

Mergers and Innovation: Evidence from the Hard Disk Drive Market

Anna Rita Bennato
Loughborough University

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

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

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

CCP Working Paper 18-4 v3
(updated March 2019)

Abstract

This case study is a relatively rare ex-post evaluation of how the level of innovation changed after the 5-to-3 consolidation of 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 user costs as indicators of innovative output. This allows us to distinguish the magnitude of the merging parties' innovative efforts from the productivity of those efforts. Of the remaining HDD manufacturers, for Seagate we found an increase in all our innovation measures following the mergers, but for Western Digital the evidence is mixed. Methodologically, the paper draws light on some of the challenges of conducting similar case-specific retrospective studies on the impact of mergers on innovation.

Contact Details:

Peter Ormosi

p.ormosi@uea.ac.uk

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

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

March 14, 2019

^{*}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 how the level of innovation changed after the 5-to-3 consolidation of 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 user costs as indicators of innovative output. This allows us to distinguish the magnitude of the merging parties' innovative efforts from the productivity of those efforts. Of the remaining HDD manufacturers, for Seagate we found an increase in all our innovation measures following the mergers, but for Western Digital the evidence is mixed. Methodologically, the paper draws light on some of the challenges of conducting similar case-specific retrospective studies on the impact of mergers on innovation.

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 a considerable amount of literature on measuring how competition affects innovation. This includes a number of studies on the effect of mergers 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 across 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. In 2011/12, three mergers (Seagate/Samsung, Western Digital (WD)/Hitachi, and Toshiba/Hitachi) reduced the number of HDD manufacturers from 5 to 3 firms. To analyse how the level of HDD innovation changed after these mergers, we assembled a rich set of data, which we used to approximate Schumpeter's innovation trichotomy and measure innovation in its entirety, as opposed to looking at its component parts (for example patents only) in isolation. First, we look at how our measures of innovation output (the number of new products, and the unit user cost of HDD capacity) changed after the mergers. This is followed by an analysis of R&D spending and patent activity. 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 approach we are able to make contributions to the innovation research literature in general, most importantly by offering evidence on whether in the HDD market R&D expenditure or patent measures are more likely to correlate with measures of innovation output.

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 crucial 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 than now, when there appears to be an increase in interest from both practitioners and academics.

To headline our key results, we find mixed evidence on how the level of innovation changed

after the 2011/12 consolidation of the HDD market. For one of the HDD manufacturers, Seagate, we find an increase in the company’s R&D spending, patent activity, number of new products marketed, and a decrease in the unit user cost of HDDs. For WD the evidence is more equivocal, indicating increase in some areas and drop in others. Our results also give indication that - at least in this specific market - R&D spending data is a better predictor of the number of new products and of unit user cost than patent data. These findings are robust to a large number of empirical models, research designs, and model specifications. The paper discusses in detail whether these changes can be causally attributed (at least partially) to the mergers.

The paper is structured as follows. We commence with a brief survey of literature, followed by an introduction of the HDD market and a description of the regulatory approval process. Section 3 presents a simple analytical framework, introduces the data and offers simple before and after statistics. In Section 4 we introduce our empirical strategy, and Section 5 offers our results. In Section 6 we look at whether the changes in innovation can be attributed to the mergers, which is followed by a discussion of the key findings of the paper. Throughout this study we have conducted a large number of econometric tests and sensitivity checks. Some of these are reported in our Appendix, and in Ormosi, Bennato, Davies, and Mariuzzo (2017).¹

1.1 Literature review

Our paper draws on various literatures, most of which originates from 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).² 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”

¹The data and our main script files are available at: https://github.com/PeterOrmosi/hdd_innovation

²Some 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).

(resonating with Arrowian arguments), but “Schumpeter was also quite correct: the prospect of obtaining market power is a necessary reward to innovation”.

From this background, a distinct body of literature has looked at how specific instances of mergers (transactions that increase market concentration) have affected innovation. An important upshot of this literature, is that sweeping generalisations are not justified, especially if one recognises 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.

On the theoretical side, the results run both ways on the impact of mergers on innovation, whilst recognising the specific circumstances in which mergers contribute to higher levels of 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 is likely to lead to lower innovation incentives, absent cost efficiencies and spillovers. However, Denicolò and Polo (2018) show that in given circumstances 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 posits 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. Letina (2016) finds that mergers decrease the variety of developed projects and decrease the amount of duplication of research. 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 again, there is no unanimity, 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 Entezarkheir and Moshiri (2018) reports increased patenting activity following mergers.³ Finally, in a recent

³In 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

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), 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”.⁴

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 finding ways of denying innovations to others.⁵ 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 to build our own work are as follows. First, there are arguments both ways on whether, in general, mergers encourage or discourage innovation. Some might argue that the balance of theory and empirical evidence points to a generally negative effect.⁶ Here, however, we are more interested in the detail of a specific case study:

at the aggregate industry level in concentrated markets in OECD countries, 2002-14.

⁴See also Janger, Schubert, Andries, Rammer, and Hoskens (2017).

⁵A more recent paper, Blind, Cremers, and Mueller (2009) shows the importance of strategic patenting in improving a firm’s reputation, giving it greater bargaining power in negotiations with other firms. Moreover, strategic patents can be used to create internal incentives for their R&D employees, and to measure their performance.

⁶The European Commission, in its recent policy brief, offers an objective assessment of mergers and

what happened in the wake of consolidations from 5 to 3 firms in the HDD market? And in general, how feasible is it to estimate the innovation impact of specific mergers? We prefer not to bring strong priors on whether or not these acquisitions stimulated innovation.

Second, we turn to a reduced-form Difference-in-Differences (DiD) methodology. As the paper will demonstrate, similar to structural methods, even with this approach it is inevitable that one draws on economic theory in order to inform assumptions on the type of competition across different technologies - something that will prove to be central to the interpretation of our results (for example, whether a control group is independent of the treatment firms depends largely on our assumptions about the type of strategic interaction between the control and treatment firms).

Third, we do 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 in the production of better HDDs. 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, we believe that the paper provides a practical blueprint for future, much needed, policy evaluations in other innovative industries.⁷

2 The Hard Disk Drive and Solid State Drive markets

We look at three mergers - Seagate/Samsung, Western Digital/Hitachi, and Toshiba/Hitachi (3.5 inch) - 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

innovation through their case law (http://ec.europa.eu/competition/publications/cpn/2016_001_en.pdf).

⁷For 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).

assemblies to store data, which records, stores and retrieves digital data without any moving parts. Solid state drives (SSD) and USB Flash drives (Flash Memory based data storage device with integrated USB interface) are Flash memory based storage. SSDs are built on semiconductor memory arranged as a disk instead of magnetic or optical storage support. Because no mechanical components are involved, SSDs are fast in comparison to rotating media (HDD), providing access to data in microseconds, instead of the several milliseconds requested by HDDs.

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 user costs 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, which is especially true regarding their performance (speed).⁸ In this respect it appears safe to conclude that HDD and SSD are likely to be vertically differentiated product (i.e. if prices were the same most would choose SSD).

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. It can be concluded that the main reason for these tendencies is the evolution of the different applications that currently rely on HDD (e.g. large archives, data servers) or SSD (e.g. portable devices). Our reading of these technologies is that HDD and SSD are complement products in (at least) many applications.

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 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

⁸The CERAM T800 SSD in 1994 had a 10MBps reading speed; today's SSDs manage over 500MBps. For HDDs, the performance improvement for the same period was only around 10 fold.

technology. The major players in SSD are Samsung, Toshiba, SandDisk, Micron, SKHynix, and Intel.

2.2 Regulatory approval

On 7 March 2011, WD and Hitachi, Ltd announced and executed a share purchase agreement for the sale of all issued and outstanding capital stock of Hitachi Global Storage Technologies (HGST), a wholly owned subsidiary of Hitachi Ltd. On 19 April 2011 Seagate announced their intention to buy Samsung's HDD business. Seagate notified the European Commission on the day of the announcement, whereas WD submitted the notification on the following day, on 20 April 2011. The Seagate/Samsung merger was unconditionally approved in every jurisdiction,⁹ 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 pressure in the HDD market, and therefore its elimination from the HDD market was not expected to affect the level of competition. The WD/Hitachi merger was approved by the European Commission and the US authorities, subject to a divestiture of the 3.5" desktop HDD manufacturing lines to Toshiba. MOFCOM, again, took a different stance and imposed a set of behavioural remedies. The divested HGST assets were acquired by Toshiba (in a transaction announced in February 2012).

