

Directed technological change and general purpose technologies: can AI accelerate clean energy innovation?

Pia Andres, Eugenie Dugoua and Marion Dumas

September 2022

Centre for Climate Change Economics
and Policy Working Paper No. 403
ISSN 2515-5709 (Online)

Grantham Research Institute on
Climate Change and the Environment
Working Paper No. 378
ISSN 2515-5717 (Online)

The Centre for Climate Change Economics and Policy (CCCEP) was established by the University of Leeds and the London School of Economics and Political Science in 2008 to advance public and private action on climate change through innovative, rigorous research. The Centre is funded by the UK Economic and Social Research Council. Its third phase started in October 2018 with seven projects:

1. Low-carbon, climate-resilient cities
2. Sustainable infrastructure finance
3. Low-carbon industrial strategies in challenging contexts
4. Integrating climate and development policies for 'climate compatible development'
5. Competitiveness in the low-carbon economy
6. Incentives for behaviour change
7. Climate information for adaptation

More information about CCCEP is available at www.cccep.ac.uk

The Grantham Research Institute on Climate Change and the Environment was established by the London School of Economics and Political Science in 2008 to bring together international expertise on economics, finance, geography, the environment, international development and political economy to create a world-leading centre for policy-relevant research and training. The Institute is funded by the Grantham Foundation for the Protection of the Environment and a number of other sources. It has 13 broad research areas:

1. Biodiversity
2. Climate change adaptation and resilience
3. Climate change governance, legislation and litigation
4. Climate, health and environment
5. Environmental behaviour
6. Environmental economic theory
7. Environmental policy evaluation
8. International climate politics
9. Science and impacts of climate change
10. Sustainable public and private finance
11. Sustainable natural resources
12. Transition to zero emissions growth
13. UK national and local climate policies

More information about the Grantham Research Institute is available at www.lse.ac.uk/GranthamInstitute

Suggested citation:

Andres, Dugoua and Dumas (2022) *Directed technological change and general purpose technologies: can AI accelerate clean energy innovation?* Centre for Climate Change Economics and Policy Working Paper 403/Grantham Research Institute on Climate Change and the Environment Working Paper 378. London: London School of Economics and Political Science

Directed Technological Change and General Purpose Technologies: Can AI Accelerate Clean Energy Innovation?*

Pia Andres[†]

Eugenie Dugoua[‡]

Marion Dumas[§]

September 14, 2022

Abstract

Transitioning away from dirty and towards clean technologies is critical to reduce carbon emissions, but the race between clean and dirty technologies is taking place against the backdrop of improvements in general-purpose technologies (GPT) such as information and communication technologies (ICT) and artificial intelligence (AI). We show how, in theory, a GPT can affect the direction of technological change and, in particular, the competition between clean and dirty technologies. Second, we use patent data to show that clean technologies absorb more spillovers from AI and ICT than dirty technologies and that energy patenting firms with higher AI knowledge stocks are more likely to absorb AI spillovers for their energy inventions. We conclude that ICT and AI have the potential to accelerate clean energy innovation.

*We thank Joëlle Noailly, Elena Verdolini, Roger Fouquet, Misato Sato and Simona Iammarino for their helpful comments at various stages of the paper's development, as well as the participants of the seminars and workshops where this paper was presented and discussed. We also acknowledge administrative support from the Grantham Research Institute, thereby acknowledging its founders: the Grantham Foundation for the Protection of the Environment and the ESRC through the Centre for Climate Change Economics and Policy.

[†]Department of Geography and Environment, London School of Economics and Oxford Martin School, University of Oxford. Email: p.andres@lse.ac.uk

[‡]Department of Geography and Environment, Center for Economic Performance and Grantham Research Institute, London School of Economics. Email: e.dugoua@lse.ac.uk.

[§]Grantham Research Institute, London School of Economics. Email: m.dumas1@lse.ac.uk.

1 INTRODUCTION

Directing technological change away from polluting technologies and towards cleaner options is central to addressing global environmental problems such as climate change. Prior work has shown that a combination of taxes and research subsidies can effectively level the playing field between clean and dirty technologies (Acemoglu et al. 2012; Aghion et al. 2016). Those policies incentivize the development and adoption of environmentally friendly technologies, which allows clean sectors’ productivity to catch up to their dirty counterparts in the longer term.

However, the race between clean and dirty technologies is taking place against a backdrop of improvements in information and communication technologies (ICT) and artificial intelligence (AI). Some highlight the positive impact those technological developments may have in helping solve environmental problems.¹ But are low-carbon technologies surfing the AI wave better than dirty technologies? AI and ICT, in some respect, resemble the textbook case of general-purpose technologies in that they have the potential to be applied in many, if not most, areas of the economy, including in high-carbon energy industries (Brynjolfsson et al. 2021; Crafts 2021; Trajtenberg 2018). Thus, *a priori*, there is no reason to believe that they can drive the low-carbon transition, as they may just as well help incumbent technologies continue to gain productivity.

This paper investigates how a new general-purpose technology (GPT) affects the direction of technological change and, in particular, the competition between clean and dirty technologies. We do so, first, theoretically and then empirically by examining the extent to which energy patents rely on AI and ICT inventions. In line with the literature on directed technological change and the environment, we consider low-carbon electricity and transport to be competing in a race with the incumbent fossil fuel-based technologies, where the latter have an advantage due to their greater maturity (i.e., in the absence of corrective policy, they will attract more talent and R&D resources). But we recast this race as happening against the backdrop of advances in AI.

Our theory shows that the arrival of a GPT opens new opportunities to shift to a clean technology equilibrium because it disrupts the path dependence mechanisms that otherwise entrench dirty incumbent technologies. In addition, the shift to the clean equilibrium is made easier if clean technologies have a higher capacity to absorb the GPT than dirty technologies. The absorptive capacity of a technology is shaped both by characteristics intrinsic to the technology and previous exposure to the GPT: both can make it easier to apply the GPT in that particular technological field.

We then study the absorptive capacities of clean and dirty technologies empirically. To do so, we analyze citations between energy patents and AI (or ICT) patents and show that clean energy technologies absorb digital technologies much more than dirty energy technologies do.

1. See, for example, Rolnick et al. (2019) or private sector initiatives such as Microsoft’s “AI for the Planet”.

This is true both across and within individual firms’ patent portfolios. We interpret this as an indication that the differences between clean and dirty technologies arise both from firm-level capacities and characteristics intrinsic to the technologies (i.e. technological reasons why there is more potential to apply AI and ICT in clean technologies than in dirty ones). At the firm level, we then find that a firm’s stock of knowledge in AI increases the extent to which it applies AI to its energy innovations, and the effect is much stronger for clean technologies. Interestingly, having a lot of prior experience in energy technologies seems to be a barrier to the use of AI, which suggests that new entrants to clean transport and electricity who have strong AI capabilities are critical to accelerating the diffusion of AI into low-carbon technologies.

In summary, this paper argues on theoretical grounds that it is critical for the low-carbon transition that clean technologies be more successful in “riding the AI wave” (i.e. applying the GPT) than dirty ones. Empirically, we find early evidence that this is the case, both because these technologies are intrinsically more able to use AI and because this, in turn, encourages firms with AI knowledge to invest in those technologies. However, compared to other technological fields, the rate at which AI is entering clean transport and electricity technologies remains low compared to other areas, such as medical technologies or telecommunications. This suggests that there are good reasons for innovation policy to deliberately target applications of AI (and digital technologies more broadly) to clean technologies.

This paper contributes to both the theoretical and empirical literatures on directed technological change and the environment (Acemoglu et al. 2012; Aghion et al. 2016; Dechezleprêtre et al. 2017; Johnstone et al. 2010; Popp et al. 2020). This literature tends to analyze environmentally beneficial innovations in isolation from other technological developments. Here, we extend it by studying the interaction with a general-purpose technology. In doing so, we also contribute to the economic literature on GPTs (Helpman et al. 1994, 1996; Lipsey et al. 2005; Rosenberg et al. 2010). This literature is mainly concerned with understanding the contribution of GPTs to growth and has not investigated how GPTs can modify the direction of technological change (except for the literature on digital technologies and skill-biased technical change).

Economic history, however, has provided detailed accounts of how specific GPTs have created new technological eras by reconfiguring technological systems, creating new complementarities between technologies, and between new technologies and infrastructure, production methods, lifestyles and consumption habits (Dosi 1982; Fouquet 2008; Grübler et al. 1999; Perez 2009; Rosenberg 1979). This qualitative strand of literature is complemented by recent empirical work that aims to quantify technological interdependencies, for the most part using patent data (Acemoglu et al. 2016; Napolitano et al. 2018; Pichler et al. 2020). These papers find that the patterns of technological interdependencies predict future rates of innovation. This underscores the importance of understanding the complementarity between clean innovation and other fast-improving fields of innovation.

The remainder of this paper proceeds as follows. Section 2 provides background on general-purpose technologies, in particular their role in economic transformations, and on ICT and AI

technologies with a focus on their potential applications to the low-carbon transition. Section 3 analyzes a model of green directed technological change in which we add a GPT. Section 4 describes the construction of our global dataset of 2,545,063 electricity and transport patent families and the extent to which they have absorbed AI and ICT knowledge. Section 5 presents our key result about clean technologies’ greater ability to absorb the GPT as compared to dirty technologies. Section 6 presents the results of the firm-level analysis, while section 7 discusses the implications of our results for the low-carbon transition.

2 BACKGROUND

Artificial Intelligence as the next General Purpose Technology Artificial Intelligence (AI) – defined by Miriam-Webster as “the capability of a machine to imitate intelligent human behaviour” – is widely thought to be the next game-changing technology about to unleash large productivity gains and a wave of automation by optimists and pessimists alike (Trajtenberg 2018). AI includes several techniques and functional applications in computer science, such as deep learning, symbolic systems and reasoning, speech processing, and computer vision, all of which are key to advancing optimization, prediction and robotics, which can be deployed in many sectors. According to Cockburn et al. (2018), deep learning has the potential to change the research process itself, thus qualifying as the “invention of a method of invention”. There is, therefore, significant evidence that AI qualifies as a general-purpose technology, and an emergent literature aspires to model its potential effects on growth and knowledge creation. For example, Aghion et al. (2018) model AI as a process of automation of goods and services, as well as the production of ideas. Agrawal et al. (2018) integrate AI breakthroughs into a knowledge production function as enabling faster discoveries in combinatorial knowledge creation.

Applications of AI in Energy Sectors Some ICT and AI technologies may have applications essential for the transition to clean energy. For example, smart grids facilitate the integration of distributed renewable energy with bulk power generation plants and bulk energy storage systems (Bose 2017), and smart buildings can benefit from effective load demand forecasting (Raza et al. 2015) and better monitoring through smart meters (Fouquet 2017). AI techniques can also be used to plan, optimize, and manage renewable energy technologies, including solar and wind systems and hydro power (Jha et al. 2017). For example, fuzzy logic controllers can adjust turbine speeds to optimize aerodynamic efficiency and extract maximum power, while neural networks can carry out automatic performance checks (Bose 2017). Lee (2020) has analyzed patent citations and found that AI has contributed to improving battery performance and optimizing cars’ energy management systems and charging systems.

