Mergers and innovation: Evidence from the hard disk drive market

https://doi.org/10.1016/j.ijindorg.2021.102755

PUBLISHER

Elsevier BV

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

This paper was accepted for publication in the journal International Journal of Industrial Organization and the definitive published version is available at https://doi.org/10.1016/j.ijindorg.2021.102755

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

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

May 12, 2021

Abstract

This paper is a retrospective evaluation of how innovation changed following mergers and subsequent policy interventions after the 5-to-3 consolidation of the worldwide hard disk drive industry in 2012. It adopts a holistic view of innovation, employing four different measures: R&D, patents, the number of new models, and their unit prices. This allows us to distinguish the magnitude of the merging parties' innovative efforts from the productivity of those efforts. Our firm-level approach confirms that there is important heterogeneity across the players, which we attribute to differences in the severity of remedies required by competition authorities.

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

*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. The paper benefited from comments received at the 2019 IIOC conference, the 2018 CCP Seminar Series, and seminars at the CMA and Ofcom. We are also grateful for the comments received from Adriaan Dierx, Niklas Duerr, Tomaso Duso, Giulio Federico, Fabienne Ilzkovitz, Carmine Ornaghi, and Mark Roberts. The data and our main codes are available at: `https://github.com/PeterOrmosi/hdd_innovation`

[†]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: s.w.davies@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, Centre for Competition Policy, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom, email: p.ormosi@uea.ac.uk

1 Introduction

This paper examines how the consolidation in the worldwide hard disk drive (HDD) market affected innovation. This involved three mergers in 2011/12 (Seagate/Samsung, Western Digital (WD)/Hitachi, and Toshiba/Hitachi's 3.5-inch production), which reduced the number of competitors from 5 to 3 firms. The first two were agreed with the competition authorities subject to remedies, one of which involved a divestment of certain assets to Toshiba, and we count the acquisition of these divested assets as the third merger.

The paper contributes to two different strands of literature. The first is the policy-driven literature on retrospective evaluations of the competitive impact of mergers. Although that literature is voluminous, it almost entirely focuses on just the price effects of mergers: two studies, Ashenfelter et al. (2014) and Kwoka (2014) identified and reviewed more than 60 such case studies, and Mariuzzo and Ormosi (2019) further disaggregated these retrospectives and reviewed over 600 market-level price-impact estimates. In stark contrast, however, we find much less on how mergers affected innovation. This paper addresses that gap. Second, it adds to the wider literature on the relationship between competition and innovation. Here, the seminal works by Arrow (1962) and Schumpeter (1942) have generated an enormous body of subsequent research on whether competition stimulates or discourages innovation, but few writers have examined the impact of mergers. As Gilbert notes in a contemporary comprehensive review: "Only a few empirical studies apply sophisticated statistical techniques to uncover the effects of mergers on R&D and innovation" (Gilbert, 2020, p.6). Furthermore, amongst the few that there have been, most have provided only aggregate, and sometimes rough, evidence, summarising the average effect across large samples of markets.

Remarkably, few have focussed on specific markets, and again this paper hopes to address this gap with an in-depth case study.

This particular case itself is of interest from a variety of perspectives. In policy terms, a reduction from 5 to 3 in any market is sufficiently concentrating to alert the interest of competition authorities: a common policy question is whether 3 is sufficient to maintain competition? Analytically, it provides an unusual chance to assess comparatively three mergers occurring more or less simultaneously, within the same market. On the other hand, it poses problems when seeking for a viable comparator. Geographically, it is a worldwide market and therefore spatially unique; in terms of product space, while there are other industries in the same broad sector of data storage, they are at different stages of their product cycle, which makes supply-side comparability difficult. In this respect, the case is representative of many others in a world of globalised markets.

We attempt to contribute to both academic methodology and merger policy debates. Methodologically, we depict "innovation" as multifaceted, involving measures of inputs (R&D spending and patent activity) and "outputs" of innovations (the number of new models and the quality-adjusted unit price of HDD capacity). For this purpose, the paper assembles a unique dataset of four different measures of innovation in a single case study. This is a more holistic approach than is usual in the previous empirical literature, and the greater breadth takes us back closer to some of the typologies to be found in the founding writers in the technology literature - Schumpeter's distinctions between invention and innovation, and process versus product innovation. It also allows us to distinguish conceptually between the effects of the merger on the amount of R&D (say) and the productivity of that R&D.

To make causal inferences from this data we use a matrix competition method, which generates the counterfactual scenario (the no-treatment outcome for the treated units) from the observed elements of the matrix of control outcomes corresponding to untreated units/periods. Matrix completion methods relax on the parallel trend assumption, which would be required in a difference-in-difference setting and which is something we cannot assume away on at least some of our outcome measures.

The paper provides significant evidence of increased R&D but falling patent activity following the mergers. We argue that the fall in patent numbers is most likely a reflection of the removal of duplication and defensive patenting. Regarding the innovation outputs, there is a heterogeneous response across firms, with Seagate and Toshiba, increasing, but Western Digital reducing the number of their new models. To examine how innovation input drives innovation output, we bring our measures together in a fixed effect model, to show that the merger contributed to improved R&D productivity for Seagate. On the other hand, for Western Digital there is no evidence that the increased R&D contributed to increased innovative outputs - they significantly reduced the range of marketed new products. In the case of Seagate, we interpret this as a sign of important synergies between Seagate and Samsung, which were substantiated by the merger. We believe that the heterogeneous innovation response to the mergers could be a sign of the difference in the intensity of merger interventions, as Western Digital was effectively held separate from Hitachi for years after the merger, whereas Seagate was objected to much less severe remedies.

As a retrospective merger study, we hope the case is of interest in its own right - enabling us to assess the competition authorities' decisions on whether or not to permit the mergers, albeit subject to remedies. As just mentioned, these mergers combined to reduce the number of participants from 5 to just 3 in a strategically important hi-tech industry - in such a case, it is important to scrutinise the decisions of competition authorities. More generally, and looking forward, innovation-focused ex-post evaluations are likely to become ever more important in the competition policy debate on information-intensive markets. Hopefully, this paper offers some insights into some of the conceptual and methodological which will confront similar future case studies.

1.1 Literature review

Our paper draws on various branches in the previous literature, beginning with the age-old, debate about the relationship between competition and innovation.[1] Much of this was stimulated by the seminal contributions of Arrow (1962) and Schumpeter (1942), and their two conflicting insights continue to run through current debates. On the one hand, where firms already enjoy market power, this can deaden their incentive to innovate because the innovation will largely displace their existing profits (Arrow). On the other, competition can discourage innovation if it constrains the potential innovator's ability to appropriate the profits from its innovation (Schumpeter). Two themes in particular from this wider competition-innovation literature are directly relevant to the current paper. First, while the balance of the literature, both theoretical and empirical, points in favour of competition, there is by no means unanimity. There is perhaps a current consensus around the inverted U, whereby competition does enhance innovation, but only up to some point, after which further increases in competition destroy the ability of potential innovators to appropriate the fruits of their innovation (Aghion et al., 2005, 2009). However, even if one accepts this more

[1] This review benefits from the appearance of Gilbert's (2020) book. This has helped structure our own thinking, and has obviated the need to include an extensive literature survey here – see his chapters 3 ("Competition and innovation basics: Arrow versus Schumpeter") and chapter 6 ("Competition and innovation: empirical evidence").

nuanced view, the immediate implications for case studies such as the current one are not clear cut since it all depends on how much is 'too much'. As will be seen, we will approach our case without strong priors about whether the mergers would have positive or negative impacts on innovation.

Second, as has long been recognised, 'innovation' is singularly difficult to measure, or even proxy, in an entirely convincing way (Gilbert, 2020, pp.110-1). The two most common measures, R&D intensity and patents, can both be dated back to the very early studies of, inter alia, (Griliches, 1979, 1990), and Scherer (1983). However, care is needed in how we interpret data and results on both patents and R&D. R&D is essentially an input, rather than an output, of innovation, and, given inevitable technical uncertainty, much R&D will fail to generate any innovation. On the other hand, improved efficiency in R&D labs may be reflected in a falling R&D input but without loss in innovative output. It should also be recognised that many innovations are not the result of formal R&D.[2]

Perhaps less obviously, similar concerns also relate to patents. In many contexts, the patent is better thought of as an invention, which may, or may not, lead to a commercially viable innovation, and which will often require considerable development expenditures. In that sense, patent data are also, at least in part, a measure of input as well as the output of innovation. Citation-weighted patents are sometimes used in an attempt to quality-adjust patent counts, but it remains the case that many even promising patents never lead to measurable innovation. Moreover, patents are often used negatively to protect an incumbent's market power (p.109 Gilbert, 2020; Cohen et al., 2002; Gilbert and Newbery, 1982), and

[2] This was first established 60 years ago by Jewkes et al.'s (1958) seminal book on twentieth-century innovations, which found 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".

likewise, some R&D is therefore essentially defensive - devoted to finding ways of denying innovations to others, rather than generating innovation themselves.

Turning specifically to the mergers-innovation literature, it should not be assumed that, when assessing the impact of a merger, we need only to refer back to the general competition-innovation literature. A merger is not necessarily synonymous with a reduction in competition, and, in the context of a research and innovation, it may lead to efficiency gains from synergies and reduced duplication which might increase rivalry and competition. With this in mind, we now consider mainly the literature on product innovation, which will be the point of interest in the HDD case.[3] Some theoretical insights resonate with results from the parallel, but much larger, the literature on the price-effects of merger: in much the same way as a merger can reduce the constraint on price rises, by reducing diversion, it can also reduce the incentive to innovate (for Arrowian reasons). In other respects, the analysis is more complex. This is because, while the price has a direct impact on the demand curve, here the firm's policy variable (invariably R&D) impacts only indirectly on the demand, via its effect on product innovation. The nature of the "innovation production function" then becomes a key issue: are there diminishing returns? Are there spill-overs, both within and across firms? Is there scope for heterogeneity in research paths? We take a more detailed look at this in Section 1.2.

On the empirical side, there have been remarkably few studies, and virtually all of these are aggregate, in the sense that they look typically at the relationship between the number of mergers and aggregate R&D spends and/or patent counts. Most have been in the

[3]A good illustration of the recent literature on process innovation is Motta and Tarantino (2017) who employ a model with simultaneous price and cost-reducing investment choices and find that, absent efficiency gains, the merger lowers total investment.

pharmaceutical sector: (Danzon et al., 2007, 165 mergers), (Ornaghi, 2009, 27 mergers), and (Haucap et al., 2019, 65 mergers), all find robustly significant average negative impact of mergers on innovation. Szücs (2014) observes 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. However, other studies find increases in R&D activity after mergers, including Bertrand (2009), and Entezarkheir and Moshiri (2018) who report increased patenting activity following mergers.[4] Stiebale (2013) focuses on acquirers (324 firms) and finds that their R&D intensity significantly increases after mergers Valentini (2012) focuses on a sample based on all mergers in the US 'medical devices and photographic equipment industry, 1988-1996, and finds that mergers had a positive effect on patenting output, but decreased patent impact, originality, and generality.

A work which is apparently close to our own, because it also concerns mergers in the HDD market, is Igami and Uetake (2020). The authors simulate a dynamic oligopoly model, calibrated with data for 1996-2016. The purpose is to endogenise merger decisions as the result of discrete choice for the firm between exit, investment in productivity and merger. Exit and merger are observed directly (there were 6 mergers, 5 exits and 1 entry over this period). 'Investment in productivity' is the measure of innovation, and this is imputed from estimates of marginal costs backed out in a familiar way from the estimated demand function and firms' output decisions. Using these estimates they then use counterfactual simulations which suggest that any innovation-dampening from mergers is minimal so long as there are at least 3 players in the market. The relevance for our own paper is obvious: following the

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

3 mergers we examine, the market still had 3 survivors, and the implication of Igami and Uetake is that this should not have negatively affected innovation.

However, as will be seen, there are three very important differences between their model and ours. First, their estimates relate to process, rather than the product, innovation; second, their estimates of innovation are only indirect, conditional on the assumptions of their structural model, including the nature of aggregate demand curve, the Cournot behaviour of the players, and the sequential ordering across firms when they choose between exit, mergers and innovation. In our work, we are concerned with product innovation, i.e. improvements in the hard drives themselves, rather than the processes which the players use to produce the drives. Also, we employ directly observed measures of innovation. Methodologically, we employ reduced form, and we believe that this provides a useful complement to Igami and Uetake's structural approach.

We draw three main conclusions from this review. First, as noted in our introduction, on the empirical side, Gilbert's contemporary survey bemoans the scarcity of studies specifically address(ing) the effects of mergers on innovation. Indeed, viewed from the antitrust perspective of retrospective case studies, assessing the impact of merger and merger control, there is a stark contrast between the enormous number of quantitative evaluations of the price impact of mergers, and little or none on the innovation impact.

Second, it is potentially misleading to equate "innovation" simplistically with either R&D or patents. R&D is an input measure which, in itself, tells us little about the efficacy of research. Indeed, it is not necessarily the case that 'more is better', if the merger was motivated by, and succeeds in, removing wasteful duplication. For patents the distinction between inputs and outputs is less clearcut – sometimes a key patent may be immediately translatable into a commercially viable innovation, but in many other cases, patents turn out to have no commercial value. Even worse, defensive patenting (designed to thwart competitors' innovation) can have a negative impact on innovation at the industry level. This suggests that a holistic approach is called for. We attempt to do this by drawing together different proxies for innovation which allow us to disentangle the quantity of innovative effort (R&D and patents) from its efficacy, as indicated by some more direct measures of output (improved product characteristics).

Third, within the theoretical literature, there are arguments both ways on whether, in general, mergers encourage or discourage innovation, and the same is true for the aggregate multi-merger empirical studies. The balance of theory and empirical evidence probably points to a generally negative effect, but this is by no means conclusive. In our view, a more important lesson to draw from the theory is that mergers are heterogeneous in numerous dimensions (across markets, firms, and measures of innovation), and this heterogeneity, both within and across mergers, could explain the mixed results. The following section further exposes this point.

1.2 Heterogeneity in the innovation response to mergers

To motivate our work, and to facilitate the understanding of the heterogeneity in the innovation response to mergers, we draw on Carl Shapiro's innovation principles: contestability, appropriability, and synergy and organise previous literature along this trichotomy. Shapiro (2011) himself also argues that this framework is a useful tool to use in merger control to predict the innovation impact of a merger. Moreover, it will enable us to structure our understanding of what happened in the HDD market and facilitate the interpretation of our results.

Heterogeneity driven by synergy

By synergies here we refer to two things, innovation related synergies, and cost savings realised by the merger. Innovation efficiencies can overturn the reduction in innovation due to the internalisation of negative externalities, and could offset the negative impact of a merger on consumer welfare (Federico et al., 2018), and in general Motta and Tarantino (2017) come to the conclusion that a merger may lead to increased investment (including R&D investment) if there are significant synergies. A related concept is the variety of innovation and duplications, which are both likely to drop after the merger (Letina, 2016). Whereas contestability and appropriability are linked to the incentive to innovate, synergies are related to the ability to do so (Shapiro, 2011). Both the Seagate/Samsung, and the WD/HGST mergers were expected by the merging firms to lead to significant cost savings,[5] although we have no information on specific innovation synergies. Nevertheless, the rationalisation of production could free up resources to invest in R&D. Without the realisation of these synergies, it was less likely that firms would have been able to increase their innovation activity post-merger. If the merging firms differed in how much synergy they realised from the mergers, it would have created some heterogeneity in their ability to increase innovation post-merger.

Heterogeneity driven by contestability

One source of heterogeneity in how firms' innovation activities respond to a merger is the

[5] `https://www.forbes.com/sites/greatspeculations/2014/12/29/how-will-hgst-integration-impact-western-` `?sh=3de85bb42fce`.

difference in the contestability of firms' demand. Contestability (often denoted as innovation externality) refers to the effect of one firm's innovation on another firm's demand. With a merger, this effect is internalised. Most works on innovation and mergers highlight the importance of how the merged entity internalises pre-merger R&D externalities (Farrell and Shapiro, 2010; Whinston, 2011; Motta and Tarantino, 2017; Federico et al., 2018; Jullien and Lefouili, 2018). If innovation in a technology produces mainly positive externalities, a merger could enhance innovation by internalising these externalities. On the other hand, if innovation is more likely to generate negative externalities, the opposite would be true. Our own intuition is that there is reasonably high contestability in the HDD market – an innovation by one firm is likely to eat into the demand of their rivals (negative externality), which, *ceteris paribus*, would predict no merger-induced increase in innovation if this effect dominates. A priori we see no reason that the three post-merger firms would largely differ in terms of contestability.

Heterogeneity driven by appropriability

Appropriability refers to the ability of the innovating firms to capture the benefits resulting from the innovation. Mergers, where appropriability was high pre-merger, may add little incentive to further innovation spending. On the other hand, where innovation spill-over (or imitation) is high, merging parties are more likely to internalise positive knowledge spill-overs, thus enhancing their incentives to innovate (Motta and Tarantino, 2017). Moreover, a merger may bring together complementary R&D assets, which would make it more difficult for rivals to successfully imitate and easier for the merged firms to appropriate the social benefits (Motta and Tarantino, 2017; Jullien and Lefouili, 2018; Federico et al., 2018). Appropriability was likely to be high in the HDD market pre-merger, as we will show for example through the high patenting rate. In this respect, the mergers are unlikely to contribute to internalising innovation spill-overs. On the other hand, the mergers do bring together complementary products and technologies, which would suggest increased innovation. For example, Seagate desperately needed a partner like Samsung to secure NAND shipments in the future in order to make Seagate more competitive in the enterprise segment.

Other potential sources of heterogeneity

A number of other sources are worth mentioning here, which could enhance or counter-balance the above effects. If the dominating effect of the mergers is due to the increased demand of the post-merger entity, R&D investment could lead to enhanced profits, which would make it more likely for innovation to increase post-merger (Bourreau et al., 2018; Jullien and Lefouili, 2018; Shapiro, 2011). Related to this is the market power effect of mergers, which could lead to higher prices, and more return to investment and hence boost post-merger innovation (Federico et al., 2017). What also matters is the type of innovation. Innovations that boost demand are more likely to increase than innovations that increase the profit margin (process innovation).