We argue that although the MOFCOM conditions were restrictive, they did not nullify the effect of the 2012 events, for the following reasons. Firstly, there is no evidence that suggests that restrictions were placed on the transfer of intellectual property rights (restrictions were only placed on the R&D activities), for example Seagate and WD were each able to access the acquired businesses' stock of intellectual property (patents). This is important, because before the merger Seagate and WD would have needed a license to use Samsung and HGST patents. With the transfer of the already existing stock of intellectual property, this was no longer required. Whether this indeed happened, we looked at the transfer of HDD-related patents between the relevant firms after the mergers. The top row of Table 1 below shows the HDD-related patent stock for the merging firms just before the merger. The bottom rows of Table 1 show the number of these patents for which ownership was transferred with the mergers in 2011/12. Seagate became the owner of around 20% of Samsung's HDD-related

⁹Following the priority principle, the Commission held that a party that is the first to notify a concentration which (i.e. the Seagate/Samsung merger), assessed on its own merits, would not significantly impede effective competition in the internal market or in a substantial part thereof, is entitled to have its operation declared compatible with the internal market within the applicable time limits. For the same reason the Commission considered the WD/Hitachi merger with a market structure that reflected conditions *after* the Seagate/Samsung merger.

patents with the mergers. The patents that Samsung did not transfer were typically not strictly on HDDs, but on complementary products that relate to HDDs. It is therefore safe to expect that the Samsung/Seagate merger, as it happened in 2011/12, had the potential to affect Seagate’s innovation activities, if not least, through the synergies resulting from shared access to some key HDD patents. Regarding WD, we did not find any ownership transfer from HGST to WD at the time of the 2011/12 conditional approval of the merger. Neither did we find any patent transfers from HGST to Toshiba at the time of the merger. This is despite the requirement that relevant HGST IP rights should be transferred to Toshiba upon their purchase of the divested 3.5-in HDD operations.¹⁰

Table 1: HDD patent stock before the merger and patents transferred with the mergers


	Samsung	Seagate	HGST	WD	Toshiba
HDD patent stock before the merger	951	194	1260	74	486
Number of patents transferred with mergers					
to Seagate	173		0		
to WD	0		0		
to Toshiba	0		0		

The MOFCOM remedies also required the parties to hold R&D activities separate and to not reduce them below a specified level. Seagate was required to spend at least \$800 million annually on R&D. However, Seagate’s average annual R&D spending pre-merger was above \$900 million). Therefore Seagate even had a leeway to somewhat drop their R&D spending and still comply with the MOFCOM requirements. This is important because it tells us that Seagate was free to increase, hold constant, or even reduce (to some extent) their R&D expenditure post-merger, which makes our exercise of testing the sign of the change in R&D a meaningful one. Western Digital was required to maintain the pre-merger level of R&D expenditure. This still allowed WD to decide between increasing or not increasing R&D expenditure.

From the MOFCOM case announcements it is clear that restrictions were much more stringent on WD than on Seagate. WD was practically forced to operate with inefficiently duplicated production, marketing and sales operations for WD and HGST. This circumvented any possibility of increased efficiency. Figure 1 shows some of the differential effects of the hold separate conditions. WD and HGST combined, having to run with duplicated units, had 80,767 employees, compared to Seagate’s 53,602. At the same time there was only

¹⁰Note that 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.

Figure 1: Number of full time employees for Seagate and WD



8 per cent difference in capacity shipped between the two firms. Moreover, the hold separate conditions were imposed for 2 years on WD and 1 year on Seagate.

3 Model and data

In modelling the impact of the mergers on the innovative output of the post-merger entities, we allow for three not mutually exclusive possibilities. The mergers might have been followed by (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 two-stage 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

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

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

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} .¹¹

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

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

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. In our framework patents are also thought of as input in our innovation production function.

Our preference for this approach recognises the traditional critiques as introduced in Section 1: (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 Data

3.2.1 R&D intensity data

R&D intensity is defined as 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). This is from firms' financial statements, as downloaded from S&P's Capital IQ database.

The R&D data draws light on two methodological issues. First, when evaluating how

¹¹When the explanatory variables include lagged variables, we assume sequential exogeneity, on unobserved heterogeneity

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 merging the two financial statements.¹²

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 if they are diversified. 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.

3.2.2 Patent activity data

We extracted patent data for each technology (HDD, Flash), and, only subsequently, grouped the data by firms.¹³ 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 for the period Q1 2007 - Q4 2014. 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. Instead of arbitrarily relying on a specific variable, we set out to create a factor variable, similar to the multiple-

¹²We remove Q4 2011 and Q1 2012 from our analysis, and we also disregard the growth in R&D intensity in this period.

¹³Relevant 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. Details of the data collection are provided in Ormosi et al.(2017).

indicator factor model in Lanjouw and Schankerman (2004). This variable brings together the richness of patent data into one variable, we construct a patent indicator. We make use of a complete set of variables collating information on patent counts, patent citations (distinguishing citations from attorneys and from the literature), patent inventors (number), patent claims (number), patent applications (number) and application countries (number). However, in contrast to their paper we choose to utilise factor analysis as the methodology in order to reduce the number of patent-related correlated variables. The justification for using this methodology (instead of principal component analysis) is that we have a set of original variables that together contribute in explaining innovation, while all those variables on their own would have limited contribution and be subject to criticism.

Through this factor analysis we 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 is the best measure of patent activity. For our introductory before and after discussion we use an easy-to-interpret measure of patents, patent citations. In the main analysis however we draw on this patent factor.

3.2.3 Product-level 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 user costs 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 2016.¹⁴ 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

¹⁴The sample accounts for the mergers that happened before 2012, for example Fujitsu is recorded as Toshiba as a result of their 2009 merger.

innovations.¹⁵ The sample consists of 33 SSD and 5 HDD brands.¹⁶

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 capacity (expressed gigabytes). This way we are also able to make comparisons with SSD. 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.

Unit user cost: This is the unit capacity retail price of HDD products. Our terminology of unit user cost reflects the fact that for the parties to the mergers, a new HDD might reflect a product innovation, but for the firms buying the HDD, any improvements in technology are a process innovation – reductions in retail price reduce the using firms costs.

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 2012, we had the price as it appeared in 2017. This might seem to go against intuition, as one could expect prices to gradually fall, and therefore the price of older products will always be smaller than the price for newer products. This however does not seem to be the case for HDDs. Archive.org takes a snapshot of 'the Internet' on a regular basis (multiple times a day). Not everything is recorded on every snapshot, the idea is to capture changes. For this reason, scraping data from Archive.org has always been a challenge, because one does not know a priori when a change is recorded

¹⁵For 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.

¹⁶This reflects the relative maturity of these two technologies. 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. A frequency table of brands is given in Table 8 in the Appendix.

on Archive.org. Moreover, it is very difficult to search for specific products through these archives, which means matching products by their product numbers almost an impossible task. For this reason we scraped price data as they were in 2017, but in Table 2 we provide a random sample of 10 HDDs and their Amazon.com retail prices over 6 years as acquired from Archive.org. The purpose of this table is to demonstrate how prices of new HDDs change after their introduction. The table shows, for 10 HDD models (product number) how their prices evolved from when they were first made available to 2017. Take the first row for example, a Seagate HDD, marketed on Amazon in 2012 for \$59.99, in 2014 its price was \$54.00, and in 2017 it was \$58.95. The gaps in the table imply that we found no information on Archive.org.

Table 2: Historical prices of a sample of HDDs (USD)

Product number	2012	2013	2014	2015	2016	2017
ST9500325AS	59.99	60.55	54.00		49.95	58.95
WD10EZEX	76.99	59.99	52.99	54.99		63.80
ST2000DM001	107.99		79.99	79.99	108.99	141.83
WD5000AAK	64.99	59.26	52.49	49.99	65.98	70.00
ST1000DM003				50.92	49.99	51.99
ST2000DM001	109.99	99.99	79.99	79.99	71.99	108.99
ST1500DL003	78.00	116.99	98.00	85.00	56.80	92.69
HGST 0F12115	121.17	134.99	108.95	117.00		
ST3000VN000		146.89	139.99	114.67		149.99
HDTC607XK3A1	72.22	59.98	54.99	63.99	123.00	

Figure 2: R&D intensity for Seagate, WD, and Toshiba




Table 2 shows that prices do recover to near their original level - after a temporary drop (the same is plotted on Figure 2 for the price averaged over our sample of 10 drives). Our interpretation of this interesting pattern is that the pace of introducing new HDDs is fast.