Potential applications of AI in the energy sector are not limited to clean technologies. AI can enhance productivity in many application sectors by automating some tasks and freeing up labor to complete other, more complex ones. It is also valuable for planning the maintenance

and deployment of physical capital or inventories. More broadly, and not specific to clean or dirty energy, Lyu et al. (2021) analyze online job postings data from 2010-2019, and find that among emerging digital technologies (among which they include Artificial Intelligence, Big data, Internet of Things, Robotics, Blockchain technology, and Cloud Computing), AI is the most widely applied in the energy sector (as measured by the extent to which new hires are asked to provide expertise in AI). AI-related knowledge also carries the highest wage premium compared to average wages and contributes most to energy firms' performance. Crucially, there are numerous potential applications for AI not just in clean but also in dirty energy. For example, AI can increase the efficiency of fossil fuel exploration (such as through well logging or geological mapping), field development and engineering, and other parts of the value chain (cf. Koroteev et al. 2021). In combustion technologies, AI can be used to monitor and optimize combustion processes. Thus, AI could accelerate innovation in clean technologies, but given its wide range of applications, it could also help the productivity of dirty technologies and prolong their attractiveness.

The economics of GPTs Our analysis is informed by several key contributions from the economics literature on GPTs and innovation spillovers. First, as emphasized by Helpman et al. (1994) and Helpman et al. (1996), the economic benefits from a new GPT may accrue only after a lag because advances in the GPT do not diffuse spontaneously: adoption requires complementary co-invention in application sectors typically happening via R&D investments. Helpman et al. (1996) model the diffusion of a new GPT, allowing for both early and late adopters, and the extent to which application sectors innovate to make use of the GPT depends on four key factors: their capacity to absorb the GPT (that is to learn from it to create large productivity gains in their sector); their market size; the historical stock of components developed for the old GPT; and the cost of developing new components. Our theoretical analysis will build on those factors.

We also follow Cohen et al. (1990) in considering absorptive capacity to be endogenous, meaning that it is the result of deliberate investments in an area of knowledge to be better able to learn from other inventors and inventions (thus, knowledge spillovers are not “free” or spontaneous).

Finally, while prior literature mainly focused on the effects of GPTs on growth, we examine how GPTs shape the race between two competing technologies and may catalyze the creative destruction of a (dirty) incumbent technology by a newer (clean) challenger. Indeed, in modeling the diffusion of a GPT, Helpman et al. (1996), for example, assume that all viable application sectors will eventually adopt the GPT. In the race between clean and dirty, however, enhancing welfare requires that the dirty sector declines and disappears.

3 THEORY

How should we expect a GPT to affect the direction of technological change? Specifically, under what conditions can a GPT accelerate the pace of innovation more in dirty rather than clean technologies? We build on the seminal model of directed technological change and the environment put forth by Acemoglu et al. (2012) by adding a general-purpose technology and letting clean and dirty sectors have potentially differing capacities to absorb this GPT.

We first consider the case where absorptive capacity is entirely exogenous, and then we partially endogenize it by allowing firms or scientists to invest in it. In both cases, we solve for the equilibrium level of innovation in the clean and dirty sectors. Endogenizing absorptive capacity also yields comparative statics that we use as hypotheses to explain the observed empirical variation in the extent to which different technologies and firms draw on the GPT.

Baseline model

Let there be an aggregate final good competitively produced from the combination of dirty and clean inputs Y_d and Y_c (e.g., energy source or material):

$$Y = (Y_{ct}^{\frac{\varepsilon-1}{\varepsilon}} + Y_{dt}^{\frac{\varepsilon-1}{\varepsilon}})^{\frac{\varepsilon}{1-\varepsilon}} \quad (1)$$

We assume that clean and dirty inputs are highly substitutable ($\varepsilon > 1$)². Sector $j \in \{c, d\}$ produces input Y_j competitively using a combination of labor and sector-specific machines:

$$Y_{jt} = L_{jt}^{1-\alpha} \int_0^1 A_{jit}^{1-\alpha} x_{jit}^\alpha di \quad (2)$$

For example, if the input is electricity and $j = c$, the machines may be wind turbines and solar panels, and for $j = d$, gas-fired power plants. The machines form a continuum; machine i has productivity A_{jit} and is consumed by the intermediate producer of input Y_j in quantity x_{jit} .

Meanwhile, scientists choose whether to work on clean or dirty technology. Having made this choice, each scientist is randomly allocated to a single machine in the sector of choice. In the standard model, scientists successfully innovate on machine i with probability η_j . If successful, the machine's productivity gets an incremental increase, denoted γ . Formally:

$$A_{jit} = (1 + \gamma)A_{jit} \quad (3)$$

The scientist then obtains a one-period patent and becomes the monopolistic producer of that machine for that period (producing each machine at a cost of τ units of the final good).

2. This is a key assumption in Acemoglu et al. (2012), which is arguably plausible in the sectors we analyze. Electricity from renewable energy sources can be used in much the same way as electricity from a coal power plant, just as an electric vehicle is a good substitute for an internal combustion engine powered one. For a more detailed discussion of this assumption and its justification, please refer to Acemoglu et al. (2012), page 135, footnote 6.

Adding Spillovers from a GPT

We modify the dynamic equation governing the change in productivity of machines (Equation 3) by introducing a stock of knowledge in a GPT (GPT_t) and an exogenous absorptive capacity β_j for scientists working on technologies of sector j . Here, we consider that spillovers from the GPT increase the *value* of an innovation by boosting the machines' productivity.³ Formally, we write:

$$A_{jit} = (1 + \gamma + \beta_j GPT_t) A_{jit} \quad (4)$$

This modeling choice is supported by Table 4 in Section 5, which shows that the value of an energy patent (as measured by the citations it receives) is greater for those patents that draw on the GPT. Equation 4 also implies that the spillovers from the GPT depend only on absorptive capacity β_j , and is therefore the same for any machine within the sector j .

Next, we focus on characterising the profitability of research in each sector to understand how the GPT affects the direction of technological change.⁴ The average productivity of sector j is:

$$A_{jt} = \int_0^1 A_{jit} di \quad (5)$$

It evolves over time according to the following equation:

$$A_{jt} = (1 + (\gamma + \beta_j GPT_t) \eta_j s_j) A_{j,t-1}, \quad (6)$$

where s_j is the share of scientists who choose to work in sector j (where market clearing of R&D labor requires $s_c + s_d = 1$). The equilibrium profits of a producer of machine with productivity A_{jit} is:

$$\pi_{jit} = (1 - \alpha) \alpha p_{jt}^{\frac{1}{1-\alpha}} L_{jt} A_{jit} \quad (7)$$

Ex-ante, the expected profit from choosing to work in sector j is:

$$\Pi_{jt} = \eta_j (1 + \gamma + \beta_j GPT_t) (1 - \alpha) \alpha p_{jt}^{\frac{1}{1-\alpha}} L_{jt} A_{j,t-1} \quad (8)$$

Solving for equilibrium values of p_{jt} and L_{jt} , and substituting, we obtain the following ratio of R&D profits for working in the clean versus dirty sector:

$$\frac{\Pi_{ct}}{\Pi_{dt}} \equiv f(s_c, s_d) = \frac{\eta_c}{\eta_d} \frac{1 + \gamma + \beta_c GPT_t}{1 + \gamma + \beta_d GPT_t} \left(\frac{1 + (\gamma + \beta_c GPT_t) \eta_c s_c}{1 + (\gamma + \beta_d GPT_t) \eta_d s_d} \right)^{-\phi-1} \left(\frac{A_{ct-1}}{A_{dt-1}} \right)^{-\phi} \quad (9)$$

Equation 9 allows us to study how the GPT affects the direction of technological change.

3. Alternatively, we could model the idea that spillovers from the GPT increase the rate of innovation (as in the notion that AI may accelerate discovery of solutions), such that $\eta_j \propto \beta_j GPT_t$. But this does not change the results of our analysis.

4. Appendix A provides the step by step derivation of the equilibrium equations.

If $f(1,0) > 1$, then $(s_c = 1, s_d = 0)$ is an equilibrium, and technological change is directed towards the clean sector. If $f(0,1) < 1$, then $(s_c = 0, s_d = 1)$ is an equilibrium, and technological change is directed towards the dirty sector. If $f(1,0) > 1$ and $f(0,1) < 1$ simultaneously, then we obtain multiple equilibria, meaning that either the dirty or the clean equilibrium is possible, and some coordination device is required to select one equilibrium.

Let's denote $\bar{A}_{c,t-1}(A_{d,t-1})$ the value of $A_{c,t-1}$ where $f(1,0) = 1$. This is the minimum value that $A_{c,t-1}$ must take, given $A_{d,t-1}$, so that a clean equilibrium becomes possible. Conversely, denote $\bar{A}_{d,t-1}(A_{c,t-1})$ the value of $A_{d,t-1}$ where $f(0,1) = 1$. This is the minimum value that $A_{d,t-1}$ must take, given $A_{c,t-1}$, so that a dirty equilibrium becomes possible. These two functions, depicted in Figure 1, delineate the area in the $(A_{c,t-1}, A_{d,t-1})$ space where we obtain a clean equilibrium, a dirty equilibrium, or multiple equilibria. Result 1 below summarizes the impact of the GPT on the direction of technological change.⁵

Result 1.

- (a) *An increase in GPT_t causes both $\bar{A}_{c,t-1}(A_{d,t-1})$ and $\bar{A}_{d,t-1}(A_{c,t-1})$ to decrease, which means that we obtain multiple equilibria for a wider set of historical states $(A_{c,t-1}, A_{d,t-1})$.*
- (b) *An increase in β_j causes $\bar{A}_{j,t-1}(A_{-j,t-1})$ to decrease and $\bar{A}_{-j,t-1}(A_{j,t-1})$ to increase, thus expanding the range of histories in which all scientists engage in innovation of type j .*

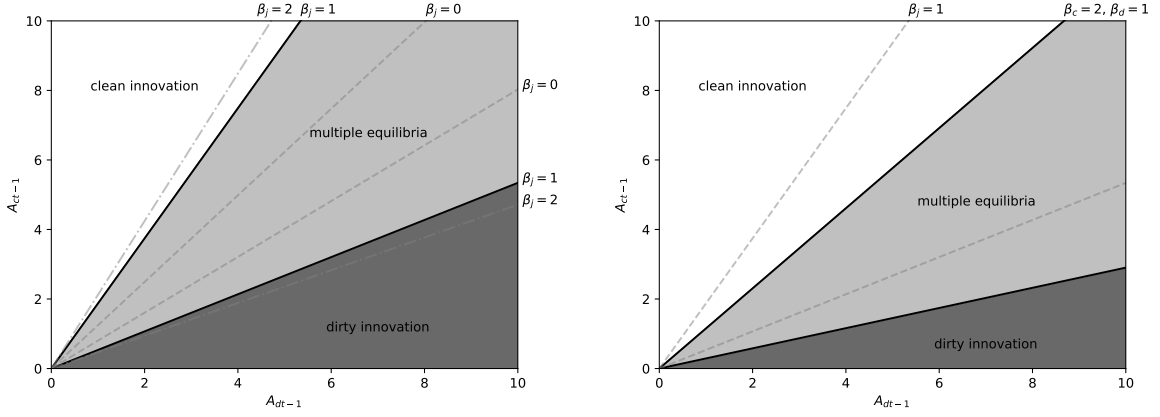
Figure 1 illustrates Result 1. As a baseline, consider $\beta_c = \beta_d = 0$ which corresponds to the case where neither sector can absorb the GPT, making it irrelevant and equivalent to the original model by Acemoglu et al. (2012). In Figure 1a, we see that, in this case, technological change is geared towards the sector that is already the most productive. There is a narrow area, when $A_{c,t-1}$ is close to $A_{d,t-1}$, where multiple equilibria are possible: actors have to coordinate on the clean or the dirty equilibrium. But, for most of the state space, the equilibria are path dependent and reflect what was done in the past. For example, an initial advantage in the dirty sector would lead to a unique equilibrium in which scientists work on dirty innovations.