The main point this systematic review of the theoretical works conveys is that the trade-offs in the merger-innovation question run across numerous dimensions. Which effect dominates depends on the specific circumstances of a merger. For this reason, individual case studies are desirable to help flesh out more of the detail on the mechanisms of how innovation changes, and on how these may differ between firms within the same industry. This was a motivation for our choice of this particular case study, which involves five different insiders in three different, but closely related, mergers within the same industry.

2 The market and the mergers

This paper examines three related mergers in the world Hard Disk Drive (HDD) market in 2011/2012. This section briefly describes the data storage market, including Hard Disk Drives, and then the mergers and judgements of the relevant CAs.

2.1 The data storage market

There are two main storage technologies, Hard Disk Drives (HDD), and Flash-based (NAND) storage. An HDD is a device that uses one or more rotating disks with magnetic surfaces (media) to store and allow access to data, whereas Flash storage uses integrated circuit assemblies to store data, which records, stores and retrieves digital data without any moving parts. Solid-state drives (SSD) and USB Flash drives are Flash memory-based storage. SSDs are built on semiconductor memory arranged as a disk instead of magnetic or optical storage support.

Because SSDs have no mechanical components, they are faster than HDD, providing access to data in microseconds, instead of the several milliseconds requested by HHDs. Their other main advantages include their smaller size, lower power consumption, increased resistance to shock, and reduced noise and heat generation. However, a major disadvantage of SSDs is their price, although SSD capacity size has been rapidly increasing and unit price falling, HDDs remain the primary technology for archiving, with SSDs mainly employed in portable devices (laptops, smartphones, tablets). So, notwithstanding the major technological improvements in SSD, the rapid growth in demand for data archives and cloud storage, has continued to favour HDDs. Storage in mobile devices, using flash-based technologies,

Table 1: The mergers and remedies

Merger/acquisition	Notification date	Interventions by CAs EU&USA	MOFCOM
Seagate/Samsung	19th April 2011	Cleared without remedy	Behavioural remedies
WD/HGST	20th April 2011	Divest 3.5in desktop HDD manufacturing to Toshiba	Behavioural remedies
Toshiba/HGST 3.5-inch	28th Feb 2012	none	none

remains only a small fraction of all storage capacity, despite its wide dissemination.

2.2 The mergers

The market structure of the HHD market had witnessed continuous consolidation since the late 1980s, and at the date of these 2011/12 mergers, there were only five players left in the market: Seagate, Western Digital (WD), Toshiba, HGST, and Samsung. Following the mergers, the market shares of Seagate, Western Digital, and Toshiba were close to a 40-40-20 split. The SSD market, on the other hand, was more fragmented, unsurprisingly, as it was a less mature technology. The major players in SSD were Samsung, Toshiba, SanDisk, Micron, SKHynix, and Intel.

Table 1 summarises the main regulatory information regarding the three mergers. On 7 March 2011, WD and Hitachi 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. The following month, on 19 April 2011, Seagate announced its 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 by European and US authorities,[6] the prevailing view was that Samsung had not exerted any appreciable competitive pressure in the market, and that its elimination would not affect the level of competition. The WD/Hitachi merger was also approved by the European Commission and the US authorities, but in this case subject to WD committing to divest the 3.5-inch desktop HDD manufacturing lines to Toshiba. Thus the third of our mergers was the result of a remedy agreed in connection with the second merger. The divested HGST assets were acquired by Toshiba in a transaction announced in February 2012. However, the Chinese CA (MOFCOM), took a different stance towards the mergers and required a set of behavioural remedies. While we found no evidence that these involved any restrictions on the transfer of intellectual property rights, it did impose remedies on R&D activities. Specifically, it required the parties to hold R&D activities separate post-merger, and not to reduce them below a specified level. In the case of Seagate, it was required to spend at least $800 million annually on R&D. However, this was relatively soft, in that Seagate's average annual R&D spending pre-merger had been above $900 million. On the other hand, the remedy was harsher for Western Digital, which was required to maintain its pre-merger level of R&D expenditure – it had no scope to rationalise on R&D.

More generally, 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

[6] Following its priority principle, the Commission held that a party that is the first to notify a concentration (Seagate) which, 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, in the context of a market structure that reflected conditions after the Seagate/Samsung merger.

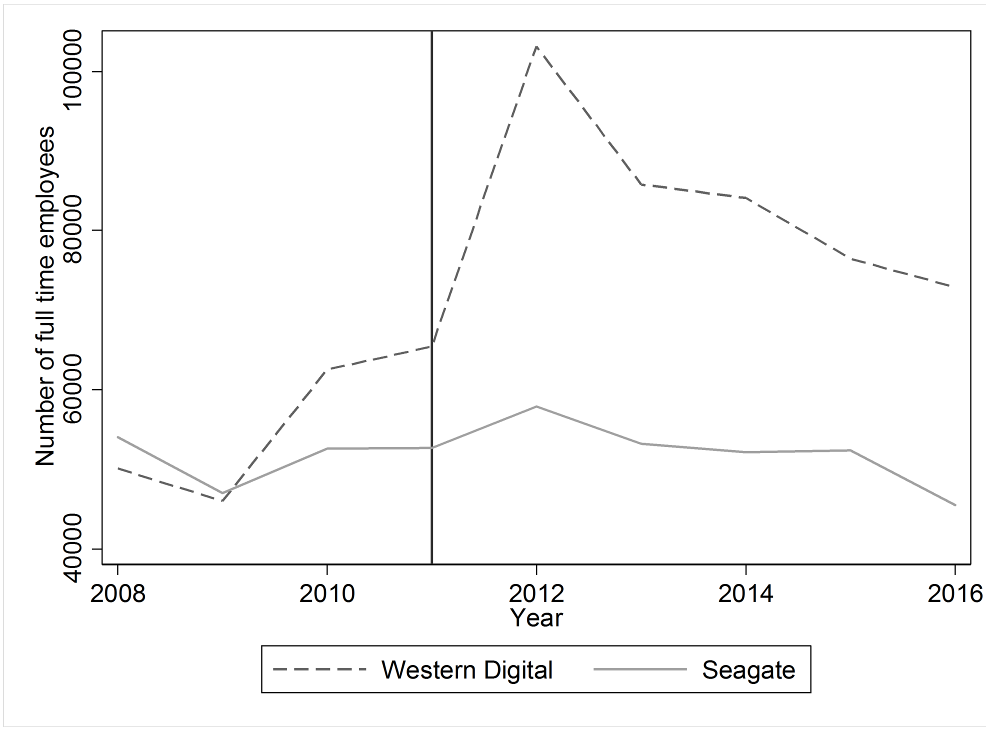

Figure 1: Number of full time employees, before and after merger

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 – this is despite there being only an 8 per cent difference in capacities shipped. Moreover, the hold separate conditions were imposed for 2 years on WD and 1 year on Seagate.

3 The theoretical framework

3.1 Mergers and innovation

In modelling how innovation might change following merger, we are guided by the lessons drawn above from the previous literature: innovation is a multidimensional concept and no one measure convincingly captures all of the dimensions. In terms of Schumpeter's trilogy (invention, innovation, and diffusion), R&D and patents are closest to invention; the emergence of marketed new products and processes to innovation. Empirically, none of the usual measures are without limitation: R&D does not always lead to fruitful outcomes; not all patents ultimately convert into innovation brought to market; both R&D and patents may often be used as a strategic defensive device to close down foreclose or hinder rival innovation. For this reason, we want to allow for three not mutually exclusive outcomes following mergers: (1) changes in the magnitudes of innovation inputs, (2) changes in the productivity of those inputs, and (3) changes due to other firm-specific reasons, unrelated to the inputs, such as changes in X-efficiency in the firm's research programme.

The above possibilities are captured by a simple two-stage framework in which, for firm i at time period t, an innovation production function relates innovative output (Y_{it}) to research effort or innovation inputs (X_{it}): $Y_{it} = f(X_{it})$. Then, the merger may enhance innovation through an increase in (X_{it}), an increase in the marginal productivity of (X_{it}), and/or a shift in the function.

In this framework, innovative output might be measured either in terms of process or product innovation (specific measures to be discussed below) and innovative inputs by the firm's R&D expenditure and/or its patenting performance. This two-stage structure represents a change in emphasis compared to nearly all 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.

3.2 Data

3.2.1 Measuring innovation

The initial intention when assembling the database for this paper was that it should include information on R&D, patenting and measures of innovative output, at the firm level for the 5 treated firms and as many flash (SSD) firms as feasible (for control purposes) over a period from the early 2000s up to the date of our research. In the event, we confined the period to 2007Q1-2016Q2 because data on the early 2000s was either incomplete or entirely unavailable (for SSD), while a major acquisition by WD of SanDisk, a key player in the SSD market, would have confused any analysis post-2016Q2.[7]

So that the magnitudes for the merged firms post-merger are comparable with magnitudes pre-merger, in each case the two merging parties are consolidated throughout the period, so for example, hereafter Seagate's patents or new models denotes the sum of Seagate and Samsung's even before the merger. The four variables, collected for each firm, are as follows.[8]

- **R&D intensity:** The ratio of R&D expenditure to total revenue, quarterly observations (2007Q1–2016Q2); source: firms' financial statements, downloaded from S&P's Capital IQ database.

- **Number of new patents:** The number of relevant patents granted to the firm, quar-

[7]There have been other acquisitions in the SSD/Flash market during our study period: `http://www.storagesearch.com/enterssdmarket.html`. However, these only led to incremental changes (at most) in market structure. The main underlying reason is that the SSD market is highly fragmented, with a large number of fringe players. These acquisitions are about buying fringe players in a competitive market, whilst some of these are acquiring products not directly on the SSD/Flash market with equally small market shares. As such, we argue that these mergers are unlikely to have substantial impact on market outcomes. In any case, in the Appendix, in Table C.7 we present results for a sample where we exclude flash/SSD firms that engaged in these mergers. This drops the number of control firms, which is a problem for the number of new models and unit price, as these models are left with only 4 firms. Nevertheless, the results qualitatively hold.

[8]More detail on how we constructed our dataset can be found in Section B in the Appendix.

terly observations (2007Q1-2016Q2). Constructed by the authors from 13515 patents (5863 HDD for the treatment and 7652 flash memory for the control), downloaded from US Patent Office (USPTO's) Patentsview API.

- **Number of new models:** Number of newly marketed HDD products (SSD products for the control) for each firm. The authors created this by scraping from Amazon's website data on 1931 HDDs and 1353 SSDs, arranged into quarterly observations (2009Q1-2016Q2).

- **Unit price:** The retail price of the HDD ($) relative to its size measured by the formatted capacity (Gb) of HDDs (SSDs for the control) first marketed in the given calendar quarter (2009Q1-2016Q2).

For R&D for Toshiba, we faced a problem not unusual for large conglomerate firms: the only data available is for the entire firm, and not for just its operations in the specific market concerned (HDD in this case).[9] Therefore Toshiba is excluded from all analysis of R&D below. Fortunately, Seagate and WD at this time were largely specialised in HDD and so this problem does not occur, and the same was true for many of our control firms (e.g. Sandisk, Kingston, Micron, Hynix).

3.2.2 Descriptive statistics: Before and after

Figure 2 shows how these four measures changed after the mergers (the vertical line indicates the period of the merger). This first descriptive step is not designed to impute causality, so at this point, it makes no allowance for any lags between the date of the merger and the time

[9] Toshiba is active in many different areas, and storage only constitutes around a quarter of its total operating revenue and R&D expenditure.

at which any innovation effect might be revealed in the data. From purely visual observation, it is clear that R&D intensity increased post-merger for both Seagate and WD (recall that there are no comparable data for Toshiba), while patenting appears to have declined for all three firms. There is a mixed picture for the two measures of innovative output: Seagate increased, WD reduced its quarterly number of new models, while no clear trend is apparent for Toshiba; unit price appears to have declined to varying extents for all three firms, apart from the early years, where information on new drives was sporadic.

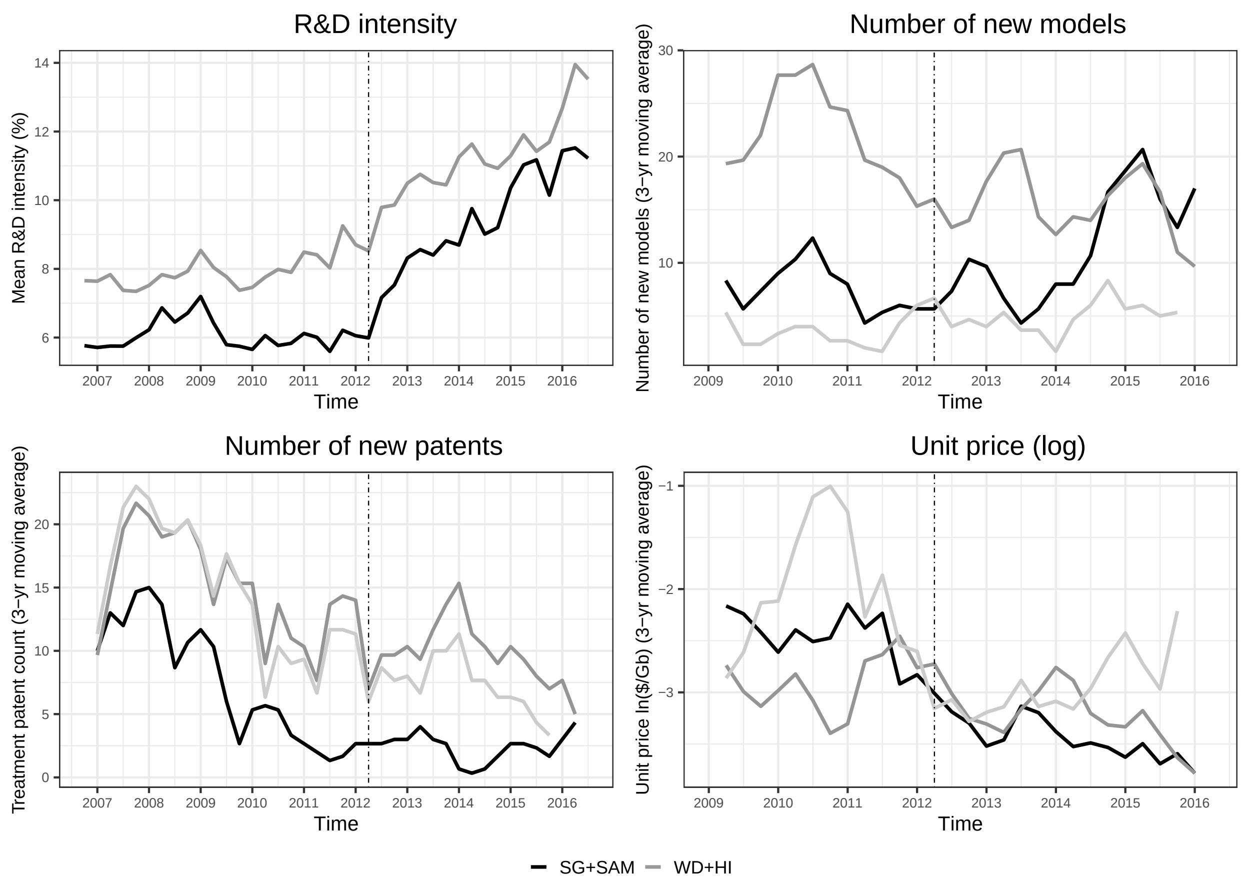

Figure 2: Four measures of innovation: before & after

Against this background, the following sections attempt to identify how far these changes can be attributed to the mergers.

3.3 Econometric methodology: a causal panel data model

To attribute any causal inference to the mergers, we need an appropriate method for estimating the counterfactuals of what would have happened in the absence of the mergers. Our approach is to construct a panel of firms observed over time, where the panel includes the "treated" (merging) firms and a control group of firms, which is similar but independent of the treatment. Then estimate causal effects in a model in which the treatment firms are exposed to the merger treatment for a subset of periods. To do this, we construct the sample which includes a sample of flash memory firms as control, and a matrix completion methodology is used to estimate the counterfactual.

3.3.1 Flash memory technology as the control

Identifying a control group is especially problematic in a case such as this because geographically there can be no comparator for a market which is genuinely worldwide, and in product space, it is difficult to identify another similar hi-tech product which is at a similar stage in its learning curve. Even the obvious, but anyway questionable, option of using other HDD manufacturers as controls is unavailable because there are no other manufacturers beyond the five parties to these mergers. Our choice for the control group is firms drawn from flash memory-based storage technologies. This choice is guided predominantly by the European Commission's (DG COMP) contemporary judgement that firms in the HDD and alternative technology (such as SSD, which uses flash memory) industries were not competing in the same market; this is reinforced by industry opinion at the time (see section 2.1 above).[10]

[10] As additional supporting evidence for the independence assumption, the Appendix also reports an indirect statistical test of our own, based on comparing trends in R&D between SSD firms and a broader sample of firms taken from the IT sector as a whole. It finds that, on this basis, SSD is no different from

In selecting the firms for the control group, we need to make a distinction between technology and product. In our control group, flash memory is the technology used in SSD products. Although flash memory is used for purposes other than data storage (for example RAM), innovations in flash memory potentially manifest in the data storage related applications of the technology. Therefore for the purposes of studying the innovation impact of mergers (impact on R&D and patent activity), we are interested in how the treatment technology (HDD) differs from the control technology (flash memory). When looking at differences in innovation output (number of new models and unit price), our focus shifts to the physical product (which incorporates the technology), and we compare HDD and SSD storage units. We found 41 firms, that had at least 10 flash memory-related patents in the our study period, and also had R&D data available. We use these as our control group for our analysis of R&D and patents. Only 8 of these were marketing (retail) SSD products in our study period, therefore we limit our control group to these 8 firms when looking at innovation output (new models and unit price).[11]

3.3.2 The counterfactual estimation

The two-way fixed effects estimator is a commonly used method to estimate causal relationships in studies that employ panel data. This method assumes linearity of the treatment effect and lack of time-varying confounding factors. In the presence of a period before the treatment and a period after the treatment, as is the case of pre- and post-merger in this work, these assumptions map to a parallel trend between treatment and control groups. If these assumptions do not hold or are too restrictive, the estimated causal effect is problem-

the IT market as a whole.

