Technological Disruption and Long-Term Investors: Managing Risk and Opportunities

Alistair Barker[1,2], Ashby Monk[1], and Dane Rook[1]

[1] *Stanford Research Initiative on Long-Term Investing (SLTI)*

[2] *AustralianSuper*

February 2024

Abstract: The current wave of uncertainty - and indeed, fear - over the possible consequences of artificial intelligence for the global economy and human society is, paradoxically, normal. Technological disruption (TD) is a perpetual feature of modern human life, occurring at various scales and presenting both risks and opportunities to investors. This paper investigates LTI approaches to TD management and proposes improvements. Through case studies and interviews with senior decision-makers from diverse LTIs worldwide, it is found that while some long-term investors (LTIs) recognize TD as a source of risk and opportunity, few have systematic approaches for managing it effectively. However, we argue that systematic TD management is feasible, and we present a toolkit for assessing and responding to TD risks and opportunities at both asset and portfolio levels.

Keywords: Alternative Data, Technological Disruption, Institutional Investment, Innovation

Contact: abarker@australiansuper.com

Acknowledgements: The authors wish to thank our interviewees and case-study subjects for their time and candor. We are also grateful to Gordon Clark, Rajiv Sharma, and Carter Casady for valuable insights, critiques, and comments. Finally, we thank the members of the Affiliate Program at the Stanford Research Initiative on Long-Term Investing (SLTI), as well as AustralianSuper for their assistance and resources to complete this project. This research would not have happened without their support. That said, none of the above are responsible for any errors or omissions herein.

1. Introduction

Technology is not exclusive to humans. Plenty of other animals (e.g., various species of primates, corvids, dolphins) apply knowledge to invent new tools - which is the definition of technology. Instead, it is the reflexive relationship between technology and society that makes *Homo sapiens* unique. A new tool can raise the odds of survival for small groups of non-humans, but it does not fundamentally transform their social structure. However, for humankind, technology can radically *disrupt* social and economic systems for our entire species; and it reliably does so at least once each generation. A person born in the 1930s will have lived through three such episodes of mass-scale technological disruption, or 'TD' [Rook et al., 2017].[1,2] And while these TDs are societal, they have knock-on consequences for the distribution of economic gains within/across industries, within/across countries, and the overall path of growth. Moreover, TD manifests at even smaller scales as well. For example, the advent of smartphones has not by itself triggered a new industrial revolution, but it has nevertheless radically impacted many industries. This multi-scale nature of TD means that, in one form or another, TD is now a normal and perpetual feature of modern human life.[3]

In addition to manifesting over various scales, TD can also develop over multiple time horizons, and with differing velocities [Christensen 1997].[4] For example, there is a major difference between 'smooth' TD and discontinuous TD. Smooth TD entails gradual adoption of new technologies, with affected industries taking several years (or even decades) to become reconstituted (e.g., electrification, the rise of automobiles, telegraphs being replaced by landline telephones). Discontinuous TD, on the other hand, involves

[1] Mass-scale tech disruptions are also referred to as industrial revolutions, and we use those terms interchangeably in this paper. Historians often refer to 'The Industrial Revolution' (always capitalized), a period of major TD that lasted from the 1760s to the 1840s. However, there have been at least three industrial revolutions since then (depending on what criteria one uses to identify an industrial revolution or mass TD). See Rook et al. [2017].

[2] Rapid acceleration of advanced artificial intelligence - typified by large language models, or LLMs - may instigate the next global-scale TD.

[3] For those requiring a more exact definition of TD, we consider it to be: major alterations in established industries, as a result of the emergence of a new technology. This is an 'investment-centric' definition; it does not consider a technology that causes changes in social/individual behavior to be disruptive unless it has economic consequences at industrial scales. Likewise, a technology that births a new industry without significantly affecting existing industries is not 'disruptive' under our definition (however, we struggle to conjure examples of such a technology).

[4] See also Danneels [2004], Markides [2005], and Teece et al. [1998], and references therein.

rapid uptake and upheaval, whereby incumbent companies undergo sudden, pronounced losses of market share to new entrants (e.g., cellular phones supplanting landlines, the advent of online retail); and these upheavals can have drastic effects on financial markets, through creation - and eventual collapse - of bubbles in asset prices. So whilst both forms of TD present risk and opportunity to investors, discontinuous TD poses a significant, material risk to investors.[5]

The community of long-term investors, which are fiduciarily bound to manage the risks and threats to delivering on stated goals, are attempting to unravel TD along its multiple different scales, horizons, and velocities to understand the possible opportunities and consequences for investment performance. That said, most long-term investors ('LTIs', including public pension funds, endowments, foundations, and sovereign wealth funds) treat TD risk as being non-standard: it has no consistent role in their core processes for assessing, monitoring, and handling risks to their portfolios.[6,7] This ad hoc approach to TD can be hazardous, and it likely distorts LTIs' estimates of how their portfolios will behave. For instance, a LTI might take undue comfort in a portfolio that is well-diversified across economic sectors, even though that same portfolio might be poorly diversified across technologies - making it highly susceptible to TD-related risks. Given the regularity with which TDs occur, multiple TD events are likely to occur over an LTI's lifespan. Therefore, it seems clear that LTIs should understand the technological makeup of their portfolios, in terms of what risks and opportunities TD poses to their invested capital and ability to meet their objectives. This understanding is a key component of what we call TD management (TDM).

[5] Throughout the rest of this paper, we focus on discontinuous tech disruption: hereafter, any mention of TD refers to discontinuous TD (any reference to smooth or gradual TD will be made explicit).

[6] LTIs are also frequently called *institutional investors*. We use the term LTI to underscore the fact that almost all of these types of investment organizations have long-term liabilities and are therefore obligated to adopt long-horizon outlooks for managing the capital under their care.

[7] This finding is consistent across not only our research subjects in the case studies for this paper, but from our work with many dozens of LTIs on a broader spectrum of topics.

This paper investigates two open questions on TDM. First, what approaches are LTIs using for TDM?[8] Second, how can these approaches be improved? To date, published research on these issues has been essentially non-existent. To help remedy that gap, we have relied primarily on case-study methods.[9] Specifically, we conducted a series of 20 firsthand *elite interviews* - i.e., a combination of structured and unstructured interviews with senior decision-makers employed by world-class LTIs. Our sample of 20 LTIs is geographically diverse (four different continents are represented, in terms of locations of domicile), and includes some of the largest LTIs worldwide (with respect to assets under management).[10] Our research subjects, technique, and findings are described extensively in later sections of this paper, but four essential realizations are that:

- Even though LTIs generally treat TD as a non-traditional risk, some leading LTIs do use organized approaches for handling it. These approaches differ substantially across investors.
- Relatively few LTIs are pursuing significant upside opportunities from TD (apart from investing in venture capital funds). Relatedly, most TDM strategies seem to lack 'coherence' overall, and are piecemeal rather than whole-portfolio responses to TD.
- LTIs' *networks* appear to be their chief sources of information on TD-related risks and opportunities (far more so than their internal capabilities for monitoring and analyzing TD information). However, few LTIs systematically leverage their networks for actionable information.
- Successful, coherent TDM requires systematic processes for analyzing TD-related risks and opportunities, and the integration of that information into investment processes and decision-making of LTIs (or their immediate agents, such as external managers).

[8] While most LTIs treat TD as a non-standard form of risk (and therefore exclude it from their core approaches to risk management), some do have processes for integrating TD into their decision-making, whether on an ad hoc or systematic basis. See Section 3 of this paper.
[9] Regarding the suitability of these methods, see (e.g.): Unseem [1995]; Strauss and Corbin [1998]; Davies [2001].
[10] Further details on our empirical approach can be found in Section 3 of this paper.