When the two sectors can both absorb the GPT, as stated in Result 1a), the window of multiple equilibria expands. In other words, thanks to the GPT, the innovation system has more opportunities to break free from the determinism of the past, even when both the incumbent and the challenger technology have the same absorptive capacity. Actors can use the GPT to move either technology sufficiently ahead of the other to make it competitive. The direction of technological change then depends on which technology actors coordinate on.

On Figure 1b, we consider a case when clean and dirty have different absorption capacities to illustrate Result 1b). We see that a higher β_c increases the area where we get multiple equilibria, and, most importantly, shrinks the area where the dirty technology dominates.

Result 1 and Figure 1 highlight what is at stake in studying the GPT's influence on clean and dirty technologies: the GPT can upend the path dependence of technology. In the absence

5. The proof is shown in Appendix A.2.



(a) Clean and dirty have equal absorptive capacity

(b) Clean has higher absorptive capacity

Note: Direction of technological change in equilibrium, that is the allocation of scientists in the clean or dirty sectors, given each sector's past stock of knowledge ($A_{j,t-1}$) and absorptive capacity (β_j).

FIGURE 1

Direction of Technological Change with a GPT

of a GPT, the more mature technology attracts more effort because, being more productive, it has a larger market. Thanks to the GPT, however, the less mature technology can catch up. A GPT can therefore fundamentally change the nature of the race between the newer clean technologies and the more mature dirty technologies. It reduces the weight of the past, by providing an opportunity to coordinate on the new clean equilibrium, especially if the clean technology has a higher absorptive capacity than the dirty.

Endogenizing the Spillovers from a GPT

We now endogenize absorptive capacity, allowing scientists to invest in their capacity to absorb the GPT. This allows us to derive comparative statics for the level of effort in absorbing the GPT. To do this, let's decompose β_j into an exogenous and an endogenous component, such that:

$$\beta_j = b_j B_j, \quad (10)$$

where b_j is exogenous (coming from the characteristics of the technology) and B_j is an endogenous investment in absorption which comes at a cost of ψB_j^2 . Scientists first choose which sector to work on (i.e. clean or dirty), and then decide how much to invest in their capacity to absorb the GPT.

The expected profit from working on technology j is now:

$$\Pi_{jt} = \eta_j(1 + \gamma + b_j B_j GPT_t) \alpha (1 - \alpha) p_{jt}^{1/(1-\alpha)} L_{jt} A_{j,t-1} - \psi B_j^2 \quad (11)$$

Hence, a scientist working in sector j would optimally invest in their absorptive capacity as

follows:

$$B_j^* = (\eta_j b_j GPT_t) \frac{(1-\alpha)\alpha}{2\psi} p_{jt}^{1/(1-\alpha)} L_{jt} A_{j,t-1} \quad (12)$$

Using Equation 12 in combination with the other equations that characterize the equilibrium, we obtain Result 2 below.⁶

Result 2. *In equilibrium, investments in absorbing the GPT in a given sector increase with the existing accessible stock of the GPT and with the intrinsic absorptive capacity of the application sector:*

$$(a) \frac{dB_j^*}{dGPT_t} > 0$$

$$(b) \frac{dB_j^*}{db_j} > 0$$

$$(c) \text{ At the equilibrium for type } j: \frac{d^2 B_j^*}{dGPT_t db_j} > 0$$

Result 2 tells us that efforts in absorbing the GPT increase with the accessible stock of the GPT and with the intrinsic absorptive capacity of the technology. In other words, the potential for spillovers encourages innovation investments in applying the GPT. We expect the extent of potential spillovers to vary by technology (due to the intrinsic absorptive capacity), but also across firms, regions or innovation systems (due to variation in the stock of the GPT across these social units).

Furthermore, as Result 2c) indicates, there is a positive interaction between intrinsic absorptive capacity and the stock of knowledge in the GPT for the technology chosen in equilibrium. In the empirical section, we will bring these comparative statics to firm-level data.

Technological Lock-In

In this part, we consider the role of technological maturity in absorbing the GPT. Specifically, we allow absorptive capacity to decay with the application sector's productivity A_{jt} . We now write the absorptive capacity β_j as a function of A_{jt} :

$$\beta_j = b_j B_j A_{jt-1}^{-\delta}, \quad (13)$$

where $\delta \geq 0$ represents an aging factor. The idea is that more mature technologies are less able to undergo radical changes, or in other words, aging causes lock-in.

Result 3.

$$(a) \frac{dB_j^*}{dA_{jt-1}} < 0 \text{ if } \delta > 1$$

$$(b) \frac{dB_j^*}{dA_{jt-1}} > 0 \text{ if } \delta < 1$$

6. The proof is shown in Appendix A.3.

Result 3 shows that the maturity of the technology in the application sector can impact the endogenous part of absorptive capacity.⁷ If the aging factor is large ($\delta > 1$), then when the technology matures and becomes more productive, fewer investments are made which leads to lower absorptive capacity. On the contrary, when the impact of aging is minimal ($\delta < 1$), an increase in the productivity of technology j leads to more investment and higher absorptive capacity.

4 DATA

Patent Data. Our next steps focus on measuring the extent to which clean and dirty technologies absorb spillovers from AI and ICT. To do so, we use data on patent applications from PATSTAT and obtain a full coverage of patents filed around the world up until 2018.⁸ To avoid double-counting, we aggregate patent applications at the level of DOCDB families, which are groups of patents that have been identified as covering the same invention.⁹ To place patent families over time, we use the priority year, that is the year when the earliest application in the family was filed.

Energy Inventions. We use technology codes from the International Patent Classification (IPC) and from the Cooperative Patent Classification (CPC) to identify inventions related to energy technologies for electricity and transportation.¹⁰ The codes are assigned by patent examiners and are often used to classify patents as either clean, grey or dirty (Acemoglu et al. 2012; Aghion et al. 2016; Dechezleprêtre et al. 2017; Johnstone et al. 2010; Lanzi et al. 2011; OECD 2016; Popp et al. 2020). Table 1 summarizes how we classify technologies. “Dirty” refers to conventional, highly polluting technologies, while “clean” includes the least polluting alternatives. The category “grey” captures increased efficiency of dirty technologies. A full list of the codes used is shown in Online Appendix Table SI1 and SI2.

We keep all energy families with a priority year between 1990 and 2018. For this period, we find a total of 1,674,751 electricity families (809,327 clean, 257,490 grey, 607,934 dirty) and 1,300,651 transport patent families (795,408 clean, 298,645 grey, 206,598 dirty). Figure

7. The proof is shown in Appendix A.4.

8. We use the 2021 Spring edition of PATSTAT. Since there is a delay between when applications are filed and when the data is transferred to the database, the years 2018 onwards are severely truncated.

9. Several patents are typically filed about the same invention because the different applications may cover slightly different claims (about the same invention) or may contain exactly the same claim but are filed in different countries. We include patent families of all sizes (i.e., including size 1) and with patent applications filed in any jurisdictions. This approach allows us to capture global trends in clean and dirty innovation without restrictions on where the invention happened and how many jurisdictions the assignees deemed interesting to file in. Some regressions will only use triadic granted families as a way to narrow down the analysis to potentially more valuable inventions.

10. We use both classifications in order to capture as many relevant families as possible. Using IPC codes is necessary to capture many families from the Chinese, Japanese and Russian patent offices which do not use the CPC. A family is assigned to a category if at least one patent within it has been assigned a relevant technology code.

TABLE 1
Technology Categories

| | Electricity | Transport |
|-------|--|--|
| Clean | Renewable Energy (Wind, Solar, Geothermal, Hydro, Marine), Nuclear Energy, Enabling technologies (e.g., smart grids) | Electric, Hybrid, or Hydrogen vehicles, Fuel cells, Batteries, Enabling technologies (e.g., charging stations) |
| Grey | Efficiency, Biomass and waste | Efficiency of internal combustion engines |
| Dirty | Combustion of traditional fossil fuels (Oil, Natural Gas and Coal), Hydrofracturing | Internal combustion engines |

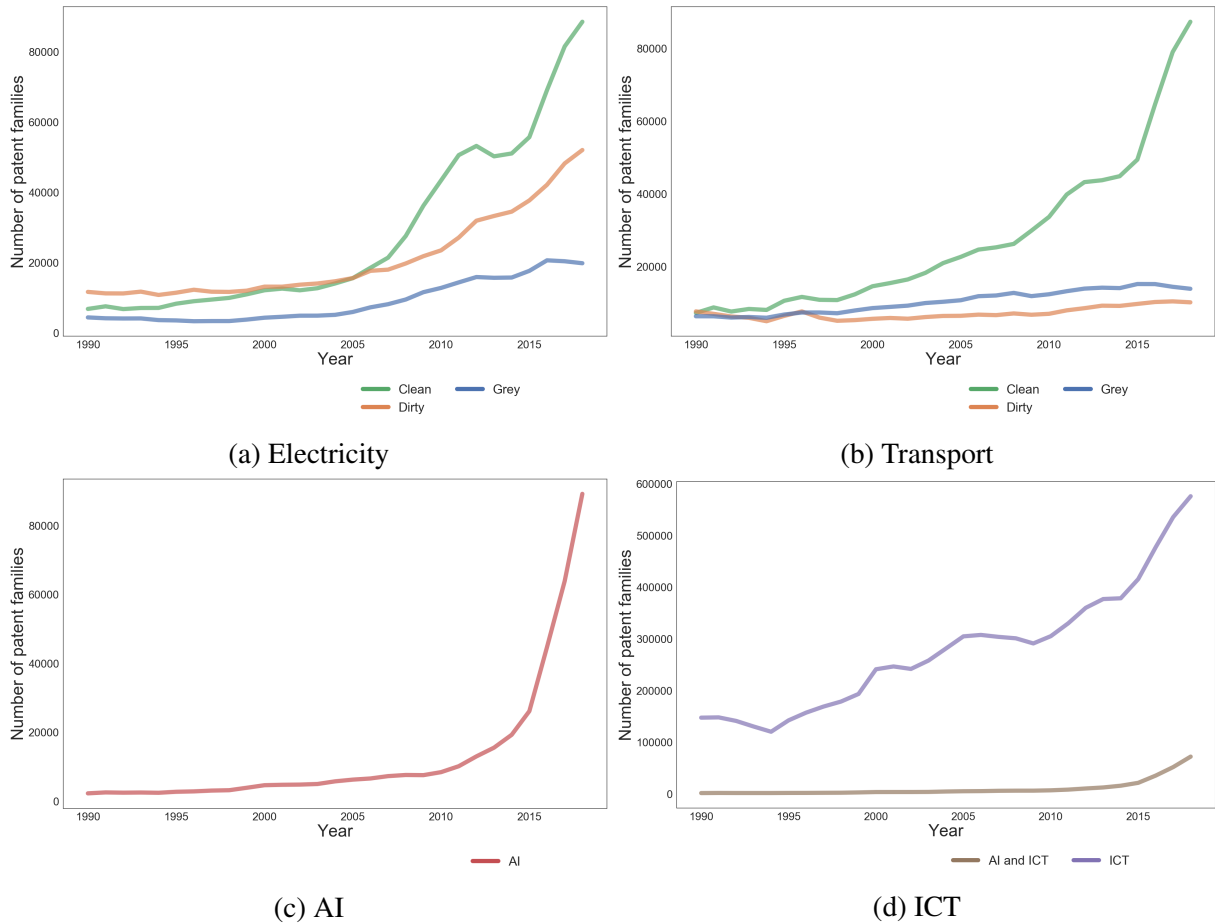
Note: The table shows the technologies we include as clean, grey or dirty electricity and/or transport. We identify patent families related to those technologies based on codes from the Cooperative Patent Classification (CPC) and the International Patent Classification (IPC). We use both classifications in order to capture as many relevant families as possible. Using IPC codes is necessary to capture many families from the Chinese, Japanese and Russian patent offices which do not use the CPC. A family is assigned to a category if at least one patent within it has been assigned a relevant technology code.