[11]For the list of firms, see Tables C.1 and C.2 in the Appendix.

atic.

Our task, to estimate causal inferences left us with a choice between different methodologies. We experimented with three methods nested in the machine learning literature, that relax these two assumptions: synthetic control (SC), the interactive fixed-effects (IFE) counterfactual (Gobillon and Magnac, 2016; Xu, 2017), and the matrix completion (MC) approach (Athey et al., 2018; Liu et al., 2020).[12] These methods do not require the same strict assumptions (parallel trend) for accurate estimation as difference-in-differences or panel data methods. Because the counterfactual is estimated in all our three potential methods, the parallel trend is imposed on the estimated counterfactual, rather than assumed from the data. We used a synthetic control method in an early version of this work (Ormosi et al., 2017). The reason we later moved on to MC and IFE methods, was the fact that they attempt to exploit both stable patterns over time (vertical) and stable patterns between units (horizontal) in imputing the missing values and therefore can deal directly with more complex missing data patterns (which we found to be the case for our innovation output measures).

MC and IFE are similar in that the estimators hold the treatment observations as missing data and estimate the treatment counterfactual using only information on non-treated observations (of treatment and control units). Time-varying confounding effects are accounted for by using a latent factor procedure. The two methods differ in the way they regularise the latent factor model (see Xu, 2017, both for a comparison between the two methods and for

[12]The latter two methods are generalizations of original solutions in the literature to the problem of unobserved time-varying confounders. The synthetic control method was proposed by Abadie et al. (2010, 2015), and the interactive fixed effects solution - which foresees unit-specific intercepts (factor loading) interacted with time-varying coefficients (latent factors) - was put forward by Bai (2009). The method of adding unit-specific time trends proposed by Mora and Reggio (2019) is also an alternative, but its impact on degrees of freedom makes it often less desirable.

a description of `gsynth` - a package that he has implemented in `R` to estimate both methods and which we used in our empirical analysis). The choice of which methodology works best can be left to the data, using the diagnostic tools and package, `fect`, implemented in `R` by Liu et al. (2020). In our empirical analysis, we run both methodologies but end-up preferring MC over IFE based on the estimated mean predicted square error. Given this preference, below we only outline how the matrix completion approach works.

For notational convenience we assume the panel to be balanced and comprising both of a large number of units N and of T periods, but this method also works with unbalanced panels and with variants of length of units and periods. We denote by D_{it} the treatment dummy, taking value one if the unit is treated at time t and zero otherwise. The set of treated observations is $\mathcal{T}$, and that of untreated observations is $\overline{\mathcal{T}}$. The outcome variable is Y_{it}. Let x_{it} be the row vector of K unit and time-varying observable control variables, whereas α_i and μ_t capture unit and time intercepts, respectively. The parameter L_{it} represents unobserved time-varying confounders. The term ε_{it} is the idiosyncratic error term. The econometric equation of interest is:

$$Y_{it} = \delta_{it} D_{it} + L_{it} + \mathbf{x}_{it}\boldsymbol{\beta} + \alpha_i + \mu_t + \varepsilon_{it}, \tag{1}$$

where the individual treatment effect is $\delta_{it} = Y_{it}(D_{it} = 1) - Y_{it}(D_{it} = 0) = Y_{it}(1) - Y_{it}(0)$. The parameters to be estimated are $\underset{N\times T}{\mathbf{L}}, \underset{K\times 1}{\boldsymbol{\beta}}, \underset{N\times 1}{\boldsymbol{\alpha}}, \underset{T\times 1}{\boldsymbol{\delta}}$ along with the penalty parameter λ. The estimation procedure requires minimizing the function:

$$\min_{\mathbf{L},\boldsymbol{\beta},\boldsymbol{\delta},\boldsymbol{\gamma},\lambda} \frac{1}{|\overline{\mathcal{T}}|} \sum_{(i,t)\in\overline{\mathcal{T}}} \left((Y_{it} - L_{it} - \mathbf{x}_{it}\boldsymbol{\beta} - \mu_t - \gamma_i)^2 + \lambda ||\mathbf{L}|| \right). \tag{2}$$

Athey et al. (2018) suggest using the nuclear norm for $||\mathbf{L}||$. The estimation procedure follows the following steps: (1) the parameters $\boldsymbol{\delta}$, $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ can be partialed out of the objective function (note that the first order conditions of equation (2) for these parameters can be solved linearly for a given value of $\mathbf{L}$); (2) the $N \times T$ parameters $\mathbf{L}$ can be estimated using an iterative procedure that relies on singular value decomposition.[13] (3) The penalty parameter λ is estimated via cross-validation of minimized mean-squared prediction error.[14]

Upon convergence of the estimator, the predicted outcome variable under no treatment is $\hat{Y}_{it}(0)$. Therefore, the estimated effect of the individual treatment effect is $\widehat{\delta}_{it} = Y_{it}(1) - \widehat{Y}_{it}(0)$, and the average treatment on the treated is given by:

$$\widehat{ATT}_t = \frac{1}{|\mathcal{T}_t|} \sum_{i \in \mathcal{T}_t} \left(Y_{it} - \widehat{Y}_{it}(0) \right) \qquad \widehat{ATT} = \frac{1}{|\mathcal{T}|} \sum_{(i,t) \in \mathcal{T}} \left(Y_{it} - \widehat{Y}_{it}(0) \right). \tag{4}$$

Standard errors can be bootstrapped as in Liu et al. (2020).

3.3.3 Identification assumptions

Inherent in the method are a number of assumptions that need further exposition here. Our task was to derive causal inferences on four different measure of innovation. Each of these

[13]The procedure can be outlined in this way. Define:

$$P_{\overline{\mathcal{T}}}(\mathbf{A}) = \begin{cases} A_{it}, & \text{if } (i,t) \in \overline{\mathcal{T}}, \\ 0, & \text{if } (i,t) \notin \overline{\mathcal{T}}. \end{cases} \quad \text{and} \quad P_{\mathcal{T}}(\mathbf{A}) = \begin{cases} 0, & \text{if } (i,t) \in \overline{\mathcal{T}} \\ A_{it}, & \text{if } (i,t) \notin \overline{\mathcal{T}}. \end{cases} \tag{3}$$

Express the singular value decomposition on a matrix $\underset{N\times T}{\mathbf{A}}$, as $\mathbf{A} = \underset{N\times N}{\mathbf{S}} \underset{N\times T}{\mathbf{\Sigma}} \underset{T\times T}{\mathbf{R}'}$. Apply the shrinkage operator to $\mathbf{A}$, so to have $\text{shrink}_\lambda \mathbf{A} = \mathbf{S}\widetilde{\mathbf{\Sigma}}\mathbf{R}'$, with $\widetilde{\mathbf{\Sigma}}$ be $\mathbf{\Sigma}$ with the $i - th$ singular value replaced by either $\max(\sigma_i(\mathbf{A}) - \lambda, 0)$ if $\sigma_i(\mathbf{A}) < \lambda$ or $\sigma_i(\mathbf{A})$ otherwise. For a given value of the penalty parameter λ, begin with $\mathbf{L}_0(\lambda) = P_{\overline{\mathcal{T}}}(\mathbf{Y})$, which corresponds to the left panel of equation 3, in which L_{it} for the treated observations is initially set to zero. Next, for any integer $h \geq 0$ calculate $\mathbf{L}_{h+1} = \text{shrink}_\lambda = (P_{\overline{\mathcal{N}}}(\mathbf{Y}) + P_{\mathcal{T}}(\mathbf{L_h}))$. Repeat the process until the sequence converges.

[14]For a complete discussion of the method (see Athey et al., 2018; Liu et al., 2020).

probably could have justified a different approach. For example, synthetic control methods were likely to work well on our R&D data, but not so well on the patent or product count data. Different methods however would have made comparisons very difficult, and it would have transpired a sense of arbitrariness. This was one of the reasons we chose the matrix completion method, which we thought performed well on model simplicity and robustness, and the reliability of the identifying assumptions across all four measures of innovation. Moreover, the matrix competition method relaxes on the parallel trends assumption, and on the fixed treatment effects assumption, both could have potentially been violated at least for part of our data.

Firstly, we assume unconfoundedness (or exogenous treatment, or conditional independence assumption), which means that treatment exposure (the merger) is (statistically) independent of the untreated potential outcome, conditional on the unit and time-specific coefficients, firm-level characteristics, and the latent factors estimated in our matrix completion method. A way to address the potential violation of this assumption (i.e. to deal with endogeneity) would be to instrument the incidence of a merger with some measure of distance (geographical or physical) from competitors.[15] 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 impossible to apply in our case, where the treatment only affects one or a small number of firms, i.e. where the merger instance is a single datapoint.

As our paper follows in a long line of merger retrospective studies, we do not have random assignment to the treatment, so our question on unconfoundedness is: having constructed

[15] This is a technique used in Hastings (2004), Dafny (2008), and Haucap et al. (2019).

our control group, can we say that once we control for all observed and latent factors, the merger was equally likely to have happened in the control group as in the treatment? This is a question that is especially hard to answer in a case where our control group is not an observable set of individuals, but an artificially created reality. Nevertheless, as we used SSD/Flash firms to estimate this control group, it is worth examining at least the intuition whether this choice satisfies the exogeneity assumption.

It appears from industry reports that innovation was only one of the many reasons why the respective parties merged. For example the Seagate/Samsung merger was part of a larger deal between the two companies in order to expand their strategic partnership, which included an agreement under which Seagate would supply disk drives to Samsung for PCs, notebooks and consumer electronics, to expand cooperation between the companies to co-develop enterprise storage solutions, to accelerate time-to-market for new products and position the companies, and also to extend and enhance the existing patent cross-license agreement between the companies.[16] This would indicate that there was some element in the motivations to merge that were driven by innovation - it was one of many reasons. We do not claim that it does not affect our results, we simply do not have evidence to test how much the exogeneity assumption is violated. This is a weakness of merger retrospectives and this paper is no exception.

The other key assumption that our results hinge on is the independence of our control group (SSD/Flash market) from the treatment firm. At the time of the mergers, the European competition authority (DG COMP) was very explicit on this in the market definitions

[16] `https://www.engadget.com/2011-04-19-samsung-sells-hdd-division-to-seagate-for-1-375-billion.html`.

it adopted in its merger control decisions for these mergers: HDD and SDD were deemed not to be in the same market. Appendix A includes a lengthy quotation from the EC's decision document (EC, 2011, paras 256-261).[17] In brief, this argues that there was little evidence on the demand side that users were able to substitute SSD for HDD for many purposes, and no evidence of effective substitutability on the supply side. This judgment seems to be in accord with general industry opinion, in which there is an understanding that, although SSD was growing, it is unlikely to be able to keep pace with the exponential growth in the quantities of data currently becoming available.

This means that many applications 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, SSDs are used largely as a complement in caching and restoring to speed up data transfers. SSDs also speed up access to store metadata or are used to boot storage pods. It appears, then, that depending on the application, one or the other of HDD or SSD is the preferred solution. In most applications they are more likely to be complements than substitutes (e.g. their relative roles in data centres). Consider also, the growing importance of hybrid HDD/SSD drives,[18] if the technologies were substitutes rather than complements then there would be no need for manufacturing products that combine the two.[19]

But even if the control group is not independent of the treatment, it should not negate

[17] EC (2011), Case No COMP/M.6214, - Seagate/ HDD Business of Samsung, Regulation (EC) No 139/2004 Merger Procedure, Article 8 (1) Date: 19/10/2011.

[18] Hybrid solutions today feature in more than a fifth of all new laptops.

[19] See Jo et al. (2009).

its use for our analysis. What would matter in this case is whether the two technologies are strategic substitutes or complements. If they are strategic complements and innovation increases in HDD, it should mean increasing innovation in SSD/Flash as well. If they are strategic substitutes, then an increase in HDD innovation would be matched by a drop in SSD/Flash. One could then use this information together with the study findings to interpret the results. Moreover, as Whinston (2011) argues, this countervailing effect is unlikely to overwhelm the direct effect, therefore "*one only needs to look at the direct effect on the merging firms' R&D holding rivals' R&D efforts fixed to discern the overall effect on R&D*".

Equally central to our research design is whether the two technologies were at a different stage in their product cycle, and how this may have influenced their innovation activities. None of the evidence we found by industry experts at the time of the merger suggested that innovation was dead in HDD. Far from it. With around 40ZB of data produced annually, a lot of this was/is inevitably stored on HDD, and therefore innovation was still happening even in the years following the mergers. An illustration of this is given by a presentation by a Senior Manager at Toshiba Electronics Europe GmbH, which highlights the importance of continuous HDD innovation in areas like unit capacity cost, and power consumption.[20]

The conclusion we take forward at this stage is that, as judged by DG COMP, HDD and SSD are largely independent.[21] In Section A.2 in the Appendix we propose a way to formally test this assumption, and arrive to the same conclusion.

[20] `https://www.cloudfest.com/wp-content/uploads/2018/03/15_Toshiba_Rainer-W.-Kaese.pdf`. See for example p.5, showing the $/capacity developments in SSD and HDD, which does not suggest to us that SSD is on an altogether different innovation trajectory).

[21] This is also the view of Igami and Uetake (2020), who exclude SSD firms from their model of competition, on the grounds that there are at least some applications where HDD does not compete with SSD.

4 Results

This section examines whether the descriptive before-after trends depicted in Figure 2 carry over into a causal setting, when compared with the counterfactuals generated by the matrix completion methodology. With causality now involved, the lag structure becomes important. There are two sources of lag: a *decision lag*, representing how long it takes for resource allocation decisions to become apparent in the data, and an *innovation lag*, representing how long before input expenditures impact on observed innovation output. So, with the mergers substantiating in 2012Q1, when will any merger effect become apparent in observed changes in R&D and patenting, and when might these impact, in turn, on innovation in the marketplace? Although there is a rich previous literature on the issue of lags in innovation,[22] there is no sufficient consensus for us to form strong priors on the lengths of these lags in this case. Instead, we will be guided by the data and report results with three alternative dates for when the impact of the mergers may have first been observed: immediately after the conditional approval in 2012Q2, or later in 2013Q2 or 2014Q2. We have chosen one-year gaps to avoid quarterly seasonalities. Arguably, an impact on input decisions might be observed as early as 2012, but an impact of innovation outputs might not have impacted fully until 2014. However, we prefer results which are robust across all three alternatives.

Figure 3 and Table 2 show the results for the matrix completion methodology. Figure 3 depicts what is, in effect, the difference between the actual and counterfactual for the pooled average of the three treatment firms for each of the four innovation measures (except for R&D where Toshiba is excluded for reasons explained above). Under our methodology,

[22]See, Pakes and Griliches (1980, 1984), Hall et al. (1986), and Wang and Hagedoorn (2014).

by construction, these graphs will be free of any trend pre-merger,[23] but systematic trends post-merger will reveal the nature of the merger impact. The aggregate ATT (average treatment effect on the treated firms) provides an overall significance test, and the shaded 95% confidence intervals reveal whether any systematic trends are becoming more pronounced over the post-merger period.[24]

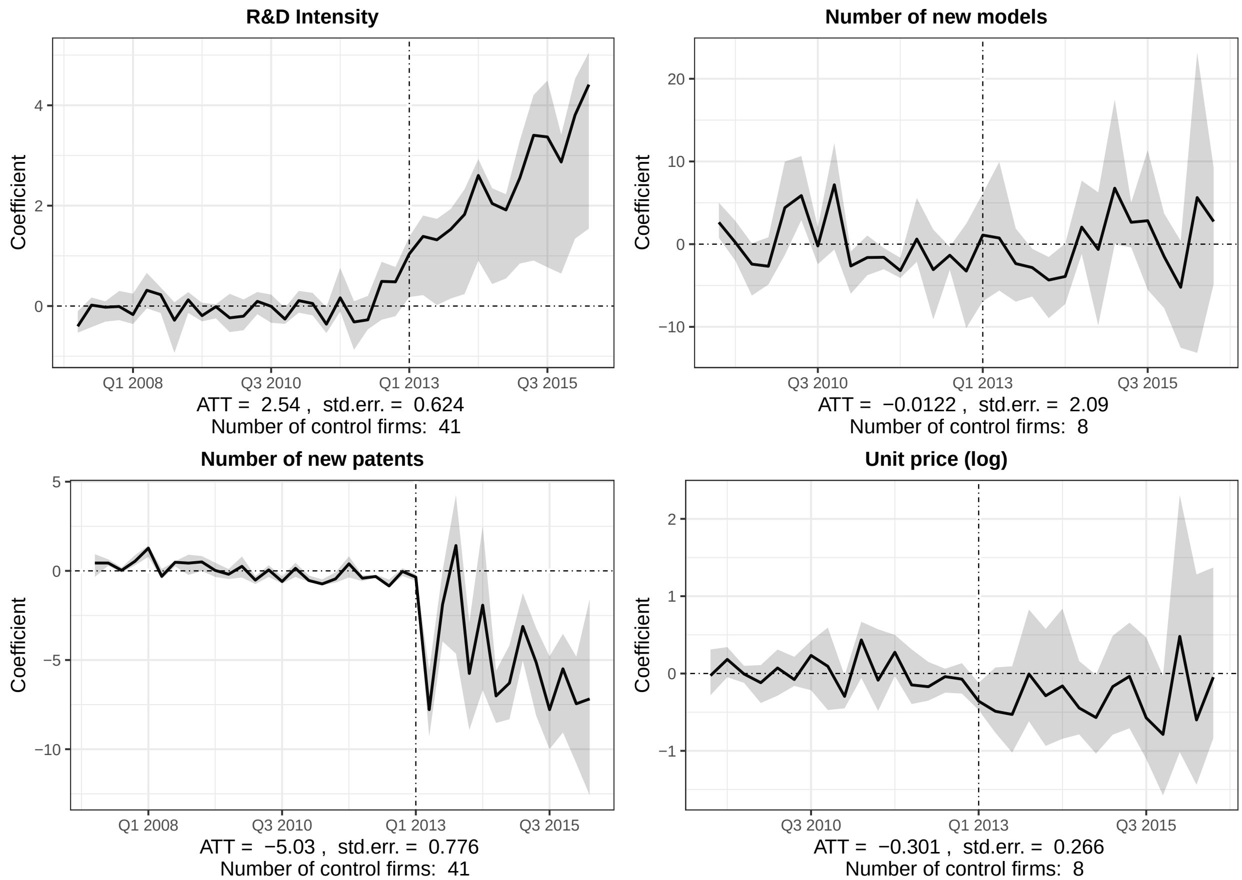

Figure 3: Average Treatment Effects on the Pooled Treated Firms (post-merger begins 2013Q2)

Results are robust across the three alternatives for dating the first effects of the merger.