For some readers, this last finding might appear paradoxical, as it is frequently assumed that TD is too idiosyncratic to be managed in any coherent, systematic manner. Yet our work finds that systematic TDM is feasible, and we present (and offer further references to) frameworks that give investors a general-purpose toolkit for conceptualizing TD risks and opportunities, as well as analyzing asset-level investment decisions - and how those decisions aggregate into portfolio-level impacts. We believe that these tools are the foundations of coherent TDM and could serve as a new standard of best-practice for long-term investors.

Some readers may find the following material especially pertinent in light of the surge in generative artificial intelligence, aided by advances in large language models (LLMs). These developments could be disruptive at market-scale, with implications for unemployment, productivity, and the reworking of corporate research-and-development processes. Indeed, various LTIs with whom we have interacted are expressing concern - and confusion - about how the continued evolution of generative AI is likely to hit their investment portfolios. While this paper concerns TDM generally (rather than the consequences of a single disruption), we believe the tools herein will help investors to gainfully tackle such questions themselves - and discover ways to translate answers into beneficial investment decisions.

The rest of this paper proceeds as follows. Section 2 discusses the nature of technological disruptions as market-scale (albeit non-traditional) risks, and implications for how TDs are analyzed. Section 3 distills findings from interviews of expert practitioners, on the topic of how their organizations are undertaking TD management. Section 4 offers a new framework for TDM. Section 5 discusses special considerations that LTIs should weigh in designing and implementing a TDM strategy. Section 6 concludes and suggests areas for further research.

2. Tech Disruption as a Non-Traditional Risk

A focal concern of this paper (and our research program, generally) is whether LTIs are currently situated to handle TD risk. Evidence from our case studies (which we cover in

Section 3) suggests that many are not. The immediate question then becomes: can LTIs' existing resources - in risk management, governance, portfolio construction, etc. - be reconfigured to better deal with TD-related risks and opportunities? The answer is largely context-specific (i.e., varies from one LTI to the next), and hinges on whether the LTI in question has adequate capacity to cope with *non-traditional market risk factors*, which we now briefly discuss.

Broadly, *market risk* refers to disadvantageous price movements across large numbers of assets. Causes of market risk (often referred to as market risk 'factors') are plentiful and diverse, but their chief point of commonality is that they impact a significant number of assets simultaneously (or nearly so) - i.e., market risk factors drive market-wide price patterns.[11] Market risk factors are divisible into *traditional* and *non-traditional* factors. Traditional factors include economic variables (e.g., unemployment levels, interest rates, inflation) as well as financial conditions (e.g., market-wide trading volumes and volatility, average profitability of listed companies).[12] In short, traditional factors are those that most investors use as the main basis for decisions related to market risk.

For most traditional factors, large volumes of granular, high-quality data are available. Also, for many traditional factors, there is well-developed financial theory connecting those factors to changes in asset values or prices. It is thus no surprise that many LTIs' risk-management systems - and organizational and governance structures for overseeing those systems - emphasize the monitoring and processing of traditional factors (in terms of information infrastructure, professional skill sets, etc.).[13]

At the most basic level (and essentially by definition), non-traditional market risk factors differ from traditional factors due to their lesser popularity: a smaller fraction of investors use non-traditional factors to drive their decisions related to market risk. Some of the most

[11] LTIs typically consider market risk to be distinct from other types of risk exposures, e.g., operational risk, counterparty risk, or strategy-related risk.

[12] The self-influencing nature of some market risk factors is worth noting. For example, recent volatility in market prices tends to strongly affect future volatility.

[13] Further, a large proportion of financial regulation centers on how financial and investment institutions monitor and account for traditional factors.

well-publicized non-traditional factors today are environmental, social, and governance (ESG) factors. Typically, there are far fewer good-quality datasets available on non-traditional factors than there are on traditional factors, and the connections between non-traditional factors and asset prices is less well-studied (from both theoretical and empirical standpoints). Alongside relative popularity and data availability, other meaningful differences between traditional and non-traditional market risk factors include:

- *Timeframes*: data on traditional risk factors often (but not always) emerge at more predictable cadences (e.g., quarterly reports, monthly economic surveys, meetings of the U.S. Federal Reserve Board), whereas data on non-traditional factors tends to emerge more idiosyncratically. Relatedly, traditional factors frequently have a measurable impact on market prices - an impact that is often more immediate and reliable than is the case for non-traditional factors (i.e., prices may take longer to absorb data on non-traditional factors, because investors are uncertain or disagree on connections between non-traditional factors and prices).
- *Standardization*: datasets on traditional factors are often available in structured, immediately usable formats, whereas non-traditional data is mostly unstructured, 'alternative' data (see Monk et al. [2019]).
- *Objectivity*: data on traditional factors is widely considered to be more objective (in terms of relevance) than that for non-traditional factors - partly because traditional factors tend to be more straightforward and readily quantifiable in single values (e.g., consider measurements of inflation, as compared with social justice).[14]

TD possesses all the hallmarks of a non-traditional market risk factor: its connection with market behavior can be unclear over short horizons; it resists simple quantification; and it suffers from issues of data availability.[15] Still, some readers might question whether TD is

[14] This difference has become a harmful stall-point for ESG factors, whether at the market or company level. Data on many ESG factors come through simplified proxy variables (e.g., CO2 is often taken to be synonymous with environmental impact, even though such impacts are more expansive than carbon emissions alone). The alternative to simplified proxies is often 'battery' scores, in which analysts attempt to translate broad factor behavior (such as governance) into single values - an approach that entails significant discretion and is subject to argument.

[15] Readers may notice that the foregoing distinctions between non-traditional and traditional market risks are chiefly made with respect to differences in data. This is not coincidental. Investing is fundamentally about converting signals of value - which are necessarily embodied in data - into returns. Thus, to an extent, differences between traditional and non-traditional factors simplify down to popular judgments about what signals matter most. But this produces something of a chicken-and-egg problem regarding data: while it is certainly true that data on traditional factors tend

genuinely a *market-level* risk factor. The fact that technology is a risk at the level of industries and companies is inarguable (after all, most public companies explicitly list it as such in their annual filings). And examples abound of technology disrupting individual industries (e.g.: streaming media and the music industry; bookings websites and the travel industry). But is TD truly a risk factor at *market-wide* scales?

The answer is: sometimes. Widespread electrification and the rise of the world-wide web are classic examples. And the continued ascension of artificial intelligence and quantum computing are current contenders as TDs that may soon achieve market-scale impact.[16] Hence, not all TDs manifest into market-level risk; however, those that have the potential to do so should be a foremost concern for LTIs. That is because, although most LTIs (attempt to) diversify their portfolios against company- and industry-level risks, not many are prepared to cope with widespread market risk (notably many passive-oriented LTIs emphasize simply trying to 'hold the market portfolio').

A key to effectively handling market-level TD risk stems from the fact that large-scale TDs affect different segments of the economy to different degrees and at differing speeds. This opens the possibility of LTIs not only defending themselves against TD risks but also capturing new investment opportunities that arise from TD. The question we thereby pose is: how (and how well) are LTIs managing these TD risks and opportunities? In the next section, we uncover some initial answers.