On average, the same manufacturer introduced a new product of exactly the *same* capacity every 6 months (5 months when only looking at the Seagate, WD, or Toshiba), 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 by too much, they would cannibalise 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.¹⁷ 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. Eventually prices all seem to start an increase, which, in our interpretation, is due to the fact that these older products are discontinued and become scarce for people who (for whatever reason) still want to buy them. For the above reasons we believe that using 2017 data is not far-fetched at all for our purposes (especially given that our goal is not to discuss the absolute magnitude of prices, rather their relative level in comparison to previous products).

From the above, we derived our two variables used for measuring technological progress. The first one, the number of new products, includes all newly marketed products. The reason we do not just focus on products with higher capacity is simply that innovation happens across many dimensions, capacity is one of them. But a new product with the same capacity can have higher speed, lower seek time, more cache, higher reliability, better transfer rates, just to mention a few.¹⁸ Our other outcome measure is the the unit retail price of new products (\$/Gb). Both variables are recorded by firm j in period t .

3.2.4 Firm characteristics

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

Firm size: There are numerous studies linking 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.)¹⁹

¹⁷See for example Douglas and Pavcnik (2001).

¹⁸For example two Seagate HDDs with the same capacity (2Tb) were marketed around a year apart, where the first one ST2000DL003 came with 5900rpm and the second one (ST2000DM001) with a higher 7200rpm performance.

¹⁹Gross profit is the difference between total revenue and the cost of revenue. In our regressions we include total revenue and gross profit, which together determine the cost of revenue. Net income includes various earnings on the firms' operations. Total debt refers to various interest bearing obligations. Total operating expenses reflects expenses not directly associated with the production of goods or services. These firm characteristics are closely correlated with each other (larger businesses will have high values, etc). We 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.

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 for constructing an unbiased Control group.

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 selecting the Control and not in the DiD estimations (where we 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.

4 The empirical strategy

4.1 Before and after the mergers

As a first step we looked at how our main measures of innovation input and output changed after the mergers. We use data on innovative output (the unit user 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.²⁰ Of these $T = 40$ time periods, there are $T_0 - 1$ time

²⁰WD acquired Sandisk to boost its SSD/Flash portfolio in 2016, which is another reason why we excluded post-2017 data.

periods measured prior to the mergers that take place in period T_0 , implying that $t \in \{1, \dots, T_0 - 1, T_0, T_0 + 1, \dots, T\}$.

Table 3 shows the difference (estimated using OLS without controlling for any covariates²¹) between the average before and after values of R&D intensity, our patent factor variable, quarterly number of new products, and quarterly unit user cost.²² The table shows somewhat different trends in the three HDD manufacturers. R&D intensity increased significantly for all three firms after the merger. Patent activity seems increased for Seagate and dropped for the others (but only strongly significantly for Toshiba). For the number of new products we see an increase for Seagate and Toshiba but a fall for WD. Finally, unit user cost has fallen for all three firms.

Table 3: Innovation input and output measures - before and after

	Seagate	WD	Toshiba
R&D	0.201***	0.415***	0.125***
p-val	(0.000)	(0.000)	(0.003)
n	32	32	30
Patent factor	0.151*	-0.155	-0.754***
p-val	(0.091)	(0.165)	(0.000)
n	32	32	32
Numbers	0.544**	-0.695***	0.399*
p-val	(0.014)	(0.001)	(0.059)
n	50	54	37
Unitprice	-1.346***	-0.532**	-1.346***
p-val	(0.000)	(0.011)	(0.000)
n	48	49	48
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$			


With the weak exception of patent activity for WD and Toshiba, these numbers do not suggest a drop in innovation after the consolidation; but whether any increase can be attributed to the mergers is a different matter. It is possible that the increase in innovation input and output factors (and the fall in patent activity) are part of a general trend, which could have happened anyway, and are not (or not fully) affected by the mergers. It is also possible that innovation would have increased at an even faster pace in the absence of the mergers. Finally, there is a possibility that the mergers triggered or contributed to the increase. Figure 3 helps illustrate this point. It plots R&D intensity for Seagate, Western Digital and Toshiba between 2007 and 2016. The two vertical lines show the start and the

²¹For better comparison with Table 4 we should have included the same covariates as in Table 4 but that would have lead to dimensionality problems due to the small number of observations we have for the individual merging firms.

²²The number of new drives and the unit user cost measures include observations for Hitachi and Samsung drives marketed under these brand names after the merger.

closure of the merger approval process. 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.²³ It appears therefore that the increase in WD's R&D intensity has started before the merger, and the merger - at best - only partially contributed to this increase. Seagate's R&D intensity follows an increasing trend after the merger,²⁴ and in this case the increase started at the time of the merger. Whether it was triggered by the merger or something else that happened in 2012 is impossible to tell from this diagram. 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. For Toshiba, the post-merger larger R&D intensity average seems to be driven by the 2009 leap, which again had nothing to do with our two analysed mergers.

Figure 3: R&D intensity for Seagate, WD, and Toshiba



In the following section we explain how far it is possible to separate out how much of these changes can be attributed to the merger. This is not simply about providing evidence for this particular case, but also (and probably more importantly) to demonstrate more generally how far acquiring unbiased robust estimates of the effect of the mergers is feasible.

²³In 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).

²⁴Note that Seagate acquired Maxtor, another HDD manufacturer in 2006, the effect of which merger might still be seen in our sample period starting in Q1 2007.

4.2 Can we capture the impact of the mergers? The econometric model

To determine whether we can make any type of causal inferences about the effect of the mergers, we apply a difference-in-differences (DiD) approach to Eqs.(1a), (1b) and (2). This requires a carefully constructed Control group. 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 $j \in \{1, \dots, J_0, \dots, J_0 + J_1\}$.

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

In our notation $D_j = 0$ if $j = \{1, \dots, J_0\} = \{\text{Control group}\}$, and $D_j = 1$ if $j \in \{J_0 + 1, \dots, J_0 + J_1\} = \{\text{Seagate, Western Digital, Toshiba}\}$. It is important to point out that in the analysis of R&D data the Treatment group only contains the three acquiring firms, i.e. we are excluding Samsung and Hitachi. As we are studying how R&D expenditure (which is firm, rather than market specific) changed for Seagate, Western Digital, and Toshiba, we are uninterested in how R&D spending 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 does not form consensus on the lags between R&D and patents when looking at their impact on company valuation.²⁵ We turn to data to find which number of distributed lags offers the best fitting model (based on Akaike Information Criterion). 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:

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

²⁵(Hall, Griliches, and Hausman, 1984), Pakes (1981), Pakes and Griliches (1980, 1984a), Wang and Hagedoorn (2014).

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:

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

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

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

4.3 Finding a Control: potential sources of bias

To estimate the impact of the merger on our measures of innovation, we need to model what would have happened 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, but 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.

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 the Control is a strategic substitute or complement (e.g. if Seagate innovates more, will the firms in the Control group follow suit).

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). Both the independence and the similarity assumptions will be examined in detail in Section 6.

For unbiased DiD estimates one would also have to assume that the treatment was exogenous. In merger retrospectives that look at the impact of specific mergers this is a

common assumption. However, the assumption of exogeneity is not always defensible, as stressed in Ormosi et al. (2017, section 3.1.1.3). There are, indeed, different reasons why the decision to merge can be endogenous. A first source of endogeneity is reverse causality, which occurs when a drastic change in the outcome variable (innovation activity in our case) is one of the reasons for the merger to take place. A second source of endogeneity is simultaneity caused by self-selection. Gugler and Siebert (2007) discuss the possibility that the firms that merge could be the more productive ones, i.e. those that would have gained higher market shares (innovation activities) even without a merger. Potentially there are simple econometric solutions to both sources of endogeneity. Feedback from the outcome variable to the decision to merge can be brought under control with lagged variables. Similarly, firm selection can be addressed using panel data techniques that control for firm unobserved effects. In our work, we account for both these sources of endogeneity as much as our observed firm characteristics can eliminate these confoundedness concerns.

