NBER WORKING PAPER SERIES

A SYNTHETIC MODEL OF DISRUPTION AND EXPERIMENTATION

Joshua S. Gans

Working Paper 29091

<http://www.nber.org/papers/w29091>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

July 2021

Thanks to Steven Callander, Scott Stern and seminar participants at the University of Toronto and Harvard Business School for useful feedback. All errors remain my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Joshua S. Gans. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Synthetic Model of Disruption and Experimentation

Joshua S. Gans

NBER Working Paper No. 29091

July 2021

JEL No. L26,O32

ABSTRACT

This paper examines how a firm's choice of the type of experiment impacts on its potential exploitation of new technological opportunities. It does so in the context of the failure of successful firms (or disruption) where the literature has informally suggested that firms undertake errors in experimental choice (in particular, choosing experiments that involved biased signals). It is shown that firms will generically choose biased over unbiased experiments even when there are no differences in their relative costs. This is done to better inform decisions regarding the exploitation of technological opportunities. It is shown that these choices can differ between incumbents and entrants based on their fundamentals as well as because of the anticipation of competition between them.

Joshua S. Gans

Rotman School of Management

University of Toronto

105 St. George Street

Toronto ON M5S 3E6

CANADA

and NBER

joshua.gans@rotman.utoronto.ca

1 Introduction

How should firms learn about new technological opportunities in order to evaluate whether they should exploit or invest in them? Being able to learn about a technology and evaluate it prior to costly investments has benefits (e.g., Nelson (1982)). Better information reduces errors in exploitation decisions both in the direction of investing in technologies that prove to have limited potential and in the direction of missed opportunities to exploit technologies that turn out to be significant. Those missed opportunities have been a particular focus on the literature regarding the loss of market leadership by firms (otherwise known as “disruption”). While the level of informativeness has been explored in the literature on innovation (e.g., Kerr et al. (2014), Agrawal et al. (2021)), there has been little focus on the type of information that should be gathered and the nature of the experiments that should be performed.

This is despite such issues being informally raised in the disruption literature. To take an example, one of the more compelling stories from Christensen (1997) was that, for disruptive innovations (that is, innovations with initially poor performance but a trajectory of faster improvement), when incumbents listen to their own customers whether they would want a product based on those technologies, the negative signals lead them astray. Christensen argued that this was because a firm’s own customers are less likely to believe they are under-served by the existing technology.

Precisely because these firms listened to their customers, invested aggressively in new technologies that would provide their customers more and better products of the sort they wanted, and because they carefully studied market trends and systematically allocated investment capital to innovations that promised the best returns, they lost their positions of leadership. (Christensen (1997))

A specific way of formally considering this story is that firms had a choice of the sample of customers they drew upon to evaluate a new technological opportunity. Implicitly, Christensen was arguing that the sample – focusing on existing customers – would give a biased signal of the technology’s potential when it is was disruptive precisely because those technologies were of greater value to under-served customers or people who may not be presently customers of the firm.¹

It is well-known that non-random sampling can give rise to biased signals. But Christensen took from this that firms were making an error in relying on experiments focused on existing customers. That error caused them to miss the technological opportunity and allow others to do fill the gap causing a loss in market leadership.

¹Azevedo et al. (2020) show that even with very large customer data-sets, there are trade-offs in the sample selection of the type that means that signal-bias is an issue.

The purpose of this paper is to formally explore this contention. Is it the case that firms exploring technological opportunities should avoid biased signals? That is, should they choose experiments with greater or lower potential for false negatives/positives (or Type I versus Type II errors)? And if not, is the bias that established firms choose in their experiments, sub-optimal or otherwise mistaken? In so doing, I show that, even when firms can choose unbiased experiments, they will choose to bias their experiments to allow for better decisions regarding whether to exploit technological opportunities. Moreover, there is a difference between established firms and startup entrants, in their preferences over experiments that is consistent with observations of their behaviour in the disruption literature but with a different normative conclusion.

1.1 The Disruption Literature

There is a large literature on why successful, established firms fail. Thanks to Clay Christensen (1997) this falls under the rubric of the term ‘disruption.’ Disruption is an event whereby an established firm fails not because of mistakes or complacency but instead because it continued to be managed according to the practices that made it successful in the first place. Christensen identified technologies (aka disruptive technologies) that, for the firm’s consumers, were of inferior current performance but made trade-offs both at the time and inter-temporally that could appeal to potential or under-served consumers. He argued that it was this characteristic whereby firms, who were oriented towards satisfying existing customers, would fail to take advantage of technological opportunities that, in retrospect, undermined their competitive advantage.

While Christensen’s argument was focused on the demand-side, Henderson and Clark (1990) presented a contemporaneous supply-side argument regarding the same phenomenon. They too argued that successful, established firms are vulnerable to losing market leadership because they continue to focus on practices that made them successful. For Henderson and Clark, the culprit was established firms’ internal organisation for innovations that evolved to stress specialisation and modularity amongst product development teams. The problem, they found, was that a new technological opportunity could present itself that had inferior current performance to the existing technology but had higher potential as it involved a new architecture by which components fit together. That potential required a reorganisation of the firm’s innovation system around a new architecture but the existing firm was structured instead to evaluate component level improvements. This created an opportunity for new entrants who could emerge that were not hamstrung by an existing organisational structure. In other words, their costs of developing the new architectural technological opportunity

were lower, perhaps considerably so, than incumbent firms.

Examples abound of both demand and supply-side forces harming established firms. For instance, Blockbuster failed to leverage its existing relationships with movie studios and customers to transition to on-demand video seeming to believe that customers enjoyed the process of physically selecting videos in store. Similarly, Blackberry, the smart-phone market leader, failed to understand that Apple's iPhone was architecturally distinct from its own offerings with a different user-interface and application development process at a time when its own hardware was more robust, both physically and for the cell phone networks it operated on. In cases such as these, and with the benefit of hindsight, many wonder why they did not see it coming.

The answer, of course, as Christensen, Henderson and Clark point out, is that they did see it coming in the sense of knowing there was a new technological opportunity that required evaluation and understanding. For Blockbuster, in the early 2000s, before Netflix was even founded, they launched a world-first movies on-demand system to supply consumers videos directly to their home over the Internet. That product failed due to a lack of home broadband in the US at the time. For Blackberry, their sales increased following the iPhone's launch in 2007 and they were able to develop a touch-screen version of their own phone relatively quickly. However, within a few years, their performance providing a device beyond calls, messages and email did not match the wealth of apps both Apple and then Android could support. While Blackberry could match or exceed hardware component performance it was and stayed behind on its integration with software; the core of Apple's architectural innovation.

This shifts the mechanism of successful firm failure from a lack of attention to, instead, the methods by which those firms evaluated new technological opportunities. Moreover, it was not the case that those opportunities could not be evaluated. Indeed, for technologies that were sustaining or modular, one of the reasons for their past success was that they were able to evaluate and exploit these technologies and maintain market leadership. Instead, it was argued that there was a failure to evaluate those new opportunities differently that sowed the seeds of their demise. Along with that came a call to recognise opportunities that were disruptive/architectural and deploy different evaluation methods to avoid squandering those opportunities to preserve their market leadership.

The research question this paper asks is: was it incorrect for established firms to continue with certain evaluation and experimental methods to understand new disruptive/architectural technologies? In so doing, we can evaluate the normative prescription for those firms that they should have adjusted their experiments accordingly and that this would have been advantageous to them. The goal here is to take the informal arguments of Christensen,

Henderson, Clark and many of those who came after them seriously and assess their efficacy. This is done using a formal economic model of experimentation and competition between an incumbent and an entrant firm.

The approach used to answer this question is as follows:

1. When do firms, evaluating whether to exploit a technological opportunity, choose certain types of experiments?
2. How do those choices relate to the type of technological opportunity (disruptive/architectural versus sustaining/modular) and to the type of firm (incumbent versus entrant)?
3. Do the answers to these questions change when incumbents and entrants face each other in potential competition based on the technological opportunity?

In answering these questions, the focus is on the nature of experiments identified in the literature and whether the optimal choices of firms match with empirical choices and proposed normative choices. Therefore, before turning to the body of the paper, it is useful to review the types of experiments discussed in the literature and how they might be classified.

1.2 Experimental Types

An experiment is a signal of an underlying state variable. Here we focus on a new technological opportunity that either “works” (W) or “does not work” (N). A technology that works allows any firm exploiting that technology to earn a positive return while one that does not work results in a lower or no return.

This paper takes an abstract view of experiments. Following a recent literature in economics, we treat experiments as distinct information sources (Che and Mierendorff (2019), Liang et al. (2018), Mayskaya (2019)). As will be argued here, this is a natural interpretation in light of experimental choices that often are, informally, advocated in studies of innovation and entrepreneurship.²

Earlier, we noted the discussion by Christensen (1997) of established firms receiving feedback on new technological opportunities from their existing customers. While Christensen expressed and argued it informally, we can express it in information-theoretic terms. Specifically, your own customers are an information source (e.g., the subjects of a survey; i.e., the experiment design) that gives you a signal of the magnitude of any increment to profits you might achieve from exploiting a new technological opportunity. Survey results are potentially a rich source of information. However, a good design would allow for a mapping from that

²See Agrawal et al. (2021) for a review of this literature.

feedback into a signal of whether the opportunity is likely to be positive or not. In other words it maps into a binary signal (with positive news above a threshold and negative news below it). If there is a positive signal, you then expect those incremental profits to be high enough to proceed with implementation. If there is a negative signal, you do not expect any incremental profits. Your customers can provide a signal as to those increments.

Seen in this light what Christensen (1997) is saying is that the signal you receive from your own customers is biased. If a technology happens to be disruptive then, in Christensen's conception, that means that it involves features that your own customers may not want because they are coupled with a reduction in quality of features they value: "they typically present a different package of performance attributes—ones that, at least at the outset, are not valued by existing customers." (Bower and Christensen (1995), p.3) Christensen argues that those customers, precisely because they are satisfied with a firm's current products, cannot evaluate the trade-offs properly. In other words, when you ask your own customers about a technology that happens to be disruptive, it is difficult to get a clear signal that the underlying state is W . Because even if W is the true state, your own customers *may* not realise that it is true. This implies that a negative signal may or may not indicate that the true state is N . By contrast, a positive signal from your own customers would be a clear indication to pursue the technology. That is, if you ask your existing customers to make a trade-off regarding features and they are happy with the trade-off, it is a clear signal that there are incremental profits to be had.