3. Present Approaches to Managing TD

Here, we explore empirical evidence on the various approaches that long-term investors are taking to manage technological disruption. As noted above, the non-uniform ways in

to be more standardized and ready for use than data on non-traditional factors (because more resources and infrastructure have been developed around the former, since more investors agree on the meaningfulness of traditional factors), it is also true that some traditional factors have become 'traditional' for the simple fact that data on them have long been easily accessible (relative to many non-traditional factors). All of this makes it unsurprising that most investors orient their decision-making processes and infrastructure more around traditional than non-traditional market risks. However, none of the foregoing makes non-traditional risks immaterial: new environmental regulations and social movements can greatly affect asset valuations - even if investors may have trouble quantifying these impacts straightaway.

[16] It is important to note here that market-level risks associated with TD are not confined to changes in companies' and industries' market shares: such risks also come from the instabilities that come from new technology.

which TDs affect markets (in terms of impacting different industries to different extents and at different rates) presents LTIs with both risks and opportunities related to TD. In what follows, we pay close attention to how LTIs' management strategies account for both these upside and downside elements of TD.

In terms of methodology, we adopted a qualitative, case-based approach that made use of elite interviews (for discussion of this technique, see Strauss and Corbin [1998] and Davies [2001]).[17] The 20 organizations in our sample are domiciled across 4 continents, and are among the largest (in terms of assets under management) sovereign wealth funds, public pension plans, and endowments on the planet. For each subject, we conducted interviews with top-level decision-makers (e.g., CIO, CEO). We cross-checked our interview findings with publicly available data through a triangulation process (see Strauss and Corbin [1998]). Following the research guidance established by Clark [1998], and Clark and Urwin [2008], we preserve our subjects' anonymity at both the individual and organizational levels. To ensure this anonymity (and to respect the confidentiality of some of the information that underpins our conclusions), we focus on synoptic findings, rather than the specific details of any one fund's approach to TDM.[18]

Our semi-structured interviews focused on the following questions:

1. How do funds think about TD: for example, are they using specific frameworks to analyze TD risk?
2. What are funds doing to help them handle TD-related risks and opportunities?
3. What hurdles do funds face with respect to TDM?

The logical interconnection between these lines of inquiry is straightforward: any single fund's perspective on TD (in terms of how they conceptualize and analyze it) inevitably shapes how that fund chooses to manage TD-related risks and opportunities; yet funds are

[17] Given the nature of elite interviews, our sample of subjects prioritizes best-of-breed (i.e., industry leaders) rather than representativeness. In essence, this leads to findings that might be considered a 'high-water mark' with respect to how some of the most respected LTIs are handling TDM.

[18] This practice allows subjects to disclose information with greater candor and depth than they might otherwise be willing to do.

not unrestricted in how they handle TD (given the constraints and inertias that every LTI experiences); therefore, how a given fund currently handles TD may fall short of its ideal TDM strategy. Below, we discuss our central findings according to this sequence of logic.

3.1. Perspectives

Among the LTIs we studied, there was significant diversity in perspectives on how TD should be monitored and analyzed. Universally, our subjects agreed that TD should be treated as a nontraditional risk factor (see Section 2), and therefore many of the tools that are used to assess traditional risks (and opportunities) are challenging to apply to TD. Many subjects stated that this challenge mirrors the one they encounter in analyzing environmental, social, and governance (ESG) factors.

A majority of subjects acknowledged that, for specific technologies, TD can be a market risk factor and affect valuations across many industries. However, some subjects voiced hesitancy about attempting to *manage* TD as a market risk factor, primarily due to their observations that:

- Technology-specific TD can affect different industries at significantly different rates and to significantly different extents (more so than many traditional market-risk factors); and
- Not all disruptive technology has market-wide impacts, and technologies that upend one industry may have negligible impact on other industries. Moreover, predicting which industries a potentially disruptive technology might affect (and to what extent) can be very difficult.

Our subjects' viewpoints on whether TD tends to be a market-level risk factor seemed to be roughly correlated with their opinions about how much effort LTIs should expend on managing TD risks and opportunities - and these opinions were diverse: some subjects asserted that the effects of TD tend to materialize over such long horizons, and with such substantial uncertainty, that allocating significant resources to TDM is rarely worthwhile; whereas other subjects considered TD risks to be one of their focal investment concerns, and argued that building strong capabilities in TDM is prudent. These diverging opinions

on the value of TDM ostensibly drive the wide range of TDM strategies that our subjects have adopted (which we discuss in the next subsection).

A point on which subjects uniformly agreed was that TDM is difficult - largely because analyzing technological disruption is such a challenge. This analytical difficulty seems to come from two sources: scarcity of trustworthy TD data ('signals'); and lack of cogent frameworks for translating TD data into assessments of TD-related risk and opportunity. Together, these problems make TD resistant to quantification, which our subjects cited as a key impediment to sound TDM (in specific, the difficulty of quantifying TD risks is a main reason why they are not factored into many of our subjects' investment decisions in any systematic, universal way - and any inclusion tends to be on an ad hoc basis).

Our subjects highlighted various causes for the weakness of TD data. One of these was the "out-of-sample" nature of technological disruption: that is, differences in historical and market contexts means that past TD events may not provide direct guidance for the future.[19] Another inhibitor of TD data quality appears to be its granularity: many subjects noted that much of the TD data available to them is industry-level data (or else data about supply-chains), which restricts the level of detail at which TDM can operate.[20]

Subjects also cited fragmentation as an issue with TD data, and this problem applies to data that is both internal and external to the organization. Internal fragmentation stems from the fact that separate teams within the organization may have unequal access to TD-relevant data. For example, an in-house VC team might not share some of its data with other parts of the organization (in many cases, this withholding would probably be unintentional - the VC team may well assume that they are the only ones with any use for such data - if they indeed make use of it at all). Such internal fragmentation appears to be a reason why some LTIs don't engage in meaningful TDM: it is challenging enough to integrate TD data from multiple sources into a unified view of TD risk - having to hunt for that data across sources

[19] However, this perspective may stem from the absence of a suitable framework for analyzing TD. In the next section, we cover an approach to TD analysis that proves there are lessons which are portable across all TD events.
[20] This also helps explain why most attempts at managing technological risk are addressed through industry-based diversification.

worsens the challenge. However, organizations that choose to ignore the need to integrate TD information into a unified, organization-level viewpoint run the risk of incoherence, whereby different teams make TD-related decisions that conflict with or contradict one another. Some subjects suggested that improvements to knowledge management (KM) practices can be a path to combatting the fragmentation problem in TD data.

Many subjects indicated that, when it comes to assessing TD, they rely predominantly on externally derived data and information, rather than data and information that gets generated internally (e.g., through primary research or in-house diligence). Specifically, we heard that external asset managers (namely, venture capital funds) tend to be the chief source of TD information for subjects in our study. However, accessing the TD-relevant information possessed by external managers is often problematic. First, there is the reality that (typically) one must be a client of such managers in order to access any valuable or differentiated insights they have on TD trends, risks, etc. (this in itself places an upper bound on how much TD information is available to a LTI - given the finite nature of assets under management). Second, there is the question of whether the external manager has an incentive or obligation to share its TD-relevant information with its clients: some managers might argue that clients have no right to such information – i.e., clients pay for net returns, not the underlying ingredients to those returns. Internal KM systems are of little help in situations where external managers are unwilling to share TD-related insights, which suggests that willingness to share information may be a valuable filtering criterion when selecting external managers.

The importance of external managers as sources for TD data and information proved to be linked to a broader observation: the quality of a LTI's network, and its ability to utilize that network in an organized way, is a crucial determinant for how well it manages TD risks and opportunities. Networks in this sense consist of the relationships - contractual and otherwise - that the organization has with external entities. A LTI's network includes its external managers, its consultants, relationships that its employees have with experts outside the organization, and any privileged *contextual connections* that the organization has. Examples of useful contextual connections here might be stakeholders, for instance:

- If a given LTI is a pension fund for construction workers, its own stakeholders might have some special perspective on how technological change might affect infrastructure assets; or
- A university endowment might have privileged access to expertise within its engineering, physical sciences, or other relevant departments that are closely aligned with cutting-edge technologies.