A third reason for the merger decision to be endogenous is that the decision to merge can be triggered by unobservable time varying factors, such as technology and/or demand shocks, which affect both the decision to merge and the outcome variable (innovation activity). Besley and Case (2000) discuss this source of endogeneity for DiD models. A way to address this type of endogeneity is to instrument the resolution to merge with some measure of distance (geographical or physical) from competitors. This is a technique used in Hastings (2004), Dafny (2008), and Haucaup and Stiebale (2016). This method provides convincing results in aggregate studies, where data is available to model what triggers the merger (e.g. data across many mergers on the technological distance between all merging firms). The problem is that these methods are probably impossible to apply in a case study such as ours, where the treatment only affects one firm, where therefore the merger instance is a single datapoint. For this reason, the credibility of any causal claim regarding our results depends on one’s belief whether our observed firm characteristics are sufficient to dispel any endogeneity concerns. In our case, it is equally possible that the merger was not innovation driven, but instead were motivated by changes in price, or costs, or some other factor that is exogenous to the analysis of our innovation outcome variable.

4.4 SSD/Flash as Control group

There were no firms in the HDD market that did not partake in the mergers, and firms in the HDD market would have likely been affected by the merger, therefore we had to look at potential controls outside the HDD market. We explore product differentiation (HDD, SSD, Flash drives) to find our preferred Control group, SSD products and SSD

manufacturing firms. To eliminate issues regarding pre-merger parallel trend, our preferred method would be to use a synthetic control group, as described in Abadie and Gardeazabal (2003), Abadie, Diamond, and Hainmueller (2010), and Abadie et al. (2015), making use of the matrix completion extension presented in Athey, Bayati, Doudchenko, Imbens, and Khosravi (2018). The idea behind this method is to create a weighted sample of firms that are most similar to the Treatment firm based on a set of observable characteristics. Synthetic control groups are data rather than theory driven, therefore the weighted combination of firms that are selected as Control are the ones where based on the observable variables (total revenue, gross profit, total assets, net income, total debt, expenses, pre-sample R&D expenditure, and the proportion of relevant business segments) the algorithm can achieve the best fit. The synthetic control method relaxes the parallel trends assumption and permits the creation of an optimal synthetic control, a weighted average of the control units which best approximates the treated unit in the pre-intervention period. This produces a synthetic group that follows the same trend as the Treatment units, which by definition should be parallel.

However, missing observations and the lack of balanced panel data for innovation output and for patents meant that we only used a synthetic control in the R&D analysis, and OLS for the rest. When using OLS, we conducted the conventional parallel trend tests.

For innovation output we use SSD drives as Control group, as listed in Table 8 in Appendix D. For the R&D analysis we used other storage (SSD and Flash storage) manufacturers as a pool for potential Controls to generate a number of synthetic control groups.²⁶ Table 7 in the Appendix lists the weights used for each control firm. Finally, for the patent analysis our preferred control group - similarly to above, is other storage technologies (Flash memory/storage) (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 7 in the Appendix) suggests a broad similarity of this Control to the Treatment group. Because there were a number of firms with very few patents, we only used the top 10 firms based on the number of their patents.²⁷

We discuss the suitability of our Control groups below in Section 6.

²⁶We used the R package *gsynth*.

²⁷Cypress Semiconductor Corp., LG Electronics, Macronix International, Micron Technology, Panasonic, SK Hynix, Sony, Taiwan, and Winbond.

5 Results

5.1 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 4 summarises these results for our main variables of innovation output. There are three panels, one for each Treatment firm (Seagate, WD, Toshiba). The columns *user cost* and *numbers* indicate our two measures of innovation output, the unit user cost of storage, and the number of new HDD models respectively.²⁸ All the variables have been standardised, therefore coefficients denote the standard deviation change in response to 1 standard deviation change in the explanatory variable.

We did not have a priori knowledge on how quickly the merger effect would trickle through 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 4 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’ $W \in \{\text{Q1 2012, Q1 2013, and Q1 2014}\}$. 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 4 also indicates for which models the pre-treatment parallel trends were rejected, and we focus on the results where the parallel trends assumption was not violated.²⁹ In our estimation we standardised all non-binary dependent and independent variables - the coefficients can be interpreted accordingly.

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, we found an increase in Seagate’s productivity of patent activity but only regarding

²⁸Note that user cost is an inverse indicator of innovation, a lower unit cost implies higher innovation.

²⁹Throughout this paper for the parallel trend tests we assumed a linear pre-merger trend for both the Treatment and the Control groups and test if these linear trends are parallel.

Table 4: Effect of mergers on innovation output

	user cost			numbers		
Treatment time	Q1 2012	Q1 2013	Q1 2014	Q1 2012	Q1 2013	Q1 2014
Seagate						
DD	0.500	0.703	0.709**	-0.279	-0.129	0.326
std.err.	(0.260)	(0.124)	(0.026)	(0.233)	(0.795)	(0.240)
Sum of R&D lags	0.187	0.156	-0.009	0.883***	0.706*	0.528
pval	(0.556)	(0.740)	(0.988)	(0.006)	(0.053)	(0.237)
Sum of R&D lags x treatment	-0.735*	-1.611**	-2.35***	1.748***	1.565**	1.439*
pval	(0.073)	(0.046)	(0.000)	(0.000)	(0.040)	(0.075)
Sum of patent lags	0.125	0.041	0.06	0.028	0.035	0.091
pval	(0.306)	(0.628)	(0.268)	(0.608)	(0.768)	(0.372)
Sum of patent lags x treatment	-0.374**	-0.489*	-0.7***	0.148	-0.111	0.061
pval	(0.039)	(0.089)	(0.002)	(0.139)	(0.603)	(0.878)
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						
DD	1.015**	-0.536	0.112	-0.706***	1.489***	-3.096***
std.err.	(0.043)	(0.263)	(0.921)	(0.008)	(0.007)	(0.004)
Sum of R&D lags	0.434	0.300	-0.178	0.522	0.175	0.353
pval	(0.160)	(0.563)	(0.779)	(0.011)	(0.595)	(0.412)
Sum of R&D lags x treatment	0.031	1.612*	0.480	-2.520***	-5.595***	4.950**
pval	(0.932)	(0.070)	(0.804)	(0.000)	(0.000)	(0.014)
Sum of patent lags	0.177*	0.109*	0.055	-0.035	-0.085	-0.014
pval	(0.072)	(0.089)	(0.420)	(0.406)	(0.345)	(0.833)
Sum of patent lags x treatment	-0.056	0.528***	0.522	0.554***	-0.885***	3.720***
pval	(0.650)	(0.008)	(0.355)	(0.000)	(0.009)	(0.000)
observations	153	146	139	155	148	141
parallel trend rejected?	Y	N	N	Y	Y	Y
parallel test (p-val)	(0.002)	(0.275)	(0.775)	(0.019)	(0.000)	(0.000)
Toshiba						
DD	5.187***	5.659**	21.63***	-1.878***	-8.810***	14.06***
std.err.	(0.000)	(0.012)	(0.000)	(0.009)	(0.000)	(0.001)
Sum of R&D lags	0.364	0.314	0.037	0.471**	0.403	0.335
pval	(0.178)	(0.567)	(0.954)	(0.026)	(0.154)	(0.323)
Sum of R&D lags x treatment	9.301***	10.244**	46.291***	-3.355	-15.973***	16.839**
pval	(0.000)	(0.016)	(0.000)	(0.001)	(0.000)	(0.012)
Sum of patent lags	0.112	0.126*	0.089	0.091	-0.012	-0.077
pval	(0.373)	(0.063)	(0.215)	(0.521)	(0.914)	(0.387)
Sum of patent lags x treatment	0.493***	0.266	3.887***	-0.134	-0.157	-4.525***
pval	(0.007)	(0.207)	(0.000)	(0.251)	(0.492)	(0.000)
observations	149	142	135	151	144	137
parallel trend rejected?	Y	Y	Y	Y	Y	Y
parallel test (p-val)	(0.001)	(0.045)	(0.001)	(0.000)	(0.076)	(0.000)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$						

one of our output measures (unit user cost).