There are other examples that incumbent firms conducted experiments that were biased towards negative news. When Blockbuster, in 2004, facing pressure from competition with Netflix, was considering abandoning late fees for video returns, it decided to conduct a trial. Nick Shepherd, the Head of International Operations, wanted late fees reduced but John Antioco, the CEO, wasn't so sure. This led to a trial that was of a particular form.

Although Shepherd balked at killing late fees altogether, Antioco urged him to test such a program in a few U.S. markets. Shepherd chose to test in Chattanooga, Tennessee, a city where nearly every test initiative went wrong. He sent a second-string team of executives to launch the no-late-fees promotion—to increase its chances of failure. (Keating (2012), p.117)

As it turned out, the experiment yielded a positive result that gave Blockbuster the confidence to end late fees.

In the terminology of this paper, Christensen's argument translates into a formal argument that an experiment based on surveying your own customers is likely to be biased in favour of generating negative news. Christensen is noting that bias and, consequently, cautioning firms to that bias as well as providing a normative recommendation to avoid those

types of experiments. The Blockbuster example though suggests that such experiments may not be flawed but reflect the decision facing an incumbent firm. As we will see below, while Christensen’s recommendation regarding the interpretation of such biased signals is sound, it is not clear that this translates into his normative recommendation that such signals should be avoided.

Christensen did not simply recommend that firms pay less attention to their own customers as experimental subjects but, instead, pursue learning strategies that shed that legacy and are ones that would be more readily available to independent start-ups. In Bower and Christensen (1995) (and more comprehensively Christensen and Raynor (2013)), they advocate that established firms set-up an independent entity located outside the firm’s main operations to explore and evaluate new technological opportunities.

It follows, then, that senior executives must first be able to spot the technologies that seem to fall into this category. Next, to commercialize and develop the new technologies, managers must protect them from the processes and incentives that are geared to serving established customers. And the *only* way to protect them is to create organizations that are completely independent from the mainstream business. (Bower and Christensen (1995), p.3, emphasis added)

Such independent units, often referred to as skunkworks, are often located in a different place to firm operations (such as IBM’s famous independent unit to explore personal computers in the 1970s and 1980s; see Bresnahan et al. (2011)). However, it can be argued that such an approach is itself a biased experiment; albeit with a different bias. The mission of such entities is avoid using a firm’s own customers as subjects. This, however, is likely to be an experiment biased towards positive news, where a positive signal may be received even if the true state is N . This is because, by ignoring the impact on a firm’s own demand, an independent unit will not receive a signal of whether there would be a negative impact on a firm’s existing product lines and the magnitude of such effects. The bias is, of course, the point of the exercise. But again this raises the question of whether that bias is appropriate for decision-making when a firm is evaluating potentially disruptive innovations.

Noting the potential bias of experiments and information sources is a theme in the innovation literature. For instance, Henderson and Clark (1990) argue that over time, engineering teams in organisations tend to have an easier time evaluating innovations to components than to architecture (i.e., how the components fit together).

Over time, engineers acquire a store of knowledge about solutions to the specific kinds of problems that have arisen in previous projects. When confronted with such a problem, the engineer does not reexamine all possible alternatives but,

rather, focuses first on those that he or she has found to be helpful in solving previous problems. (Henderson and Clark (1990), p.16)

They go on:

But information that might warn the organization that a particular innovation is architectural may be screened out by the information filters and communication channels that embody old architectural knowledge. Since radical innovation changes the core design concepts of the product, it is immediately obvious that knowledge about how the old components interact with each other is obsolete. The introduction of new linkages, however, is much harder to spot. Since the core concepts of the design remain untouched, the organization may mistakenly believe that it understands the new technology. (Henderson and Clark (1990), p.17)

The particular bias that Henderson and Clark identify is that associated with asking internal, specialised teams to evaluate the potential of new technologies. Such teams are likely to provide a biased signal of a technology's potential if that primary changes are architectural. In other words, the signal is biased towards negative news and may sometimes dismiss profitable technological opportunities. Unlike Christensen, while they acknowledge the bias, Henderson and Clark also note that established firms are likely to have challenges adopting architectural innovations precisely because it requires breaking out of existing communication and information filtering built into organisation structure. Instead, they argue that if an organisation is to adopt architectural innovations, it needs to be designed differently so that the costs of adopting such innovations are lowered.

Both Christensen and Henderson and Clark argue that the particular biases arise when evaluating particular types of technological opportunities. For instance, asking your own customers does not necessarily give you a biased signal if a technology is sustaining and internal evaluation processes can work well for component innovations. For disruptive and architectural technological opportunities, the problems arise and, in both cases, it is important that firms recognise biases in certain experiments that can arise there. Being able to act on those biases, of course, therefore, requires recognising distinct technological types as well; something that itself can be challenging (Gans (2016)).

Another form of bias in experiments come from the evaluation of technological opportunities potentially *ahead of their time*. For instance, when Blockbuster launched its on-demand video service, it was at a time when broadband was not widely available in the US. Arguably, this meant that the experiment was biased towards negative news relative to an experiment conducted later in the broadband diffusion cycle. That is, it may provide a clear signal that

on-demand video was valuable to consumers but was biased towards providing an ambiguous negative signal even if that technology would, as we now know, be successful. Thus, by analysing a firm's choice of experiment, we can also evaluate what types of firms are more likely to experiment at the beginning or later stages of underlying technology diffusion cycles.

Other forms of experimentation, often associated with entrepreneurial firms as well as established firms can also be characterised in terms of their potential bias. Moreover, unlike the ones currently described, these biases may be inherent rather than depending on the nature of the technological opportunity itself.

Take, for instance, the release of products for the evaluation of early adopters. One approach, popularised by Ries (2011) is to release *minimum viable products* to early adopters. Here is how he describes that:

[W]e build a minimum viable product, an early product that is terrible, full of bugs and crash-your-computer-yes-really stability problems. Then we ship it to customers way before it's ready. And we charge money for it. (Ries (2011), p.4)

A minimum viable product is a product that only just works and is often built in a way that is not scalable (say, relying on human labor when automation might eventually be possible). What do innovators learn from releasing these to potential customers? If, given all of its flaws, the customers accept and like that product, then that is a strong signal that the opportunity is promising. If they do not like that product, there is very little signal as the opportunity may be promising or not. The very fact that the release is flawed biases the signal. In this case, it is biased towards negative outcomes.

An alternative approach is to build a very high quality product designed to satisfy lead-users and establish a beach-head in a target market. This is advocated by Moore (2014). He argues that innovators should evaluate potential customer segments that may be amenable to products based on the opportunity and then put your best foot forward in trying to win that segment. It can be seen that this is somewhat opposed to the minimum viable product approach. This is because those customers need to be highly satisfied so as to become a base upon which to scale from. In this case, you want to invest in a product that creates that satisfaction. In other words, you want that product to be an experiment that is *w*-biased so that if it fails, you know for sure that, for instance, looking for alternative customer segments that are 'close' to the one targeted is not worthwhile. There are entrepreneurial companies that use this type of approach (for instance, Segway) but it is also pursued by established companies (most notably, Apple).

Once again, for both minimum viable product and lead-user approaches, the question

is whether the particular bias inherent in each is appropriate for the implementation decision innovative firms are evaluating. We will address this question when we examine the experimental choices of entrants below.

1.3 Preview of the main result

The formal model analyses a firm's choice of experiment and, in particular, whether it is biased towards favourable or unfavourable news regarding a technology whose value is uncertain. The key assumption is that the experiment does not yield a perfect signal regarding the technology's value. An imperfect signal gives credence to the choice of experiment type and also to how different firm characteristics might drive that choice.

Of particular note is the firm's predisposition regarding whether to exploit the technology in the absence of a technology. It is shown that a firm is more likely to choose an experiment that is biased in the same direction as their own pre-disposition to exploit the technological opportunity. If a firm would, absent an experiment, choose to exploit an opportunity, they choose an experiment biased towards positive news. This is because that experiment will be of most use to them if it can convince them to change their pre-disposition and not exploit the opportunity. This occurs if a negative signal very clearly indicates that the firm choose contrary to their disposition.

Given this, we are able to determine that it is, in fact, optimal for established firms facing disruptive/architectural opportunities to choose to evaluate those opportunities with an experiment biased towards unfavourable news (such as a survey of their own customers). This is because such opportunities invariably mean that current productive assets and profit sources will be cannabilised by the new technological opportunity (as pointed out by Arrow (1962)). This means that incumbent firms will be pre-disposed to not exploit disruptive/architectural opportunities because of low incremental profits or high organisational costs of so doing. Thus, they optimally opt for an experiment biased against a favourable signal regarding the technology. By contrast, we find that entrants, in considering whether to exploit a disruptive/architectural technological opportunity are more likely to adopt *w*-biased experiments.

This main result supports the positive observations of the past literature on disruption that successful established firms were likely to adopt experiments that were biased towards unfavourable assessments of the new technology. However, it also demonstrates that their choices in that regard were optimal. Thus, the model here challenges the normative proposition that such firms refrain from such experiments in evaluating technological opportunities that are potentially disruptive.

1.4 Plan for the paper

The paper proceeds as followed. In Section 2, the model is set-up and the initial result regarding own-biased experimentation is proven. Importantly, in that model experiments involve the same cost to the firm and thus, there is no cost advantage to the firm from choosing a biased experiment over unbiased ones that are also in the feasible set. Thus, we are not constraining firms to choose between biased experiments but allowing them to choose the degree of bias. Section 3 then turns to consider how the type of firm (incumbent versus entrant) and type of technology (disruptive/architectural versus sustaining/modular) impact on experimental choice. In addition, comparative results regarding the different probabilities of these firms in exploiting technological opportunities are established. Section 4 then turns to consider direct competition and cooperation amongst incumbent and entrant and how these impact on choices. It is demonstrated that the main result is reinforced in such strategic environments. Section 5 then examines the robustness of the main result to sequencing of actions and the observability of experimental outcomes. A final section concludes.

2 Model Setup

A new technological opportunity appears. A firm needs to decide whether to exploit it or not. The state of the world is such that the new technological opportunity either “works” (W) or “not” (N). If the technology works and the firm exploits it, their payoff is v_W while exploiting a technology that does not work earns v_N . The firm can also choose not to exploit the technology in which case it earns v_S (where S stands for status quo). It is assumed that $v_W > v_S > v_N$. The firm’s prior probability that the technological opportunity will “work” is p . Absent any other information, the firm will exploit the technology if:

$$pv_W + (1 - p)v_N \geq v_S$$

Let $\hat{p}$ be p such that $pv_W + (1 - p)v_N = v_S$ or $\hat{p} = \frac{v_S - v_N}{v_W - v_N}$.