Every LTI's network is unique. However, we observed several commonalities in how our subjects claimed to us their networks for TD-related purposes, which gives clues on what constitutes a 'good' network as far as TDM is concerned:

- Consultants: no subjects claimed any meaningful reliance on consultants for TD information or data. It was widely felt that consultants fail to give differentiated, actionable intelligence on TD.
- Weak ties: repeatedly, subjects said that many elements of their network were not necessarily valuable as direct sources of information, but rather for the power of their own networks (whether formal or informal). In network theory, the value of such indirect connections is known as the 'strength of weak ties' (see, e.g., Granovetter [1973; 1983]); in this setting, the value of weak ties manifests (for example) in knowing someone who knows someone with special TD information or knowledge (and is willing to share it). A related observation was the power of geography. Some subjects stated that their organizations keep offices in 'high-tech' locations, such as Silicon Valley, because they feel that having "boots on the ground" in such locales helps TD information to be transferred "osmotically". This phenomenon is almost certainly attributable to the power of weak ties.
- Venture access: multiple subjects stated that their involvement in venture capital (whether through in-house teams or external managers) is done as much (or more) for information access as investment returns. They claimed that direct exposure to new technologies through startups keeps them better informed on technology risks elsewhere in their portfolios.

Although these characteristics appear to be shared by all 'good' networks (as far as TDM is concerned), it is clear - as various subjects were keen to point out - that a high-quality network does not necessarily translate to high-quality TDM. To do TDM *successfully*, a LTI must have a consistent (usually formalized) process to actively mine its network for TD-related insights, as well as a coherent approach to integrating and weighting the various 'signals' that it gathers from across its network. This ability to weigh TD information was seen as especially important, because LTIs themselves must determine the "threshold for action" on any given piece of TD information: that is, the necessary conditions for information to prompt a portfolio-level action. Our subjects varied substantially on how they manage these translations of information into action, which is the topic we next explore.

3.2. Management Approaches

Taken as a whole, our interviews indicated that there are four essential approaches which LTIs adopt in managing technological disruption: *dismiss*, *defense*, *offense*, and *hybrid*. These approaches to TDM are not mutually exclusive (or even mutually incompatible) at the organizational level: it is entirely possible for a LTI to use different approaches for different segments of its portfolio at different points in time. However, from what we observed, this was rarely the case in practice, because each of these approaches appears to stem from top-level perspectives (discussed in the previous subsection) and capability sets. We now discuss each of these approaches in brief.

- *Dismiss:* While all of our research subjects were aware that TD poses risks to their portfolios, several stated that their organization still opted to dismiss those risks by electing to not directly manage them. Those subject LTIs who embraced this approach tended to have fewer organizational capabilities (often as a consequence of being smaller funds, in terms of total assets under management, or using predominantly passive strategies overall), and relatively few direct investments. In general, those who adopted this strategy voiced the opinion that responsibility for handling TD risk could be passed on to external managers, or else that holding a portfolio with very broad market exposure could diversify any specific TD risk.

Subjects who used this approach also tended to have little confidence that their organization could (or should) attempt to capitalize on TD-related opportunities.

- *Defense:* The majority of our subjects assumed a 'defensive' posture to TDM, which emphasizes risk mitigation rather than pursuit of TD-related upside. For the most part, subject LTIs with this defensive style tend to have significant exposures to assets that are 'incumbent', and more 'at-risk' from disruption (as opposed to standing to gain from TD). In specific, these LTIs have portfolios that are heavy in one or more of the following: 'value' stocks; infrastructure and real estate; leveraged buyout funds; and commercial credit/fixed income. For the most part, subjects with a defensive approach voiced a preference for outright divesting rather than portfolio tilts as a way to react to TD risk; that is, they mostly favor exiting industries (or even whole asset classes) or specific holdings once risk from TD becomes excessive, versus proactively trying to shift portfolio allocations toward assets that offset this risk.[21] Some subjects pointed to the scalability of divestment as a redeeming quality for larger investors.

- *Offense*: All subjects felt that capitalizing on upside opportunities caused by TD is challenging to do well, both consistently and at sufficiently large scale to make a meaningful impact on an LTIs portfolio returns. Said differently: it is harder to play offense than defense. The main reason given for this was that obtaining any special or differentiated insights on TD is hard, and the resources needed to even attempt to do so are typically expensive. These impressions largely explain why many of our subjects' organizations make no direct effort to play offense in their TDM strategies. Other, related reasons that our subjects gave for avoiding offense-centered approaches to TDM include: 1) a presumption that playing offense requires "picking winners" - which is akin to speculation for some investors (i.e., it is assumed to have a high failure rate, even if the payoffs for correct picks are

[21] Generally, this preference for reactive divestments over proactive tilts was attributed to the career risk associated with tilts that turned out to ultimately be unnecessary or ineffective.

high)[22]; and 2) for most TD events, historical precedents are rarely treatable as 'comparables' - they are seen as too dissimilar - and this lack of precedent convinces many funds that playing direct offense on any particular TD opportunity is imprudent.[23] Instead, many funds choose to 'outsource' offense to venture capital (VC) funds, which is seen as a more prudent approach, given its popularity; indeed, it is the most popular way that funds in our subject pool play offense (thematic investing around technological trends is a distant second place in popularity).[24] It should be noted, however, that TDM is not a primary reason why most of our subject organizations invest in VC funds: high expected returns and diversification benefits relative to public markets are usually the core reasons why LTIs participate in VC, and TDM benefits are just secondary considerations for some LTIs. Although many of our subjects expect their funds to increase portfolio allocations to VC in the coming years, they still expressed concern about the scalability of VC as a standalone solution to TDM: given the relatively small fraction of portfolio share that most funds apportion to VC *and* the fact that many VC funds attempt to diversify their holdings, the net capacity of a LTI to capitalize on TD is likely to be minimal, if investing a fraction of its capital in VC funds is the only element of its TD strategy. Relatedly, there is the matter of portfolio-level *coherence*: simply adding new assets to one's current portfolio does not imply that existing assets are any better managed with respect to TD, which can cause incoherent execution on TDM from a whole-portfolio standpoint. This problem of coherence is a serious one, and it reappears in most LTIs' TDM approaches.

[22] Our subjects also acknowledged that 'picking winners' is not a whole-portfolio approach to doing TDM, and that it is usually not possible to make bets on winners at a scale which would have any substantial portfolio-level impact (specifically, many funds' are functionally prevented from 'putting too many eggs in one basket').

[23] In the next section, we demonstrate how the logic behind this latter reason (lack of historical precedent) is flawed.

[24] It is worth mentioning that VC funds are ostensibly in the business of picking winners, and they are incentivized to configure their resources to do so more efficiently than many LTIs may be able to do by themselves (that is, VC funds enjoy economies of scale in winner-picking when it comes to TD). Moreover, due to the way that VC funds pool LTI capital across multiple deals, they can diversify risk in picking TD winners in ways that most LTIs cannot do by themselves. However, VC managers may be incentivized to over-diversify, relative to LTIs' TDM objectives: that is, by tempering their exposure to some TD risks through diversification, VC managers may be blunting LTIs' ability to play TD offense via externally managed VC funds.