For WD there is evidence of increased R&D and patent productivity but only for the number of new products. On the other hand, innovation productivity seems to have fallen when looking at the unit user cost of storage capacity. 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 4 also offers some more general findings: R&D spending (productivity of R&D) is associated with increased innovation output, when output is measured by the number of new products, and that patent activity is not crenellated with either measure of innovation output.

5.2 Impact on innovation input

Table 5 shows the DiD coefficients in our innovation input models (R&D intensity and patent factor). All the variables have been standardised, therefore coefficients denote the standard deviation change in response to 1 standard deviation change in the explanatory variable.

Table 5: Estimated changes in innovation input

Method	Seagate		Western Digital		Toshiba	
	R&D	Patent factor	R&D	Patent factor	R&D	Patent factor
	Synth	OLS	Synth	OLS	Synth	OLS
DD	0.016***	0.750**	0.043***	0.828***	0.016**	-0.273
(p-val)	(0.000)	(0.044)	(0.000)	(0.025)	(0.000)	(0.431)
Observations	78	279	78	279	74	279
Parallel trend rejected		N (0.305)		Y (0.023)		N (0.687)
p-val in parentheses						
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$						

It appears that in comparison to other storage technologies, R&D intensity increased for all three HDD firms. Regarding patent activity, there seems to be an increase for Seagate and WD and a drop for Toshiba. The positive sign in the case of Seagate suggests that although patent citation is falling in general (see Table 3), following the mergers it fell at a slower rate in HDD than in our control group of SSD/Flash. For WD's result we rejected the parallel trend assumption. As we show in our detailed results in the Appendix (Table 6) when we use other patent measures - ones where the parallel trends assumption holds - WD had a drop in its patent activity in comparison to the Control firms. Again, if we juxtapose these numbers with our before and after figures in Table 3, then this result suggests that post-

merger WD's and Toshiba's patenting activity dropped at a faster rate than in SSD/Flash patent citations. Whether these effects were caused by the merger depends on the credibility of the assumptions we made for unbiased DiD estimates are reasonably credible. Below we discuss this in more detail.

6 Examining the DiD assumptions

We now return to the question of whether we can attribute the above effects to the mergers. This depends on whether assumptions required for unbiased DiD estimates have been violated, and we now start by looking at the independence of the Control groups used, and the similarity assumption.

6.1 Independence of the Control group

For all the models above we used some combination of SSD/Flash manufacturers/products as Control. The first question we need to answer is the relationship between SSD/Flash and HDD. In their 2011/12 investigation the European Commission concluded that the two are not in the same product market. Igami and Uetake (2017) also acknowledges the difference between HDD and SSD (which may be one of the reasons why SSD firms are excluded from their model of competition). Their justification is based on the argument that there are at least some applications where HDD does not compete with SSD. This is in line with our findings.

Firstly, looking at industry opinion, there seems to be an understanding that although SSD is growing, it is unlikely to be able to keep up with the fast rise in the amount of data available.³⁰ This means that many applications, especially where unit cost dictates purchasing behaviour, are still dominated by HDD (e.g. data servers - cold and archive data, or surveillance data). On the other hand, in some applications SSD is going to be (or has already become) dominant (mobile and local user created data). In each of these areas, competition does not appear very strong between the technologies. For example in data archives, where HDDs are overwhelmingly used, SSD is not considered for the same use as HDD, but used in caching and restoring to speed up data transfers. SSDs also speed up access to storage metadata or are used to boot storage pods. It appears that depending on the application HDD or SSD is the preferred solution. The main driver of this distinction is of course the cost of SSD, which is often prohibitive for many applications. Based on

³⁰See for example: https://www.theregister.co.uk/2017/08/24/ssds_will_not_kill_off_hdds/, or https://www.cloudfest.com/wp-content/uploads/2018/03/15_Toshiba_Rainer-W.-Kaese.pdf.

the above, we would argue, that SSD and HDD are vertically differentiated, where SSD is superior in every characteristics apart from price (if prices were the same, i.e. SSD/Flash was available to the same extent as HDD, most, if not all, users would choose SSD/Flash.

Moreover, the different applications of the two products would imply that in many applications the two are more likely to be complements rather than substitutes - take the above examples of their relative roles in data centres, or hybrid HDD/SSD drives for example ³¹. If the two technologies were substitutes rather than complements then surely there would be no need for manufacturing products that combine the two³². This vertical relationship and product complementarity are important when we need to decide about the independence between SSD and HDD for judging the quality of our Control group.

The next question to ask is if there is dependence between the two technologies, whether the dependence is driven by strategic complementarity or substitutability across R&D decisions in the two technologies, i.e. if HDD increases R&D, will it reduce (strategic substitutes) or increase (strategic complements) R&D spending in SSD. There is rich IO literature on this question for horizontal competition (among duopolists of substitute products), but less on vertically differentiated, complementary products. ³³ Moreover, to our knowledge there is no empirical treatment of the question of strategic behaviour between vertically differentiated complement technologies. For this reason we do not have any a priori assumption about the strategic R&D relationship between HDD and SSD.

Could our data be used to imply something in this strategic behaviour? To find out, we took a few preliminary steps that gave us some tentative information. We relied on various synthetic control methods as referenced above. First, we constructed a synthetic group of firms where we were confident of independence (general IT firms excluding storage manufacturers). In our context, a group of other IT firms may sound too distant to be useful as control, but the very idea of creating a synthetic composition of these IT firms is to create an artificial, data-driven replication of the treatment group. This would suggest that purely for the purposes of replicating the outcome variable of interest (R&D spending), the distance between technologies matters less if we are confident in the scope of our observable characteristics, which are used to create the Control.

We then compared the estimated coefficients under the 'independent' (general control) and the 'potentially confounded' (storage control). The idea was that this could tell us

³¹Hybrid solutions today feature in more than a fifth of all new laptops

³²Jo, Kwon, Kim, Seo, Lee, and Maeng (2009)

³³One way to relate to theory would be to think of our scenario as one where the superior technology is an incumbent (or new entrant), for example as analysed by Cestone and Fumagalli (2005) or under a distinction between R&D as a profit appropriation or as a competitive tool, as discussed by Beath, Katsoulacos, and Ulph (1989).

about the strategic reaction in the SSD based Control group to the Treatment. These results (presented in detail in Appendix A) show no difference between the two groups, which would suggest independence of the SSD control group. Of course, we appreciate the preliminary nature of this test, which is why we take a cautious interpretation of the results. Nevertheless, we believe that this approach could be further advanced to provide a more definite answer on strategic R&D behaviour in any market.

6.2 Similarity of the Control group

Whereas the independence assumption is very difficult to formally test with data, checking pre-treatment parallel trends has become the norm in testing the viability of the similarity assumption.

6.2.1 Innovation output estimates


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 improvements in technology (unlike in typical mature industries where technology tends to be static), and therefore there is still intensive technological progress in HDDs (for example in areal density).³⁴ For both our variables of interest (number of new products and unit user cost) improvements are still ongoing both in HDD and SSD. Figure 4 compares the evolution of unit user cost of disk capacity 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. This would suggest that – at least for this particular characteristic – SSD is not an outlandish choice as Control, when it comes to similarity. As reported in Table 4 the parallel trend assumption was rejected for all models for Toshiba, and for much of WD’s models. Only for Seagate could we reject the null hypothesis of non-parallel trends for both measures of innovation output.

6.2.2 Patent analysis

For patents, we estimated our DiD models for all possible measure of patents. In Appendix B.2 we present a ‘meta-analysis’ of all these estimates. One of the advantages of this approach is that we could select only those models which satisfy the parallel trend assumption (for each patent measure we tested whether the over-time pre-merger slopes were significantly

³⁴<https://www.tomcoughlin.com/Techpapers/HDD%20Market%20Down%20to%20Three%20Suppliers,%20042011.pdf>

Figure 4: Unit cost of storage [$\ln(\$/\text{Gb})$]



different), and then examined the robustness of results only across this selected group of Controls. These are given in Table 6 in Appendix B.2.1.

7 Discussion of results

We presented a detailed discussion of the assumptions required for establishing causality between the mergers and our estimates of how our measures of innovation changed. The causal interpretation of the results comes down to how much the reader is convinced of our assumptions. Our personal view is that we have evidence to claim causality in the case of Seagate, but not for WD and Toshiba. Nevertheless, we decide to remain on a cautious side and refrain from causal claims, and interpret these results as what they really are, estimates of how innovation changed in the Treatment firms, in comparison to the Control firms. Against this backdrop, we offer a number of competing explanations below.