2a and 2b show that the number of energy families have been going up both for electricity and transport. In transport, clean vastly outpaces dirty and grey throughout most of the period, while in electricity clean innovation has exceeded dirty and grey since the early 2000s.

AI and ICT Inventions. To identify patents related to AI, we follow the methodology developed by World Intellectual Property Organization (2019) that uses technology codes and keyword searches in abstracts and titles. Keywords include “artificial or computational intelligence”, “neural networks” or “learning model or algorithm”. For ICT, we use a series of technological codes following Inaba et al. (2017). These codes include inventions classified as related to the “transmission of digital information”, “self-organising networks, e.g. ad hoc networks or sensor networks” or “high speed computing”. In the end, this procedure identifies 548,641 AI families and 10,883,849 ICT families. We note that, to this day, the stock of ICT knowledge is vastly greater than that of AI. Figures 2c and 2d show that the number of AI families remains relatively small and has only begun rising sharply since 2010. On the other hand, more than 150,000 ICT families have been filed each year since the early 1990s. We also find that a majority of AI families also qualify as ICT: this implies that, to some degree, AI can be thought of as a sub-field of ICT (see Online Appendix Figure SI1).

Backward Patent Citations. We use backward citations to quantify the extent to which energy inventions rely on AI and ICT. Specifically, as a measure of absorption, we calculate the percentage of backward citations that each energy family makes to AI or ICT patent families.¹¹ In our sample, the average energy family cites about 3.8 patent families with 0.3% going to AI and 4.3% to ICT. This hides considerable variation, however, since some families have 100% of their backward citations going to AI or ICT patents while others cite none. Table 2 provides

11. PATSTAT provides information about citations at the family level, meaning if two patents in the same family cite the same AI patent, that patent counts only as one citation.



Note: The figure plots the total number of patent families over time filed worldwide for each of the following categories: a) electricity, b) transport, c) AI and d) ICT. The year used is the priority year of the family. We note that AI patenting has seen a sharp increase since 2010.

FIGURE 2
Patenting Trends Over Time By Family Type

TABLE 2
Examples of Energy Innovations Citing AI Patents

| Patent Application Title | Sector | Type | Year | Citations to AI | |
|--|------------------------|-------|------|-----------------|----|
| | | | | # | % |
| Improved Flow Valve Port for a Gas Regulator | Electricity | Dirty | 2007 | 49 | 67 |
| Robotic cleaning device | Transport | Clean | 2013 | 297 | 41 |
| Virtual sensor system and method | Transport | Dirty | 2007 | 37 | 26 |
| Battery agnostic provisioning of power | Transport, Electricity | Clean | 2016 | 119 | 13 |
| System and approach for dynamic vehicle speed optimization | | Grey | 2015 | 51 | 10 |
| Dual fuel heater with selector valve | Electricity | Grey | 2011 | 38 | 9 |
| Method and apparatus for configuring a communication interface | Electricity | Clean | 2014 | 55 | 2 |

Note: The table illustrates how AI may be applied to energy technologies by showing examples of energy patent families with a high number of citations to AI.

examples of energy patents with high reliance on AI. The first patent in the table, for instance, corresponds to a dirty electricity family filed in 2007 entitled "Improved Flow Valve Port for a Gas Regulator". The patent makes 49 citations to other patents and 67% of those are citations going to AI families.

Proxies of Patent Quality. We follow prior work by using the number of citations received (a.k.a. forward citations) as a proxy of patent quality (Jaffe et al. 2017; Jaffe et al. 2000). The number of times a particular family is cited by other families, however, heavily depends on the number of years since it was first filed: the older the family, the more opportunities there have been for other families to cite it. It is therefore inappropriate to compare families filed in different years since the younger ones would mechanically have fewer citations. To avoid this problem, our main measure is the number of forward citations received within 3 years.¹² As an additional proxy of patent quality, we also use the number of countries where the patent family was filed as well as the size of the family (i.e., the total number of applications in the family).

Firm-Level Data. We use European Patent Office data obtained from the Bureau Van Dijk Orbis hard-drive to link PATSTAT patent ids to Orbis firm identifiers. We then construct firm-level innovation indicators: for each firm, we count the yearly number of families of different types (e.g., clean electricity or dirty transport). We also construct proxies of firm-level knowledge stocks by calculating cumulative discounted sum of families going back to 1980. We discount stocks by 15% each year following prior work (Hall et al. 2005). Finally, we collect financial and legal data on firms from Orbis. We follow Kalemli-Ozcan et al. (2015) when

12. We also use the number of forward citations received within 5 years as a robustness check where appropriate. Since we use forward citations here to make statements about families relative to other families within a particular time window, the particular time window used should not matter (assuming that there is not much variation in how citations appear over time across families). In any case, we find that citations peak after 4 years, and so, our robustness checks using citations received within 5 years ensure that our measures cover the majority of citations.

cleaning the data; in particular, we use multiple vintages to optimize coverage.¹³ The end result is a dataset of 1,460,034 observations covering 21,046 firms over 1990 to 2018.

5 AI AND ICT ABSORPTION INTO CLEAN AND DIRTY INVENTIONS

This section examines the extent to which energy families have absorbed knowledge spillovers from AI and ICT over the last decades. First, on Figure 3, we plot trends over time in the percentage of backward citations going to AI or ICT for different types of energy families. A key take-away is that, in the transport sector, clean patents build on AI and ICT more than dirty and grey, while, in the electricity sector, grey slightly leads clean. On Figure 3a, we see that, overall, the average percentage of backward citations going to AI is low, typically well below 1%, even though it has been increasing since 2010 which coincides with the rise of AI patenting seen on Figure 2c. We also note that AI absorption is higher in clean than in grey or dirty (especially since 2010) and that it is higher in transport than in electricity.

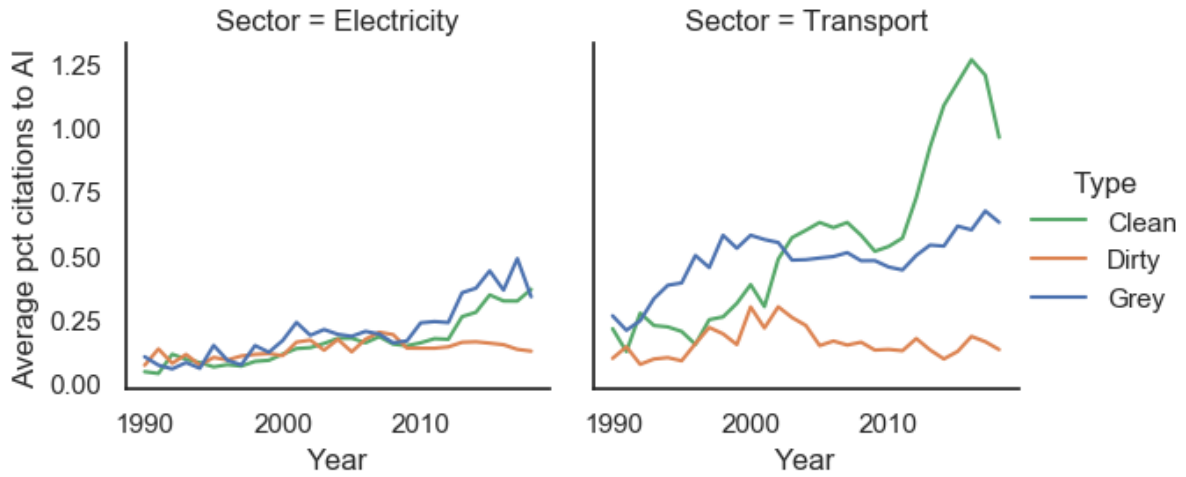
Figure 3b shows that, similar to the case of AI, the percentage of backward citations going to ICT is higher in clean than in grey or in dirty. The magnitude of ICT absorption in clean electricity is particularly high: the average percentage of citations going to ICT reached nearly 20% in the late 2000s, while other technology groups have remained below 10% throughout. We also note that the share of ICT in backward citations is overall much higher than that of AI, but this should not be surprising since ICT is more mature and constitutes a larger pool of potential patent families to be cited.

Next, we run a series of regressions to investigate how the absorption of AI and ICT for clean relative to dirty technologies varies when we include firm fixed effects and quality controls. The main specification is as follows:

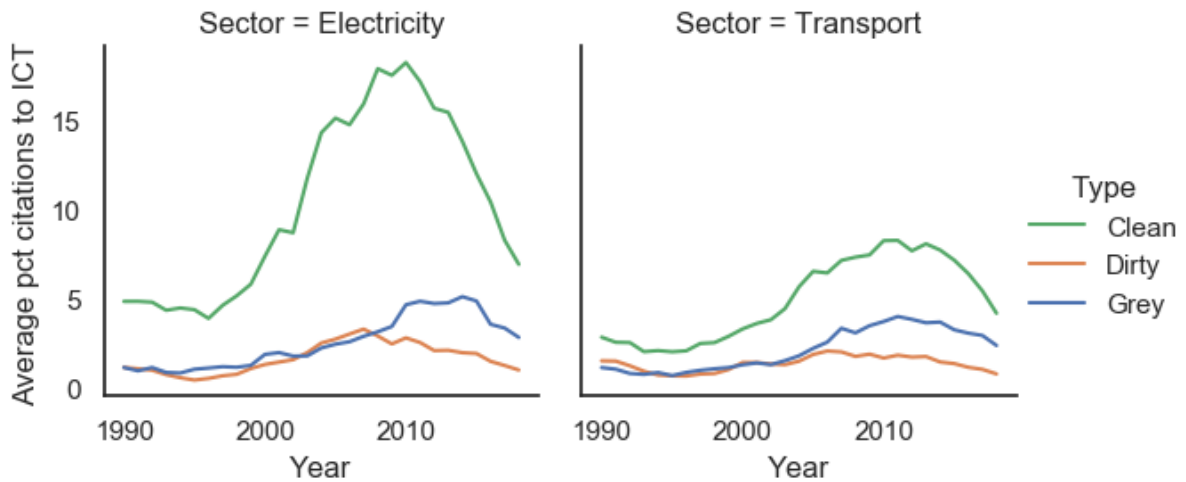
$$Absorption_{ijt} = \beta_0 + \beta_c Clean_i + \beta_g Grey_i + \mathbf{bX}_i + \delta_t + \delta_j + \varepsilon_{ijt} \quad (14)$$

$Absorption_{ijt}$ is the percentage of backward citations going to AI or ICT for patent family i filed by firm j in year t . $Clean_i$ and $Grey_i$ are binary variables that equal 1 if family i is classified as clean or grey, respectively (either in transport or in electricity). β_0 is the intercept. \mathbf{X}_i is a series of variable proxying the quality of family i which includes the number of forward citations received by family i in three first years of its filing, the size of family i and the number of countries where family i was filed. δ_t and δ_j are year and firm fixed effects, respectively. Table 3 presents the regression results. Column 1 to 4 focus on AI, Column 5 to 8 on ICT. Column 1 and 5 show specifications with year fixed effects but without firm fixed effects; this allows us to document the size of the effect in the whole sample of families without any

13. We use the following vintages: 201709, 201812, 201912, 202012, and 202106. Please refer to the Online Appendix A.3 for more details on our data cleaning process.