[23] In other words, the parallel trends test, so crucial in conventional difference-in-differences, is automatically satisfied.

[24] These plots are a comparison of the observed treatment data, and the estimated treatment data. One can think of the pre-treatment area of Figure 3 as a measure of how well the counterfactual estimation works, because pre-treatment the difference between the observed and the estimated values should be zero.

Figure 3 shows the results assuming the 2013 date of initial impact, figures C.3 and C.4 in the Appendix show the comparable figures for the two alternative dates. As can be seen, there is a significant increasingly positive impact on R&D. For patents, the downward post-merger effect becomes increasingly pronounced over the period, and the confidence interval is consistently negative after two exceptional quarters in the immediate post-merger years. Figures C.1 and C.2 in the Appendix present results for a number of different measures of patent activity. This includes using patent citation instead of patent count, and looking at the estimated treatment effects on patent activity, where only the most relevant patents are included in the analysis.[25] We also report results for measures of patent intensity (patent count and patent citation divided by the company's turnover). These results confirm a drop in patent count intensity, and patent citation numbers, but an increase in patent citation intensity. This would suggest that although there are fewer patents post-merger, these few patents receive relatively more citation. On the output side, however, there is no strong visual evidence of any effect for new models or unit price, although, for the latter, the line does lie consistently below the horizontal axis for nearly the entire post-merger period.

Table 2 presents the matrix completion results,[26] for each treatment firms, and this reveals broad symmetry between the three firms on the input side, but not for outputs.[27] Thus, R&D increases significantly, and patents decline significantly, relative to the counterfactual

[25]To choose these most relevant patents, we used the USPTO classification, where the class G06F - Electric Digital Data Processing - is the most frequent class for storage technologies.

[26]In the Appendix, in Tables C.4 and C.5 we report the results of using fixed effects and synthetic control estimators. These tables can be viewed as a verification of the robustness of our findings, as the results remain in the same ballpark as those reported in Table 2. Table C.4 confirms that the parallel trend assumption is rejected for R&D and patents, justifying our choice of method.

[27]Table C.6 in the Appendix shows the mean squared prediction errors for these models to demonstrate the quality of these estimators.

in all cases.[28] For new models, however, significant differences emerge, with Seagate and Toshiba (to a lesser extent) both significantly increasing their number of models relative to the counterfactual, while the opposite is true for WD, who significantly reduces their number.[29] On unit price, qualitatively, all three move in the same direction, recording declines relative to the counterfactuals, but not significantly so, except for Toshiba for which this is significant in two of the three alternative merger start-dates.

To interpret these results, consider first the implication of (erroneously) applying a simple checklist count, in which, conventionally, R&D, patents and new models would be positive indicators and unit price an inverse measure of innovative performance. On that basis, the top section of Table 2 reports 11 indicators for the three firms, 10 of which are significant, and of these 4 are positive and 6 negative. As a case study, while this would offer strong support for the hypothesis that mergers affect innovation, it is unclear whether the effect is beneficial or harmful.

Moreover, such a line of reasoning is over-simplistic because (i) it treats the four measures as perfectly substitutable alternatives for each other, and (ii) it fails to recognise the potential for asymmetries between firms. As we have argued earlier, different indicators are not direct substitutes because they measure different aspects of "innovation"; and part of the motive for research into this case, is that it would allow us to explore asymmetries between different players within a given market. We return below to possible explanations for the firm-asymmetries, but perhaps the most striking feature of these results is that all three firms record significantly declining post-merger patenting activity, in contrast to the

[28] Again, except for Toshiba, for whom we have no appropriate R&D data.

[29] The opposite effects for new models between Seagate and Toshiba on the one hand, and WD on the other explains why, in Figure 3, there appears to be no effect on average.

Table 2: Treatment effects by Treated Firm

	R&D Int	Number of new patents	Number of new models	Unit price (log)
	Post-merger period begins 2012Q2			
Seagate	2.870***	-5.420***	4.120***	-0.898**
s.e.	(0.824)	(0.182)	(0.475)	(0.458)
WD	2.590***	-5.230***	-7.470***	-0.492
s.e.	(0.701)	(0.139)	(0.429)	(0.399)
Toshiba		-8.070***	1.600***	-0.962***
s.e.		(0.017)	(0.189)	(0.229)
Pooled	2.730***	-5.800***	-0.815	-0.717**
s.e.	(0.761)	(0.759)	(2.92)	(0.309)
Number of control firms	41	41	8	8
	Post-merger period begins 2013Q2			
Seagate	2.740***	-4.840***	4.120***	-0.392
s.e.	(0.605)	(0.132)	(0.550)	(0.400)
WD	2.340***	-4.450***	-4.940***	-0.349
s.e.	(0.534)	(0.113)	(0.462)	(0.329)
Toshiba		-7.100***	0.845***	-0.453***
s.e.		(0.210)	(0.136)	(0.175)
Pooled	2.540***	-5.030***	-0.0122	-0.301
s.e.	(0.621)	(0.772)	(2.070)	(0.263)
Number of control firms	41	41	8	8
	Post-merger period begins 2014Q2			
Seagate	2.570***	-3.950***	7.970***	-0.348
s.e.	(0.500)	(0.189)	(0.600)	(0.360)
WD	1.960***	-4.400***	-5.670***	-0.526*
s.e.	(0.311)	(0.232)	(0.468)	(0.290)
Toshiba		-6.580***	2.120***	0.158
s.e.		(0.428)	(0.217)	(0.144)
Pooled	2.280***	-5.010***	1.380	-0.205
s.e.	(0.512)	(0.935)	(3.329)	(0.248)
Number of control firms	41	41	8	8

*= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$. One-way (unit) fixed effects model. Bootstrapped (1000 iterations) standard errors. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

significantly increasing R&D expenditures, and (admittedly weaker) evidence of improving output measures (more products, at a lower price). One interpretation of the exceptional finding for patents might lie in its potential limitations as a measure of innovation, as noted earlier here and in the previous literature. In particular, perhaps the sharp reduction from

5 to 3 rivals may have reduced duplication between rivals, and reduced defensive patenting with exclusionary motives. In Table C.3 in the Appendix we examine other measures of patenting activity (patent intensity, citation, and citation intensity), and confirm the drop in patent numbers, but find that at the same time citation intensity increased, suggesting that the fewer patents are becoming more important in other innovations.

4.1 The innovation production function

Given that Table 2 has shown positive impacts of the mergers on R&D intensity for both Seagate and WD, but negative impacts on patenting, we now examine how these changes in innovation input have impacted on subsequent changes in the innovation outputs. As argued earlier, this can be either a direct effect – the merger changes the magnitude of R&D or patents and this leads to changed outputs, and/or an indirect productivity effect – the merger affects how productive each unit of R&D or patent is, and this could lead to changes in innovative output, even absent any change in the magnitudes of R&D or patents.

The matrix completion methodology is now no longer directly applicable - it has already established that there are time-variant effects of the mergers, but it is not designed to identify what are the time-variant input causes. Instead, we now revert to a standard difference-in-differences approach, including the key explanatory variables, innovative inputs. The estimated equation is now a straightforward linear regression of innovative output Y_{it}, on innovative inputs ($z_{ji,t-L}$). The time lag L is allowed to take on alternative values and j is one of the inputs $j = (R\&D, patents)$. In our headline results, we assume a one-year (four-quarter) lag, but results with other lags are given in Tables C.8 and C.9 in the Appendix.

Equation 1 modifies to:

$$Y_{it} = \delta_0 D_{it} + \sum_j \left(\delta_{j1} + \delta_{j2} D_{it}\right) z_{ji,t-L} + x_{it}\phi + \alpha_i + \mu_t + \varepsilon_{it}, \tag{5}$$

where δ_{j1} is the non-merger specific effect of the innovation input on output, δ_{j2} is the effect of the merger on the productivity of innovative input, δ_0 gives us a residual firm-specific effect of the merger on innovation. The model is estimated separately for Seagate and Western Digital, for each of the new models and unit price; as before, Toshiba is excluded because of unavailability of HDD-related R&D data. Finally, given the limited number of SSDs available to purchase from Amazon before 2010, we limit our time period to run from Q12010 to Q22016.

Table 3 shows the impacts on Seagate. Our choice of method no longer imposes a pre-merger parallel trend between the control and the treatment. Therefore we needed to test whether this assumption holds. For this we conducted a series of tests to show whether there is a significant difference in the slopes of the pre-merger linear trends (the p-value of the test is reported in the tables below, with the null implying parallel trend). We assumed two scenarios, depending on whether the effect of the merger on innovation would first manifest in 2013Q2 or 2014Q2. First of all, it is apparent that the parallel trend assumption holds for 2013, but not always for the 2014 scenario. One possible explanation to this would be that some effects of the merger started after 2013 but before 2014Q2. For this reason, our discussion below focusses on the 2013 columns of Table 3.

Regarding new HDD models, our results for Seagate show that R&D expenditure is associated with increased numbers of new models (δ_{11}), and that the merger further contributed

Table 3: Impact of R&D and patents on innovation output - Seagate

	2013Q2		2014Q2	
	Number of new models	Unit price	Number of new models	Unit price
Post-merger ($\overline{\mu}_t$ post-merger)	1.198***	-1.15***	1.321***	-1.202***
	(0.4359)	(0.1397)	(0.4667)	(0.1498)
Treatment effect (δ_0)	-26.646***	0.996	-19.222*	0.385
	(5.827)	(1.8669)	(10.4734)	(3.3603)
R&D (δ_{11})	0.231**	-0.021	0.231**	-0.037
	(0.1036)	(0.0332)	(0.1029)	(0.033)
Patent count (δ_{21})	-0.081	0.004	-0.055	-0.008
	(0.0768)	(0.0246)	(0.0766)	(0.0246)
R&D x treatment effect (δ_{12})	3.746***	-0.063	3.085***	0.016
	(0.6237)	(0.1998)	(1.1118)	(0.3567)
Patent x treatment effect (δ_{22})	-1.424***	0.016	-1.437***	0.041
	(0.4983)	(0.1597)	(0.533)	(0.171)
P-val parallel trend	0.7475	0.756	0.0056	0.8223
N	224	224	224	224
R2	0.328	0.31	0.32	0.3
F-val	9.065	8.328	8.741	7.937

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Robust standard errors. We tested for the joint significance of time fixed effects and decided against using them, we have instead preferred a post-merger dummy. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

to this positive relationship (δ_{12}). On the other hand, the effect of patent activity (δ_{21}) is not significant, but the merger seems to have reduced the productivity of patents (δ_{22}). The residual effect (δ_0) of the merger (not due to the change in inputs) was negative. For unit price, we found no significant changes following the mergers. It appears that in this case patent activity is not a good predictor of innovation output.

Table 4 presents contrasting results for Western Digital. R&D intensity contributes to more new models (significant) and lower unit price (not significant), but the merger made Western Digital less able to transfer the benefits of increased R&D to more new models. On the other hand, patent activity did not have a significant effect. The residual effects of the merger improved the number of new models but not the unit price.

To better understand what the above findings mean in terms of our innovation function,

Table 4: Impact of R&D and patents on innovation output - Western Digital

	2013Q2		2014Q2	
	Number of new models	Unit price	Number of new models	Unit price
Post-merger ($\overline{\mu}_t$ post-merger)	1.167**	-1.134***	1.371***	-1.196***
	(0.4546)	(0.1367)	(0.4949)	(0.1463)
Treatment effect (δ_0)	33.856***	1.65	-19.21	2.781
	(10.2284)	(3.077)	(24.0032)	(7.0967)
R&D (δ_{11})	0.179*	-0.019	0.158	-0.033
	(0.1081)	(0.0325)	(0.1088)	(0.0322)
Patent count (δ_{21})	0.029	-0.016	0.036	-0.022
	(0.0697)	(0.021)	(0.0689)	(0.0204)
R&D x treatment effect (δ_{12})	-3.78***	-0.086	0.942	-0.164
	(0.9499)	(0.2858)	(2.1537)	(0.6368)
Patent x treatment effect (δ_{22})	0.089	0.005	0.201	-0.012
	(0.2005)	(0.0603)	(0.2436)	(0.072)
P-val parallel trend	0.4522	0.9729	0.000	0.9085
N	224	224	224	224
R2	0.178	0.311	0.142	0.305
F-val	4.017	8.36	3.079	8.152

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Robust standard errors. We tested for the joint significance of time fixed effects and decided against using them, we have instead preferred a post-merger dummy. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

in Table 5 we decomposed these estimates into the effect of the change in R&D and R&D productivity.[30] We assume only two periods, one period before the merger ($t = 0$) and one after the merger ($t = 1$), in which case the variable rd_{1-L} is the average value of rd (lagged L periods) over the time units post-merger and, similarly, rd_{-L} is the average value of rd (lagged L periods) over the time units before the merger. Denoting with T the treatment and C the control units, we take the first differences of equation (5) and then calculated the

[30] We only focus on R&D as the results for the effect of patents were weak.

difference in differences:

$$Y_{T1} - Y_{T0} = \delta_0 + \delta_{11}\left(rd_{T1-L} - rd_{T-L}\right) + \delta_{12} rd_{T1-L} + \mu_1 \tag{6a}$$

$$Y_{C1} - Y_{C0} = \delta_{11}\left(rd_{C1-L} - rd_{C-L}\right) + \mu_1 \tag{6b}$$

$$(Y_{T1} - Y_{T0}) - (Y_{C1} - Y_{C0}) = \underbrace{\underbrace{\delta_0 + \delta_{12} rd_{T1-L}}_{\text{R\&D productivity change}} + \underbrace{\delta_{11} rd_{DiD-L}}_{\text{R\&D change}}}_{\text{R\&D joint effect}}. \tag{6c}$$

The decomposition results in Table 5 suggest that for Seagate, the merger contributed to the increased number of new models in two way: first through the increase in R&D intensity (as shown in Table 2),[31] and second, through the improved innovation productivity of R&D. Both of these effects are significant. Regarding unit price, the merger generated R&D increase is associated with a lower unit price. However, the merger did not improve the productivity of R&D regarding unit price. When decomposing the effect of R&D on innovation output for Western Digital, we find that the increase estimated for WD's R&D intensity in Table 2 increased the number of new models, however, the merger reduced the innovation productivity of R&D.

As a robustness check, we also estimate a weighted version of these regressions, where each control firm is weighted by the weights recovered from a variant of the stage 1 results in Table 2. These weights enforce parallel trends on the pre-merger period between the treatment and the control units. The results of these estimates are reported in Tables C.10 and C.11 in the Appendix. These confirm the above findings.

[31] Paying attention to the fact that Table 2 shows average treatment effects on the treated, rather than average treatment effects.

Table 5: Decomposition of the effect of R&D

	2013Q2		2014Q2	
	Number of new models	Unit price	Number of new models	Unit price
Seagate				
R&D change	0.385**	-0.036	0.472**	-0.076
	(0.1731)	(0.0554)	(0.2102)	(0.0674)
R&D productivity change	5.868***	0.453	9.51***	0.533
	(1.6986)	(0.5442)	(1.687)	(0.5413)
R&D joint effect	6.253***	0.417	9.982***	0.457
	(1.6863)	(0.5403)	(1.6683)	(0.5353)
Western Digital				
R&D change	0.265*	-0.028	0.251	-0.052
	(0.1601)	(0.0482)	(0.1731)	(0.0512)
R&D productivity change	-6.398**	0.735	-8.764***	0.958
	(2.564)	(0.7713)	(3.0467)	(0.9008)
R&D joint effect	-6.133**	0.707	-8.513***	0.905
	(2.5558)	(0.7689)	(3.0387)	(0.8984)

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Standard errors in parentheses.

4.2 Discussion of results

This section describes what we see to be the substantive story of the case in its own right: the relative experiences of the three acquiring parties and our retrospective judgement on the decisions taken by the competition authorities.

Across the market (averaging across the three mergers) we found increasing R&D intensity, falling patent count, and no significant effect for the number of new products of unit price. The results are worth comparing with some of the previous studies on mergers and innovation that report estimates for the same variables (R&D intensity and patent count). Although these are aggregate studies, this comparison can establish how our findings in the HDD market compare with other industries. We found a 2.5 percentage point increase in R&D intensity (Figure 3). Closest to our market (HDD), Stiebale (2013) offers estimates in a similar ballpark (increase of 3.5 percentage points, s.e. 1.11) in knowledge based industries.

Szücs (2014) finds a significant 2.5% drop (s.e. 0.5%) in R&D intensity, but paired with his findings of R&D growth, he attributes this drop to the sales expansion of the merging firms. Ornaghi (2009) reports a non-significant 0.2% (s.e. 0.3%) drop. Regarding patents, we found that the number of new patent applications drops by an average 5 patents per calendar quarter. This is equivalent to an average 44% fall in patent applications in comparison to the pre-merger period. This figure is comparable to previous estimates of aggregate studies, in which Ornaghi (2009) estimated a 32.5% drop by the third year following mergers, and Haucap et al. (2019) reported a 20-25% drop and a further 45-50% drop for patents where there were overlaps pre-merger (this latter result addresses the removal of duplicate patents story). As these are aggregate studies, what would be interesting to see is what proportion of the mergers in their sample resulted in an increase comparable to ours.