For Seagate we found no evidence of falling innovation activity in Seagate following mergers. Rather we found an increase in the productivity of R&D intensity; moreover, Seagate’s innovation inputs (both R&D intensity and patents) also increased. 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.³⁵ To the extent the newly shared patent pool incorporates previously separate but complementary patents, this is a similar scenario as the one predicted by (Davidson and Ferrett (2007)), where as a result of this complementarity, post-merger innovation increases. The post-merger level of innovation increase could have also been explained by the elimination of duplications between Seagate’s and Samsung’s production and R&D lines (Denicolò and Polo (2018)). Another explanation is that there was increasing pressure from another technology, SSD. Although SSD is often used as a complement to HDD, the two products are vertically differentiated, where SSD is superior in most (if not all) characteristics apart from price, therefore if SSD continues to innovate and further reduce prices, then most consumers would undoubtedly choose SSD over HDD. This pressure could be triggering more innovation in HDD.

For WD we did not find conclusive evidence of R&D productivity for the change in the number of new products, but found that WD unit user cost increased. On the other hand we found evidence of a positive change in R&D intensity and a drop in patent activity. One possible explanation for the 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 duplicated their R&D, production, marketing, and sales operations. This was crippling for WD’s efficiency and the inability to remove duplications could have affected R&D productivity (Denicolò and Polo (2018)). This might also explain some of the increase in R&D (i.e. duplicated spending for the units held separate). Other events might have 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.

For Toshiba, most evidence appears weak, or likely biased, therefore the findings have to be treated with caution. We found evidence of increasing R&D spending after the mergers - in fact this is the only strong result. However, it is worth adding that R&D figures include Toshiba’s other segments (around 25% of Toshiba’s revenue comes from storage related operations) and the change in R&D could be picking up changes in other segments (i.e. a general, not HDD specific drop in Toshiba’s R&D spending). For our analysed time period there was no 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

³⁵Although 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).

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. This might explain why we found evidence of a drop in patent activity. The drop in Toshiba’s patent activity seems to indicate that Toshiba’s acquisition of the divested Hitachi IP rights did not have the expected effect - or simply that these effects were delayed beyond our observation period. 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.

The above results can also be linked to previous theoretical literature. For example Federico et al. (2018) describes that one dimension of merger effects is the reduction of downward pricing pressure by the internalisation of competition. With increased profits more resources become available for R&D and firms will invest in it (if post-innovation profit exceeds the level that would prevail in the absence of extra innovation). In a somewhat related scenario, Jullien and Lefouili (2018) and Bourreau, Jullien, and Lefouili (2018) show how higher post-merger margins can trigger merging firms to increase post-merger demand by increasing innovation (in mergers with product innovation). Both of these are relevant to the HDD market. As HDD firms still compete in quality, it is conceivable that this, and the corresponding drive to expand demand and post-merger profits could have played a role in Seagate. How much of this is induced by the merger, or the expectation of the merger is difficult to tell for reasons discussed above.

8 Conclusion

This paper offered a rare opportunity to simultaneously examine three levels of innovation: R&D spending, patent activity, and the characteristics of new products. On the one hand this unique dataset provided a novel evaluation of the relationship between mergers and innovation, and found mixed evidence of how innovation changed after the merger for the 3 post-merger entities.

The paper also highlighted the difficulties of conducting a case-specific merger retrospective study focusing on innovation effects. Even in our case where data was relatively easily available, causality (the effect of the mergers) was very difficult to establish. Nevertheless, these studies are important. The role of innovation in merger control is attracting increasing interest from academics and practitioners alike. This renewed interest brought the assessment of innovation effects to the forefront of merger-related discourse. The objective of this paper was to contribute to the growing body of theoretical works and aggregate studies, by offering an example of a merger retrospective study on innovation effects. In this respect

our goal was not simply to provide another datapoint of evidence. Instead, we hope that this study would start a stream of similar retrospective studies, in order to create a large pool of evidence, much the same way as is the case with the ex-post analysis of the impact of specific mergers on price.

Given the width of the data we collected, as a more general contribution to the innovation literature, we also found that innovation input (R&D intensity) positively boosts innovation output. Moreover, we found no evidence that patent activity would have an effect on our measures of innovation output in HDD manufacturing. 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 is 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. Innovation levels changed differently for the three HDD manufacturers. 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.

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.
- ATHEY, S., M. BAYATI, N. DOUDCHENKO, G. IMBENS, AND K. KHOSRAVI (2018): “Matrix completion methods for causal panel data models,” Discussion paper, National Bureau of Economic Research.
- BEATH, J., Y. KATSOULACOS, AND D. ULPH (1989): “Strategic R & D Policy,” *The Economic Journal*, 99(395), 74–83.
- 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.
- 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).
- CESTONE, G., AND C. FUMAGALLI (2005): “The strategic impact of resource flexibility in business groups,” *RAND Journal of Economics*, pp. 193–214.
- 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.
- DAVIDSON, C., AND B. FERRETT (2007): “Mergers in multidimensional competition,” *Economica*, 74(296), 695–712.

- 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.
- ENTEZARKHEIR, M., AND S. MOSHIRI (2018): “Mergers and innovation: evidence from a panel of US firms,” *Economics of Innovation and New Technology*, 27(2), 132–153.
- 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, 925–116.
- (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., Z. GRILICHES, AND J. A. HAUSMAN (1984): “Patents and R&D: Is there a lag?,” .
- 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.
- JO, H., Y. KWON, H. KIM, E. SEO, J. LEE, AND S. MAENG (2009): “SSD-HDD-hybrid virtual disk in consolidated environments,” in *European Conference on Parallel Processing*, pp. 375–384. Springer.
- LANJOUW, J. O., AND M. SCHANKERMAN (2004): “Patent quality and research productivity: Measuring innovation with multiple indicators,” *The Economic Journal*, 114(495), 441–465.
- LETINA, I. (2016): “The road not taken: competition and the R&D portfolio,” *The RAND Journal of Economics*, 47(2), 433–460.
- 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.