- *Hybrid*: Some subjects claimed that their funds aimed to use a hybrid approach to TDM - i.e., blend both defensive and offensive strategies. However, few believed that their funds were succeeding in these hybrid approaches, mainly because their defensive and offensive activities were not well coordinated (in terms of data, risk management, investment portfolio decision-making etc.), and thus offense and defense did not suitably complement one another - in short, most hybrid approaches were seen to suffer from some degree of incoherence.[25] Generally, funds that used a hybrid approach still tended to rely on VC as the crux of their TDM strategies. However, some funds noted that this reliance was not just for controlling TD-related risks and returns: they noted that participating in VC can be a valuable source of material information and insight about TD, which can be translated into decisions in other areas of the portfolio.[26] Yet, problematically, most LTIs may not be equipped to act on the information derived from their VC connections; none of our subjects' organizations had a well-defined decision-making framework for converting VC-sourced information into concrete investment actions (this problem was not unique to VC: most LTIs seem to lack a formal process for translating TD information into portfolio action, regardless of where in their networks such information originates).

Again, we hasten to point out that these four categories of TDM approaches (dismiss, defense, offense, and hybrid) are generalizations, and the boundaries between them are not necessarily crisp. Yet, regardless of how one chooses to classify LTIs' responses to TD, these organizations face a common set of hurdles to performant TDM.

3.3. Hurdles

At the highest level, across our whole subject pool (as well as organizations that we have more informally studied in recent years), the greatest hurdle to managing technological

[25] Perhaps the most developed hybrid approach pursued by any of our subjects was a "3-pronged" strategy used by one fund, which entailed: avoiding losers, picking winners, and using insights from both of these activities to inform how the organization could improve its own internal technology.

[26] Multiple subjects pointed to a parallel with corporate venture capital, whereby conglomerate companies undertake direct venture investments as much (or more) for learning value as for returns amplification.

disruption is a lack of *coherence*, in terms of how TD is conceptualized and managed. That is, LTIs tend to have a non-integrated view about what TD entails (in terms of its impacts on assets, ramifications for markets, etc.) and what their organization should do about TD - not just from the standpoint of portfolio management, but also with respect to configuring organizational resources. In the next section, we introduce a framework for handling TDM coherence overall. For now, we point out some of the forces that underpin TDM incoherence; these can all be considered smaller hurdles that collectively add up to the primary hurdle of incoherence.

- Disengaged leadership: quite simply, if a LTIs' senior leadership is not supportive of TDM, then that LTI will almost certainly fail at TDM. To be done well, TDM needs coordinated perspectives and resources across the organization - e.g.: TD-relevant information needs to flow freely; different teams must be able to reconcile their interpretations of various TD signals (and successfully integrate those signals); TDM must factor into diligence and risk management processes in a consistent way; and execution of TD strategies (whether defense, offense, or hybrid) should be complementary across the portfolio. While some of this coordination can be achieved via process architecture, it will never succeed unless the organization's leaders are willing to actively pursue it. This need for direct support from senior leaders is largely due to the existence of the following additional hurdles (which leadership must work to offset).

- Mindset: institutional investment organizations are rarely set up to be innovative; their governance and resourcing structures are geared around incremental changes to the portfolio and organization.[27] The mindset that often arises from this natural incrementalism can be difficult to reconcile with the 'step changes' that inherently come from TD (and the step changes in the portfolio and organization that might sometimes be needed as 'best responses' to TD).

[27] See Monk and Rook [2020].

- Fragmentation of investment portfolios and investing teams: most LTIs are (directly or indirectly) have some form of organizational and portfolio groupings for common types of investments or activities to manage a large and complex portfolio. The way these grouping exist varies across LTIs, including along functional lines (e.g., in-house or through external funds), along asset class lines (e.g., equities, debt, derivatives) or by market type (e.g., public or private, foreign or domestic). Without extensive communication and transparency across these separate segments of the organization, it is difficult to properly manage TD risks: for example, a portfolio might be 'well diversified' across asset classes as far as traditional market risks are concerned, but it may fail to be diversified across TD risks, especially if different parts of the organization have conflicting perspectives on how to conceptualize, measure, and react to TD risks. Moreover, fragmentation by asset classes can inhibit flows of TD-related information. Earlier, we discussed the importance of networks for monitoring TD signals; however, the effectiveness of those networks is degraded whenever information cannot flow freely inside the organization after it is extracted from the external network.

- Misaligned incentives: TD risks and opportunities manifest along timescales that are variable. Sometimes the effects of TD are quick to take hold, and other times they may take decades to be significantly felt. Typically, the effort and resources that a LTI is likely to expend in directly managing a particular TD is negatively correlated with how distant in time that disruption's impacts will be at an asset level; said differently, the sooner a disruption's effects are noticeable, the more attention it will get from LTIs. This is not entirely unreasonable, but it does not mean that TDs with far-off consequences should be ignored: there are many near-term decisions that LTIs make that affect their long-term positions - for example, investing in private equity funds with commitment periods of a decade or more. Nonetheless, the compensation structures of many LTIs' employees cause them to prioritize short-term performance, and potentially over-discount effects of TD that occur in the mid-term or long-term future (the same is true for external managers).

- Resources: almost all LTIs face tight budgetary constraints, and securing funding for various new initiatives can be extremely difficult (and even politicized). Based on reports from research subjects, TDM is still very much a "new initiative" for most LTIs, and one that is certainly not cost-free. Given that tech disruptions involve non-traditional risk, entail substantial ambiguity, and have impacts that may not be felt for several years, it can be heavily challenging to obtain dedicated resources to manage them – even though such resources are needed for successful TDM.[28]

Many of the foregoing hurdles will be familiar to readers; they are essentially identical to the hurdles that promote short-termism among LTIs in general. We caution, however, that TD should not be construed as just another problem of myopia: it is a concern that goes beyond the difficulties LTIs face in reconciling short and long horizons. That said, we are aware (based on the decades-long struggle against LTI myopia, both in academic research and professional practice) that the foregoing hurdles are unlikely to be removed any time soon…if ever. Nevertheless, an integrated framework for TDM could help mitigate them, and that is what we next present.

4. Coherent TDM

Here, we introduce a framework for LTIs to begin coherently managing tech disruption - that is, doing TDM in a way that is consistent across the portfolio, and drives asset-level decisions that are complementary (rather than being disconnected or in conflict with one another). Essentially, this framework is jointly a perspective *and* process: it yields a way of conceptualizing tech disruption and subsequently evaluating it for the sake of making investment decisions.

We begin by noting that our framework is fundamentally a risk-management framework, because any performant TDM framework must ultimately be centered around risk: quite

[28] Although many of the subjects in our research pool stated that their organizations have few resources specifically dedicated to TDM, there were several notable exceptions. For instance, one employed a "disruption analyst" whose exclusive focus was to study and monitor technological disruptions and ensure that the organization is well informed about them.

simply, the impact of TD at an asset level boils down to the success or failure of a given (potentially) disruptive technology - hence, that success or failure poses risk to the asset. The magnitude of that risk depends on two variables: 1) the extent to which the asset is influenced by *either* the new technology or any one of the existing technologies that the new technology threatens to replace; and 2) the asset's degree of *incumbency*, whereby:

- 'Incumbent' assets are those that face risk from the new technology *succeeding* (in terms of its functionality and uptake), along with their ability to react to the success of that new technology (by adopting it themselves, developing other technologies to counter it, or shifting to new markets or modes of operating).
- 'New entrant' assets are those that face risk from the new technology *failing* (in terms of its functionality and uptake), along with their ability to respond to the failure of that new technology (by embracing or developing other technologies or pivoting to other markets or modes of operating).[29]

In truth, many assets may not be exclusively incumbents or new entrants, because it is possible to reside somewhere between these two poles (e.g., a business's costs may be reduced when a disruptive technology benefits its supply chain, but that same technology could cause the business's sales to decline). Regardless, we argue that a comprehensive TDM framework must necessarily be a risk-management framework; and the most generalized template for any risk-management framework is (likely) ISO 31000:2018.[30]

ISO 31000:2018 is a standard which proposes that responsible risk management should be an *ordered* process consisting of five sequential steps: context; identification; analysis; evaluation; and treatment. This standard articulates a detailed definition for each of these steps, but when applied to the specific problem of TDM, these steps compel LTIs to:

1. Context: understand (to a suitable level) the potentially disruptive technology and the likely improvements (or dis-improvements) it might make to technologies that

[29] A new entrant might, e.g., be a startup company, or an existing company that is able to (via use of the disruptive new technology) to enter a new industry or market.
[30] See: https://www.iso.org/obp/ui/#iso:std:iso:31000:ed-2:v1:en.

it could foreseeably replace. Also, study broader social, economic, and historical backdrops within which the disruptive technology is continues to evolve.