(a) Citations to AI



(b) Citations to ICT

Note: The figure shows the percentage of citations going to AI (a) and ICT (b) for the average electricity (left) and transport (right) family over time. The year used is the priority year of the family. Since 2010, AI clearly makes up a greater share of backward citations in clean transport families than in grey and dirty families. For electricity the picture is less clear. The share of ICT in backward citations is higher for clean families throughout the period in both electricity and transport.

FIGURE 3
Percentage of Citations to AI and ICT Over Time

controls. As we move from Column 1 to Column 4, we add more restrictions on the sample and on the specifications such as firm fixed effects and quality controls. Whether the coefficients on “Clean” and “Grey” change at all from Column 1 and Column 4 is instructive in understanding what may or not be driving the effect. In particular, showing the difference between results with and without firm fixed effects illustrates that the magnitude of the overall trends can be driven, to a large extent, by differences between firms (rather than within).

Consistent with Figure 3, the coefficients on “Clean” are positive and statistically significant, indicating that clean families rely more on AI and ICT than their dirty counterparts. To allow for an easier interpretation of the main effect captured by β_c , the line “Ratio Clean/Dirty” in Table 3 expresses the magnitude of the effect in percentage term relative to dirty and can be interpreted as the relative absorptive capacity of clean vs dirty. Formally, it corresponds to $\frac{100 \times \beta_c}{mean_d}$, where $mean_d$ is the average percentage of backward citations going to AI (or ICT) in the average dirty family. For example, Column 1 indicates that the absorptive capacity for AI is 304% higher in clean than in dirty. It is 502% higher for ICT (see Column 5).

The relative absorptive capacity may be high for reasons intrinsic to the technologies (e.g., many clean technologies may simply be technologically closer to ICT or AI) or due to general equilibrium effects (e.g., because R&D is being redirected towards clean technologies across the economy). Another reason, however, could be that clean inventions are developed by firms that are better able to leverage AI and ICT technologies into their energy inventions. The high relative absorptive capacity may therefore be driven by firm-level characteristics rather than intrinsic technological differences. To investigate whether firm-level characteristics play a significant role, Column 2 and 6 include firm fixed effects. We find that the ratio changes little for AI but decreases for ICT, highlighting that firms may play a larger role for ICT than AI. The ratio remains high showing that, even within the same firm, clean inventions cite more than 300% as much AI than dirty ones. Section 6 explores the role of firm-level characteristics in more depth.

In Column 3, 4, 7 and 8, we examine whether clean inventions maintain their lead when restricting the analysis to high-quality inventions. To do so, Columns 3 and 7 run the same regressions as Columns 2 and 6 while limiting the sample to triadic patent families that have been granted.¹⁴ Columns 4 and 8 further control for a series of variables proxying for quality (forward citations, family size and number of countries). We find that AI and ICT absorption in clean remain much higher than for dirty in those specifications too.¹⁵ When running regressions separately for transport and electricity families, we find that the AI absorption gap between clean and dirty is stronger in transport than in electricity. The reverse is true for ICT: clean

14. A family is said to be triadic if it was filed at the three main patent authorities: the USPTO, the EPO and the JPO.

15. All coefficients excluded from the main tables are shown in the long version of the same table in Supplementary Online Table SI3. The number of patents in the family and the number of countries in the family are not significant. The number of citations received (within 3 years) is positive and significant at the 10% level for AI and at the 5% level for ICT. They do seem to add some explanatory power to the model since the R squared goes from 0.058 to 0.060 for AI and from 0.441 to 0.445 for ICT.

TABLE 3
Estimating the Absorptive Capacity of Clean, Grey and Dirty Technologies

| | (1) AI | (2) AI | (3) AI | (4) AI | (5) ICT | (6) ICT | (7) ICT | (8) ICT |
|----------------------------|----------------------|---------------------|---------------------|---------------------|-----------------------|---------------------|----------------------|----------------------|
| Clean Family | 0.437*** (0.024) | 0.530** (0.069) | 0.463** (0.077) | 0.420** (0.070) | 8.243*** (0.263) | 7.071** (0.943) | 10.329*** (0.951) | 9.951*** (0.947) |
| Grey Family | 0.264*** (0.001) | 0.040 (0.105) | -0.124 (0.098) | -0.151 (0.103) | 0.894** (0.142) | 0.432 (0.255) | 0.443 (0.211) | 0.196 (0.208) |
| Nbr Citations Made (1000s) | 8.134*** (0.497) | 3.081** (0.610) | 0.177 (0.302) | -0.434 (0.235) | 103.298*** (9.297) | 47.204** (7.022) | 3.524* (1.124) | -3.953 (1.487) |
| Constant | 0.121*** (0.010) | 0.245** (0.042) | 0.575*** (0.033) | 0.624*** (0.030) | 1.402*** (0.107) | 4.591*** (0.461) | 7.691*** (0.443) | 9.157** (1.172) |
| Ratio Clean/Dirty | 304.35*** (16.63) | 212.71** (27.86) | 94.15** (15.77) | 85.39** (14.24) | 501.72*** (16.00) | 239.56** (31.96) | 229.29*** (21.11) | 220.91*** (21.02) |
| Sample | | | Gr. Triadic | Gr. Triadic | | | Gr. Triadic | Gr. Triadic |
| Year FEs | X | X | X | X | X | X | X | X |
| Firm FEs | | X | X | X | | X | X | X |
| Quality Proxies | | | | X | | | | X |
| Adjusted R2 | 0.006 | 0.043 | 0.058 | 0.060 | 0.067 | 0.312 | 0.441 | 0.445 |
| Observations | 2,550,428 | 1,495,048 | 131,564 | 131,564 | 2,550,428 | 1,495,048 | 131,564 | 131,564 |

Linear Regression.

Standard Errors in Parentheses. Clustered at the type and firm level.

Dependent Variable: Percentage of backward citations going to AI or ICT

Note: To allow for an easier interpretation of the main effect captured by β_c , the line “Ratio Clean/Dirty” expresses the magnitude of the effect in percentage term relative to dirty and can be interpreted as the relative absorptive capacity of clean vs dirty. Formally, it corresponds to $\frac{100 \times \beta_c}{mean_d}$, where $mean_d$ is the average percentage of backward citations going to AI (or ICT) in the average dirty family. Quality proxies include the number of citations received within three years, the size of the family and the number of countries where the family was filed. Column 1 and 5 use observations at the family level while the other columns use observations at the family-firm level. Some families are associated with several firms, implying that those families appear multiple times in the data. For this reason, the number of observations in Columns 2 (and 6) could in theory be larger than Columns 1 (and 5). All coefficients excluded from the main tables are shown in the long version of the same table in Online Appendix Table SI3.

electricity is much more ahead in absorbing ICT than dirty.¹⁶

Next, we examine how the relative absorptive capacity of clean vs dirty has changed over time by running a similar regressions as Column 1 and 2 for AI (and Column 4 and 5 for ICT) but for each year separately. We then plot the yearly estimated “Ratio Clean/Dirty” either with or without firm fixed effects on Figure 4. The dotted lines represent a measure of relative absorption arising from intrinsic characteristics and general equilibrium effects alone, whereas the solid lines should be interpreted as a measure of relative absorption that also includes firm composition effects (e.g., changes in the number of firms with high capacity to use the GPT).

We see that the relative absorptive capacity for AI has increased over the years: it is fairly noisy up until around 2002, then becomes positive, and reaches close to 600% in 2018. For most years, it is very similar whether or not firm fixed effects are used in the estimation. Since 2008, however, the within-firm absorptive capacity is consistently higher. This means that AI did not diffuse through all firms at the same pace.¹⁷ We further explore heterogeneity along firm characteristics in the next section.

For ICT, the story differs slightly. The relative absorption is positive and significantly higher for ICT than for AI for most years. AI has recently caught up and, by the end of our sample, we see that clean inventions have a similar lead in both. The difference between the relative absorptive capacity estimated with and without firm fixed effects is larger for ICT than for AI. But, this time, the line without firm fixed effects is on top. This means that firms specialising in clean patenting have an easier time absorbing ICT, compared to other firms.

Finally, in Table 4, we explore whether inventions relying on AI or ICT generate greater value. For this purpose, we proxy “value” by the number of citations received within 3 years of the priority year.¹⁸ First, in Column 1, we see clean families receive about 66% more citations than dirty.¹⁹ This is consistent with prior work by Dechezleprêtre et al. (2017) and implies that clean inventions are more valuable than dirty. Second, Column 2 shows that families citing AI receive about 27% more citations.²⁰ The effect of citing AI declines somewhat when firm fixed effects are included, but the magnitude remains relatively high at around 14%. The interaction between being clean and citing AI is positive and significant implying that the effect of citing

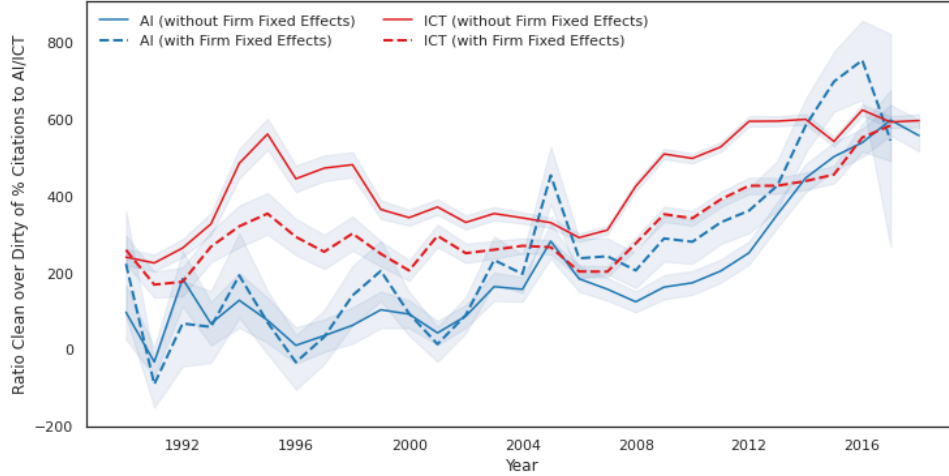
16. Those regressions are shown on Online Appendix Tables SI6 and SI7. For AI, the coefficients on Clean are lower in Electricity compared to Transport. When controlling for quality proxies, the coefficients become insignificant. This indicates that among high-quality electricity families, there is no difference between the ability of clean and dirty technologies to absorb AI. The same coefficient for ICT remains strongly significant however.