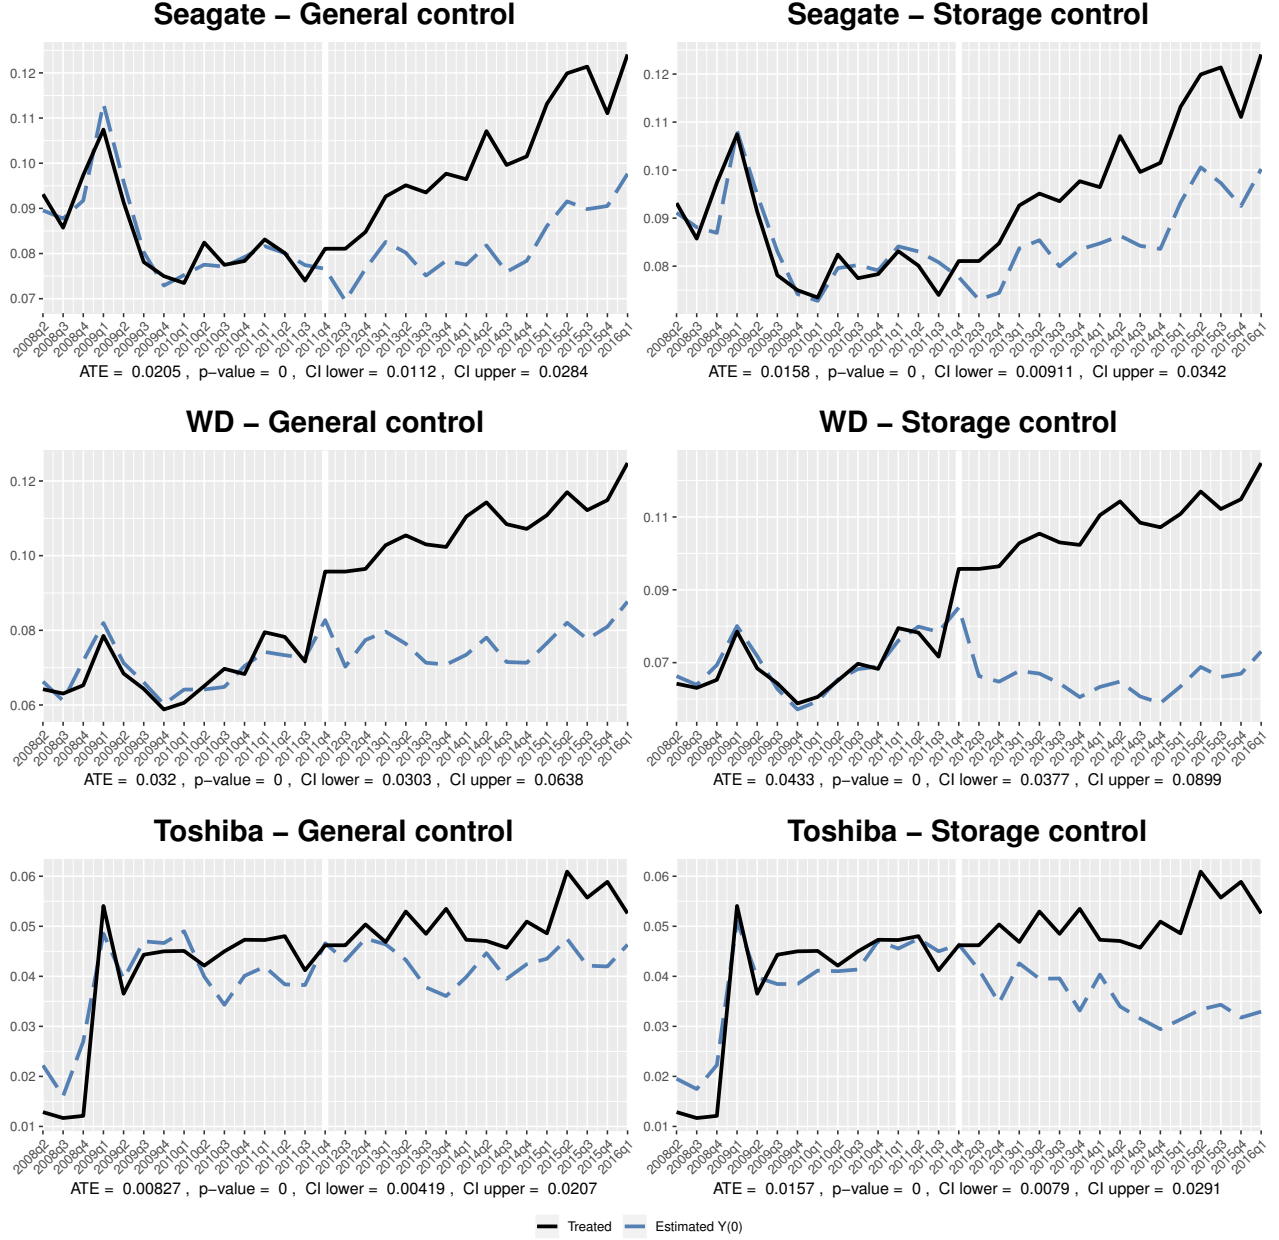
9 Appendix

A Strategic complements or substitutes?

We looked at what the data can tell us about the strategic R&D behaviour between HDD and SSD. For this we applied a simple test, in which we compared two potential Control groups, one that is very likely independent (a synthetic composition of IT firms excluding storage manufacturers, we call it ‘General control’), and one that is potentially affected by the strategic relationship, SSD manufacturers (‘Storage control’). The idea is that if SSD R&D activity is unaffected by the mergers, then post-merger the Storage control response to the merger will


be similar to the General control's response. If however the Storage control response, in comparison to the General control, is directly/inversely proportional to the Treatment group's response, that would be a sign of strategic complementarity/substitutability between HDD and SSD.

Figure 5: Change in HDD in comparison to two different Synthetic control groups



First, in Figure 5 we present the results where we compare HDD to the two types of synthetic controls, the General control (without storage firms) and the storage-based synthetic control.³⁶ Figure 5 suggests that there is a small difference between the results for the two Control groups, although the difference is not significantly different. This becomes obvious if we compare the General with the Storage control,³⁷ we find no significant post-merger difference (Figure 6). Taken at its face value this would suggest that R&D spending in the group that we suspected to be confounded (SSD/Flash) is no different from the group where we can more confidently assume independence.

Figure 6: Comparing SSD (treatment) with a general synthetic group



³⁶Below in Table 7 we show the firms with highest weights

³⁷This was achieved by creating a synthetic group of the general to match the storage group

B Robustness checks

B.1 R&D expenditure results

We run a number of OLS models using three other potential Control groups: a other storage firms (SSD and Flash Drive), and an unweighted and a PSM weighted group of other IT firms. These models produced results that led to the same conclusion as the synthetic control results on how R&D changed after the merger. A detailed discussion of these other models, the corresponding estimates, and their robustness checks can be found in Ormosi et al. (2017).

B.2 Patents

To check the robustness of our patent results, we run a large number of DiD models for various patent measures and model specifications. 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) standardised measures, and (iv) normalising the patent measures by total revenue 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 45 different measures of patent activities.

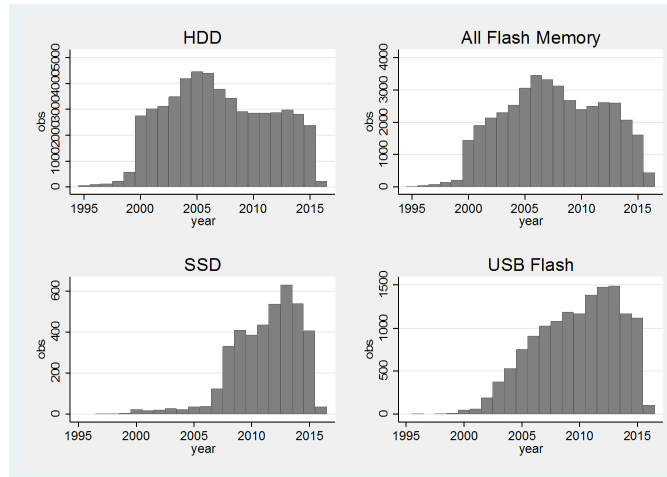
Then, separately for Seagate, Toshiba and Western Digital, we estimate the a DiD model on each of the counts of patent activities. In order to make results comparable we standardise all continuous variables, including the control variables total debts, total assets, total revenue and total R&D intensity (all up to four lags). The procedure gives, for each of the three companies, 45 standardised DiD estimates of patent activity. We combine these estimates using a meta-analysis approach and obtain the average effects and the distribution of estimates. In doing so we use three different Control groups: the one presented in the main text (Flash/SSD), and two others, introduced below.

B.2.1 Different control groups in the patent analysis

Flash/SSD patents: In our preferred results we used firms with Flash related patents as Control. Figure 7 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 can potentially satisfy the similarity assumption, which is what we used as Control in the main text.

Figure 7: Number of HDD, Flash Memory, SSD, and USB Flash patents per year



Other HDD patents: This group consists of the top 10 firms in terms of the number of HDD-related patents held in our data.³⁸ These patents are not innovations of the HDD units but innovations on something complementary to HDD.³⁹ 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.

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

³⁸The list of Control firms includes: Canon, Funai, Hon Hai Precision Industry, IBM, Inventec, Lenovo, LG Electronics, Panasonic, Ricoh, and Sony.

³⁹For 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.

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 could imply that the estimates are biased downward.

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.

We tested for parallel trends in all models. Those estimates, where the parallel trend assumption was rejected were filtered out. Table 6 provides two sets of results. The top half of the table includes all the estimates, and the bottom half only the ones where the parallel trend assumption could not be rejected.