2. Identification: recognize the key entities likely to be involved in the disruption (e.g.: incumbent companies and new entrants; industries that could be directly impacted; relevant regulatory bodies).

3. Analysis: understand and model (qualitatively and quantitatively) how the entities in the previous step are likely to interact over the time-course of the disruption, and what the anticipated consequences of these interactions are likely to be, both in the short and long terms (e.g., loss of market share or bankruptcies/defaults by incumbents, increased inflation or unemployment, creation of new legislation or regulation).

4. Evaluation: assess possible investment actions that can be taken (for the sake of loss reduction or gain-seeking) in light of the (probability-adjusted) consequences identified in the previous step.

5. Treatment: select one or more candidate actions from the evaluation step, and implement them.

We maintain that, applied to TDM, this sequence is fundamentally *complete*, in the sense that each of these steps is necessary, and they are collectively sufficient for sound TDM. Moreover, any reasonable candidates for additional steps turn out to merely be subsets or combinations of these five.

Despite the completeness of this ISO-derived framework for TDM, it suffers from a lack of detail, and it fails to give users concrete guidance on how to execute any one of the steps. This under-specification can be resolved, however, by grafting two tools into this framework: Innovation Axes [Rook et al. 2017] and the DARLing approach to assessing

Resilience [Rook et al. 2021].[31] Later in this section, we give a detailed discussion of these tools and how they intersect with TDM. For now, we observe that:

- Innovation Axes are the essential dimensions along which disruptive technologies evolve. Framing TD by means of Innovation Axes gives clarity (in ways that we discuss below) to the steps of context and identification.
- The DARLing methodology helps deconstruct how assets respond to non-BAU[32] shocks. TDs fall into this category of shock, and DARLing yields a concrete way to execute the analysis and evaluation steps of the ISO-derived framework.

From a procedural standpoint, Innovation Axes and DARLing could be used for TDM as follows:

1. Use Innovation Axes to help understand the nature and impact potential of a new, disruptive technology, and identify assets that are likely to be materially affected (directly or indirectly) by the ascendence of that technology.
2. Assess the Resilience of those identified assets through the DARLing lens.
3. Aggregate those Resilience assessments up to the portfolio level.
4. Make asset-level decisions that acknowledge the whole-portfolio impacts of the disruptive technology.

We posit that this sequence leads to a coherent approach to TDM because it appraises TD at the whole-portfolio level, and therefore most strongly encourages consistent responses across asset classes, at adequate scales. (For example, we expect it will lead to responses that go beyond simply 'upping the VC allocation', or selectively divesting from the most under-threat incumbent companies.) We also note that this proposed approach closely ties TDM to Resilience-based portfolio management, which was introduced by Rook et al. [2021] and refined by Monk and Rook [2023] and Rook and Monk [forthcoming]. We lack space here to cover the Resilience paradigm in its entirety, but simply observe that its key differentiating aspect lies in how it emphasizes the 'joint' nature of drawdowns and

[31] "Resilience" here is capitalized for reasons articulated by Rook et. al [2021]; we give a very brief explanation of that reasoning later in this section.
[32] Here, BAU stands for 'business-as-usual'.

recoveries: it frames risk in terms of how assets respond to shock events (whether of known or unknown origin), both initially (drawdown) and eventually (recovery). Plainly, TD events are shock events, which makes the Resilience paradigm applicable. Indeed, we think Resilience is the most practical lens through which TDM should be conducted, for the sake of coherence.

Returning to the more immediate topic of Innovation Axes and the DARLing approach, we observe that these two tools can fulfill the first four steps (context, identify, analyze, evaluate) of a performant TDM process, as well as give guidance on the treatment step - in terms of informing a prioritization of actions that can be taken in response to a given episode of tech disruption (they do, however, stop short of indicating *how* chosen actions should be implemented; we discuss several key considerations for implementation in the penultimate section of this paper).

In the remainder of this section, we dissect Innovation Axes and the DARLing approach as they relate to TDM. Our coverage of these two tools is, however, necessarily succinct, and we refer readers to the original papers (Rook et al. [2017], and Rook et al. [2021]) for additional detail. (Readers should be aware that these two tools apply regardless of the particular approach to TDM favored by an investor - defense, offense, or hybrid. We also expect that some smaller LTIs may lack sufficient internal resources to fully implement these tools themselves, in which case they should explore the capacity and willingness of their external managers to do so - i.e., we see potential proficiency with these tools as a due diligence criterion in manager selection.)

4.1. Innovation Axes and Patterns

The first two steps of ISO 31000:2018 (context and identification) effectively reduce to the following problem: evaluating an emerging - and potentially disruptive - technology's prospects for replacing existing technologies (these prospects are determined not only by any technical advancement, but also by users' preferences, efficiency of production, etc.); and recognizing entities that stand to materially benefit or lose out from that technology's ascendence (or failure). We suggest that Innovation Axes are an effective tool for helping

solve that problem. Innovation Axes can be understood as the dimensions along which all (potentially) disruptive technologies evolve, relative to existing ones; in essence, they are the 'lanes' by which new technologies gain a comparative advantage over existing ones.

Innovation Axes were first recognized by Rook et al. [2017], during a study that sought to isolate points of commonality across all industrial revolutions.[33] Those authors closely examined each of the four industrial revolutions in modern history, and discovered that, regardless of what specific innovations that propelled a given revolution, the dimensions along which those innovations impacted industries, the economy, and society were fixed and unchanging across revolutions. In that study, five such dimensions (Innovation Axes) were uncovered: control, integration, reconfigurability, scale, and sustainability impact:

> "**Control** concerns how the behavior of an output or process is determined, as well as how specifically it is determined (e.g.: How variable are the outputs of an assembly line? Are routes taken by delivery drivers centrally decided or do the drivers have flexibility in routes they take?) **Integration** reflects the extent to which units in an economic (sub-)system are interlinked or coordinated. **Reconfigurability** is the ease with which (or extent to) which an output or process can be adjusted to meet shifting conditions or needs. **Scale** relates to both 'size' (in terms of physical proportions) and 'flow' of processes (e.g., sequential or parallel?) (for example, innovation along this Axis might involve extending the footprint of a supply chain from dozens to thousands of miles, or creating medical therapies that target specific cells or genes, rather than entire tissues or organs). **Sustainability impact** measures resource intensity or wastefulness (i.e., whether an output or process can be produced/executed over a long horizon without needing major revision). [Rook et al., 2017, p. 2-3]

It is crucial to note that these Axes are generic modes of change, and the specific changes that take place along them vary from one industrial revolution to the next. Rook et al. used the term "Innovation Patterns" to describe these specific sets of changes along the Innovation Axes, and each of the past industrial revolutions had its own Patterns (see the figure below).