2.1 Experiment Set

Prior to choosing whether to exploit the new technology, the firm can explore its feasibility through experimentation. Let $\mathbf{E}$ be the set of experiments. There are no costs incurred by

	w	n
W	λ_w	$1 - \lambda_w$
N	$1 - \lambda_n$	λ_n

Table 1: **Experiment Maps from Signal to State**

selecting an experiment from that set.³ Experiments yield two possible signals, w and n , whose likelihood is governed by the true underlying state.

It is assumed that $\Pr[w|W] = \lambda_w$ and $\Pr[n|N] = \lambda_n$ (see Table 1). The past literature tends to think of experimentation as generating a signal to the technology's quality (Kerr et al. (2014)) that is unbiased in the sense that $p = \Pr[w] = p \Pr[w|W] + (1 - p) \Pr[w|N]$. This implies $\frac{1-\lambda_n}{1-\lambda_w} = \frac{p}{1-p}$.

Here, instead, we focus on situations where the firm can choose from a set of experiments with differing degrees of bias (including no bias at all). It is assumed that λ_w and λ_n are constrained in magnitude so that $\lambda_w + \lambda_n \leq \Lambda$ where $\Lambda < 2$. Experiments are informative so that the posterior beliefs satisfy $\Pr[W|w] \geq p \geq \Pr[W|n]$ which, in turn, implies that $\lambda_w + \lambda_n > 1$. Thus, a typical element of $\mathbf{E}$ is a pair $\{\lambda_w, \lambda_n\} \in (0, 1]^2$ such that $\lambda_w + \lambda_n \in (1, \Lambda]$.

This experimental choice set makes explicit a trade-off for the firm in choosing between the probability of a false positive ($1 - \lambda_n$) and the probability of a false negative ($1 - \lambda_w$). Or, stated another way, if the null hypothesis being tested was that the technology worked, then $1 - \lambda_w$ is the probability of a Type I error while $1 - \lambda_n$ is the probability of a Type II error. An unbiased experiment balances these probabilities relative to the prior.⁴

While there are many possible experiments, two, in particular, stand out for closer attention. These are a w -biased experiment, $\{\lambda_w, \lambda_n\} = \{1, \lambda\}$ and an n -biased experiment, $\{\lambda_w, \lambda_n\} = \{\lambda, 1\}$, where $\lambda \equiv \Lambda - 1$. An w -biased experiment yields a signal of w with probability 1 if the underlying state is W but either signal is possible if the true underlying state is N . Thus, the news is biased towards receiving ' w .' If a signal of n is generated, it is known with certainty that the true state is not W and must be N ; i.e., $\Pr[N|n] = 1$. This experiment can, therefore, generate a possible signal, n , where, regardless of their prior, any firm optimally chooses not to exploit the opportunity. This experiment eliminates the possibility of a false negative. For an n -biased experiment, the reverse is true and a signal of w implies that the true state is W ; that is, $\Pr[W|w] = 1$ eliminating the possibility of a

³We only allow one experiment to be selected, so implicitly there is an infinite cost of selecting more than one experiment.

⁴In other terminology, λ_w is the sensitivity of an experiment while λ_n is its specificity.

false positive.⁵ With this experiment, a signal of w implies the firm should always exploit the opportunity.⁶

2.2 Experiment Choice

We are now in a position to state our first result characterising when the firm chooses a w -biased experiment rather than an n -biased experiment.

Proposition 1 *If $p > \hat{p}$, then the firm optimally chooses the w -biased experiment. If $p < \hat{p}$, the firm optimally chooses the n -biased experiment.*

All proofs are in the appendix. This proposition says that if, without any experimentation, the firm would choose to exploit the new technology, then they will choose a w -biased over an n -biased experiment. The reason is that they are relatively optimistic about the technology's prospects and it is only if they receive a clear signal otherwise that the firm will change its mind and choose not to exploit it. This reflects the result of Che and Mierendorff (2019) that information seekers will choose experiments that reflect their own bias.

While Proposition 1 considers the choice between the w -biased and n -biased experiment, it is important to note that, for all but a degenerate case, these are the only two experiments that will be optimally chosen. This can be shown graphically. In Figure 1, the red line depicts the frontier of the set $\mathbf{E}$. The blue lines are indifference curves whereby on a particular curve the same expected outcome is achieved if the decision rule whereby the technology is implemented if there is a signal of w and is not implemented if there is a signal of n . That is, each indifference curve involves:

$$V = p\lambda_w v_W + (1-p)(1-\lambda_n)v_N + ((1-p)\lambda_n + p(1-\lambda_w))v_S$$

These indifference curves have a slope of $-\frac{p}{1-p} \frac{v_W - v_S}{v_S - v_N}$ which is greater than -1 if $p > \hat{p}$ and less than -1 otherwise. Thus, the optimal experimental choice is typically at a kink of the

⁵Awaya and Krishna (2021) also considers learning about a technology when such learning is imperfect. However, their analysis examines the choices of whether to experiment or not and the consequences of that. In so doing, they assume that experiments involve $\Pr[w|W] = \Pr[n|N]$. Here, however, $\Pr[w|W] = \lambda_w$ while $\Pr[n|N] = \lambda_n$ and we explicitly consider a choice of experiments where potentially $\lambda_w \neq \lambda_n$.

⁶The naming of these experiment types is challenging. I had considered calling them a Sure-No and Sure-Go experiment respectively with a nod to Savage's Sure Thing Principle. I also tried to find more interesting names that were, say, based on some ancient parable but came up short. Instead, I relied on the terminology used by Che and Mierendorff (2019) and the fact that the prior p would not equal the expected posterior for the two extreme experiments of interest. The term 'bias,' reflects the notion that news sources are slanted towards one type of news, w or n as the case may be.

Figure 1: **Choice of Experiment**

frontier of the experiment set; that is, a w -biased experiment at A if $p > \hat{p}$ and an n -biased experiment at B if $p < \hat{p}$.⁷

3 Potential Disruption and Experiment Choice

In the previous section, it was demonstrated that a firm's choice of experiment depends on their prior and its relationship to a threshold determined by the relative benefits of success or failure above a status quo outcome. Our interest in this paper is how a firm's positive choice regarding experimental bias relates to the normative claims made regarding that choice in the face of disruptive technological opportunities and whether that firm is an incumbent or an entrant. Thus, our next step is to characterise incumbent/entrant preferences regarding experimental bias. Following that, how these relate to the disruptive nature of the technological opportunity are analysed.

3.1 Incumbents versus entrants

We begin by considering the difference between incumbents and entrants and how this impacts on their experiment choices. Suppose that an entrant (E) can earn (gross) profits of π if they successfully exploit the new technological opportunity and 0 otherwise. The cost of entry is c ($\leq \pi$). Given this, $v_W^E = \pi - c$, $v_N^E = -c$ and $v_S^E = 0$. By contrast, an incumbent

⁷The experimental set puts a cap on $\lambda_w + \lambda_n$ and so the boundary of that set is a straight line. It may be that the set of experiments is such that there are diminishing opportunities to receive a clear signal in one direction or another and so the boundary of $\mathbf{E}$ is strongly concave to the origin. In this situation, the degree of bias chosen will be a function of the distance between p and $\hat{p}$ and the same qualitative conclusions that will be drawn throughout this paper will continue to hold.

(I) can earn (gross) profits of $\Pi + \Delta$ if they successfully exploit the new technological opportunity where Π are their pre-innovation profits. It is also possible that trying to exploit the new technological opportunity is disruptive internally to the incumbent and, thus, costs the C regardless of whether that technology works or not. Finally, we assume that if it either does not exploit or exploits an opportunity that does not work the incumbent earns Π . Given this, $v_W^I = \Pi + \Delta - C$, $v_N^I = \Pi - C$ and $v_S^I = \Pi$.⁸

Given this, we can now prove a result showing a relationship between the type of experiments favoured by entrants and incumbents and their differences.

Proposition 2 *The range of p whereby an incumbent chooses a w -biased experiment (i.e., $(\hat{p}, 1]$) becomes relatively larger (smaller) than the range for the entrant as $c - C$ and/or $\Delta - \pi$ increases (decreases).*

This can be seen by substituting in the relevant v_W^i , v_N^i and v_S^i ($i \in \{I, E\}$) into the equation for $\hat{p}$ to give:

$$\hat{p}_E = \frac{c}{\pi} \qquad \hat{p}_I = \frac{C}{\Delta} \qquad (1)$$

Thus, if $C < c$ and $\Delta > \pi$, then $\hat{p}_I < \hat{p}_E$. If, however, either $C < c$ and $\Delta < \pi$ or $C > c$ and $\Delta > \pi$, then the difference between incumbent and entrant experimental choices depends on the ratio of supply and demand-side factors. Note that this result only characterises a difference in tendencies to adopt distinct experiments. However, for p very high or low, both firms opt for the same experimental choice.⁹

3.2 The type of technological opportunity

The parameters underlying the results in Proposition 2 arise in the literature on disruption. In particular, for a technological opportunity where $\frac{\pi}{c} > \frac{\Delta}{C}$, it is generally thought that opportunity favours entrant rather than incumbent adoption. Consider a situation where $\pi > \Delta$ which can arise even when $\Pi \gg \pi$. What is going on here is that the technological opportunity leads to products that largely cannibalise an incumbent's own pre-innovation profits. This is a well-known constraint on incumbent innovation arising from the *replacement*

⁸A recent paper by Callander and Matouschek (2021) provides greater underpinnings consistent with these assumptions in the context of disruption but their focus is on the conditions for welfare-enhancing innovation-based entry.

⁹Note that this assumes that the incumbent and entrant have a choice from the same set of experiments, **E**. However, it may be that one firm, say, the incumbent has a more informative set than the entrants in the form of a higher Λ . Note that in this case, the result still holds and $\hat{p}_I$ and $\hat{p}_E$ drive the choices of the bias of the experiment. Obviously, if $\Lambda \geq 2$, for some firm, they will choose a fully informative experiment with no bias by definition.

effect (Arrow (1962)). A form of this also appears in Christensen (1997) where he argues that for disruptive innovations (that involve, lower initial customer satisfaction for existing customers relative to marginalised customers) provide an opportunity for entrants to profit from niche areas in a way that incumbents cannot achieve. The reverse is true for sustaining innovations that do offer increased satisfaction for existing customers.

While the distinction between a disruptive and sustaining innovation involves a consideration of demand-side factors, we can also provide a supply-side characterisation of the technological opportunity. Following Henderson and Clark (1990), the costs of exploiting technological opportunities for an incumbent vary depending upon whether a technological opportunity involves modular versus architectural innovation. A modular innovation is, by definition, easy for an incumbent to integrate into their existing technological and organisational structure. Thus, for these innovations it is likely that $C < c$. By contrast, an architectural innovation involves creating a new organisational structure to properly exploit the technological opportunity. This is difficult (and maybe impossible) for incumbent firms to do while it is relatively easy for new entrants. Thus, an architectural innovation involves $C > c$.