Table 6: The effect of the mergers on patent activity

Control	Seagate	WD	Toshiba
All regressions			
Flash/SSD	0.375*** (0.000)	0.168*** (0.000)	-0.491*** (0.000)
Other HDD	0.549*** (0.000)	0.517*** (0.000)	-0.383*** (0.000)
Top Storage	0.386*** (0.000)	0.309*** (0.000)	-0.546*** (0.000)
Parallel trends only			
Flash/SSD	0.231*** (0.001)	-0.125** (0.023)	-0.489*** (0.000)
Other HDD	0.356*** (0.000)	0.232*** (0.000)	-0.364*** (0.000)
Top Storage	0.245*** (0.000)	-0.172*** (0.000)	-0.080 (0.201)
pvals in parentheses			
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$			

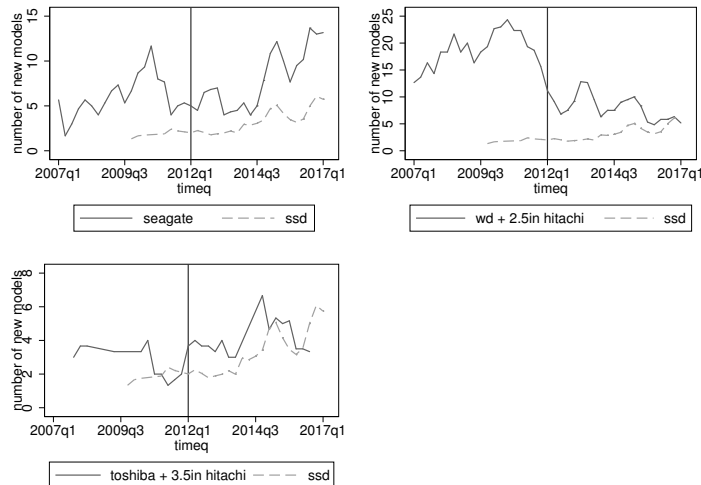
The main finding of increasing patent activity seems consistent for Seagate across the three models with alternative Control groups. For Western Digital the change depends on what the Control group is. Patent activity increased in comparison to other HDD patent holders (excluding the merging firms) but decreased in comparison to Flash. Moreover, the sign of the estimate using Flash/SSD as Control changes when we eliminate the estimates where the parallel trend assumption was likely violated. For Toshiba, the drop in patent activity seems confirmed in all three cases.

C Further tests on innovation output estimates

C.1 Assumptions required for unbiased DiD estimates

Figure 8 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 8: Per-firm quarterly average number of new products marketed on Amazon



It immediately stands out from Figure 8 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 8 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 9 shows that there has been a continuous decrease in the unit user cost of SSD capacity. HDD on the other hand displays a mixed picture. Unit user cost has been steadily falling for Seagate and WD (steeper for Seagate), and fell first then levelled out for Toshiba.

Figure 9: 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) - displayed in Table 4. However, the main story here does not hinge on our DiD estimate. Rather, it is about the effect of previous R&D spending and patent activity on product numbers and unit user cost. This is also important regarding the independence assumption required for unbiased DiD, because, strictly speaking, in this respect even the choice of our Control group is irrelevant here. To illustrate why, take the example of Seagate. We have shown how R&D spending in Seagate increased for all Control groups. Here we simply show that this increased R&D activity is associated with an increased number of new products and lower unit user cost.

In Ormosi et al. (2017) we also present 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.

D Brands/firms used as Control

D.1 Firms with largest weights in the R&D analysis

Table 7: Synthetic control weights - storage only

General		Storage	
weight	firm name	weight	firm name
Seagate			
-22.942	ShoreTel, Inc. (NasdaqGS:SHOR)	-1.1852912	Trek 2000 International Ltd
-22.510	Chicony Electronics Co., Ltd. (TSEC:2385)	-0.7966607	SanDisk LLC
-21.512	KongZhong Corporation (NasdaqGS:KZ)	-0.638343	I-O Data Device, Inc. (TSE:6916)
-18.255	TechMatrix Corporation (TSE:3762)	-0.1770924	Micron Technology, Inc. (NasdaqGS:MU)
-10.888	Soft-World International Corp. (GTSM:5478)	-0.1649102	Quantum Corporation (NYSE:QTM)
5.083	Elite Material Co., Ltd. (TSEC:2383)	0.2431215	Transcend Information, Inc. (TSEC:2451)
6.349	Taiwan Pcb Techvest Co., Ltd. (TSEC:8213)	0.2894449	Power Quotient International Co., Ltd. (TSEC:6145)
7.012	Nihon Dempa Kogyo Co., Ltd. (TSE:6779)	0.4821508	Chaintech Technology Corporation (TSEC:2425)
7.816	Samyoung Electronics Co., Ltd (KOSE:A005680)	0.7167503	Gigastorage Corporation (TSEC:2406)
16.558	Wistron NeWeb Corporation (TSEC:6285)	0.9628517	Hitachi, Ltd. (TSE:6501)
WD			
-9.989	Kona I CO., Ltd. (KOSDAQ:A052400)	-7.737355	Powerchip Technology Corp.
-9.600	TechTarget, Inc. (NasdaqGM:TTGT)	-3.4934264	Micron Technology, Inc. (NasdaqGS:MU)
-9.550	Chicony Electronics Co., Ltd. (TSEC:2385)	-2.1280186	SanDisk LLC
-7.844	Skyworks Solutions, Inc. (NasdaqGS:SWKS)	-0.7447585	Lite-on Japan, Ltd. (JASDAQ:2703)
-5.356	Ssangyong Information & Communications Corp. (KOSDAQ:A010280)	-0.4181387	NetApp, Inc. (NasdaqGS:NTAP)
2.411	Ebix, Inc. (NasdaqGS:EBIX)	0.5300235	Transcend Information, Inc. (TSEC:2451)
2.982	Taiwan Surface Mounting Technology Corp. (TSEC:6278)	0.932503	Gigastorage Corporation (TSEC:2406)
3.354	Nihon Dempa Kogyo Co., Ltd. (TSE:6779)	3.617181	Hitachi, Ltd. (TSE:6501)
3.715	Samyoung Electronics Co., Ltd (KOSE:A005680)	4.257823	Quantum Corporation (NYSE:QTM)
7.899	Wistron NeWeb Corporation (TSEC:6285)	4.4763798	Trek 2000 International Ltd
Toshiba			
-14.288	TechTarget, Inc. (NasdaqGM:TTGT)	-0.4676411	SanDisk LLC
-11.056	KongZhong Corporation (NasdaqGS:KZ)	-0.2251632	ADATA Technology Co., Ltd. (GTSM:3260)
-7.635	Chicony Electronics Co., Ltd. (TSEC:2385)	-0.1588192	I-O Data Device, Inc. (TSE:6916)
-6.908	Ssangyong Information & Communications Corp. (KOSDAQ:A010280)	-0.1566601	Chaintech Technology Corporation (TSEC:2425)
-6.132	ASM Pacific Technology Limited (SEHK:522)	-0.1396133	Power Quotient International Co., Ltd. (TSEC:6145)
2.833	Ebix, Inc. (NasdaqGS:EBIX)	-0.0358594	Trek 2000 International Ltd
3.413	Taiwan Surface Mounting Technology Corp. (TSEC:6278)	0.1290028	Lite-on Japan, Ltd. (JASDAQ:2703)
4.034	Nihon Dempa Kogyo Co., Ltd. (TSE:6779)	0.2068797	Micron Technology, Inc. (NasdaqGS:MU)
4.392	Samyoung Electronics Co., Ltd (KOSE:A005680)	0.4984574	Hitachi, Ltd. (TSE:6501)
9.431	Wistron NeWeb Corporation (TSEC:6285)	0.7731834	Quantum Corporation (NYSE:QTM)

D.2 Brands in the innovation output analysis

Table 8: Brands in the innovation output analysis

SSD			HDD		
brand	Freq.	Percent	brand	Freq.	Percent
ableconn	11	0.85	hitachi	190	11.82
adata	28	2.17	samsung	47	2.92
apple	20	1.55	seagate	438	27.24
arch memory	14	1.08	toshiba	146	9.08
axiom	51	3.94	wd	787	48.94
corsair	24	1.86			
crucial	30	2.32			
dell	25	1.93			
edge	11	0.85			
hp	65	5.03			
intel	147	11.37			
kingdian	40	3.09			
kingspec	33	2.55			
kingston	43	3.33			
lenovo	38	2.94			
micron	19	1.47			
mushkin	13	1.01			
mydigitalssd	12	0.93			
ocz	40	3.09			
other_ssd	267	20.65			
owc	61	4.72			
patriot	11	0.85			
plextor	22	1.7			
pny	10	0.77			
samsung	93	7.19			
sandisk	40	3.09			
seagate	13	1.01			
silicon power	10	0.77			
super talent	19	1.47			
systor	11	0.85			
toshiba	20	1.55			
transcend	39	3.02			
visiontek	13	1.01			
Total	1,293	100	Total	1,608	100