[33] Those authors observed that, historically, the economic and techno-social transformations brought on by industrial revolutions have occurred at distinct paces - with some resulting in market 'bubbles', and others not.

Figure 1: Mapping of Innovation Patterns to Axes of Innovation, Across Revolutions

Axis of Innovation	Industrial Revolution			
	First (1760s to mid-1800s)	***Second*** (1870s to about WWI)	***Third*** (1960s to early 2000s)	***Fourth*** (began ~ early 2010s)
Control	Mechanization	Standardization	Automation	Autonomy
Integration	Networked Transportation	Electrification	Decentralization	Hyperconnectivity
Reconfigurability	Fixed-Purpose Machines	Heavy-Duty Construction	Reprogrammability	Adaptability
Scale	Factory-Based Production	Mass-Production Systems	Digitalization	On-Demand
Sustainability Impact	Intensified Throughput	Resource Exploitation	Resource Depletion	Renewability

Source: Rook et al. [2017]

We posit that these constructs - Innovation Axes and Innovation Patterns - apply not only to industrial revolutions, but also to all episodes of technological disruption in general (noting that industrial revolutions are simply instances of TD at the largest scale). As such, Innovation Axes offer a self-consistent, reusable lens through which LTIs can study TD. For example, LTIs might analyze performance gaps along these Axes between incumbents and new entrants in an industry that is facing emergence of a new technology.

In summary, Innovation Axes offer a coherent template with which LTIs can analyze the ways in which disruptive technologies are likely to transform industries. The companies (or other investable assets) that are most susceptible to these modes of change are those that LTIs should focus on in subsequent analysis and evaluation, for the sake of TDM - and the DARLing approach seems well suited to such further assessment.

4.2. Resilience and DARLing

The Resilience paradigm was introduced by Rook et al. [2021] as a vantage point to help investors deal with uncertain (and even ambiguous) shock events.[34] The Resilience view emphasizes the need to explore not only the initial impact of shocks on financial markets (both at the asset and market levels), but also the trajectory of recovery. Many other risk paradigms in finance (e.g., value-at-risk, maximum drawdown) concentrate on the initial consequences of a shock, without considering longer-term ramifications. To illustrate the importance of taking recovery into account, imagine two assets, A and B, that experience a shock: their prices fall immediately by 10% and 15% respectively; but B recovers twice as fast as A. Not accounting for recovery may prompt a knee-jerk conclusion that B is the riskier asset, even though this is not entirely evident. The Resilience paradigm was designed to handle these sorts of analytical complications.

Technological disruption is a form of shock, which is one reason why the Resilience view is well-suited to TDM. Another reason is that Resilience is rooted in signal detection. As was suggested in Section 2, handling of non-traditional market risk factors relies on being able to *detect* the onset of (or else changes to) such factors.[35] In addition to the Resilience concept, Rook et al. [2021] created the DARLing template for analyzing the Resilience of companies and assets, which questions the company or asset's abilities in:

- *Detecting* relevant shock events from data signals
- *Absorbing* the initial impacts of relevant shocks
- *Recovering* from those initial impacts[36]
- *Learning* from its experiences in the wake of a shock

[34] We follow earlier practice and capitalize the type of "Resilience" referred to by Rook et al. [2021], in order to differentiate it from usage of the term "resilience" elsewhere in financial practice (which, as the authors note, is mostly a confusion of resilience and robustness). See that paper for a thorough justification.

[35] Detection is less fundamental for most traditional market risk factors, given the highly standardized data pipelines for such factors. For instance, few investors (apart from some very niche or well-resourced ones) spend much effort on 'detecting' inflation: it simply is not cost-efficient to do so given the abundance of low-latency data on inflation.

[36] Recovery here is a nuanced concept: it does not necessarily mean an asset or company returning to the exact same state that it was in prior to the shock (e.g., sales volume or capital expenditure). In some cases, recovery might be the achievement of some 'equally good' state that differs from the pre-shock one, and is on a separate trajectory (in terms of competitive strategy, resources, sustainability, etc.).

Our suggestion here is that, for purposes of TDM, investors should apply the DARLing template not just to individual companies and assets, but also to their whole investment portfolios.[37] Specifically:

- In analysis of TD risk (or opportunity), LTIs should pay attention to how (and how well) relevant assets, companies, and industries absorb and recover from the effects of disruption caused by particular technologies. For instance, an incumbent company that suffers from early difficulties in integrating a new technology but ultimately excels in using it is far different from an incumbent that experiences early difficulties from which it never recovers. These abilities in absorption and recovery at lower levels (company or industry) can then be aggregated up to the portfolio level for a more encompassing perspective of TD risk (or opportunity).

- For long-term investors, evaluation of feasible investment decisions - as responses to TD - should consider the full time-course of a TD event, which means viewing investment choices in light of the capacity of different components of the portfolio to absorb, recover, and learn from a TD event - and also the LTI's own organizational ability in learning and adapting to evolving risks and opportunities.

It is worth noting that any portfolio- or asset-level decisions a LTI makes must account for not only TD-related risks, but other forms of risk as well. For example, a particular asset might be highly susceptible to disruption from a particular technology, but that asset might serve as a strong buffer against other forms of risk; any decision about the level of exposure to maintain in that asset must therefore balance TD and other forms of risk. We believe that this necessity is another reason why DARLing (and the Resilience paradigm in general) is an appropriate tool for TDM, and risk management at large: DARLing can be applied to all forms and sources of risk, not just TD. Therefore, it is a mechanism that can

[37] We think that it is crucial for LTIs to apply DARLing analysis not only to their portfolios (and assets therein), but also to their own organizations. That is, they should assess their own capacity to detect, absorb, recover, and learn from TD events, and what internal resources they have (or need) to improve those capacities.

promote coherent investment decisions overall - not merely coherence in relation to technological disruption.

Readers will have noticed that the approach to TDM we describe here does not fully flesh out the 'treatment' step advised by ISO 31000:2018. The reason for this omission is that all LTIs are different, in terms of their portfolios, resources, and constraints; therefore, the best investment choices in response to TD, and paths for implementing those choices, can vary (substantially) from one LTI to the next. In future work, we hope to construct a clear mapping of general investor 'types' to archetypal TDM strategies; but that effort exceeds the scope of this paper. For now, we restrict our attention to a few key considerations for LTIs in implementing TDM - which is the focus of the next section.

5. Special Considerations in TDM

In selecting and executing a particular TDM strategy, every LTI should take account of its own unique *context* - that is, its resource set and constraints in building and managing its investment portfolio. Failure to do so can result in an ill-fitting TDM approach that will have a low probability of achieving its goals. The extent to which any single resource or constraint matters to a TDM strategy will vary from one LTI to the next. However, there are a few considerations that we think deserve special attention from LTIs when they are designing and implementing TDM strategies. These considerations relate to: 1) what is needed to undertake a bottom-up approach to TDM; 2) accounting for the variable speeds at which TD can occur; and 3) dealing with valuations, including asset price bubbles. In what follows, we provide some succinct guidance on these considerations.

5.1. Bottom-Up TDM

Previous sections have implied that performant TDM is 'bottom-up' in nature: the impact of TD must be explored at the asset level, and these impacts aggregated up to understand portfolio-level effects; TDM strategies can then be rooted in controlling these aggregated effects. This bottom-up approach necessarily relies on having suitable information and/or data. But, as noted in Section 2, conventional financial data will usually be deficient for thorough TD-related analysis. This means that 'alternative' data (alt-data) will be needed

to conduct effective TDM. Monk et al. [2019] uncover insights on how LTIs can begin to build capabilities around alt-data.