17. The specifications using firm fixed effects mechanically drop the families filed by firms that do not have both clean and dirty families. As a result, those specifications capture only the relative absorptive capacity in the context of firms that do both clean and dirty patenting. Intuitively, we can expect those to be large diversified firms. Conversely, the specifications without firm fixed effects contain families associated to any kind of firms: either firms filing both clean and dirty, firms specializing in dirty patenting only, or firms specialising in clean patenting.

18. We run a similar analysis using citations received within 5 years in Online Appendix Table SI10 and find similar results.

19. The specification is log-linear, hence we convert the coefficients in the following way: $100 * (e^{0.508} - 1) = 66.2\%$.

20. Similarly: $100 * (e^{0.240} - 1) = 27.1\%$.



Note: The figure examines how the relative absorptive capacity of clean vs dirty has changed over time. To do so, we run similar regressions as Column 1 and 2 for AI (and Column 4 and 5 for ICT) but for each year separately. We then plot the yearly estimated “Ratio Clean/Dirty” either with or without firm fixed effects. The dotted lines represent a measure of relative absorption arising from intrinsic characteristics and general equilibrium effects alone, whereas the solid lines should be interpreted as a measure of relative absorption that also includes firm composition effects (e.g., changes in the number of firms with high capacity to use the GPT). While the relative absorptive capacity of clean technologies is higher for ICT through most of the period, relative absorptive capacity for AI seems to be catching up in the most recent years. The differences between the solid and dotted lines indicate that firm-level characteristics are playing a significant role, which we further investigate in Section 6.

FIGURE 4
Relative Absorptive Capacity of Clean vs. Dirty Over Time

AI on forward citations is stronger for clean than dirty inventions. This interaction effect is much greater when firm fixed effects are included.

6 FIRM-LEVEL MECHANISMS

In this section, we examine cross-firm variation in the capacity to absorb AI and ICT spillovers into energy inventions. In the previous section, family-level analyses highlighted the role of firms’ characteristics in determining relative absorptive capacity. In addition, recall that our theoretical results show that spillovers from a GPT knowledge stock should be an important determinant of the level of absorption (see Result 2). Arguably, some firms may have access to larger GPT stocks, especially as a large number of firms in our sample patent both in energy and AI or ICT (see Online Appendix Figure SI2).

To estimate the role of GPT spillovers within firms, we construct a dataset at the firm-year-portfolio level where a “portfolio” is a group of patents of a particular type. Firms’ portfolio can be either clean electricity, clean transport, grey electricity, grey transport, dirty electricity or dirty transport. For each firm-year-portfolio observation, we count the number (and percentage) of families in the portfolio that cite at least one AI family. We construct similar measures relative to ICT.

Table 5 provides some examples of top patenting firms, together with the average annual

TABLE 4
Do Families Citing AI or ICT Receive More Forward Citations?

| | (1) AI | (2) AI | (3) AI | (4) AI | (5) ICT | (6) ICT | (7) ICT | (8) ICT |
|--------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| Clean Family | 0.508*** (0.024) | 0.497*** (0.022) | 0.413*** (0.042) | 0.394*** (0.041) | 0.508*** (0.024) | 0.480*** (0.040) | 0.413*** (0.042) | 0.377*** (0.042) |
| Grey Family | 0.324*** (0.019) | 0.322*** (0.017) | 0.265*** (0.032) | 0.262*** (0.030) | 0.324*** (0.019) | 0.342*** (0.022) | 0.265*** (0.032) | 0.262*** (0.027) |
| AI Citing | | 0.240*** (0.046) | | 0.130*** (0.026) | | | | |
| Clean X Citing AI | | 0.061*** (0.014) | | 0.119*** (0.022) | | | | |
| Grey X Citing AI | | 0.008 (0.017) | | 0.042 (0.028) | | | | |
| ICT Citing | | | | | | 0.335*** (0.047) | | 0.156*** (0.039) |
| Clean X Citing ICT | | | | | | -0.111*** (0.005) | | 0.007 (0.020) |
| Grey X Citing ICT | | | | | | -0.126*** (0.016) | | -0.022 (0.027) |
| Constant | -1.407*** (0.088) | -1.385*** (0.093) | -0.960*** (0.090) | -0.945*** (0.095) | -1.407*** (0.088) | -1.401*** (0.090) | -0.960*** (0.090) | -0.957*** (0.091) |
| Sample | | | | | | | | |
| Year FEs | X | X | X | X | X | X | X | X |
| Firm FEs | | | X | X | | | X | X |
| Quality Proxies | X | X | X | X | X | X | X | X |
| Pseudo R2 | 0.282 | 0.284 | 0.338 | 0.339 | 0.282 | 0.285 | 0.338 | 0.340 |
| Observations | 2.55e+06 | 2.55e+06 | 1.47e+06 | 1.47e+06 | 2.55e+06 | 2.55e+06 | 1.47e+06 | 1.47e+06 |

Poisson Pseudo-Likelihood Regression.

Standard Errors in Parentheses. Clustered at the type and firm level.

Dependent Variable: Citations Received Within 3 Years of Priority.

Note: Quality proxies include the size of the family, the number of countries where the family was filed, the logged number of citations made by the family, whether it is granted, and whether it is triadic. All coefficients excluded from the main tables are shown in the long version of the same table in Online Appendix Table SI9.

TABLE 5
Examples of Top Energy Patenting Firms

| Firm Type | Name | Count Energy | % Clean | % Dirty | % Clean Families Citing AI | % Dirty Families Citing AI |
|-------------|-------------------------|--------------|---------|---------|----------------------------|----------------------------|
| Electricity | Sharp Corporation | 256 | 87 | 8 | 1 | 0 |
| Electricity | GE | 115 | 8 | 45 | 14 | 5 |
| Electricity | Kobe Steel,Ltd. | 92 | 22 | 56 | 2 | 0 |
| Transport | Toyota | 3259 | 54 | 11 | 5 | 1 |
| Transport | Bosch | 1215 | 33 | 9 | 11 | 3 |
| Transport | Denso | 1108 | 30 | 27 | 9 | 1 |
| Both | Panasonic | 1096 | 85 | 10 | 2 | 0 |
| Both | Sanyo Electric Co.,Ltd. | 651 | 97 | 2 | 0 | 0 |
| Both | Toshiba | 615 | 80 | 8 | 2 | 1 |

Note: The table shows the number of energy patent families, the percentage of families which are clean or dirty, and the percentage of each which cite AI, for some of the top patenting firms. The values correspond to averages over the period 1990-2018. To classify firms, we calculate the following ratio: $\frac{CountTransport - CountElectricity}{CountTransport + CountElectricity}$, where *Count* refers to the number of family in the category. We define firms as “Transport” if this ratio is greater than 0.5; “Electricity” if it is smaller than -0.5; and “Both” if it ranges from -0.5 to 0.5.

number of clean and dirty families and the percentage citing AI. For clarity, we group firms into three types: those that mostly patent electricity-related inventions, those that patent mostly transport-related inventions and those that do both.²¹ We note that the percentage of families citing AI is always higher in clean portfolios than in dirty but the percentage can go from 3% (e.g., Panasonic) to 11% (e.g., Vestas, a leading wind energy firm).

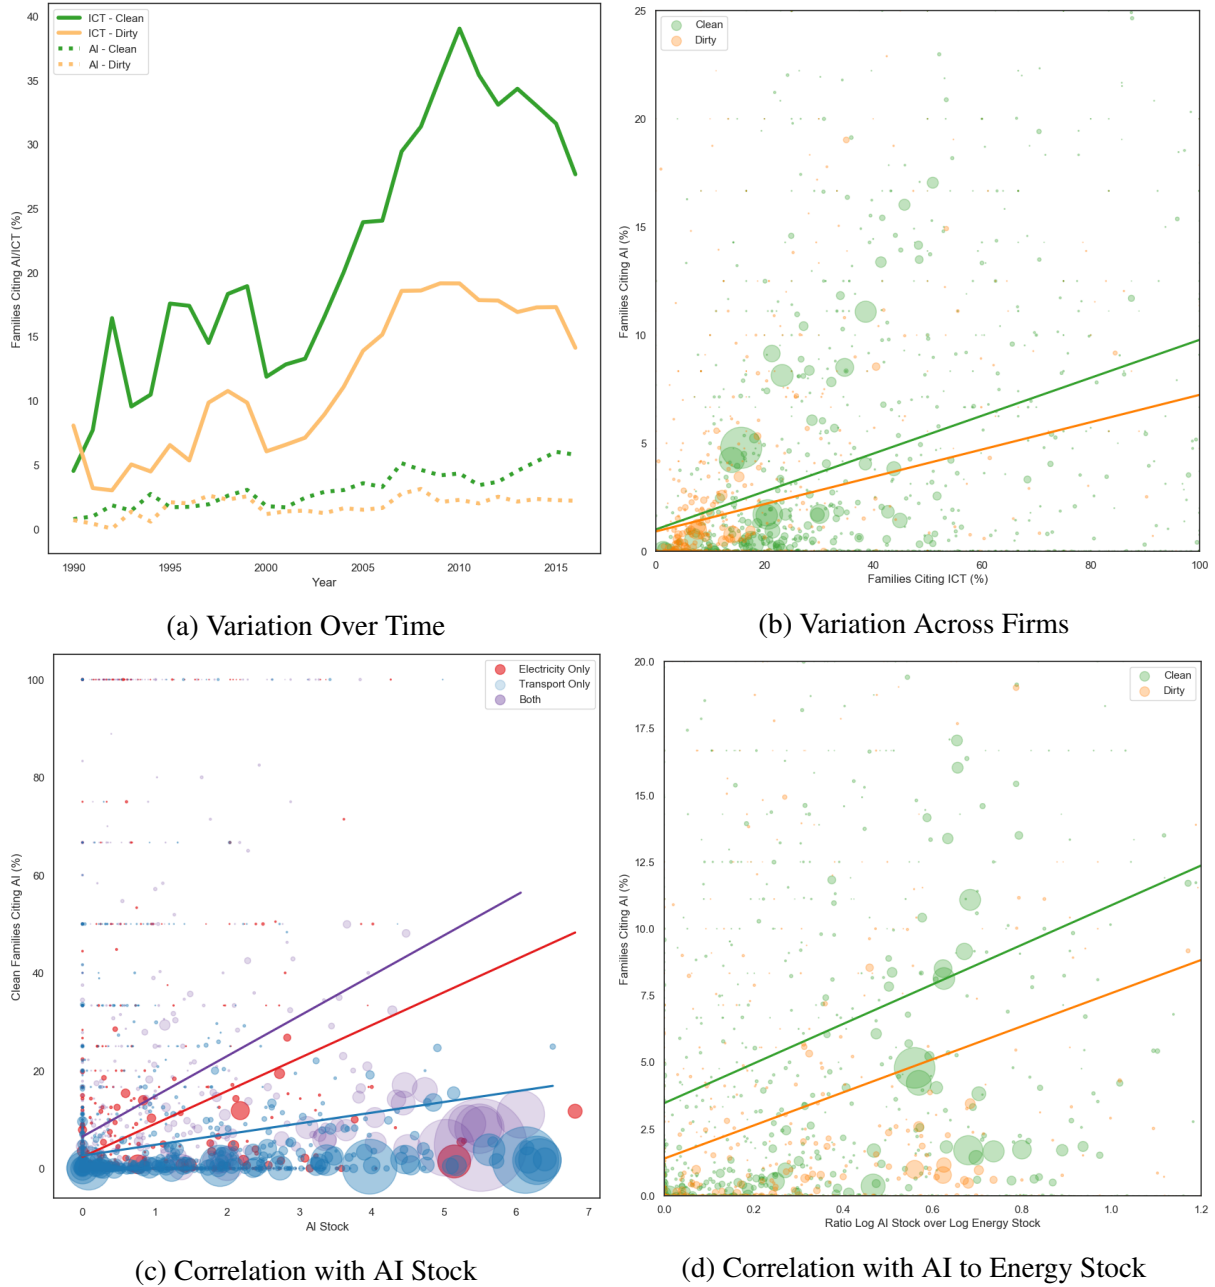
Figure 5 provides more evidence of firm-level variation in absorption capacity. First, on Figure 5a, we see that, the average firm’s clean portfolio always relies more on AI and ICT than dirty.²² For AI, we also note that the gap between clean and dirty has somewhat been widening over time, and especially since 2010. These trends are consistent with what we observed at the family level on Figure 4. We note, however, that ICT absorption has been going down since 2010. This is almost coincidental with the temporary slowdown in clean patenting observed on Figure 2a and 2b.