When looking at firm level heterogeneity, for Seagate, we found no evidence of falling innovative activity post-merger. On the contrary, the results suggest that the firm increased its R&D and there was also an increase in the productivity of the R&D. Although patenting appears to have declined, we believe this is most likely a reflection of the removal of duplication and defensive patenting. This is supported by the evidence that at the same time patent citation intensity increased (i.e. duplicate patents, that typically have lower citation, are less likely registered by Seagate). Moreover, as Figure C.5 shows, patent litigation activity involving the margin parties shows a fall after the merger, which would also support this interpretation.[32] On the output side, there is significant evidence that the firm increased the numbers of its new models and (weaker) evidence of reductions in the size adjusted prices it

[32]Patent litigation intensified pre-merger, and fell after the merger. This could provide partial explanation why the parties decided to merge, i.e. to avoid excessive litigation over access to patents.

offered in the marketplace. There were innovation synergies between Seagate and Samsung, which were substantiated 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.[33] To the extent the newly shared patent pool incorporates previously separate but complementary patents, this is conducive to increased innovation. Moreover, even if the synergies between the companies were not directly related to R&D, the savings could have also led to more investment in R&D. 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. Regarding the incentives to innovate, theory from Section 2.1 suggests that our findings of increase post-merger innovation activity, could have been a result of the merger: (1) eliminating some of the negative externalities, (2) decreasing potential innovation spill-overs, (3) combining complementary technologies, or (4) increased post-merger demand. We do not have data to verify which of these effects was the most likely driver of our findings.

For WD we found a more mixed story: as for Seagate, there was the unexpected combination of significantly increasing R&D but declining patenting activity (with increased patent citation intensity). But in this case, there is no evidence that the increased R&D led to increased innovative outputs. On the contrary, the firm significantly reduced the range of new products brought to market. One explanation 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

[33] Although there were remedies in place to ensure that the brands were kept separate and that the acquired brand did 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).

R&D, production, marketing, and sales operations of the two parts. This was crippling for WD's efficiency and the inability to remove duplications could have affected R&D productivity. 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 3.5-inch operations to Toshiba had to include all 3.5-inch related IP rights. To put this in the context of our discussion in Section 1.2, the remedies hindered WD's ability to realise any merger related synergies. Without such synergies, theory would predict that the ability to innovate also suffers, and one would need increased incentives to offset it. Our weak findings on WD suggest that this was not the case.

For Toshiba, the evidence is least conclusive – not least because no suitable data were available which could allow us to disentangle this part of the firm's R&D from its much wider conglomerate worldwide activities (only around 25% of Toshiba's revenue comes from storage-related operations). However, there is evidence that the firm significantly increased its range of new models post-merger.

As a retrospective study of the effects of mergers on innovation, we believe that the most valuable lessons to be learned are methodological as described in the final section. But in terms of the case itself, these two, simultaneous mergers, which gave rise to a third merger because of the agreed remedy, certainly created a market structure – a literal triopoly - which might have led to competition concerns. Based on our results, however, such fears are unfounded: R&D increased significantly across the board with evidence that this led to increased models and (only weak evidence) of increased process innovation as reflected in lower size-adjusted prices. The only real question mark is the reduced patenting activity of all three firms, but we have explained that in itself may merely reflect reduced duplication.

As for the decisions of the competition authorities, DG COMP's agreement seems to have been justified. First, for Seagate and Western Digital, the decision to conditionally allow the merger seems to be justified. Whether our findings are due to the MOFCOM (and in the case of WD the FTC and EC) remedy, or whether we would have observed the same outcome without remedies is not something we can establish from our work (a structural analysis might be able to estimate the impact of the WD divestiture, although we cannot think of a way the specific impact of the MOFCOM behavioural remedies could be estimated). However, when contrasting the effects on Seagate and on WD, our interpretation is that the less positive experience of WD could well have reflected the rather demanding behavioural remedies imposed by the Chinese competition authority. For Toshiba, the story is different. It was an exogenous merger, imposed by the FTC and DG COMP. We found no evidence that the divestiture to Toshiba was an ineffective remedy.

5 Conclusion

This paper is a rare attempt at an in-depth retrospective evaluation of the impact of mergers on innovation. Our main results paint a mixed, but rich, picture, in which different measures of innovation apparently tell different stories, and in which the mergers impacted asymmetrically on the parties concerned. The purely descriptive evidence shows that (for the two firms for which we have the data) R&D increased; but, for all three, patenting decreased post-merger; two firms increased the numbers of new models, while one reduced its number; and for at least two of the three firms, the unit price of the HDD memory fell, albeit moderately. These results are confirmed when compared to counterfactual estimates derived by a matrix completion methodology: broadly speaking this confirms that our descriptive results are significant. We then move on to the relationship between R&D and the two output measures – the innovation production function. This establishes a significant asymmetry between the two main players. For one (Seagate), the major impact of the merger was significantly increasing the productivity of R&D, but for the other (Western Digital), increased R&D appears to have led to declining productivity. We suggest that WD's poor performance may be a direct consequence of its reduced ability to realise merger-specific synergies due to the stringent behavioural remedies imposed by the Chinese competition authority on WD, but not Seagate.

Looking beyond the immediate details of this case, some of our findings may have more general applicability. First, for all three treatment firms, patenting declined post-merger – both in absolute terms and relative to the synthetic comparator, and there is no evidence that the amount of patent activity had any effect on the outputs of innovation. Given the positive evidence on R&D and to some extent the output measures, this calls into question how best to interpret patenting activity in mergers cases. In this instance, we suggest that the reduction in patenting might reflect a removal of duplicative and/or defensive patenting as a result of internalisation within the merged forms of previously competitive activity. If this is true, then reduced patenting post-merger is not necessarily a negative outcome. Moreover, we found little evidence that patent activity is a good predictor of changes in innovation output. In any event, the case confirms our prior expectation that no one measure of innovation is likely to tell the whole story.

Second, previous empirical studies, taking the industry as the unit of observation and then attempting to identify the causes of inter-industry differences may often miss more important intra-industry differences between firms in the same industry. Certainly these asymmetries are pronounced in this case.

Third, this case confirms the increased complexity entailed in conducting a case-specific merger retrospective study of the innovation effects of mergers. In the traditional merger retrospective, the focus is often on just one, easily measured, policy variable – price. But here, the firm's policy variable concerns an input decision, research activity, and now what also matters is how that input translates into innovation with a value in the market place. Our answer to this extra layer of complexity has been to depict "innovation" as multidimensional, with measures of inputs and outputs, however, this adds extra demands on necessary data. Especially in global markets, usually involving giant conglomerates, disaggregated data of the various dimensions to innovation will often be scarce. Equally problematic is the identification of a suitable control group to use as a comparator – given that there is no geographical comparison possible and that identifying an alternative hi-tech industry at a similar stage in its evolution. In this case study, we have been fortunate in that the two main players are both relatively specialised in the industry concerned, and that technical product data are available. In this particular case, there is a comparator product, SSD, in the same wider data-storage sector. In choosing this as an independent but similar comparator, we have drawn on the original decision of the original CA (DG COMP) report which judged that HDD and SSD are in separate product products. We have also used a current best practice econometric technology (matrix completion) to create a control which, by construction, satisfies the key parallel trends assumption.

Regarding policy implications, in our specific case there is nothing to suggest that innovation was hindered. Moreover, our findings are relevant for the design of merger remedies as they can hinder merger synergies and hence dampen the ability of the merging firms to increase their innovation efforts. But our most importantly policy message is that there are complex mechanisms that determine how innovation changes after mergers. Our study shows that this complexity can lead to different outcomes even within the same market. For this reason, overall generalisations are difficult to make. Instead, we need many more case studies similar to ours (this rhymes with the arguments of Gilbert, 2020), akin to the rich body of merger retrospectives on the price effect of mergers.

References

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

Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2):495–510.

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

Aghion, P., Blundell, R., Griffith, R., Howitt, P., and Prantl, S. (2009). The effects of entry on incumbent innovation and productivity. *The Review of Economics and Statistics*, 91(1):20–32.

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

Ashenfelter, O., Hosken, D., and Weinberg, M. (2014). Did robert bork understate the competitive impact of mergers? evidence from consummated mergers. *The Journal of Law and Economics*, 57(S3):S67–S100.

Athey, S., Bayati, M., Doudchenko, N., Imbens, G., and Khosravi, K. (2018). Matrix completion methods for causal panel data models. Technical report, National Bureau of Economic Research.

Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279.

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

Bourreau, M., Jullien, B., Lefouili, Y., et al. (2018). Mergers and demand-enhancing innovation. Technical report, Toulouse School of Economics (TSE).

Cohen, W. M. (2010). Fifty years of empirical studies of innovative activity and performance. In *Handbook of the Economics of Innovation*, volume 1, pages 129–213. Elsevier.

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

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

Entezarkheir, M. and Moshiri, S. (2018). Mergers and innovation: Evidence from a panel of US firms. *Economics of Innovation and New Technology*, 27(2):132–153.

Farrell, J. and Shapiro, C. (2010). Upward pricing pressure in horizontal merger analysis: Reply to Epstein and Rubinfeld. *The BE Journal of Theoretical Economics*, 10(1).

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

Federico, G., Langus, G., and Valletti, T. (2018). Horizontal mergers and product innovation. *International Journal of Industrial Organization*, 59:1–23.

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

Gilbert, R. J. (2020). *Innovation matters: Competition policy for the high-technology economy*. MIT Press.

Gilbert, R. J. and Newbery, D. M. (1982). Preemptive patenting and the persistence of monopoly. *The American Economic Review*, 72(3):514–526.

Gobillon, L. and Magnac, T. (2016). Regional policy evaluation: Interactive fixed effects and synthetic controls. *Review of Economics and Statistics*, 98(3):535–551.

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

Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28:4.

Hall, B. H., Griliches, Z., and Hausman, J. A. (1986). Patents and r and d: Is there a lag? *International economic review*, pages 265–283.

Hastings, J. S. (2004). Vertical relationships and competition in retail gasoline markets: Empirical evidence from contract changes in southern california. *American Economic Review*, 94(1):317–328.

Haucap, J., Rasch, A., and Stiebale, J. (2019). How mergers affect innovation: Theory and evidence. *International Journal of Industrial Organization*, 63:283–325.

Igami, M. and Uetake, K. (2020). Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016. *The Review of Economic Studies*, 87(6):2672–2702.

Irwin, D. A. and Pavcnik, N. (2004). Airbus versus Boeing revisited: international competition in the aircraft market. *Journal of International Economics*, 64(2):223–245.

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

Jo, H., Kwon, Y., Kim, H., Seo, E., Lee, J., and Maeng, S. (2009). SSD-HDD-hybrid virtual disk in consolidated environments. In *European Conference on Parallel Processing*, pages 375–384. Springer.

Jullien, B. and Lefouili, Y. (2018). Horizontal mergers and innovation. *Journal of Competition Law & Economics*, 14(3):364–392.

Kwoka, J. (2014). *Mergers, merger control, and remedies: A retrospective analysis of US Policy*. Mit Press.

Letina, I. (2016). The road not taken: competition and the R&D portfolio. *The RAND Journal of Economics*, 47(2):433–460.

Liu, L., Wang, Y., and Xu, Y. (2020). A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data. *Available at SSRN 3555463*.

Mariuzzo, F. and Ormosi, P. L. (2019). Post-merger price dynamics matters, so why do merger retrospectives ignore it? *Review of Industrial Organization*, 55(3):403–429.

Mora, R. and Reggio, I. (2019). Alternative diff-in-diffs estimators with several pretreatment periods. *Econometric Reviews*, 38(5):465–486.

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

Ormosi, P. L., Bennato, A. R., Davies, S., and Mariuzzo, F. (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.

Pakes, A. and Griliches, Z. (1980). Patents and R&D at the firm level: A first report. *Economics Letters*, 5(4):377–381.

Pakes, A. and Griliches, Z. (1984). Patents and R&D at the firm level: a first look. In *R&D, Patents, and Productivity*, pages 55–72. University of Chicago Press.

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

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

Shapiro, C. (2011). Competition and innovation: Did arrow hit the bull's eye? In *The rate and direction of inventive activity revisited*, pages 361–404. University of Chicago Press.

Spulber, D. F. (2013). Innovation economics: The interplay among technology standards, competitive conduct, and economic performance. *Journal of Competition Law and Economics*, 9(4):777–825.

Stiebale, J. (2013). The impact of cross-border mergers and acquisitions on the acquirers' r&dfirm-level evidence. *International Journal of Industrial Organization*, 31(4):307–321.

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

Valentini, G. (2012). Measuring the effect of M&A on patenting quantity and quality. *Strategic Management Journal*, 33(3):336–346.

Wang, N. and Hagedoorn, J. (2014). The lag structure of the relationship between patenting and internal R&D revisited. *Research Policy*, 43(8):1275–1285.

Whinston, M. D. (2011). Comment on" competition and innovation: Did arrow hit the bull's eye?". In *The Rate and Direction of Inventive Activity Revisited*, pages 404–410. University of Chicago Press.

Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1):57–76.

Appendix

A The validity of the Flash technology (SSD) as a control for HDD

A.1 European Commission's (DG COMP) view

Case No COMP/M.6214, - Seagate/HDD Business of Samsung, Regulation (EC) No 139/2004 Merger Procedure, Article 8 (1) Date: 19/10/2011, paragraphs 256-261:

"It may, therefore, be concluded that SSDs and HDDs are not currently substitutable due to the significant price differential between the two technologies and the limited storage capacity of SSDs compared to HDDs.

Moreover, any potential future replacement of some types of HDDs with SSDs, notably in the Mission-Critical Enterprise space and the high-end Notebook market such as ultra-portable notebooks, is likely to occur only in the longer term. The competitive pressure currently exerted by SSDs on HDDs would appear to be too limited to impose any price constraint over HDDs suppliers and the current market conditions are not expected to dramatically change in the short term. It could, therefore, be relatively easy for HDDs suppliers to raise HDDs prices in the short term without risking reducing their sales in favour of SSDs. This is the case as the price of SSDs is currently 20 times higher than the price of HDDs, therefore even a price rise of HDDs by more than 50% would not trigger a significant shift towards SSDs.

Given that respondents to the market investigation do not expect this price gap to close in the coming three years, it can be concluded that at least in the near future SSDs will not exert sufficient competitive pressure on HDDs to prevent HDDs suppliers from raising their price. Consequently, SSDs do not currently belong to the same relevant HDD product markets.

It is apparent that from a demand-side perspective customers appear unable to substitute HDDs produced for certain end-uses with other drives displaying a different form factor or other technical features required by different end-use applications.

From a supply-side perspective, the results of the market investigation could not establish sufficient supply-side substitutability in terms of effectiveness and immediacy to justify a broader market definition."

A.2 Testing the independence of the control group

Central to the validity of our assumptions for our empirical model is the strategic interaction between innovating firms. A number of previous papers have looked for evidence of strategic interaction in competing firms' innovation activities in a theoretical setting (list) or highlighted the importance of strategic innovation behaviour (Spulber, 2013; Cohen, 2010). Cohen (2010) reviews these and concludes that they converge to the finding that competing firms match their innovations (i.e that they are strategic substitutes), but warns that there are numerous explanations for positively correlated R&D or product introduction behaviors other than strategic behavior (e.g. common changes in industry-level technological opportunity and demand conditions, spillovers from leading firms which increase the marginal productivity of rival R&D, or simply a catch-up phenomenon where equally capable firms involved in similar activities all move in the same direction at roughly similar rates).

We looked at what the data can tell us about the strategic R&D behaviour between HDD and Flash/SSD. Although these two technologies are not direct competitors - such as the examples reviewed by Cohen (2010) - the two technologies have some overlap in their applications. 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, Flash/SSD manufacturers ('Storage control'). The idea is that if Flash/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's 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 Flash/SSD.

Figure A.1 presents the results comparing HDD (Seagate and Western Digital pooled) to the two types of synthetic controls, the General control and the storage-based. Figure A.1 suggests that there is a small difference, although the difference is not significantly different. This becomes obvious if we compare the General with the Storage control, we find no significant post-merger difference (Figure A.2). Taken at its face value this would suggest that R&D spending in the group where there might be a suspicion of being confounded (Flash/SSD) is no different from the group where we can more confidently assume independence.