To properly use such data, LTIs will (in most cases) need some capabilities in simulation and stress-testing (S&ST). In some instances, standard, off-the-shelf S&ST tools may be applicable. In other situations, more bespoke tools may be needed. While some LTIs may have the ability to build customized S&ST tools themselves, others will need to outsource such efforts by, e.g., relying on external asset managers or startup companies that are able to produce tailored software. Alternatively, LTIs may find it beneficial to collaboratively develop purposed S&ST tools - ones that fit their exact (collective) needs, and for which they do not need to pay fees to external parties, and over which they maintain full control. Additionally, LTIs may need to rethink conventional industry and sectoral classification systems, in light of TD. Many LTIs' allocation strategies for public equities rely (to some extent) on sector-based allocations. However, sectoral classifications may not accurately represent the technological underpinnings of companies - that is, they can easily result in groupings that are too heterogeneous within groups, and too homogeneous across groups (which is a hallmark of a poor classification system). While LTIs may not be benefitted by discarding conventional classification systems entirely, they should feel compelled to deeply question whether those systems capture asset similarity and risk as effectively as they purport to do.

5.2. Speed Variability in Tech Disruption

We have noted earlier that different technological disruptions proceed at different paces. Some are gradual, and their impacts are 'smoother', whereas others are rapid and result in 'discontinuities' (e.g., they may cause sudden bankruptcies of incumbents, and drastic 'overnight' losses/gains in market valuation). Moreover, disruption can occur at different speeds across different types of company (e.g., across different conventional industry and sector groupings). The most appropriate strategy for managing TD risks and opportunities can vary according to the pace of the disruption in question (as well as level of certainty about that pace). For investors, a key consideration about these pace differentials lies in their effects on discounting.

Relatedly, investors should also be on the lookout for disruptions that are susceptible to exogenous 'accelerants', a recent example of which is the role played by COVID-19 in boosting the popularity of remote-working technologies, as well as further degrading the market shares of many brick-and-mortar retailers.

5.3. Valuation of Disruptions and Pricing Bubbles

For investors, the end aim of any effective TDM must be to act in advance of the market valuation of a TD. The non-uniformity of rates at which disruption occurs (and the non-traditional means and data used to analyze it) means that many assets may take time to be properly discounted to reflect TD risk. This can present an opportunity for LTIs to act if their TDM is effective. For example, to buy an asset where the profit opportunities or the potential loss from disruption are not yet fully reflected in market prices.

But there is a flipside to this matter, and it relates to pricing bubbles. Whilst it can take time for a TD to be fully discounted, overshooting is also likely for some instances of disruption, where investor enthusiasm for a new technology can become sizably disconnected from that technology's actual disruptive potential, resulting in a pricing bubble that will eventually burst (as was the case with the Dot-Com crisis of the early 2000s), or slowly deflate (as was the case with the stocks of several technology companies in the 1980s and 2010s, due to hype surrounding advances in artificial intelligence that ultimately failed to meet investors' expectations). Regardless of whether a LTI's TDM strategy is more tuned toward offense or defense, protecting portfolios against TD-related bubbles should be a core concern.

6. Conclusion

The emergence of generative artificial intelligence heralds a transformational shift in the investment landscape, introducing unparalleled opportunities and unprecedented challenges. And yet, the current wave of uncertainty (and, indeed, fear) among investors over the possible consequences of technological disruption is – paradoxically - normal. In fact, the cyclical nature of technological disruption (TD), evidenced by historical patterns of speculative bubbles and transformative gains, underscores how LTIs are continuously

having to adapt to a changing technological landscape. And this adaptation is not only about managing risks but also about seizing opportunities that arise from disruptive innovations, as evidenced by the transformative impact of mobile, cloud, and AI technologies on markets and investment portfolios. In short, all LTIs should have methodologies for addressing TD risks, because TD is an inherent aspect of modern existence, manifesting routinely at different magnitudes, horizons, and even velocities.

However, as we show in this paper, many LTIs lack a structured method for TD management (TDM), despite recognizing TD as a significant hazard. Through 20 case studies around the world, we observed that LTIs' TD strategies are unstructured, disjointed, and heterogenous. Many LTIs exhibit a fragmented understanding of TD's impacts on assets and markets and lack a unified strategy for organizational adaptation to TD challenges. Most LTIs rely on external managers for sourcing TD-related data and insights, and they hope that TD risks will be addressed "osmotically" thanks to networks, overseas offices, and relationships.

We argue that addressing TD requires a comprehensive approach that not only encompasses portfolio adjustments but also involves strategic organizational realignment towards innovation and agility. Unlike traditional market risks, the timeframe for TD is often longer and more uncertain, its impacts are unevenly distributed, and reliance on historical data is less effective. This necessitates a proactive management stance – one that focuses on developing capabilities to identify, plan for, defend against, and capitalize on technological disruptions. Specifically, we advocate a structured approach to TDM that prioritizes driving *coherent*: organization-level perspectives on TD analysis; and portfolio-level responses. While we are firm believers in this approach, it is certainly not the only one that might be sensible; however, it is the only one that we know to exist. Accordingly, the entire field of TDM is an area that is ripe for (and duly deserves) future research.

REFERENCES

Christensen, C. [1997] *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business School Press.

Clark, G. [1998] "Stylized Facts and Close Dialogue: Methodology in Economic Geography," *Annals of the Association of American Geographers*, 88(1):73-87.

Clark, G., and R. Urwin [2008] "Best-Practice Pension Fund Governance," *Journal of Asset Management*, 9(1):2-21.

Danneels, E. [2004] "Disruptive Technology Reconsidered: A Critique and Research Agenda," *Journal of Product Innovation Management*, 21(4): 246-258.

Davies, P. [2001] "Spies as Informants: Triangulation and Interpretation of Elite Interview Data in the Study of Intelligence and Security Services," *Politics*, 21(1):73-80.

Granovetter, M. [1973] "The Strength of Weak Ties," *American Journal of Sociology*, 78(6):1360-1380.

Granovetter, M. [1983] "The Strength of Weak Ties: A Network Theory Revisited," *Sociological Theory*, 1:201-233.

Markides, C. [2005] "Disruptive Innovation: In Need of Better Theory," *Journal of Product Innovation Management*, 23(1):19-25.

Monk, A., M. Prins, and D. Rook [2019] "Rethinking Alternative Data in Institutional Investment," *Journal of Financial Data Science*, 1(1):14-31.

Monk, A., and D. Rook [2020] *The Technologized Investor: Innovation through Reorientation*, Stanford University Press.

Monk, A., and D. Rook [2023] “Resilience as an Analytical Filter for ESG Data,” in H. Bril, G. Kell, and A. Rasche (eds.) *Sustainability, Technology, and Finance: Rethinking How Markets Integrate ESG*, Routledge.

Rook, D., S-Y. In, and A. Monk [2021] “The Resilience Manifesto: Integrating ESG and Risk Management, in Portfolios and Organizations,” *Social Science Research Network*, available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3942909.

Rook, D., and A. Monk [forthcoming] “Resilient Portfolio Theory,” pre-publication.

Rook, D., A. Salvatori, J. van Moyland, and P. Rosa [2017] “Innovation Patterns: Upgrading Sectoral Classification for the Fourth Industrial Revolution,” *Social Science Research Network*, available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2928335.

Strauss, A., and J. Corbin [1998] *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*, Sage.

Teece, D., G. Pisano, and A. Shuen [1997] “Dynamic Capabilities and Strategic Management,” *Strategic Management Journal*, 18(7):509-533.