Figures 5b, 5c and 5d illustrate the variation across firms. On these graphs, each bubble represents a firm-year-portfolio observation where the bubble’s size is proportional to the number of families in the portfolio in that year. The values are calculated for the years 2005 to 2015. First, Figure 5b shows the variation across firms in the percentage of families in clean and dirty portfolios that cite AI or ICT. Unsurprisingly, portfolios rely on ICT in larger proportions than for AI (the y-axis’ scale is larger than that of the x-axis). The solid lines further show that, for a given level of ICT absorption, AI absorption is typically higher in clean portfolios compared to dirty. Again, this is consistent with what we saw in the previous section.

Next, we examine whether firm-level stock of AI knowledge is an important predictor of absorptive capacity. In other words, do firms that filed more AI patent families in the past rely

21. To classify firms, we calculate the following ratio: $\frac{CountTransport - CountElectricity}{CountTransport + CountElectricity}$, where *Count* refers to the number of family in the category. This ratio spans values from -1 to +1, where -1 corresponds to firms doing 100% electricity and +1 100% transport. We define firms as “Electricity” if the ratio is smaller than -0.5; “Transport” if it is greater than 0.5; and “Both” if it ranges from -0.5 to 0.5.

22. To be exact, Figure 5a plots the *weighted* mean share of families that cite AI or ICT in a given portfolio. The mean is weighted by the size of the portfolio so that firms with larger portfolios weigh more in the calculation.



Note: Figure 5a plots the weighted mean share of families that cite AI or ICT in a given portfolio. The mean is weighted by the size of the portfolio so that firms with larger portfolios weight in more in the calculation. On the other figures, each bubble represents a firm-year-portfolio observation and the bubble's size is proportional to the number of families in the portfolio in that year. The values are calculated for the years 2005 to 2015.

FIGURE 5
Variation Over Time and Across Firms in the Percentage of Families Citing AI and ICT

more on AI in their clean portfolios? According to Figure 5c, the answer is a tentative yes. The solid lines highlight that firms patenting in both sectors (purple) or mostly in transport (red) have a higher level of absorption on average. The correlation seems also stronger for those firms relative to those patenting mostly in electricity.

Finally, we explore whether energy incumbents may be at a disadvantage relative to new entrants. To do so, we calculate firms' energy stock as the sum of clean, grey and dirty electricity/transportation patent stocks. Figure 5d shows there is a positive relationship between firm-level AI absorption and the ratio of the firm's AI stock to Energy stock. This suggests that firms with a very high energy stock relative to their AI stock are less able to apply AI to energy technologies, which would be consistent with new or smaller energy firms being better able to absorb AI and ICT into their inventions.

We probe those relationships further using linear regressions. First, we check whether clean portfolios absorb more AI and ICT than dirty ones. The first specification is, therefore, as follows:

$$\begin{aligned} FamilyCountCitingGPT_{jtk} = & \beta_0 + \beta_1 FamilyCount_{jtk} + \beta_c Clean_k + \beta_g Grey_k \\ & + \mathbf{bX}_{jt} + \delta_t + \delta_j + \varepsilon_{jtk} \end{aligned} \quad (15)$$

$FamilyCountCitingGPT_{jtk}$ is the count of families in portfolio k filed in year t by firm j that cite some AI or ICT patents. $FamilyCount_{jtk}$ is the total number of families in portfolio k (filed in year t by firm j). $Clean_k$ and $Grey_k$ are binary variables equal to 1 if the portfolio is clean or grey. We run separate regressions for the transport and electricity portfolios. \mathbf{X}_i is a series of firm-level controls that include total assets, number of employees and years since incorporation. δ_t and δ_j are year and firm fixed effects.

To examine more closely how absorption varies by type of firms, we also include two other control variables interacted with Clean and Grey. The first, "Firm Sectoral Focus", is a variable with values between -1 and 1 that captures the degree of sectoral specialization. It equals to -1 when the firm's energy families are all in electricity, and to 1 when they are all in transport. Specifically, it is equal to $\frac{CountTransport - CountElectricity}{CountTransport + CountElectricity}$, where $Count$ refers to the number of families in the category (for firm j in year t). Second, we include variables that capture specialization within the clean-grey-dirty space. In particular, "Firm Clean Focus" is the percentage of clean families out of all energy families (in year t) and allows us to explore whether the propensity to absorb AI and ICT is higher or lower in firms that specialize in clean energy inventions.

Next, we examine 1) whether higher stocks of AI or ICT facilitate AI and ICT absorption, and 2) whether there is a negative incumbent effect such that a higher stock of energy patents correlates with lower levels of absorption. To do so, we study a regression similar to the one above, but adding terms for the AI, ICT and energy stocks:

$$\begin{aligned}
FamilyCountCitingGPT_{jtk} = & \beta_0 + \beta_1 FamilyCount_{jtk} + \beta_c Clean_k + \beta_g Grey_k \\
& + \beta_2 StockGPT_{jt-1} + \beta_3 StockEnergy_{jt-1} \\
& + \mathbf{bX}_{jt} + \delta_t + \delta_j + \varepsilon_{jtk}
\end{aligned} \tag{16}$$

$StockGPT_{jt-1}$ is the discounted cumulative count of AI or ICT families firm j filed up to time $t - 1$. $StockEnergy_{jt-1}$ is the discounted cumulative count of energy families (of any type) firm j filed up to time $t - 1$. We also include interactions between the stock variables and $Clean_k$ and $Grey_k$. Table 6 and 7 present the regression results. In both tables, Column 1 to 4 focus on AI, Column 5 to 8 on ICT. Columns 1-2 and 5-6 examine “Transport” portfolios, while Columns 3-4 and 7-8 examine “Electricity” portfolios.

First, Columns 1, 3, 5 and 7 in Table 6 show that clean portfolios typically absorb more AI or ICT than dirty: in all columns, the coefficient on Clean is positive and significant at the 1% level, except for AI in electricity.²³ This is indicative that clean technologies have a greater *intrinsic* capacity to use AI in transport, and ICT in both transport and electricity. On the other hand, it seems unlikely that clean technologies in electricity have a much higher intrinsic absorptive capacity for AI than dirty.

Results shown in Columns 2, 4, 6, and 8, however, highlight that the lead of clean over dirty is significantly different for firms with different specializations. For transport portfolios, the lead of clean over dirty in absorbing AI and ICT is mostly present when firms’ patenting concentrates on clean transport.

In electricity, the story is different. First, clean portfolios do not appear to absorb significantly more AI, and, in fact, dirty may lead slightly when firms concentrate on dirty electricity. Indeed, in Column 4, the coefficient on “Clean Portfolio” is negative and significant at the 10% level.

ICT absorption in electricity also presents a mixed picture. Although Column 7 shows that clean indeed leads over dirty, the effect mostly disappears (but remains positive) when controlling for firms’ sectoral and clean specializations. The coefficients are not significant for those variables, but qualitatively, their signs indicate that the gap between clean and dirty is stronger when firms specialize in electricity.

Next, on Table 7, we find that firms with higher AI (resp. ICT) stocks cite more AI (resp. ICT) patents in their inventions (Columns 1, 3, 5 and 7). This is consistent with the earlier theoretical result 2a) which stated that GPT absorption will increase with the existing accessible GPT stock.

When adding firm fixed effects, however, the coefficients on AI stock is no longer significant (Columns 2). In other words, the variation over time within firms in the size of the AI stock explains little of the variation in absorptive capacity. The coefficient on the interaction, however, remains significant at the 10% level.

23. This is consistent with our family-level results presented in Section 5.

TABLE 6
Do Firms' Clean Portfolios Rely more on AI and ICT than Dirty Portfolios?

| | (1) AI | (2) AI | (3) AI | (4) AI | (5) ICT | (6) ICT | (7) ICT | (8) ICT |
|-----------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| Family Count (log) | 0.944*** (0.049) | 0.974*** (0.050) | 0.979*** (0.040) | 0.997*** (0.055) | 0.888*** (0.038) | 0.921*** (0.042) | 0.979*** (0.026) | 0.970*** (0.030) |
| Clean Portfolio | 1.471*** (0.098) | -0.020 (0.297) | 0.156 (0.123) | -0.599* (0.363) | 0.904*** (0.062) | 0.176 (0.182) | 0.495*** (0.057) | 0.188 (0.154) |
| Firm Sectoral Focus | | -0.028 (0.194) | | -0.101 (0.167) | | -0.046 (0.117) | | -0.091 (0.099) |
| Firm Clean Focus | | -0.004 (0.004) | | -0.005 (0.003) | | -0.001 (0.002) | | -0.003* (0.002) |
| Clean X Firm Sectoral Focus | | 0.523*** (0.195) | | 0.179 (0.162) | | 0.201* (0.115) | | -0.142 (0.111) |
| Clean X Firm Clean Focus | | 0.012*** (0.005) | | 0.012** (0.005) | | 0.006** (0.003) | | 0.003 (0.002) |
| Portfolio Type | Transport | Transport | Electricity | Electricity | Transport | Transport | Electricity | Electricity |
| Portfolio FEs | X | X | X | X | X | X | X | X |
| Year FEs | X | X | X | X | X | X | X | X |
| Firm FEs | X | X | X | X | X | X | X | X |
| Firm level controls | X | X | X | X | X | X | X | X |
| Observations | 10,733 | 10,733 | 10,082 | 10,082 | 17,310 | 17,310 | 22,476 | 22,476 |
| R2 | 0.738 | 0.740 | 0.450 | 0.455 | 0.835 | 0.836 | 0.732 | 0.733 |

Poisson pseudo-maximum likelihood regression. Standard errors in parentheses, Clustered at firm level.