Given this, we will say that a technological opportunity is disruptive/architectural if $\frac{\pi}{c} > \frac{\Delta}{C}$ and is sustaining/modular if $\frac{\pi}{c} < \frac{\Delta}{C}$. Seen in this way, Proposition 2 provides a new perspective on incumbent and entrant experimental choices and how they relate to the nature of technological opportunity. An n -biased experiment is more likely to be chosen by incumbents than entrants if the nature of the technological opportunity is more associated with disruptive/architectural innovation than sustaining/modular innovation (see Figure 2). This is because, in this situation, the entrant, in the absence of an experiment, is more predisposed to exploit the innovation than the incumbent would be. Such experiments are designed to convince a pre-disposed innovator to refrain from innovating. An incumbent, for a disruptive/architectural technological opportunity, is more likely to choose an n -biased experiment. This is because they are potentially are a reluctant innovator and are looking to experiments to convince them that the technological opportunity should be exploited.

What Proposition 2 shows is that the normative claim of Christensen (1997) that incumbents facing disruptive technological opportunities is incorrect. Instead, if, as argued earlier, listening to your own customers is n -biased then that is likely to be a more informative experiment for established firms than a w -biased experiment like an independent skunkworks. This is because disruptive innovations, being disruptive, are less likely to be profitable for incumbents than entrants and so incumbents will be pre-disposed not to pursue them. That pre-disposition means that an n -biased experiment is more valuable as it is more likely to cause them to change that decision. Specifically, if $p < \frac{C}{\Delta}$, the expected payoff to the incum-

Figure 2: **Experimental Choice and Technology Opportunity**

bent of running an n -biased experiment is $\Pi + p\lambda(\Delta - C)$ and is only $\Pi + p\Delta - (1 - (1 - p)\lambda)C$ for a w -biased experiment; a difference of $(1 - \lambda)(C - p\Delta) > 0$.

Given the relationship derived in Proposition 2, what does this imply about the relative likelihoods for the incumbent and entrant in terms of being observed to exploit the technological opportunity?

Proposition 3 *If the technology opportunity is disruptive/architectural, the exploitation probability is (weakly) lower for the incumbent than the entrant. If the technology opportunity is sustaining/modular, the exploitation probability is (weakly) higher for the incumbent than the entrant*

Proposition 3 says that, under conditions where the technological opportunity is disruptive/architectural, the entrant is *more likely* to exploit the opportunity than the incumbent whereas the reverse is true if the technological opportunity is sustaining/modular.

This proposition shows that, even when given an opportunity to learn about a new technological opportunity *and* when such learning may lead an incumbent to exploit the technology, the types of experiments chosen by incumbent and entrant end up, in equilibrium, continuing to reinforce their pre-dispositions. Intuitively, in the absence of an experiment, in the range of $p \in (\hat{p}_I, \hat{p}_E)$ where the incumbent and entrant choose different experiment types, without experimentation the entrant will exploit the opportunity and the incumbent will not. Outside of this range, without an experiment, they each either exploit or do not exploit the opportunity. A firm chooses an experiment with a bias designed to increase the chances that they change their default exploitation incentive. For the entrant, this is towards non-exploitation. For the incumbent, this is towards exploitation. If each choose

the same experiment $\{\lambda_w, \lambda_n\}$ the probability of exploitation for the entrant (incumbent) would fall (rise) to $p\lambda_w + (1-p)(1-\lambda_n)$; the probability that a firm receives a w signal. This probability is rising in λ_w and falling in λ_n . Thus, choosing an n -biased experiment over a w -biased one will lower the probability you exploit and vice versa. An incumbent chooses an experiment that causes them to switch only if the evidence is overwhelming. This means that they will not risk an error of exploiting when they should have not exploited; i.e., they minimise $\Pr[N|w] = 1 - \lambda_w$. This involves favouring a higher λ_w .

In terms of observations, this result suggests (a) that incumbents will be less likely to exploit new technological opportunities when the conditions are what is commonly associated with disruptive/architectural innovation ($\frac{\pi}{c} > \frac{\Delta}{C}$) and (b) conditional on exploitation, incumbents are more likely to exploit successfully in that situation. The latter result arises because the probability of successful exploitation for an entrant with a w -biased experiment is $\Pr[W|w] = \frac{p}{p+(1-\lambda)(1-p)} < 1$ and for an incumbent with an n -biased experiment is $\Pr[W|w] = 1$. Thus, the observed failure rate will be *higher* for entrants than incumbents conditional on exploiting disruptive/architectural opportunities.

Interestingly, while the incumbent experimental choices and how they match with disruptive/architectural tendencies confirms the observations in the literature (namely, that they tend to be n -biased), the results for entrants might be regarded as counter-intuitive. Earlier we argued that the minimum viable product approach of Ries (2011) was n -biased while the lead user approach was w -biased. Our analysis here suggests that the minimum viable product approach will be optimally pursued by start-ups in a less disruptive environment where their own costs and potential profits are relatively low. While this accords with the ‘lean’ approach that recommends cheap and frequent experimentation, that approach is often associated with start-ups taking advantage of disruptive opportunities (Gans et al. (2019), Ching et al. (2019)). However, in those situations, the optimal experimental strategy for start-ups is likely to be pursuing a w -biased experiment such as focusing on establishing a beachhead with lead users. This suggests that care must be taken in mapping experimental approaches to underlying technological and market parameters.¹⁰

Proposition 3 can shed some light on past observed incumbent responses to potentially disruptive technological opportunities. As Gans (2016) documents, it was often the case that incumbents often ran experiments on new technological opportunities before entrants. For instance, Blockbuster launched a video on demand product in the early 2000s well before Netflix was even founded. This product failed due to the lack of widespread broadband

¹⁰Gans (2021) provides a resolution of this apparent puzzle with respect to entrants showing that when such entrants have to persuade others to devote resources to the venture, the optimal persuasive experiment is n -biased regardless of how optimistic an entrepreneur is regarding the value of the technological opportunity.

in the US. Given that we know now that video on demand was a workable technology with sufficiently broadband, launching it without that capability to see if there was low bandwidth interest was an n -biased experiment yielding a signal that it was not worth pursuing even if the technology was ultimately workable. For an established firm with a significant infrastructure that would be replaced by this new technology, they would be pre-disposed not to adopt it, so running an n -biased experiment was optimal. However, it also meant that it was less likely to pursue that opportunity than an entrant who ran a w -biased experiment. Arguably, Netflix, by establishing a beachhead with its DVD by mail service, was doing just that.

3.3 Confirmation bias

Thusfar, the analysis here assumes that the incumbent and entrant have a common prior on whether the technological opportunity will work. The disruption literature often posits potential decision-making biases that impact on subjective probabilities. When there is fundamental uncertainty, these are more likely to manifest themselves.

To explore this issue, it is assumed here that, when, in the absence of an experiment, a decision-maker is pre-disposed to exploit the technology, their subjective belief, p_E or p_I as the case may be, is greater than the objective prior, p . Thus, an entrant who is pre-disposed to exploit the technology is partially pre-disposed because $p > \hat{p}_E \implies p_E > p$. The reverse is true if the decision-maker is reluctant to exploit the technology, in which case, say, for an entrant $p < \hat{p}_E \implies p_E < p$. Confirmation bias of this kind means that a given decision-maker has an increased incentive to choose an experiment that is designed to run counter to that bias as that is even more informative. For instance, a reluctant incumbent with $p_I < p$, will always choose a n -biased experiment as $p_I < \hat{p}_I$.

Given this, note that the range of p such that the entrant and the incumbent choose different experiment types is broader than the model considered without differing priors. In particular, when $\hat{p}_I \geq p_E > p > p_I \geq \hat{p}_E$, then entrant chooses the w -biased experiment and the incumbent chooses the n -biased experiment. In this case, the true or observed exploitation is still based on p and, thus, the incumbent is less likely to exploit than the entrant. In other words, while the range of p whereby the incumbent and entrant have different exploitation rates is wider, when they are different, the rankings of those rates remains the same. The same is true in situations where $\hat{p}_I \geq p_E > p > p_I \geq \hat{p}_E$.

4 Competition

The analysis, thusfar, has analysed incumbent and entrant experiment independently from one another. The purpose was to identify how their different starting positions impacted on their optimal experimental choice. This implicitly assumes, however, that experimental outcomes themselves do not mould or are moulded by the anticipation of competition. Certainly, the informal statements of those advocating for incumbents to perform w -biased rather than n -biased experiments when facing potential disruption, do often state that it is the existential risk that incumbents face from competition by would-be disruptors that should cause them to undertake investigations to better uncover those threats and react to them.

Here we add a competitive element to our model. This allows us to analyse, say, the impact of both firms obtaining a signal suggesting that the technology is workable with the implication that there will be more competitive outcomes ex post than if only one firm received such a signal. There are also nuances that arise as the exploitation actions of the firms themselves provide a signal of the outcomes of experiments and also that ex post competition might be avoided by the incumbent and entrant reaching a cooperative deal (e.g., licensing or acquisition). In this section, we explore these outcomes.

4.1 Competition without Signaling

We begin by considering ex post competition in the absence of the exploitation actions providing a signal to competitors. Let σ_i^I and σ_i^E denote the experimental choice of the incumbent (I) and entrant (E) respectively between $i \in \{w, n\}$ -biased. Let Π and Δ be I 's profits and increment to profits if they face no competition from E and Π_c and Δ_c be their profits if there is competition. For E , let π be their profits if the incumbent does not innovate and π_c be the profits if the incumbent innovates. It is assumed that $\Pi + \Delta > \Pi_c + \Delta_c + \pi_c$. We assume that the entrant, if they enter but the technology does not work, immediately exits. In addition, we assume that I and E make their exploitation choices on the basis of their own experimental outcomes and are not in a position to draw inferences from their rival's exploitation decisions.