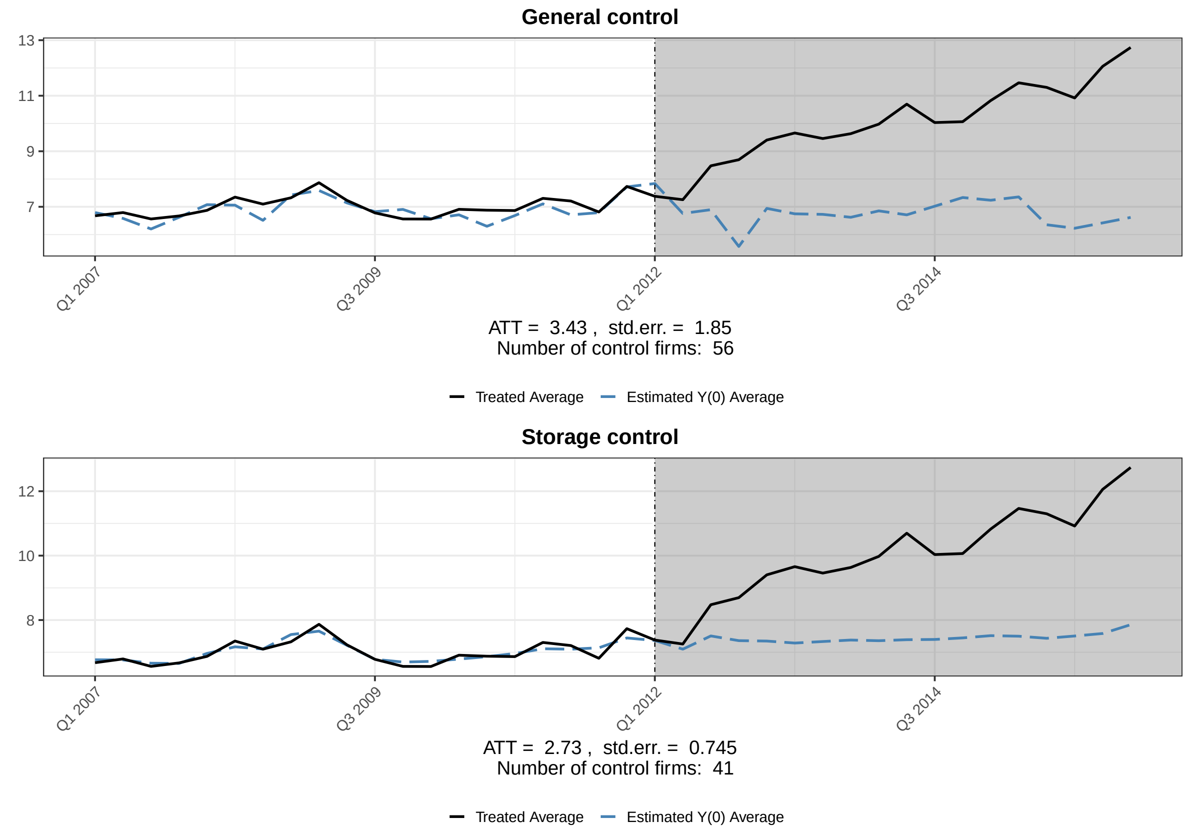

Figure A.1: Change in R&D in HDD in comparison to two different Synthetic control groups

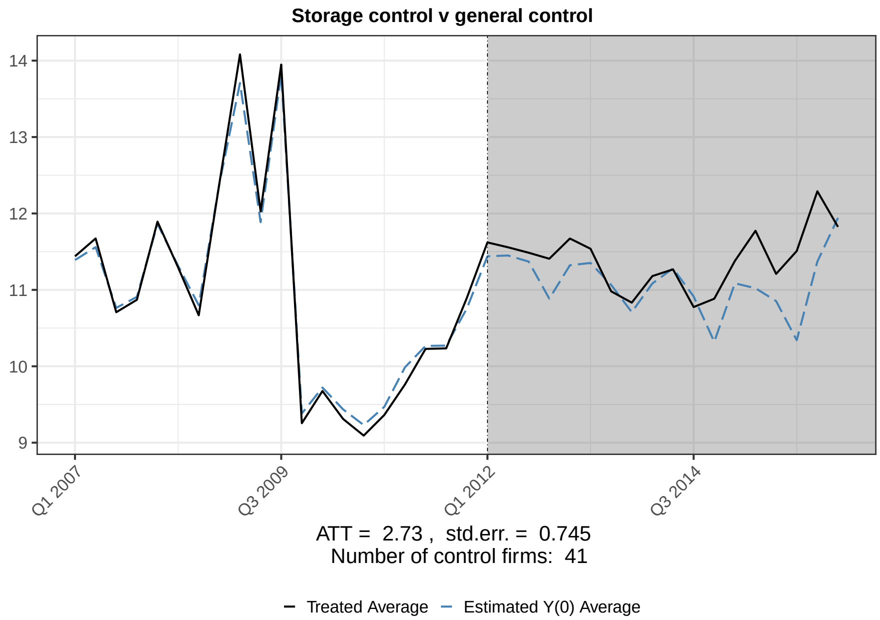

Figure A.2: Comparing SSD (treatment) with a general synthetic group

B Data

B.1 R&D intensity

The R&D data draws light on two methodological issues. First, when evaluating how R&D intensity changes after a merger, one must not ignore an important artefact of this type of data, that is, following a merger, elements of the financial statement of the acquired company are added to the corresponding elements of the financial statement of the acquiring company. This means that for simple arithmetic reasons R&D expenditure and total revenue will be higher in the post-merger period even if the merger does not increase the R&D intensity of the relevant businesses. For this reason, we decided to merge the relevant firm-level characteristics (including R&D expenditure and turnover) for the merging firms. For example, to consolidate the R&D figures of Seagate/Samsung, and WD/Hitachi we made the following assumption. We looked at the leap in the R&D figures in the quarter when the merger was substantiated and assumed that this hike was due to adding the figures from the acquired firm's balance sheet to the acquiring firms. We assumed that this difference is the HDD relevant R&D expenditure of the acquired business. Figure A.1 gives a visual illustration for Seagate's total revenue and R&D expenditure.

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

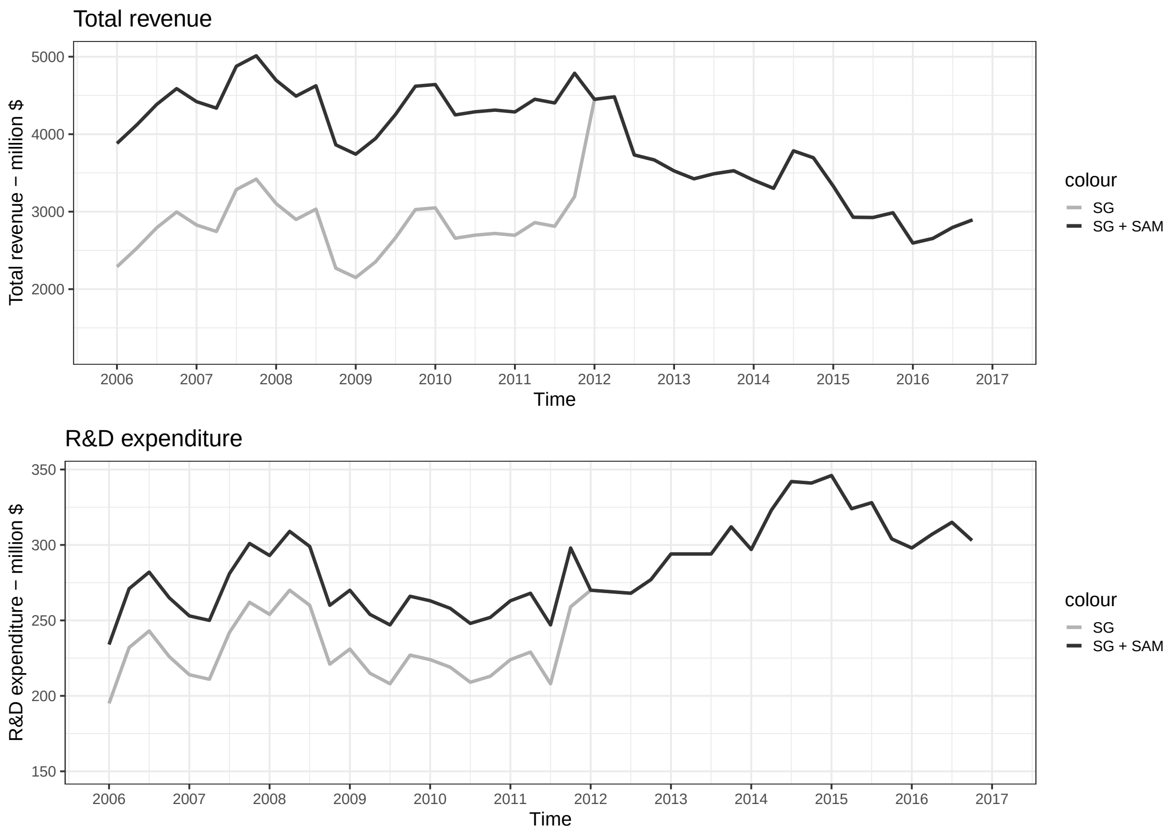

Figure B.1: Consolidation of data - an example from Seagate

different areas, and storage only constitutes around a quarter of its total operating revenue and R&D expenditure. For this reason, we dropped Toshiba from our R&D related analysis.

B.2 Patent activity

We extracted patent data for each technology (HDD, Flash) using PatentsView API, and subsequently grouped the data by firms. The date of each patent application refers to the quarter when a first application is registered at the USPTO. 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. For our analysis, we decided to use a simple patent count. We offer the results using other measures: patent citation, patent intensity, and patent citation intensity, as well as a smaller sample of patents in the most closely relevant patent class (Electric digital data processing). In earlier versions of this paper, we experimented with numerous other measures, including a patent factor, in which 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).

B.3 Innovation output

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

We collected information on 1931 HDDs and on 1353 SSDs that were sold on Amazon between 2001 and 2016.[34] 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

[34]We account for the mergers that happened before 2012, for example, Fujitsu is recorded as Toshiba as a result of their 2009 merger.

in HDD are likely to have a 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 are representative of the whole industry in terms of technological innovations.[35] The sample consists of 33 SSD and 5 HDD brands.[36]

We have access to the following product characteristics for HDDs and SSDs:

B.3.1 Date first appeared 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).

B.3.2 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.

[35] For 98 HDDs and 54 SSDs we could not identify a brand from the scraped data and these were removed from the sample. We removed brands with fewer than 10 products, and we also removed hybrid drives as they represent a combination of the two technologies.

[36] This 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 C.12 in the Appendix.

B.3.3 Storage capacity

Ideally, one would have looked at areal density. However using retail data we had limited access to technical 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.

B.3.4 Retail price

This is the unit capacity retail price of HDD products. Our thinking is that this retail price potentially reflects the cost of HDD (reflecting changes in process and product innovation). Moreover, this can also be thought of as unit user cost as 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 necessarily know when a change happened. Moreover, it is very difficult to search for specific products through these archives, which means matching products by their product numbers was not possible for us. For this reason, we scraped price data as they were in 2017, but in Table B.1 we provide a random sample of 26 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 these 26 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.

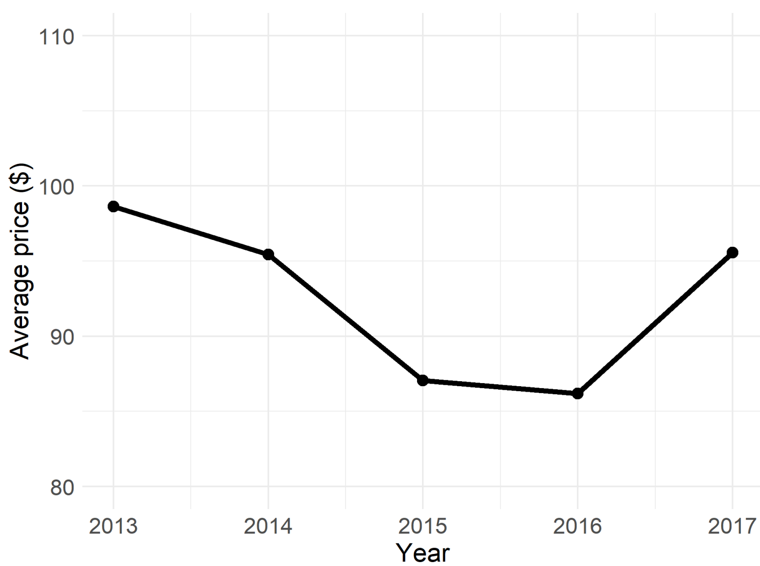

Figure B.2: Historical prices of a sample of HDDs (USD)

Table B.1 shows that prices do recover to near their original level - after a temporary drop (the same is plotted on Figure B.2 for the price averaged over our sample of 26 drives).

Table B.1: Historical prices of a sample of HDDs (USD)

Product number	**2013**	**2014**	**2015**	**2016**	**2017**
WD10EZEX	59.99	52.99	54.99		63.80
ST2000DM001		79.99	79.99	108.99	141.83
WD5000AAK	59.26	52.49	49.99	65.98	70.00
ST1000DM003			50.92	49.99	51.99
ST2000DM001	99.99	79.99	79.99	71.99	108.99
ST1500DL003	116.99	98.00	85.00	56.80	92.69
ST9500325AS	60.55	54.00	87.45	49.95	58.95
HGST 0F12115	134.99	108.95	117.00	134.33	162.53
ST3000VN000	146.89	139.99	114.67	112.37	149.99
HDTC607XK3A1	59.98	54.99	63.99	123.00	
WD10EZEX	67.99	59.99	54.63	49.99	49.99
WD30EFRX	144.44	129	176.56	105.99	109.99
ST1000DM003	79.18	63.14	53.72	48.99	60.24
WD1002FAEX	91.33	116.05	102.23	99.51	92.98
ST31000524AS	65.39	71.7	66.16	61.04	84.99
ST500DM002	59.99	30.11	30.11	30.11	55.54
0F12115	180.01	161.54	147.04	136.66	161.25
ST2000DM001	89.99	79.99	82.96	77.11	77.47
ST95005620AS	102.24	114.88	104.56	97.18	154.98
WD2500BEVT	58.32	63.62	57.92	53.82	58.97
ST9750420AS	77.39	106.09	96.56	89.74	57.2
ST2000DL003	119.42	130.27	92.86	86.3	98.81
WD2002FAEX	155.7	188.09	171.21	159.13	128.78
WD15EARS	131.05	125.28	92.16	95.18	86.04
WD10EARS	92.82	101.25	58.88	85.66	92.84
ST31000528AS	113.5	123.78	92.16	104.73	118.07

Our interpretation of this 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 (Irwin and Pavcnik, 2004). 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).[37]

From the above, we derived our two variables used for measuring innovation output, both recorded by firm j in period t:

- **Number of new models:** This variable includes all newly marketed products. The reason we do not just focus on products with higher capacity is simply that innovation

[37] In the paper we found that the unit price of older products is still cheaper than the unit price of newer products. Therefore even if we found that the retail price of older products as collected in 2018 is lower than the original retail price, it would only mean that our finding of a falling unit price is likely to be an underestimate.

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.[38]

- **Unit price:** The ratio of retail price to storage capacity. We observe for each firm j and time period t the lowest unit price ($/Gb).

C Additional figures and tables

Table C.1: Control firms used in the R&D and patent regressions

Advanced Micro Devices, Inc.	Microsemi Corporation
Alcatel-Lucent	Microsoft Corporation
Apple Inc.	NEC Corporation
Applied Materials, Inc.	NXP Semiconductors N.V.
Broadcom Limited	Nanya Technology Corporation
Canon Inc.	NetApp, Inc.
Casio Computer Co., Ltd.	Oki Electric Industry Co., Ltd.
Cisco Systems, Inc.	Phison Electronics Corporation
Cypress Semiconductor Corporation	Powerchip Technology Corp.
Dell Technologies Inc.	QUALCOMM Incorporated
Dongbu HiTek Co., Ltd.	SK Hynix Inc.
HP Inc.	STMicroelectronics N.V.
Hon Hai Precision Industry Co., Ltd.	SanDisk LLC
Infineon Technologies AG	Seiko Epson Corporation
Intel Corporation	Sharp Corporation
International Business Machines Corporation	Silicon Motion Technology Corporation
LG Electronics Inc.	Taiwan Semiconductor Manufacturing Company Limited
Lenovo Group Limited	Texas Instruments Incorporated
Lite-On Technology Corporation	United Microelectronics Corporation
Macronix International Co., Ltd.	Winbond Electronics Corporation
Micron Technology, Inc.	

[38]For 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.

Table C.2: Control firms used in the innovation output regressions

Apple Inc.
Dell Technologies Inc.
HP Inc.
Intel Corporation
Lenovo Group Limited
Micron Technology, Inc.
SanDisk LLC
Silicon Motion Technology Corporation

Table C.3: Treatment effect estimates for different patent measures - Matrix completion

	Patent count	Patent intensity	Patent citation	Patent citation intensity
	Post-merger period begins 2012Q2			
Seagate	-5.47***	-0.000777***	-27.9***	0.383***
	(0.123)	(0.000173)	(5.58)	(0.0535)
WD	-5.22***	-0.000286**	248***	0.664***
	(0.12)	(0.000131)	(3.33)	(0.0539)
Toshiba	-8.18***	-0.00571***	-165***	0.068
	(0.0144)	(0.000805)	(1.8)	(0.0437)
	Post-merger period begins 2013Q2			
Seagate	-4.9***	-0.000788***	-35.5***	0.384***
	(0.124)	(0.000191)	(3.33)	(0.055)
WD	-4.1***	-0.000362	353***	0.635***
	(0.247)	(0.000247)	(4.68)	(0.0649)
Toshiba	-7.48***	-0.00815***	-145***	0.101***
	(0.0145)	(0.000304)	(8.68)	(0.0294)
	Post-merger period begins 2014Q2			
Seagate	-3.96***	-0.000472***	-17.6**	0.359***
	(0.232)	(0.00018)	(7.14)	(0.0803)
WD	-4.39***	-0.000162	137***	0.585***
	(0.227)	(0.000141)	(34.4)	(0.116)
Toshiba	-6.89***	-0.00522***	-153***	0.336***
	(0.485)	(0.000844)	(10.8)	(0.043)

Table C.4: Treatment effect estimates for all innovation measures - OLS

	R&D Int	Number of new patents	Number of new models	Unit price (log)
Post-merger period begins 2012Q2				
Seagate	2.672***	-5.481***	2.896***	0.468**
	(0.7054)	(0.2517)	(0.5542)	(0.2029)
N	1589	1589	215	215
parallel p-val	0.0701	0	0.7495	0.3318
WD	2.435***	-5.426***	-4.238***	0.878***
	(0.6141)	(0.2285)	(0.5809)	(0.1901)
N	1589	1589	215	215
parallel p-val	0.0967	0	0.2821	0.3718
Toshiba		-7.784***	-0.377	0.728***
		(0.2535)	(0.37)	(0.1277)
N		1588	215	215
parallel p-val		0	0.7385	0.1936
Post-merger period begins 2013Q2				
Seagate	3.124***	-4.878***	2.896***	0.468**
	(0.6594)	(0.2342)	(0.5542)	(0.2029)
N	1589	1589	215	215
parallel p-val	0.0701	0	0.7495	0.3318
WD	2.656***	-4.611***	-4.238***	0.878***
	(0.5822)	(0.218)	(0.5809)	(0.1901)
N	1589	1589	215	215
parallel p-val	0.0967	0	0.2821	0.3718
Toshiba		-6.966***	-0.377	0.728***
		(0.2428)	(0.37)	(0.1277)
N		1588	215	215
parallel p-val		0	0.7385	0.1936
Post-merger period begins 2014Q2				
Seagate	3.206***	-4.703***	6.267***	0.449*
	(0.6141)	(0.2309)	(0.6674)	(0.2623)
N	1589	1589	215	215
parallel p-val	0.0701	0	0.9555	0.2461
WD	2.664***	-4.897***	-6.431***	0.891***
	(0.5467)	(0.2156)	(0.637)	(0.2303)
N	1589	1589	215	215
parallel p-val	0.0967	0	0.4053	0.2888
Toshiba		-7.656***	-0.14	0.937***
		(0.2396)	(0.4172)	(0.1318)
N		1588	215	215
parallel p-val		0	0.4177	0.1466

*= $p < 0.10$, **= $p < 0.05$, ***= $p < 0.01$. Robust standard errors. We tested for the joint significance of time fixed effects and decided against using them, we have instead preferred a post-merger dummy. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

Table C.5: Treatment effect estimates for all innovation measures - Generalised Synthetic Control

