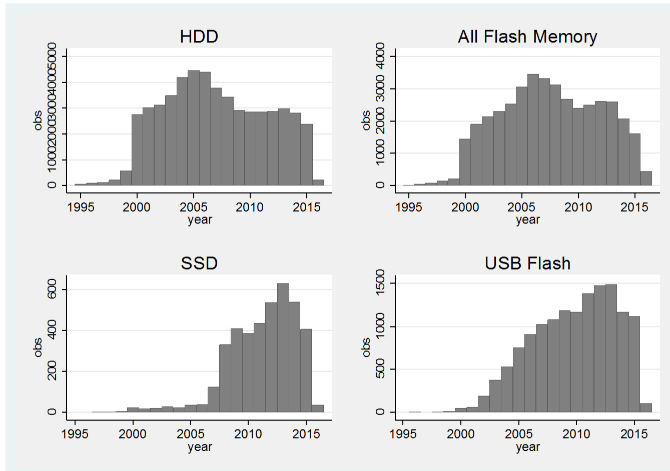

factor analysis as the methodology in order to reduce the number of patent-related correlated variables. The justification for using this methodology (instead of principal component analysis) is that we have a set of original variables that together contribute in explaining innovation, while all those variables on their own would have limited contribution and be subject to criticism.

Table 9 shows the DD estimates for each Treatment firm, using a factor variable. The results are qualitatively the same as in our headline table.

Table 9: 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

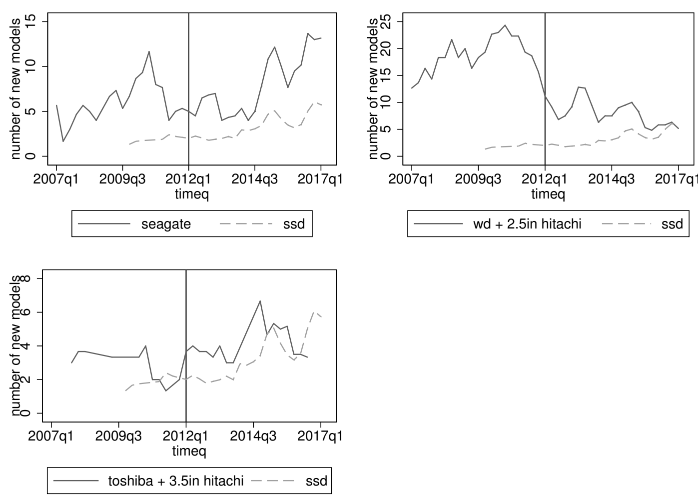

It immediately stands out from Figure 7 that SSD's only appear in our sample from 2009, and especially at the beginning the per-firm average number of SSDs was very low. Figure 7 shows that for Seagate there was an increase in the number of new drives marketed roughly 2 years after the merger. This is 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.

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

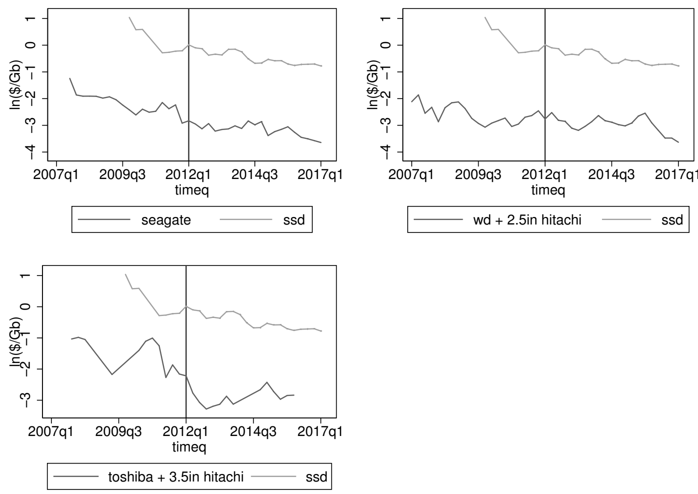

We tested separately whether the Treatment and Control follow a parallel trend before the treatment(s). We found that out of our 5 treatment events, the first two estimates are likely to be biased because pre-treatment trends were not parallel (this is visually confirmed on Figure X and Y). This would 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 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.