There are four potential exploitation outcomes depending on the signal received:

1. (Both receive n) Entrant does not exploit and incumbent is able to earn their ex ante monopoly profits of Π .
2. (Incumbent receives w while entrant receives n) In this case, the entrant does not

exploit while the incumbent earns an expected payoff of

$$P(W|\{w, \sigma_i^I\}, \{n, \sigma_i^E\})(\Pi + \Delta - C) + P(N|\{w, \sigma_i^I\}, \{n, \sigma_i^E\})(\Pi - C)$$

3. (Entrant receives w while incumbent receives n) In this case, the entrant exploits and receives an expected payoff of

$$P(W|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})(\pi - c) + P(N|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})(-c)$$

while the incumbent receives an expected payoff of

$$P(W|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})\Pi_c + P(N|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})\Pi$$

4. (Both receive w) Both exploit with the entrant's expected payoff

$$P(W|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(\pi_c - c) + P(N|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(-c)$$

and the incumbent's expected payoff of

$$P(W|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(\Pi_c + \Delta_c - C) + P(N|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(\Pi - C)$$

These outcomes can be quite varied as the incumbent and entrant may be conducting experiments with different biases even though they are trying to learn about the same underlying state. That said, if, say, the incumbent conducts an n -biased experiment and receives a signal of w , then they know that E also received that signal even if E conducted an n -biased experiment and, even if E did not, then the underlying state that determines their competitive threat will be W . Thus, determining which experiments each firm might choose could be complex. Nonetheless, we will continue to assume that each has a common prior, p .¹¹ The probabilities that each of these signal combinations are generated (contingent upon the experiments being chosen) are listed in Table 2.

Despite these potential complexities, the following proposition shows that the equilibrium choices are fundamentally driven by the factors as in Proposition 1. That is, for this game, we can write:

$$v_W^I = \begin{cases} \Pi_c + \Delta_c - C & \text{if } \sigma_w^E \\ \lambda(\Pi_c + \Delta_c) + (1 - \lambda)(\Pi + \Delta) - C & \text{if } \sigma_n^E \end{cases} \quad v_W^E = \begin{cases} \pi_c - c & \text{if } \sigma_w^I \\ \lambda\pi_c + (1 - \lambda)\pi - c & \text{if } \sigma_n^I \end{cases}$$

¹¹Because we are not a masochist.

	σ_w^I, σ_w^E	σ_w^I, σ_n^E	σ_n^I, σ_w^E	σ_n^I, σ_n^E
$\{\sigma_i^I = w, \sigma_i^E = w\}$	$p + (1-p)(1-\lambda)^2$	$p\lambda$	$p\lambda$	$p\lambda^2$
$\{\sigma_i^I = w, \sigma_i^E = n\}$	$(1-p)\lambda(1-\lambda)$	$1-\lambda$	0	$p\lambda(1-\lambda)$
$\{\sigma_i^I = n, \sigma_i^E = w\}$	$(1-p)\lambda(1-\lambda)$	0	$1-\lambda$	$p\lambda(1-\lambda)$
$\{\sigma_i^I = n, \sigma_i^E = n\}$	$(1-p)\lambda^2$	$(1-p)\lambda$	$(1-p)\lambda$	$1-p+p(1-\lambda)$

Table 2: **Signal Probabilities**

$$v_N^I = \begin{cases} \Pi_c - C & \text{if } \sigma_w^E \\ \lambda\Pi_c + (1-\lambda)\Pi - C & \text{if } \sigma_n^E \end{cases} \quad v_N^E = -c$$

$$v_S^I = \begin{cases} \Pi_c & \text{if } \sigma_w^E \\ \lambda\Pi_c + (1-\lambda)\Pi & \text{if } \sigma_n^E \end{cases} \quad v_S^E = 0$$

Thus, it can be seen that the threshold values $\hat{p}_E$ and $\hat{p}_I$ will, for the most part, depend on their rival's experimental choice. In particular, both firms have higher expected payoffs in their rival is pursuing an N rather than a w -biased strategy. That is because, in that case, there is a chance that they will be the only firm who exploits the opportunity.

Given this, we can now characterise the pure strategy Nash equilibrium of the game where the incumbent and entrant simultaneously choose their experiments.

Proposition 4 *Suppose that I and E simultaneously choose σ_i^I and σ_i^E . Then, the pure strategy Nash equilibrium choices are: (i) $\{\sigma_w^I, \sigma_w^E\}$ if $p > \max\{\frac{C}{\Delta_c}, \frac{c}{\pi_c}\}$; (ii) $\{\sigma_w^I, \sigma_n^E\}$ if $p \in [\frac{C}{(1-\lambda)\Delta + \lambda\Delta_c}, \frac{c}{\pi_c}]$; (iii) $\{\sigma_n^I, \sigma_w^E\}$ if $p \in [\frac{c}{\lambda\pi_c + (1-\lambda)\pi}, \frac{C}{\Delta_c}]$; and (iv) $\{\sigma_n^I, \sigma_n^E\}$ if $p \leq \min\{\frac{C}{(1-\lambda)\Delta + \lambda\Delta_c}, \frac{c}{\lambda\pi_c + (1-\lambda)\pi}\}$.*

The proof (in the appendix) shows that the level of p drives the equilibrium outcomes. The particular thresholds of p are driven by the same considerations as Proposition 1 in that they mark the level of p whereby a firm would choose, in the absence of an experiment, to exploit

if p is above that threshold and otherwise not exploit. As noted earlier, the difference here is that those thresholds are influenced by the experimental choices of the rival firm.¹²

Importantly, note that the non-experiment incentives to exploit decrease as a rival moves from an N to a w -biased experiment. This is because the experimental choices of each firm are strategic substitutes. For instance, $v_I(\sigma_w^I, \sigma_w^E) - v_I(\sigma_n^I, \sigma_w^E) \leq v_I(\sigma_w^I, \sigma_n^E) - v_I(\sigma_n^I, \sigma_n^E)$ implies that: $(1 - \lambda)(p\Delta_c - C) \geq (1 - \lambda)p((1 - \lambda)\Delta + \lambda\Delta_c - C)$ which always holds. The same qualitative relationship holds for E . Hence, for intermediate levels of p , Proposition 4 shows that the two firms choose experiments of different types. In particular, this reinforces the differing preferences over experiments in the stand-alone case whereby the incumbent firm is more (less) likely to choose a w -biased experiment than then entrant for a sustaining/incremental (disruptive/architectural) technological opportunity.

This provides more insight into the potential for disruption whereby an entrant chooses to exploit an opportunity while the incumbent does not. When $p \in [\frac{C}{(1-\lambda)\Delta + \lambda\Delta_c}, \frac{c}{\pi_c}]$, something that is more likely to arise as $c - C$ or $\pi_c - \Delta_c$ rises (i.e., the opportunity is disruptive/architectural), I chooses the n -biased experiment and, therefore, exploits the opportunity with probability $p\lambda$ while the entrant exploits the opportunity with probability $p\lambda + 1 - \lambda$. While the entrant is more likely to exploit than the incumbent, the probability that the entrant successfully exploits is p . Thus, compared with the stand-alone outcome in Proposition 3, the entrant has a higher probability of exploitation with the same probability of successful exploitation while the incumbent's exploitation and successful exploitation probabilities are unchanged. That is, in the presence of an incumbent simultaneously conducting a w -biased experiment, the entrant's expected payoff falls to $v_E(\sigma_n^I, \sigma_w^E) = p\lambda(\pi_c - c) + (1 - \lambda)(p\pi - c)$ from $v_E(0, \sigma_w^E) = p\pi - (p + (1 - p)(1 - \lambda))c$. From the incumbent's perspective, the presence of competition from the entrant reduces their expected payoff but only reinforces their incentives to undertake an n -biased experiment relative to the stand-alone case.

4.2 Ex post Cooperation

The above model assumes that the incumbent and entrant become competitors if both should successfully exploit the technological opportunity. However, as Gans and Stern (2000) note, there are virtues to the incumbent and entrant cooperating rather than competing. Moreover,

¹²Let $x = 1$ indicate that a firm always exploits and $x = 0$ indicate that a firm never exploits. Thus, for example, if I exploits while E chooses an experiment, I 's expected payoff is $v_I(1, \sigma_i^E)$ while if E chooses never to exploit this payoff is $v_I(1, 0)$. Thus, we have, to derive the threshold for which p above that threshold the firm chooses a w -biased experiment, $v_I(1, \sigma_w^E) = p(\Pi_c + \Delta_c) + (1 - p)\Pi - C > v_I(0, \sigma_w^E) = p\Pi_c + (1 - p)\Pi \implies p > \frac{C}{\Delta_c}$ and $v_I(1, \sigma_n^E) = p\lambda(\Pi_c + \Delta_c) + (1 - p\lambda)\Pi + p(1 - \lambda)\Delta - C > v_I(0, \sigma_n^E) = p\lambda\Pi_c + (1 - p\lambda)\Pi \implies p > \frac{C}{\lambda\Delta_c + (1 - \lambda)\Delta}$. For E , $v_E(1, \sigma_w^I) = p\pi_c - c > 0 \implies p > \frac{c}{\pi_c}$ and $v_E(1, \sigma_n^I) = p(\lambda\pi_c + (1 - \lambda)\pi) - c > 0 \implies p > \frac{c}{\lambda\pi_c + (1 - \lambda)\pi}$. You can see that these threshold levels correspond to those in Proposition 4.

as Marx et al. (2014) document, it is often a feature of competition in disruptive technologies that entrants enter first and then switch to cooperation thereafter.

Here we consider that possibility by assuming that, if it chooses to enter, E can, if the technology is workable, come to a cooperate deal with I . In that deal, I pays E a lump sum about τ that neutralises competition between them and allows the incumbent to earn $\pi + \Delta$. It is assumed that this deal can occur either if the incumbent has not chosen to exploit the opportunity (in which case, C is never incurred) or even if it has.¹³ For E to be part of a cooperative deal, they must enter and incur c .¹⁴

The first step is to determine the outcome of bargaining over the price of any cooperative arrangement should the opportunity arise. There are two cases to consider. In the first, I exploits the technology alongside E . If they do not arrive at a deal, I earns $\Pi_c + \Delta_c$ and E earns π_c . If they arrive at a deal the joint profits become $\Pi + \Delta$. If we use the Nash bargaining solution with equal bargaining power, then $\tau_c = \frac{\Pi + \Delta - \Pi_c - \Delta_c + \pi_c}{2}$. In the second case, I does not exploit the technology. In this case, absent a deal, I earns Π_c and E earns π . This implicitly assumes that I cannot exploit at a later stage. Here, $\tau = \frac{\Pi + \Delta - \Pi_c + \pi}{2}$.