	R&D Int	Number of new patents	Number of new models	Unit price (log)
	Post merger period begins 2012-Q2			
Seagate	2.20*	-8.01***	6.10***	-3.21***
s.e.	(1.27)	(1.59)	(1.56)	(1.03)
WD	2.38***	-1.87	7.39	-3.34***
s.e.	(0.872)	(1.77)	(7.01)	(1.22)
Toshiba		-3.54**	0.361	-2.34***
s.e.		(1.72)	(0.492)	(0.815)
Number of control firms	41	41	8	8
	Post merger period begins 2013-Q2			
Seagate	2.63**	-4.68***	8.35***	-2.68***
s.e.	(1.19)	(1.06)	(2.31)	(0.636)
WD	2.62***	-1.08	-1.72	-2.71***
s.e.	(0.93)	(1.67)	(4.29)	(0.699)
Toshiba		-2.19	3.42**	-1.64**
s.e.		(1.82)	(1.6)	(0.775)
Number of control firms	41	41	8	8
	Post merger period begins 2014-Q2			
Seagate	2.57***	-3.95***	7.97***	-0.348
s.e.	-0.5	-0.189	-0.6	-0.36
WD	1.96***	-4.4***	-5.67***	-0.526*
s.e.	-0.311	-0.232	-0.468	-0.29
Toshiba		-6.58***	2.12***	0.158
s.e.		-0.428	-0.217	-0.144
Number of control firms	41	41	8	8

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Generalised synthetic control estimates, as in Xu (2017). The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

Table C.6: Treatment effects by Treated Firm (with MSPE values)

	R&D Int	Number of new patents	Number of new models	Unit price (log)
	Post-merger period begins 2012Q2			
Seagate	2.870***	-5.420***	4.120***	-0.898**
s.e.	(0.824)	(0.182)	(0.475)	(0.458)
mspe	6.59	4.39	2.07	0.596
WD	2.590***	-5.230***	-7.470***	-0.492
s.e.	(0.701)	(0.139)	(0.429)	(0.399)
mspe	6.59	4.29	2.22	0.582
Toshiba		-8.070***	1.600***	-0.962***
s.e.		(0.017)	(0.189)	(0.229)
mspe	6.54	4.2	2.01	0.623
Pooled	2.730***	-5.800***	-0.815	-0.717**
s.e.	(0.761)	(0.759)	(2.92)	(0.309)
mspe	4.94	4.44	5.02	0.705
Number of control firms	41	41	8	8
	Post-merger period begins 2013Q2			
Seagate	2.740***	-4.840***	4.120***	-0.392
s.e.	(0.605)	(0.132)	(0.550)	(0.400)
mspe	6.65	5.47	2.24	0.587
WD	2.340***	-4.450***	-4.940***	-0.349
s.e.	(0.534)	(0.113)	(0.462)	(0.329)
mspe	6.66	6.02	2.7	0.625
Toshiba		-7.100***	0.845***	-0.453***
s.e.		(0.210)	(0.136)	(0.175)
mspe	6.65	5.64	2.08	0.705
Pooled	2.540***	-5.030***	-0.0122	-0.301
s.e.	(0.621)	(0.772)	(2.070)	(0.263)
mspe	5.12	4.24	8.15	0.698
Number of control firms	41	41	8	8
	Post-merger period begins 2014Q2			
Seagate	2.570***	-3.950***	7.970***	-0.348
s.e.	(0.500)	(0.189)	(0.600)	(0.360)
mspe	8.57	7.3	2.2	0.536
WD	1.960***	-4.400***	-5.670***	-0.526*
s.e.	(0.311)	(0.232)	(0.468)	(0.290)
mspe	8.59	7.25	3.48	0.464
Toshiba		-6.580***	2.120***	0.158
s.e.		(0.428)	(0.217)	(0.144)
mspe	8.52	6.99	1.98	0.559
Pooled	2.280***	-5.010***	1.380	-0.205
s.e.	(0.512)	(0.935)	(3.329)	(0.248)
mspe	6.47	4.76	11.7	0.566
Number of control firms	41	41	8	8

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. One-way (unit) fixed effects model. Bootstrapped (1000 iterations) standard errors. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

Table C.7: Treatment effect estimates for all innovation measures - excluding Flash/SSD firms involved in mergers

	R&D Int	Number of new patents	Number of new models	Unit price (log)
	Post merger period begins 2012-Q2			
Seagate	2.99***	-4.81***	4.63***	-1.13
s.e.	(1.1)	(0.304)	(0.838)	(1.19)
WD	2.79***	-4.73***	-6.94***	-0.619
s.e.	(0.558)	(0.352)	(1.5)	(0.996)
Toshiba	0.401***	-7.07***	0.934***	-0.676***
s.e.	(0.134)	(0.38)	(0.0756)	(0.138)
Pooled	2.89***	-5.93***	-0.527	-0.829
s.e.	(0.981)	(0.78)	(2.75)	(0.705)
Number of control firms	30	30	4	4
	Post merger period begins 2013-Q2			
Seagate	2.93***	-4.15***	5.05***	-0.714
s.e.	(0.734)	(0.363)	(1.12)	(0.984)
WD	2.55***	-3.93***	-4.02	-0.478
s.e.	(0.632)	(0.328)	(2.7)	(5.25)
Toshiba	0.261	-6.32***	0.422***	-0.242**
s.e.	(0.25)	(0.407)	(0.0911)	(0.0966)
Pooled	2.74***	-5.2***	0.542	-0.502
s.e.	(0.827)	(0.811)	(2)	(0.526)
Number of control firms	30	30	4	4
	Post merger period begins 2014-Q2			
Seagate	2.62***	-3.22***	8.86***	-0.458
s.e.	(0.442)	(0.315)	(0.476)	(0.768)
WD	2.04***	-3.97***	-4.95**	-0.492
s.e.	(0.504)	(0.344)	(2.45)	(0.581)
Toshiba	0.526*	-6.07***	1.62***	0.159
s.e.	(0.308)	(0.382)	(0.0836)	(0.104)
Pooled	2.36***	-4.96***	2.04	-0.264
s.e.	(0.636)	(0.942)	(3.15)	(0.343)
Number of control firms	30	30	4	4

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. The table presents matrix completion results for a subset, where we excluded those Flash/SSD firms that engaged of a merger of any type. All of these mergers were either buying fringe players or conglomerate mergers. One-way (unit) fixed effects model. Bootstrapped (1000 iterations) standard errors. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

Table C.8: Impact on innovation output, merger effect at 2013Q2 (R&D and patent lags) - Seagate

	R&D (lag 4) Patent (lag 2)		R&D (lag 4) Patent (lag 4)		R&D (lag 6) Patent (lag 3)		R&D (lag 6) Patent (lag 6)		R&D (lag 8) Patent (lag 4)		R&D (lag 8) Patent (lag 8)	
	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price
Post-merger ($\overline{\mu}_t$)	1.09**	-1.146***	1.198***	-1.15***	1.277***	-1.159***	1.262***	-1.168***	1.403***	-1.161***	1.362***	-1.161***
	(0.4362)	(0.1392)	(0.4359)	(0.1397)	(0.4419)	(0.1371)	(0.421)	(0.1367)	(0.4114)	(0.1363)	(0.4012)	(0.1349)
Treatment effect (δ_0)	-25.608***	1.143	-26.646***	0.996	-23.954***	1.934	-46.022***	1.957	-28.002***	1.63	-36.26***	1.686
	(5.9795)	(1.9077)	(5.827)	(1.8669)	(6.1867)	(1.9187)	(6.0305)	(1.9587)	(5.0061)	(1.6582)	(4.4727)	(1.5038)
R&D (δ_{11})	0.216**	-0.021	0.231**	-0.021	0.112	0.028	0.112	0.029	0.236**	0.059*	0.241**	0.058*
	(0.1043)	(0.0333)	(0.1036)	(0.0332)	(0.1053)	(0.0327)	(0.1007)	(0.0327)	(0.0978)	(0.0324)	(0.0964)	(0.0324)
Patent count (δ_{21})	0.07	-0.002	-0.081	0.004	-0.006	-0.019	0.017	-0.01	-0.068	0.004	-0.051	0.011
	(0.0809)	(0.0258)	(0.0768)	(0.0246)	(0.0821)	(0.0255)	(0.0733)	(0.0238)	(0.0736)	(0.0244)	(0.0748)	(0.0252)
R&D × treatment effect (δ_{12})	3.67***	-0.073	3.746***	-0.063	3.71***	-0.201	5.445***	-0.202	4.374***	-0.168	4.751***	-0.169
	(0.6349)	(0.2026)	(0.6237)	(0.1998)	(0.6992)	(0.2169)	(0.6771)	(0.2199)	(0.5923)	(0.1962)	(0.5572)	(0.1874)
Patent × treatment effect (δ_{22})	-1.487***	-0.016	-1.424***	0.016	-1.176**	0.015	2.546***	0.006	-0.84*	0.001	1.537***	-0.002
	(0.5114)	(0.1632)	(0.4983)	(0.1597)	(0.5394)	(0.1673)	(0.5218)	(0.1695)	(0.4934)	(0.1634)	(0.4677)	(0.1573)
P-val parallel trend	0.7475	0.756	0.7475	0.756	0.1323	0.8857	0.1323	0.8857	0.3662	0.6617	0.3662	0.6617
N	224	224	224	224	224	224	224	224	224	224	224	224
Decomposition of R&D effects												
R&D change	0.361**	-0.035	0.385**	-0.036	0.158	0.039	0.158	0.04	0.268**	0.067*	0.274**	0.066*
	(0.1742)	(0.0556)	(0.1731)	(0.0554)	(0.1482)	(0.046)	(0.1418)	(0.046)	(0.1112)	(0.0368)	(0.1096)	(0.0369)
R&D Productivity change	6.249***	0.51	4.868***	0.453	5.644***	0.331	-2.584	0.347	4.671***	0.376	-0.776	0.426
	(1.6841)	(0.5373)	(1.6986)	(0.5442)	(1.7437)	(0.5408)	(1.7114)	(0.5558)	(1.6465)	(0.5454)	(1.6224)	(0.5455)
R&D joint effect	6.609***	0.475	6.253***	0.417	5.802***	0.37	-2.426	0.387	4.939***	0.444	-0.503	0.492
	(1.6708)	(0.533)	(1.6863)	(0.5403)	(1.7369)	(0.5386)	(1.7029)	(0.5531)	(1.6423)	(0.544)	(1.6162)	(0.5434)
Decomposition of patent effects												
Patent change	-0.091	0.003	0.184	-0.01	0.009	0.027	-0.063	0.035	0.156	-0.01	0.256	-0.055
	(0.1045)	(0.0333)	(0.1747)	(0.056)	(0.1151)	(0.0357)	(0.2664)	(0.0865)	(0.1675)	(0.0555)	(0.3722)	(0.1251)
Patent productivity change	-28.925***	1.108	-29.932***	1.033	-26.577***	1.968	-40.342***	1.97	-29.94***	1.632	-32.833***	1.683
	(5.6116)	(1.7904)	(5.5172)	(1.7676)	(5.6911)	(1.765)	(5.5133)	(1.7907)	(4.5546)	(1.5086)	(4.3066)	(1.448)
Patent joint effect	-29.015***	1.11	-29.749***	1.023	-26.569***	1.996	-40.405***	2.005	-29.784***	1.622	-32.577***	1.628
	(5.6113)	(1.7903)	(5.5151)	(1.7669)	(5.6893)	(1.7644)	(5.5072)	(1.7887)	(4.5504)	(1.5073)	(4.2926)	(1.4433)

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Robust standard errors. We tested for the joint significance of time fixed effects and decided against using them, we have instead preferred a post-merger dummy. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

Table C.9: Impact on innovation output, merger effect at 2013Q2 (R&D and patent lags) - Western Digital

	R&D (lag 4) Patent (lag 2)		R&D (lag 4) Patent (lag 4)		R&D (lag 6) Patent (lag 3)		R&D (lag 6) Patent (lag 6)		R&D (lag 8) Patent (lag 4)		R&D (lag 8) Patent (lag 8)	
	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price	Number of new models	Unit price
Post-merger ($\overline{\mu}_t$)	1.127**	-1.152***	1.167**	-1.134***	1.284***	-1.162***	1.233***	-1.161***	1.333***	-1.146***	1.325***	-1.154***
	(0.4503)	(0.1364)	(0.4546)	(0.1367)	(0.4551)	(0.1343)	(0.4178)	(0.1341)	(0.4507)	(0.1333)	(0.4315)	(0.1323)
Treatment effect (δ_0)	40.029***	1.724	33.856***	1.65	23.581**	3.184	11.883	3.062	16.705**	2.897	18.798***	2.884
	(11.7817)	(3.5695)	(10.2284)	(3.077)	(9.363)	(2.7638)	(8.7081)	(2.7939)	(7.4349)	(2.1982)	(7.1885)	(2.204)
R&D (δ_{11})	0.168	-0.021	0.179*	-0.019	0.105	0.025	0.102	0.025	0.235**	0.056*	0.231**	0.057*
	(0.1076)	(0.0326)	(0.1081)	(0.0325)	(0.1089)	(0.0321)	(0.1001)	(0.0321)	(0.1076)	(0.0318)	(0.1037)	(0.0318)
Patent count (δ_{21})	0.107	0.008	0.029	-0.016	0.042	-0.013	0.107	-0.017	0.038	-0.015	0.112	-0.015
	(0.0682)	(0.0207)	(0.0697)	(0.021)	(0.071)	(0.021)	(0.0652)	(0.0209)	(0.0703)	(0.0208)	(0.072)	(0.0221)
R&D × treatment effect (δ_{12})	-3.988***	-0.079	-3.78***	-0.086	-2.67***	-0.257	-2.42***	-0.255	-2.503***	-0.25	-1.682**	-0.241
	(1.0052)	(0.3046)	(0.9499)	(0.2858)	(0.8887)	(0.2623)	(0.8181)	(0.2625)	(0.7878)	(0.2329)	(0.7248)	(0.2222)
Patent × treatment effect (δ_{22})	-0.268	-0.002	0.089	0.005	-0.149	0.016	0.685***	0.022	0.213	0.02	-0.653***	0.011
	(0.2207)	(0.0669)	(0.2005)	(0.0603)	(0.1931)	(0.057)	(0.1479)	(0.0474)	(0.2132)	(0.063)	(0.1578)	(0.0484)
P-val parallel trend	0.4522	0.9729	0.4522	0.9729	0	0.2431	0	0.2431	0	0.4186	0	0.4186
N	224	224	224	224	224	224	224	224	224	224	224	224
Decomposition of R&D effects												
R&D change	0.25	-0.031	0.265*	-0.028	0.186	0.045	0.181	0.045	0.425**	0.101*	0.417**	0.102*
	(0.1594)	(0.0483)	(0.1601)	(0.0482)	(0.1936)	(0.0571)	(0.178)	(0.0571)	(0.1942)	(0.0574)	(0.1872)	(0.0574)
R&D productivity	-2.447	0.878	-6.398**	0.735	-3.78	0.553	-12.912***	0.453	-7.878***	0.438	2.285	0.516
	(2.7195)	(0.8239)	(2.564)	(0.7713)	(2.4391)	(0.72)	(2.1288)	(0.683)	(2.6834)	(0.7933)	(2.3124)	(0.709)
R&D joint effect	-2.198	0.847	-6.133**	0.707	-3.594	0.598	-12.731***	0.498	-7.453***	0.539	2.702	0.619
	(2.7101)	(0.8211)	(2.5558)	(0.7689)	(2.4315)	(0.7177)	(2.12)	(0.6802)	(2.6777)	(0.7917)	(2.302)	(0.7058)
Decomposition of patent effects												
Patent change	-0.144	-0.011	-0.075	0.04	-0.118	0.038	-0.291	0.046	-0.098	0.039	-0.394	0.052
	(0.0913)	(0.0277)	(0.1787)	(0.0538)	(0.2007)	(0.0592)	(0.178)	(0.0571)	(0.1803)	(0.0533)	(0.2527)	(0.0775)
Patent productivity	37.144***	1.7	34.809***	1.707	22.047**	3.347	19.676**	3.317	18.995**	3.113	11.217	3.012
	(10.7645)	(3.2613)	(10.1809)	(3.0627)	(9.1761)	(2.7086)	(8.4461)	(2.7098)	(7.8217)	(2.3125)	(7.2037)	(2.2086)
Patent joint effect	37***	1.688	34.735***	1.748	21.929**	3.384	19.385**	3.363	18.897**	3.152	10.822	3.063
	(10.7643)	(3.2613)	(10.1793)	(3.0622)	(9.1732)	(2.7077)	(8.4438)	(2.7091)	(7.8184)	(2.3116)	(7.1988)	(2.2071)

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Robust standard errors. We tested for the joint significance of time fixed effects and decided against using them, we have instead preferred a post-merger dummy. The control variables are: total revenue, gross profit, net income, expenses, total assets, and total debt.

Table C.10: Impact on innovation output, weighted regression - Seagate

	2013Q2		2014Q2	
	Number of new models	Unit price	Number of new models	Unit price
Post-merger ($\overline{\mu}_t$ post-merger)	0.32	-2.246***	-0.061	-2.258***
	(0.6832)	(0.2696)	(0.7037)	(0.2831)
Treatment effect (δ_0)	-27.207***	1.921	-16.065	1.37
	(8.8893)	(3.6783)	(17.0797)	(7.1159)
R&D (δ_{11})	1.012***	-0.068**	0.98***	-0.073**
	(0.1257)	(0.0306)	(0.1241)	(0.0308)
R&D $\times$ treatment effect (δ_{12})	3.3***	-0.027	2.348	0.044
	(1.0056)	(0.414)	(1.8238)	(0.7588)
N	225	225	225	225
Decomposition of R&D effects				
R&D change	4.543***	-0.115**	4.538***	-0.154**
	(0.5645)	(0.0517)	(0.5746)	(0.0647)
R&D productivity change	1.433	1.691**	5.81***	1.778**
	(2.0209)	(0.8047)	(2.1072)	(0.8472)
R&D joint effect	5.976***	1.576**	10.348***	1.624*
	(1.9405)	(0.803)	(2.0273)	(0.8447)

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Robust standard errors. We tested for the joint significance of time fixed effects and decided against using them, we have instead preferred a post-merger dummy.