Given this the four outcomes become:

1. (Both receive n) Entrant does not exploit and incumbent is able to earn their ex ante monopoly profits of Π .
2. (Incumbent receives w while entrant receives n) In this case, the entrant does not exploit while the incumbent earns an expected payoff of

$$P(W|\{w, \sigma_i^I\}, \{n, \sigma_i^E\})(\Pi + \Delta - C) + P(N|\{w, \sigma_i^I\}, \{n, \sigma_i^E\})(\Pi - C)$$

3. (Entrant receives w while incumbent receives n) In this case, $\tau = \frac{\Pi + \Delta - \Pi_c + \pi}{2}$ and the entrant exploits and receives an expected payoff of

$$P(W|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})(\tau - c) + P(N|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})(-c)$$

while the incumbent receives an expected payoff of

$$P(W|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})(\Pi + \Delta - \tau) + P(N|\{n, \sigma_i^I\}, \{w, \sigma_i^E\})\Pi$$

4. (Both receive w) Both exploit and $\tau_c = \frac{\Pi + \Delta - \Pi_c - \Delta_c + \pi_c}{2}$ giving the entrant's expected

¹³Each of these situations might run afoul of anti-trust laws but we assume those laws either do not or are not applied.

¹⁴It is only by doing this that they can signal to the incumbent that the technology is workable and so be in a position to make a cooperative deal.

	σ_w^E	σ_n^E
σ_w^I	$v_I = p(\Delta - \tau_c)$ $+ \Pi - (1 - (1 - p)\lambda)C$ $v_E = p\tau_c - (1 - (1 - p)\lambda)c$	$v_I = p\lambda(\Delta - \tau_c)$ $+ \Pi + p(1 - \lambda)\Delta - (1 - (1 - p)\lambda)C$ $v_E = p\lambda\tau_c - p\lambda c$
σ_n^I	$v_I = p(\Delta - \lambda\tau_c - (1 - \lambda)\tau)$ $+ \Pi - p\lambda C$ $v_E = p(\lambda\tau_c + (1 - \lambda)\tau) - (1 - (1 - p)\lambda)c$	$v_I = p\lambda(\Delta - \lambda\tau_c - (1 - \lambda)\tau)$ $+ \Pi + p\lambda(1 - \lambda)\Delta - p\lambda C$ $v_E = p\lambda(\lambda\tau_c + (1 - \lambda)\tau) - p\lambda c$

Table 3: **Payoffs with Ex Post Cooperation**

payoff

$$P(W|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(\tau_c - c) + P(N|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(-c)$$

and the incumbent's expected payoff of

$$P(W|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(\Pi + \Delta - \tau_c - C) + P(N|\{w, \sigma_i^I\}, \{w, \sigma_i^E\})(\Pi - C)$$

Given this, we can prove the following.

Proposition 5 *Suppose that I and E simultaneously choose σ_i^I and σ_i^E . If ex post cooperation is possible, then, the pure strategy Nash equilibrium choices are: (i) $\{\sigma_w^I, \sigma_w^E\}$ if $p > \max\{\frac{C}{\tau - \tau_c}, \frac{c}{\tau_c}\}$; (ii) $\{\sigma_w^I, \sigma_n^E\}$ if $p \in [\frac{C}{(1-\lambda)\Delta + \lambda(\tau - \tau_c)}, \frac{c}{\tau_c}]$; (iii) $\{\sigma_n^I, \sigma_w^E\}$ if $p \in [\frac{c}{\lambda\tau_c + (1-\lambda)\tau}, \frac{C}{\tau - \tau_c}]$; and (iv) $\{\sigma_n^I, \sigma_n^E\}$ if $p \leq \min\{\frac{C}{(1-\lambda)\Delta + \lambda(\tau - \tau_c)}, \frac{c}{\lambda\tau_c + (1-\lambda)\tau}\}$.*

What this shows is that the cooperation case follows the contours of the competition case. The experimental choices of both firms remain strategic substitutes and moving from an n -biased to a w -biased experiment is driven by the incremental profits from generating a workable innovation (that is, a payment for E and a reduced payment for I) as well as what might be earned if a firm can innovate successfully while the other does not (that is, an increased payment for E and incremental profits without a payment for I).

Note that both I and E choices depend on $\tau - \tau_c = \frac{1}{2}(\pi - \pi_c + \Delta_c)$ but in opposite directions. An increase in $\tau - \tau_c$ causes I to favour a w -biased experiment whereas the opposite is true for E . In addition, both firms are pushed towards a w -biased experiment as Δ increases. Thus, on the demand-side, as $\Delta - \pi$ increases, E favours a w -biased experiment while I 's response is more muted and potentially ambiguous (depending on λ). Thus, we can see that, for the entrant, the pressures to undertake a w -biased experiment actually fall as the technological opportunity becomes potentially more disruptive.

5 Experiments with market signals

One view of disruption is that it represents a squandered opportunity. In this view, established firms are in a position to exploit a technology opportunity ahead of others but, because they fail to evaluate the technology's potential appropriately, they miss the opportunity to preempt competitors. There is a literature in economics that complements this view, showing that Arrow's replacement effect may not arise if, in the absence of innovation by incumbents, competition is inevitable (Gilbert and Newbery (1982)). Translated into the language of disruption, it is argued that incumbents should disrupt themselves before they are disrupted.

This paper has focused on the choice of experiments as the mechanism by which technological opportunities are evaluated for the purpose of exploitation. To evaluate the squandered opportunity argument, here we consider a situation where I is able to conduct an experiment and exploit the technology, if they so choose, ahead of any actions by the entrant. Given this timeline, there are two broad cases of interest. In the first, we maintain our assumption from the previous section that the outcomes of experiments are not observable but the exploitation decision is. In the second, we consider a situation where experimental outcomes themselves are observable to the entrant.

5.1 Unobservable experimental outcomes

Suppose that I is able to experiment and exploit the technological opportunity prior to E . Our interest is in whether, under conditions we have already identified where it is optimal for I to choose an n -biased experiment, whether being able to experiment and exploit ahead of E changes the incumbent's choice. From Proposition 4, $\{\sigma_n^I, \sigma_w^E\}$ is an equilibrium outcome in the simultaneous move game if $p \in [\frac{c}{\lambda\pi_c + (1-\lambda)\pi}, \frac{C}{\Delta_c}]$. Note that this continues to hold if the opportunity is also disruptive/architectural in the sense that $\frac{c}{\pi} < \frac{C}{\Delta}$ as I would still choose an n -biased innovation. Thus, we will assume that both of these conditions hold in what follows.

To evaluate this choice, we work backwards by considering E 's response to I 's actions. As we are looking for equilibrium outcomes, in deriving this we presume that E correctly infers the experimental choice (if any) of I and infers information from their subsequent exploitation choice.

First, note that if I chooses an n -biased experiment, their best outcome is if E does not enter regardless. However, if I receives a signal of w , I will exploit as it earns a payoff of $\Pi_c + \Delta_c - C$ because E exploits with certainty in this case. By contrast, if the signal is n , E may not enter. This happens if the posterior probability $\Pr[W|\{n, \sigma_n^I\}]$ is less than $\frac{c}{\pi_c}$.

However, if this condition did not hold, then it would be optimal for E to conduct a w -biased experiment. Thus, there is a weak improvement in I 's return to a n -biased experiment in this case as a negative signal reduces the probability that E ends up exploiting the technology.

Second, note that if I chooses a w -biased experiment, E will only enter if I exploits the technology (that is, I receives a signal of w). The problem is that, under the conditions above, it is always optimal for E to conduct a w -biased experiment if I chooses to exploit the technology. Thus, a preemptive w -biased experiment does not improve the probability that I can exploit the technology in the absence of entry by E . Thus, the returns to the w -biased experiment are no higher than before.

Thus, we are left to conclude that because E has the option to experiment following I 's exploitation decision that the relative return to a w -biased experiment relative to an n -biased experiment for I are not improved if I can preemptively experiment and exploit prior to the entrant.

5.2 Observable experimental outcomes

The experimental choices analysed in the previous section assumed that, while firms could infer the bias of experiments being conducted by their rival, the experimental outcomes themselves were private. This could be a reasonable assumption if the experiments involved direct feedback from groups about product prototypes rather than the launch of a minimum viable product or any type of product as part of an experiment to test its market performance. In such cases, experimental outcomes could be observed by the rivals. In this section, we relax that assumption for an incumbent firm who is able to experiment and exploit preemptively.

Consider the case, akin to the previous subsection, where only the incumbent conducts an experiment but the outcome of this experiment can be observed by a potential entrant. This scenario aligns with a situation where the technological opportunity is first presented to the incumbent who can then choose to run an experiment to test it. This approximates, for instance, the experiment by Blockbuster in on-demand video in the early 2000s before others had done so. As argued earlier, because broadband technology was nascent at the time, this represented an n -biased experiment. Thus, having an observed n signal caused Blockbuster to not pursue that opportunity while not precluding others from inferring that on-demand products may be successful in the future and entering on that basis.

One factor that will guide the incumbent's choice is how the entrant reacts to different signals. Thusfar, we have assumed that a w signal results in exploitation while an n does not. In this case, the incumbent can expect the following:

1. With an n -biased experiment, a w signal results in I receiving $\Pi_c + \Delta_c - C$ while an

n signal results in Π .

2. With a w -biased experiment, a w signal results in I receiving $\frac{p}{p+(1-p)(1-\lambda)}(\Pi_c + \Delta_c) + \frac{(1-p)(1-\lambda)}{p+(1-p)(1-\lambda)}\Pi - C$ and an n signal results in Π

Given this, I will favour an w -biased experiment over an n -biased one if:

$$p(\Pi_c + \Delta_c) + (1-p)\Pi - (p + (1-p)(1-\lambda))C \geq p\lambda(\Pi_c + \Delta_c - C) + (1-p\lambda)\Pi$$

This never holds if $\Pi_c + \Delta_c < \Pi$ as p cannot be less than $\frac{C}{\Pi_c + \Delta_c - \Pi}$. If $\Pi_c + \Delta_c > \Pi$, then a w -biased experiment will be chosen if $p \geq \frac{C}{\Pi_c + \Delta_c - \Pi}$. Nonetheless, we can conclude that when the incumbent's experimental outcome is observable this makes it more likely that the incumbent will choose an n -biased experiment. This, however, is a result of the payoff from a favourable experimental outcome being lower than a strategic effect to impact on entrant behaviour.

This raises the natural question of whether there are ways for the experimental choice to impact on entrant choices? Specifically, is it possible for an experimental choice to lead to a situation where an entrant does not enter following a w signal? Clearly, if $\pi_c \geq c$, this is only possible if I conducts a w -biased experiment. In this case, following a w signal, E 's expected payoff is $\frac{p}{p+(1-p)(1-\lambda)}\pi_c - c$. If $p\pi_c < (1 - (1-p)\lambda)c$ or $p < \frac{(1-\lambda)c}{\pi_c - \lambda c}$, then E will never exploit if I conducts a w -biased experiment.¹⁵ Given this, I will favour an w -biased experiment over an n -biased one if:

$$p(\Pi + \Delta) + (1-p)\Pi - (p + (1-p)(1-\lambda))C \geq p\lambda(\Pi_c + \Delta_c - C) + (1-p\lambda)\Pi$$

or $p \geq \frac{(1-\lambda)C}{\Delta + \lambda(\Pi - \Pi_c - \Delta_c)}$. In this case, we see that there can be strategic value to the incumbent choosing a w -biased experiment because it presents an unclear signal to the entrant and this is sufficient to deter their entry.

Interestingly, however, this effect is strongest if the technological opportunity is more sustaining/modular than disruptive/architectural. In other words, if the opportunity is disruptive/architectural, choosing a w -biased experiment even if it potentially deterred E 's entry, would not be optimal for the incumbent.¹⁶ What this means is that even if there is a scenario whereby I choosing a w -biased experiment can facilitate in preempting en-

¹⁵This also means that E themselves would never conduct a w -biased experiment themselves in this situation.

¹⁶In addition, having received a w signal from the incumbent's experiment, if experiments are costless, it is optimal for the entrant to pursue an n -biased experiment in this case as it would otherwise not exploit the technology. If this occurred, then the incumbent's upside from deterring entry would not occur and it is easy to see that the incumbent would find it optimal to initially choose an n -biased experiment.

try by E , this is not likely to be optimal precisely when the technological opportunity is disruptive/architectural.

6 Conclusion

This paper has put the nature (in particular, the bias) of experiments at the forefront of studying how new technologies are evaluated. In so doing, it addresses the (ahem) implicit bias that many informal discussions have towards unbiasedness. It has been shown that firms will favour biased experiments when evaluating new technological opportunities because these offer a way of obtaining a clearer signal when the state of the world is in opposition to their predispositions regarding exploitation in the absence of experimentation. Thus, the paper shows that many of the biases that exist in experiments including choices of sample selection, timing with respect to technology evolution or prototype design can be optimal given the nature of the decision facing firms.

Nonetheless, this paper is still at a high level on several dimensions. First of all, bias was modelled generically rather than derived from more specific information regarding experiments being run. Future work could actually measure the bias of such experiments based on their context and methodology. Indeed, the results here suggest that assessing bias more carefully is warranted. A recent example of the literature moving in this direction is Cao et al. (2020) who relate bias to the gender-composition of a product’s beta testers and find that such bias impacts on the assessment and launch of new product designs.¹⁷

Second, the role of experiments presented here is simply to inform the exploit or not choice by a single decision-maker. This is not uncommon but in the literature on disruption, decisions can often be driven by internal frictions within incumbent firms (see, for example, Bresnahan et al. (2011)) and dealing with outside resource providers for entrant firms (see, for example, Kerr et al. (2014), Agrawal et al. (2021)). This introduces an additional role for experiments as a means of persuading other agents as to the efficacy of exploitation choices in order to reduce or manage such frictions.¹⁸

Finally, the model here considers a firm’s choice of a single experiment. In principle, there is nothing preventing firms from running multiple experiments each with different biases spanning the available set. Moving in this direction, however, would require making assumptions regarding the relative costs of different experiments; something that I have avoided doing here for clarity. How to manage a portfolio of experiments is, therefore, a

¹⁷Another recent paper by Camuffo et al. (2020) shows how experimental methods training can improve business outcomes for innovators.

¹⁸These issues are explored in a companion paper (see Gans (2021)).

fruitful direction for future research.

7 Appendix: Proofs of Propositions

7.1 Proof of Proposition 1

Let σ_w and σ_n denote the w - and n -biased experiments respectively. First, if the firm chooses the w -biased experiment, their payoff upon receiving a signal of n is v_S as they optimally choose not to exploit. If they receive a signal of w and exploits, their expected payoff is $P(W|w, \sigma_w)v_W + (1 - P(W|w, \sigma_w))v_N$ where $P(W|w, \sigma_w) = \frac{p}{p+(1-\lambda)(1-p)}$. This will be an optimal response if this exceeds v_S or $p > \frac{(1-\lambda)(v_S-v_N)}{v_W-v_S+(1-\lambda)(v_S-v_N)}$. The RHS is less than $\hat{p}$ so this holds for $p > \hat{p}$.

Second, if the firm chooses the n -biased experiment and exploits contingent on receiving a signal of w , they earn v_W as they optimally choose to exploit. If they receive a signal of n , they will choose to exploit if $(1 - P(N|n, \sigma_n))v_W + P(N|n, \sigma_n)v_N > v_S$ where $P(N|n, \sigma_n) = \frac{1-p}{p(1-\lambda)+1-p}$. This is an optimal response if $p > \frac{v_S-v_N}{(1-\lambda)(v_W-v_S)+v_S-v_N}$. The RHS is greater than $\hat{p}$ so this may or may not hold if $p > \hat{p}$.

There are, therefore, two cases to consider regarding the firm's experiment choice if $p > \hat{p}$: (1) $p \in [\hat{p}, \frac{v_S-v_N}{(1-\lambda)(v_W-v_S)+v_S-v_N}]$ and (2) $p > \frac{v_S-v_N}{(1-\lambda)(v_W-v_S)+v_S-v_N}$. In case (1), the firm chooses σ_w if

$$(1 - (1 - p)\lambda)(P(W|w, \sigma_w)v_W + (1 - P(W|w, \sigma_w))v_N) + (1 - p)\lambda v_S \geq p\lambda v_W + (1 - p\lambda)v_S$$

which holds for all $p > \hat{p}$. In case (2), the firm will exploit regardless of the signal from σ_L . Thus, the firm chooses σ_w if

$$(1 - (1 - p)\lambda)(P(W|w, \sigma_w)v_W + (1 - P(W|w, \sigma_w))v_N) + (1 - p)\lambda v_S \geq pv_W + (1 - p)v_N$$

which holds for all p which proves the first part of the proposition. The second part can be proved in the same way.

7.2 Proof of Proposition 3

The exploitation probability is the probability that a firm is observed to exploit a technological opportunity; that is, it is $p\lambda_w + (1 - p)(1 - \lambda_n)$. Note that when $p \geq \max\{\hat{p}_E, \hat{p}_I\}$, both incumbent and entrant choose the w -biased experiment and their exploitation probability is $p + (1 - p)(1 - \lambda)$ for both. When $p \leq \min\{\hat{p}_E, \hat{p}_I\}$, both incumbent and entrant choose the n -biased experiment and their exploitation probability is $p\lambda$ for both. If $\frac{\pi}{c} > \frac{\Delta}{C}$, then it is possible that $\hat{p}_E < p < \hat{p}_I$ and the incumbent chooses the n -biased experiment while the entrant chooses a w -biased experiment. This implies that the probability that the entrant

exploits the opportunity is $p + (1-p)(1-\lambda)$ while the probability that the incumbent exploits the opportunity is $p\lambda$. It is easy to see that $p + (1-p)(1-\lambda) > p\lambda$ for $\lambda < 1$. The proof for $\frac{\pi}{c} < \frac{\Delta}{C}$ follows analogously.

7.3 Proof of Proposition 4

Let's first focus on the incumbent's experimental choice. If the incumbent chooses σ_w^I , then their expected payoff if the entrant also chooses σ_w^E is:

$$\begin{aligned} v_I(\sigma_w^I, \sigma_w^E) = & (p + (1-p)(1-\lambda)^2)(P(W|\{w, \sigma_w^I\}, \{w, \sigma_w^E\})(\Pi_c + \Delta_c) \\ & + P(N|\{w, \sigma_w^I\}, \{w, \sigma_w^E\})\Pi - C) \\ & + (1-p)\lambda(1-\lambda)(\Pi - C) + (1-p)\lambda(1-\lambda)\Pi + (1-p)\lambda^2\Pi \end{aligned} \quad (2)$$

where $P(W|\{w, \sigma_w^I\}, \{w, \sigma_w^E\}) = \frac{p}{p+(1-p)(1-\lambda)^2}$ and $P(N|\{w, \sigma_w^I\}, \{w, \sigma_w^E\}) = \frac{(1-p)(1-\lambda)^2}{p+(1-p)(1-\lambda)^2}$. If the entrant chooses σ_n^E , the incumbent's expected payoff becomes:

$$\begin{aligned} v_I(\sigma_w^I, \sigma_n^E) = & p\lambda(\Pi_c + \Delta_c - C) + (1-\lambda)(P(W|\{w, \sigma_w^I\}, \{n, \sigma_n^E\})\Delta + \Pi - C) \\ & + 0\Pi_c + (1-p)\lambda\Pi \end{aligned} \quad (3)$$

where $P(W|\{w, \sigma_w^I\}, \{n, \sigma_n^E\}) = \frac{p(1-\lambda)}{1-\lambda}$. Notice that it is only in situations where a signal confirms the bias of the signal chosen by *both* firms that there is ambiguity regarding the eventual state. For instance, $P(W|\{w, \sigma_w^I\}, \{n, \sigma_n^E\}) = 0$ because E 's choice of a w -biased experiment means that if they receive a signal of n , then the true state of the world is N with certainty. This simplifies the expressions considerably. We can write similar expected payoffs for I should they choose σ_n^I while the entrant chooses σ_w^E :

$$\begin{aligned} v_I(\sigma_n^I, \sigma_w^E) = & p\lambda(\Pi_c + \Delta_c - C) + 0(\Pi + \Delta - C) + (1-\lambda)(P(W|\{n, \sigma_n^I\}, \{w, \sigma_w^E\})\Pi_c \\ & + P(N|\{n, \sigma_n^I\}, \{w, \sigma_w^E\})\Pi + (1-p)\lambda\Pi \end{aligned} \quad (4)$$

where $P(W|\{n, \sigma_n^I\}, \{w, \sigma_w^E\}) = \frac{p(1-\lambda)}{1-\lambda}$ and $P(N|\{n, \sigma_n^I\}, \{w, \sigma_w^E\}) = \frac{(1-p)(1-\lambda)}{1-\lambda}$. If the entrant chooses σ_n^E , the incumbent's expected payoff becomes:

$$\begin{aligned} v_I(\sigma_n^I, \sigma_n^E) = & p\lambda^2(\Pi_c + \Delta_c - C) + p\lambda(1-\lambda)(\Pi + \Delta - C) \\ & + p(1-\lambda)\lambda\Pi_c + (1-p+p(1-\lambda)^2)\Pi \end{aligned} \quad (5)$$

This last case is interesting because both have chosen n -biased experiments. This means that if either one gets a signal w and innovates, there is no residual uncertainty. Of course,

	σ_w^E	σ_n^E
σ_w^I	$v_I = p(\Pi_c + \Delta_c) + (1-p)\Pi$ $-(1-(1-p)\lambda)C$ $v_E = p\pi_c - (1-(1-p)\lambda)c$	$v_I = p\lambda(\Pi_c + \Delta_c) + p(1-\lambda)\Delta$ $+(1-p\lambda)\Pi - (1-(1-p)\lambda)C$ $v_E = p\lambda\pi_c - p\lambda c$
σ_n^I	$v_I = p(\Pi_c + \lambda\Delta_c) + (1-p)\Pi - p\lambda C$ $v_E = p(\lambda\pi_c + (1-\lambda)\pi)$ $-(1-(1-p)\lambda)c$	$v_I = p\lambda(\Pi_c + \lambda\Delta_c) + p\lambda(1-\lambda)\Delta$ $+(1-p\lambda)\Pi - p\lambda C$ $v_E = p\lambda(\lambda\pi_c + (1-\lambda)\pi) - p\lambda c$

Table 4: **Payoffs with Ex Post Competition**

the same is true if neither innovate; something that arises in all four cases considered above.