Table C.11: Impact on innovation output, weighted regression - Western Digital

	2013Q2		2014Q2	
	Number of new models	Unit price	Number of new models	Unit price
Post-merger ($\overline{\mu}_t$ post-merger)	-0.02 (0.3951)	-0.297*** (0.0415)	0.174 (0.4219)	-0.3*** (0.0432)
Treatment effect (δ_0)	40.356*** (9.4097)	0.685 (1.0503)	-11.742 (22.3406)	1.142 (2.3613)
R&D (δ_{11})	0.842*** (0.1856)	-0.057 (0.0349)	0.677*** (0.1866)	-0.055* (0.0328)
R&D × treatment effect (δ_{12})	-4.353*** (0.8835)	-0.055 (0.1004)	0.453 (2.0148)	-0.1 (0.2137)
N	225	225	225	225
Decomposition of R&D effects				
R&D change	1.626*** (0.3584)	-0.139* (0.0843)	1.379*** (0.38)	-0.136* (0.0812)
R&D productivity change	-6.007*** (1.1889)	0.104 (0.1496)	-6.725*** (1.2837)	0.03 (0.1525)
R&D joint effect	-4.381*** (1.1336)	-0.034 (0.1235)	-5.345*** (1.2261)	-0.106 (0.129)

*=$p < 0.10$, **=$p < 0.05$, ***=$p < 0.01$. Robust standard errors. We tested for the joint significance of time fixed effects and decided against using them, we have instead preferred a post-merger dummy.

Table C.12: HDD and SSD brands in the sample

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

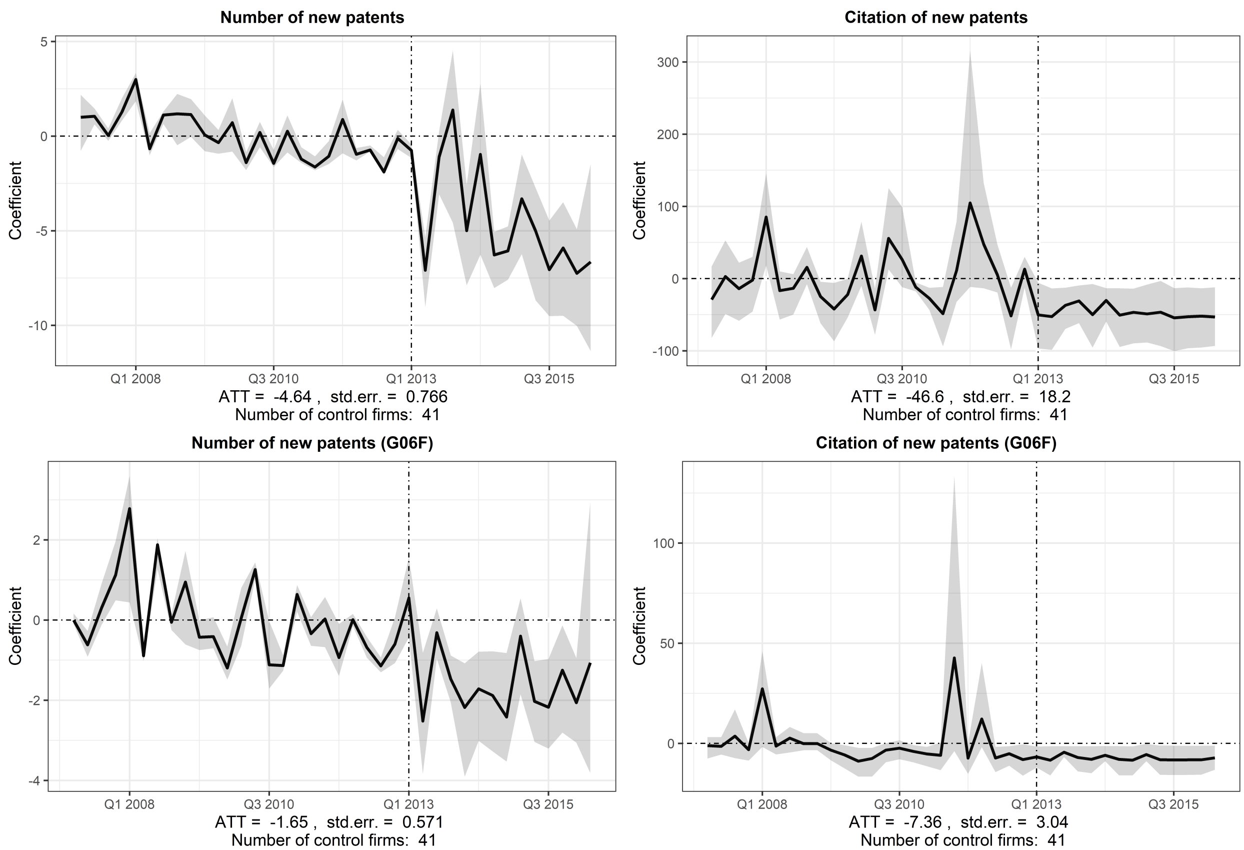

Figure C.1: Treatment effect estimates for various patent measures

Note: The number of new patents is the measure we used as our headline result in Figure 3. Citations measure the number of times a patent is cited by other USPTO patents. G06F represents the Electrical digital data processing class (G06F) of USPTO patents.

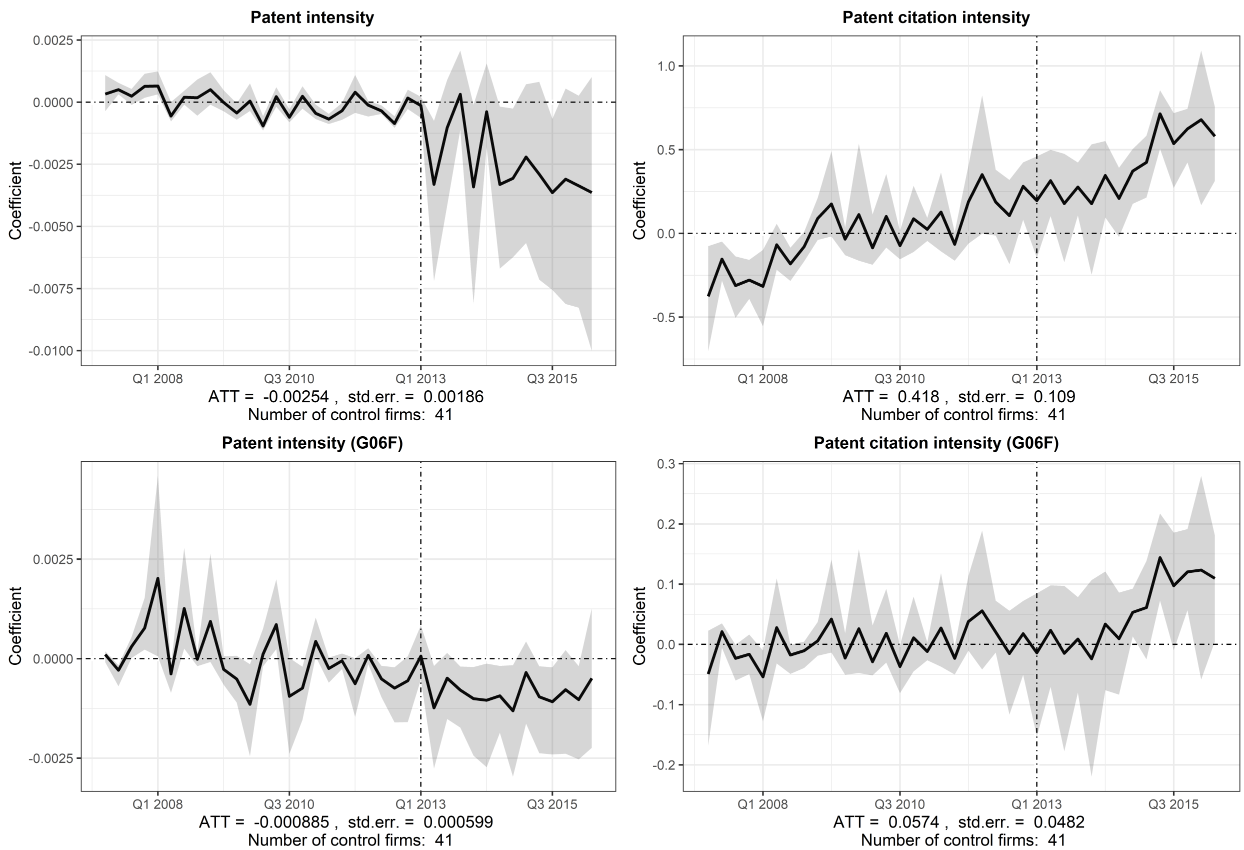

Figure C.2: Treatment effect estimates for patent count and citation

Note: Patent intensity measures the number of new patents divided by the company's turnover. Patent citation intensity is citation count divided by the company's turnover. G06F represents patents that are in the Electrical digital data processing class (G06F) of USPTO patents.

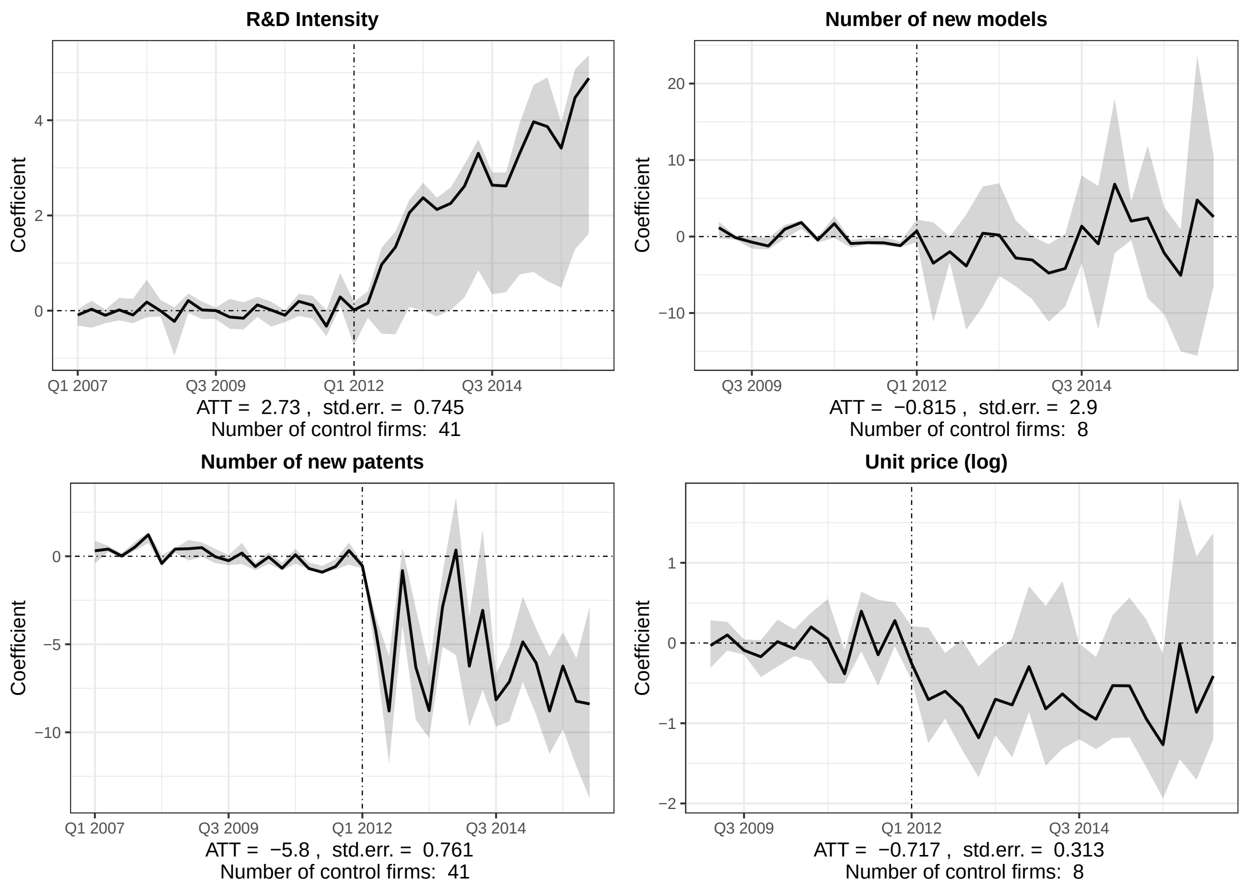

Figure C.3: Average Treatment Effects on the Pooled Treated Firms (post-merger begins 2012Q2)

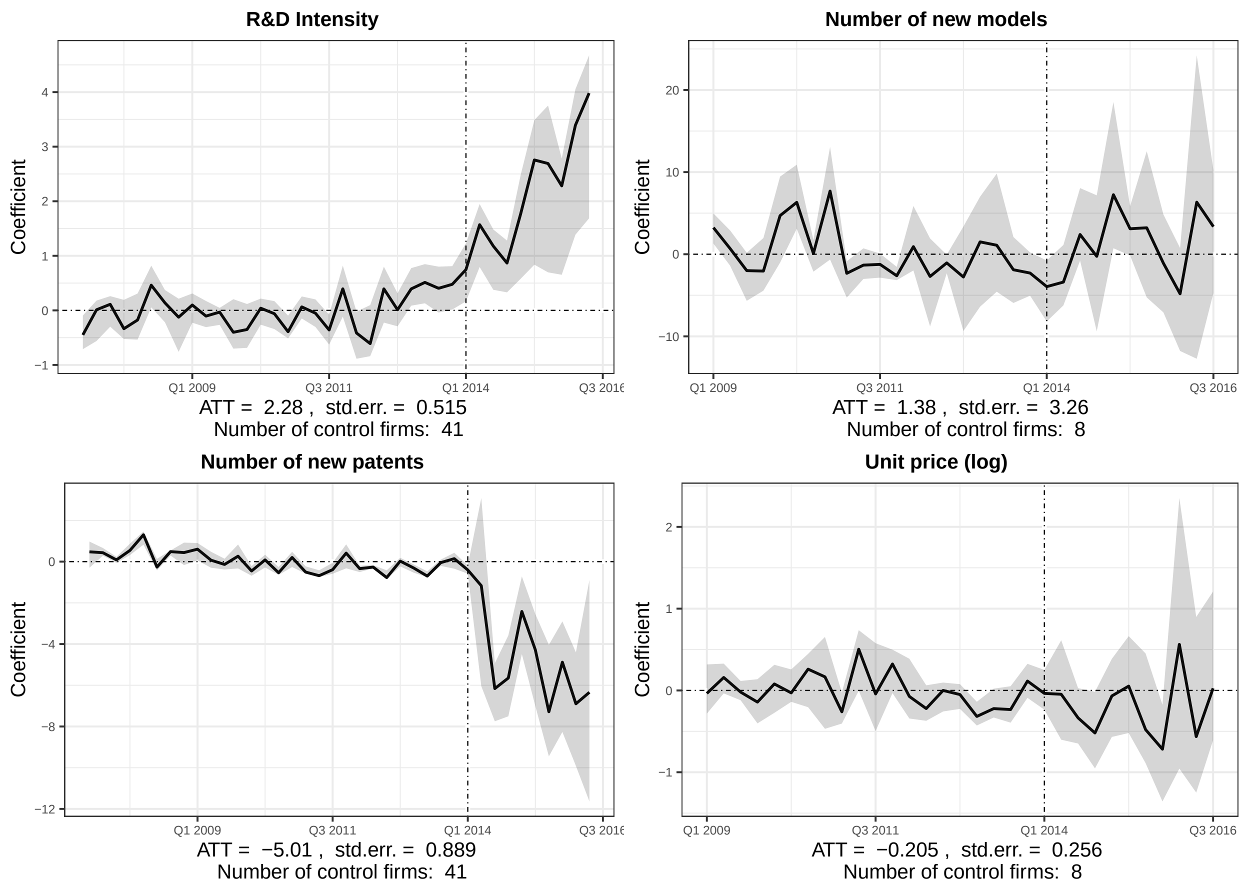

Figure C.4: Average Treatment Effects on the Pooled Treated Firms (post-merger begins 2014Q2)

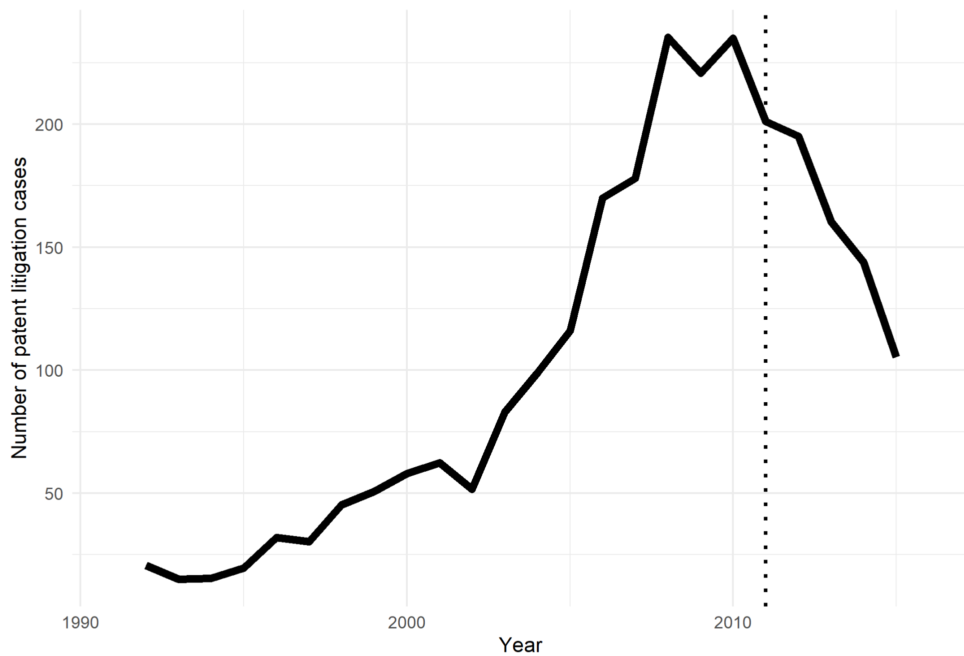

Figure C.5: The number of patent litigation cases related to HDD

Source: `https://www.uspto.gov/learning-and-resources/electronic-data-products/patent-litigation-docket-reports-data`