I prefers the w -biased experiment if:

$$v_I(\sigma_w^I, \sigma_w^E) \geq v_I(\sigma_n^I, \sigma_w^E) \implies (1-\lambda)(p\Delta_c - C) \geq 0 \implies p \geq \frac{C}{\Delta_c} \equiv \hat{p}_I(\sigma_w^E)$$

I prefers the n -biased experiment if:

$$v_I(\sigma_w^I, \sigma_n^E) \leq v_I(\sigma_n^I, \sigma_n^E) \implies p(1-\lambda)((1-\lambda)\Delta + \lambda\Delta_c) \leq (1-\lambda)C \implies p \leq \frac{C}{(1-\lambda)\Delta + \lambda\Delta_c} \equiv \hat{p}_I(\sigma_n^E)$$

Turning now to E 's choice of experiment, the four cases for I translate into the following payoffs for E :

$$v_E(\sigma_w^I, \sigma_w^E) = p(\pi_c - c) + (1-p)(1-\lambda)^2(-c) + (1-p)\lambda(1-\lambda)(-c) \quad (6)$$

$$v_E(\sigma_w^I, \sigma_n^E) = p\lambda(\pi_c - c) \quad (7)$$

$$v_E(\sigma_n^I, \sigma_w^E) = p\lambda(\pi_c - c) + (1-\lambda)(p(\pi - c) + (1-p)(-c)) \quad (8)$$

$$v_E(\sigma_n^I, \sigma_n^E) = p\lambda^2(\pi_c - c) + p(1-\lambda)\lambda(\pi - c) \quad (9)$$

Table 4 summarises the payoff matrix. E prefers an w -biased experiment if:

$$v_E(\sigma_w^I, \sigma_w^E) \geq v_E(\sigma_w^I, \sigma_n^E) \implies (1-\lambda)p\pi_c \geq c \implies p \geq \frac{c}{\pi_c} \equiv \hat{p}_E(\sigma_w^I)$$

$$v_E(\sigma_n^I, \sigma_w^E) \leq v_E(\sigma_n^I, \sigma_n^E) \implies p(1-\lambda)(\lambda\pi_c + (1-\lambda)\pi) \leq (1-\lambda)c \implies p \leq \frac{c}{\lambda\pi_c + (1-\lambda)\pi} \equiv \hat{p}_E(\sigma_n^I)$$

We are now in a position to characterise the Nash equilibrium outcomes regarding the simultaneous experimental choice by I and E .

1. $p > \max\{\hat{p}_I(\sigma_w^E), \hat{p}_E(\sigma_w^I)\}$: it is optimal for both firms to choose the w -biased experiment given that the other firm is doing so.
2. $p \in [\hat{p}_E(\sigma_n^I), \hat{p}_I(\sigma_w^E)]$: it is optimal for E to choose the w -biased experiment and I to choose the n -biased experiment.
3. $p \in [\hat{p}_I(\sigma_n^E), \hat{p}_E(\sigma_w^I)]$: it is optimal for E to choose the n -biased experiment and I to choose the w -biased experiment.
4. $p < \min\{\hat{p}_I(\sigma_n^E), \hat{p}_E(\sigma_n^I)\}$: it is optimal for both firms to choose the n -biased experiment given that the other firm is doing so.

If, say, $p \in (\hat{p}_E(\sigma_n^I), \hat{p}_E(\sigma_w^I))$, then E would want to choose a w -biased experiment only if I chose an n -biased one. However, if $p > \hat{p}_I(\sigma_w^E)$, then I will want to choose a w -biased experiment. In this case, there would be no pure strategy Nash equilibrium.

7.4 Proof of Proposition 5

First, note that should E choose to exploit the technology in the absence of a cooperative deal, then, under Nash bargaining, a cooperative outcome will emerge because joint surplus, $\Pi + \Delta \geq \Pi_c + \pi$ or $\Pi + \Delta \geq \Pi_c + \Delta_c + \pi_c$ as the case may be. Given this, the expected payoffs for E and I are as depicted in Table 3. Since $\pi \geq \pi_c - \Delta_c$, then $\tau > \tau_c$.

Second, using those payoffs it is easy to see that, for E , $v_E(\sigma_w^I, \sigma_w^E) - v_E(\sigma_w^I, \sigma_n^E) = (1-\lambda)(p\tau_c - c) > 0 \implies p > \frac{c}{\tau_c}$ while $v_E(\sigma_n^I, \sigma_w^E) - v_E(\sigma_n^I, \sigma_n^E) = (1-\lambda)p((\lambda\tau_c + (1-\lambda)\tau) - c) > 0 \implies p > \frac{c}{\lambda\tau_c + (1-\lambda)\tau}$. Moreover, $v_I(\sigma_w^I, \sigma_w^E) - v_I(\sigma_n^I, \sigma_w^E) = p(1-\lambda)(\tau - \tau_c) > C \implies p > \frac{C}{\tau - \tau_c}$ and $v_I(\sigma_w^I, \sigma_n^E) - v_I(\sigma_n^I, \sigma_n^E) = p(1-\lambda)((1-\lambda)\Delta + \lambda(\tau - \tau_c)) > C \implies p > \frac{C}{(1-\lambda)\Delta + \lambda(\tau - \tau_c)}$ as the condition for I making a w -biased choice. The proof is completed by the same method as the proof of Proposition 4.

References

- Agrawal, Ajay K, Joshua S Gans, and Scott Stern**, “Enabling Entrepreneurial Choice,” *Management Science*, 2021.
- Arrow, Kenneth**, “Economic Welfare and The Allocation of Resources for Invention,” in “The Rate and Direction of Inventive Activity,” Princeton University Press and NBER, 1962.
- Awaya, Yu and Vijay Krishna**, “Startups and Upstarts: Disadvantageous Information in R&D,” *Journal of Political Economy*, 2021, 129 (2), 000–000.
- Azevedo, Eduardo M, Alex Deng, José Luis Montiel Olea, Justin Rao, and E Glen Weyl**, “A/b testing with fat tails,” *Journal of Political Economy*, 2020, 128 (12), 4614–000.
- Bower, Joseph L and Clayton M Christensen**, “Disruptive technologies: Catching the wave,” *Harvard Business Review*, 1995, pp. 1–12.
- Bresnahan, Timothy F, Shane Greenstein, and Rebecca M Henderson**, “Schumpeterian competition and diseconomies of scope: illustrations from the histories of Microsoft and IBM,” in “The rate and direction of inventive activity revisited,” University of Chicago Press, 2011, pp. 203–271.
- Callander, Steven and Niko Matouschek**, “The Novelty of Innovation: Competition, Disruption, and Antitrust Policy,” *Management Science*, 2021.
- Camuffo, Arnaldo, Alessandro Cordova, Alfonso Gambardella, and Chiara Spina**, “A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial,” *Management Science*, 2020, 66 (2), 564–586.
- Cao, Ruiqing, Rembrand Koning, and Ramana Nanda**, “Biased sampling of early users and the direction of startup innovation,” *Harvard Business School Entrepreneurial Management Working Paper*, 2020, (21-059), 21–059.
- Che, Yeon-Koo and Konrad Mierendorff**, “Optimal dynamic allocation of attention,” *American Economic Review*, 2019, 109 (8), 2993–3029.
- Ching, Kenny, Joshua S Gans, and Scott Stern**, “Control versus execution: endogenous appropriability and entrepreneurial strategy,” *Industrial and Corporate Change*, 2019, 28 (2), 389–408.

- Christensen, Clayton and Michael Raynor**, *The innovator's solution: Creating and sustaining successful growth*, Harvard Business Review Press, 2013.
- Christensen, Clayton M**, *The Innovator's Dilemma: when new technologies cause great firms to fail*, Harvard Business Review Press, 1997.
- Gans, Joshua S**, *The Disruption Dilemma*, MIT press, 2016.
- , “A Theory of Disruptive Vision,” mimeo., Toronto 2021.
- **and Scott Stern**, “Incumbency and R&D incentives: Licensing the gale of creative destruction,” *Journal of Economics & Management Strategy*, 2000, 9 (4), 485–511.
- , —, **and Jane Wu**, “Foundations of entrepreneurial strategy,” *Strategic Management Journal*, 2019, 40 (5), 736–756.
- Gilbert, Richard J and David MG Newbery**, “Preemptive patenting and the persistence of monopoly,” *The American Economic Review*, 1982, pp. 514–526.
- Henderson, Rebecca M and Kim B Clark**, “Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms,” *Administrative science quarterly*, 1990, pp. 9–30.
- Keating, Gina**, *Netfixed: The epic battle for America's eyeballs*, Penguin, 2012.
- Kerr, William R, Ramana Nanda, and Matthew Rhodes-Kropf**, “Entrepreneurship as Experimentation,” *Journal of Economic Perspectives*, 2014, 28 (3), 25–48.
- Liang, Annie, Xiaosheng Mu, and Vasilis Syrgkanis**, “Optimal and myopic information acquisition,” in “Proceedings of the 2018 ACM Conference on Economics and Computation” 2018, pp. 45–46.
- Marx, Matt, Joshua S Gans, and David H Hsu**, “Dynamic commercialization strategies for disruptive technologies: Evidence from the speech recognition industry,” *Management Science*, 2014, 60 (12), 3103–3123.
- Mayskaya, Tatiana**, “Dynamic choice of information sources,” 2019.
- Moore, Geoffrey A**, *Crossing the chasm*, Harper business, 2014.
- Nelson, Richard R**, “The role of knowledge in R&D efficiency,” *The quarterly journal of economics*, 1982, 97 (3), 453–470.

Ries, Eric, *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*, Currency, 2011.