Dependent variable: Count of Families citing AI or ICT.

Firm level controls include total assets (log), number of employees (log), and years since incorporation (log).

The label (log) refers to the natural logarithm of 1 + the variable in question.

Note: Firm Sectoral Focus is a variable from -1 to 1 that captures the degree of specialization (in year t). It equals to -1 when the firm's energy families are all in electricity; it equals to 1 when they are all in transport. Specifically, it is equal to $\frac{CountTransport - CountElectricity}{CountTransport + CountElectricity}$, where *Count* refers to the number of families in the category. Firm Clean Focus is the percentage of clean families out of all energy families (in year t). "Grey" and interactions with "Grey" are included in the regressions but left out of the table for clarity. All coefficients shown in the longer version of the same table in Online Appendix Table SI11.

The interaction between the AI stock and Clean is also positive and significant for transport (at the 10% level). It indicates that a higher AI stock facilitates AI absorption more so for clean than dirty portfolios. This is consistent with our theoretical result 2c), which predicted a positive interaction between the intrinsic absorption capacity of a technology and the GPT stock (for the technology which is chosen as the direction of technological change).

The story of AI absorption in “Electricity” portfolios is similar but weaker. Column 3 indicates that firms with a higher AI stock cite more AI in both their clean and dirty portfolios. Coefficients lose significance once adding firm fixed effects (Column 4) highlighting that the correlations in Column 3 were driven by variation across firms.

ICT absorption also seems to increase when firms have a higher stock of ICT patents. However, now the coefficient on the interaction between Clean and ICT Stock is negative and significant, highlighting that the facilitating effect of the ICT stock is stronger for dirty than clean portfolios. Interestingly, those results hold also when including firm fixed effects, thus using only variation over time within firms. This indicates that, as firms grow their stock of ICT, they also increase the proportion of their energy patents absorbing ICT.

Last but not least, we explore whether experience in energy patenting accelerates or slows down absorption in the GPT. Here, the coefficients on “Stock Energy” are negative and almost always strongly significant. This indicates that firms with many energy patents (e.g., incumbents) tend to absorb the GPT less. The coefficients on the interaction “Clean X Stock Energy” are generally not significant, highlighting that this effect is similar for both clean and dirty portfolios. Importantly, this negative incumbent effect is consistent with our theoretical result 3 assuming an aging parameter δ greater than 1; in other words, mature application sectors are less able to absorb the GPT.

TABLE 7
Does Experience in AI, ICT and/or Energy Patenting Facilitate AI/ICT Absorption?

| | (1) AI | (2) AI | (3) AI | (4) AI | (5) ICT | (6) ICT | (7) ICT | (8) ICT |
|---------------------------------|----------------------|---------------------|----------------------|---------------------|----------------------|----------------------|----------------------|----------------------|
| Family Count (log) | 0.982*** (0.044) | 0.922*** (0.045) | 1.064*** (0.126) | 1.017*** (0.042) | 0.939*** (0.030) | 0.887*** (0.044) | 1.032*** (0.036) | 1.016*** (0.026) |
| Clean Portfolio | 0.750*** (0.147) | 1.014*** (0.273) | 0.350*** (0.112) | 0.006 (0.184) | 0.680*** (0.075) | 0.697*** (0.126) | 0.903*** (0.064) | 0.476*** (0.102) |
| Stock AI (log, t-1) | 0.273*** (0.066) | 0.020 (0.096) | 0.333*** (0.082) | -0.067 (0.087) | | | | |
| Clean X Stock AI (log, t-1) | 0.138* (0.073) | 0.137* (0.074) | -0.030 (0.101) | -0.014 (0.047) | | | | |
| Stock Energy (log, t-1) | -0.199*** (0.045) | -0.186** (0.083) | -0.136*** (0.051) | -0.048 (0.063) | -0.169*** (0.026) | -0.244*** (0.049) | -0.250*** (0.023) | -0.183*** (0.045) |
| Clean X Energy Stock (log, t-1) | -0.029 (0.046) | -0.007 (0.065) | -0.112* (0.068) | 0.033 (0.042) | 0.109*** (0.025) | 0.093*** (0.026) | 0.017 (0.033) | 0.101*** (0.032) |
| Stock ICT (log, t-1) | | | | | 0.231*** (0.022) | 0.197*** (0.060) | 0.305*** (0.019) | 0.083* (0.049) |
| Clean X Stock ICT (log, t-1) | | | | | -0.121*** (0.027) | -0.072*** (0.024) | -0.099*** (0.023) | -0.084*** (0.023) |
| Portfolio Type | Transport | Transport | Electricity | Electricity | Transport | Transport | Electricity | Electricity |
| Portfolio FEs | X | X | X | X | X | X | X | X |
| Year FEs | X | X | X | X | X | X | X | X |
| Firm FEs | | X | | X | | X | | X |
| Firm level controls | | X | | X | | X | | X |
| Observations | 26,810 | 9,610 | 41,591 | 9,097 | 26,810 | 15,604 | 41,591 | 20,266 |
| R2 | 0.660 | 0.742 | 0.335 | 0.449 | 0.769 | 0.836 | 0.639 | 0.726 |

Poisson pseudo-maximum likelihood regression. Standard errors in parentheses, Clustered at firm level.

Dependent variable: Count of Families citing AI or ICT.

Firm level controls include total assets (log), number of employees (log), and years since incorporation (log).

The label (log) refers to the natural logarithm of 1 + the variable in question.

Note: Stock Energy corresponds to the firm's total stock of energy patents (i.e., the sum of clean, grey and dirty electricity/transportation patent stocks). We add interactions between the portfolio type and the GPT stock but only show the interaction for "Clean". "Grey" is included in the regressions but left out of the table for clarity. All coefficients excluded from the main tables are shown in longer versions of the same table in Online Appendix Table SI18 and SI19.

7 DISCUSSION AND CONCLUSION

This paper explores theoretically and empirically whether AI has the potential to accelerate clean energy innovation. We first examine how a GPT can affect the race between clean and dirty in a model of directed technological change. We find that, depending on the relative absorptive capacity of clean and dirty, the GPT can break path dependency and help clean technologies compete with dirty. We then use patent data to develop empirical proxies of absorptive capacity and examine how clean and dirty technologies compared over the last two decades. We find evidence, both at the patent family and firm levels, that clean inventions consistently absorb more AI and ICT spillovers than dirty ones. Moreover, this trend has been particularly clear since 2010 for AI.

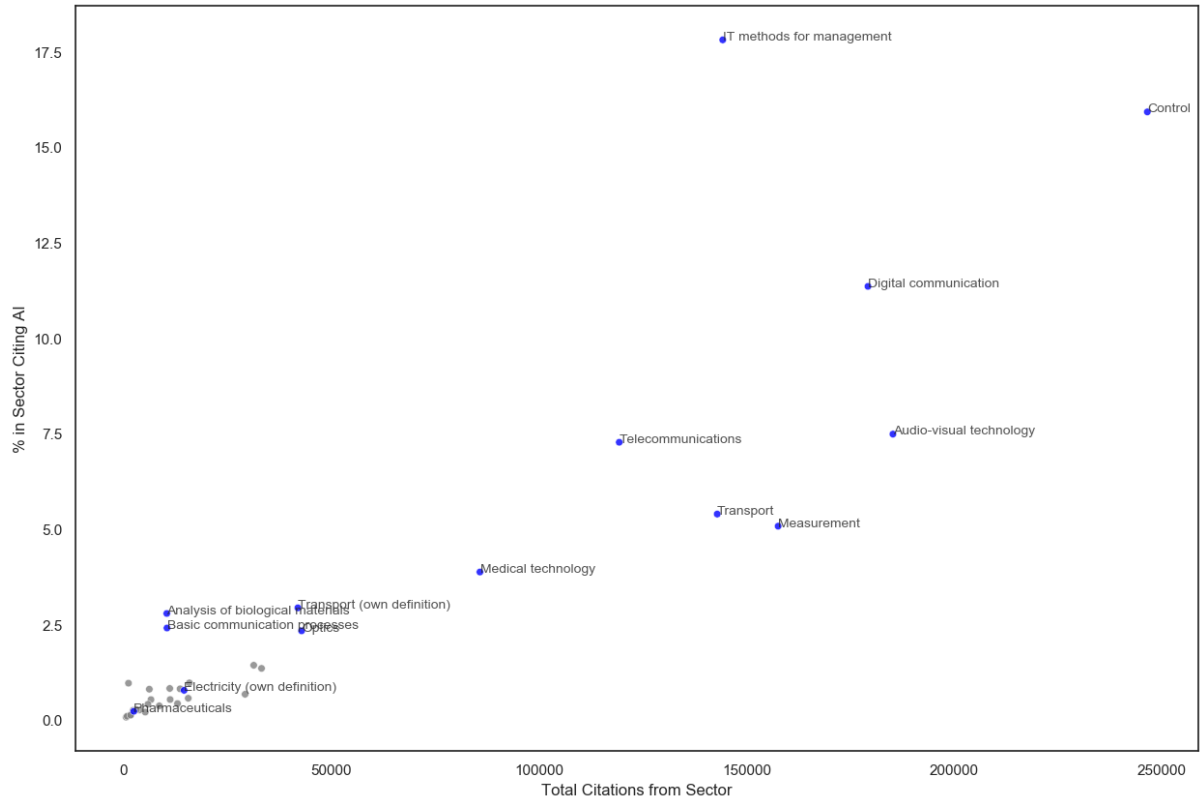
These results provide grounds for cautious optimism regarding the potential for AI to accelerate the transition to clean energy. Indeed, our theory highlights that a GPT can make new technologies more attractive for R&D investments, especially if they more effectively absorb the GPT than incumbent technologies. The theory also shows that this can generate a virtuous feedback. If inventors start preferring clean, they will put more effort into applying the GPT to it, which in turn increases the technology’s productivity, further encouraging innovators to focus on it.

Our firm-level empirical results provide supporting evidence for this process. First, clean technologies’ advantage over dirty ones not only holds, but increases within firms, suggesting that clean tech has a higher intrinsic absorptive capacity and is now the preferred direction of technological change.²⁴ If this is the case, our theory predicts that a higher stock of the GPT leads innovators to put more effort into applying the GPT, *especially to the clean sector*. We find evidence of this in the data. We further find that a firm’s prior focus on energy hinders absorption, in line with the idea that the GPT helps break path dependence and open new opportunities.

Our optimism, however, is cautious. Indeed, the rate of AI absorption is still low. On average, only about 0.3% of backward citations that energy patents make go to AI inventions. Similarly, only about 9% of firms’ patents cite any AI invention. These figures are much lower than the trends for ICT between 1990 and 2010. Figure 6 also puts these statistics into a broader context by plotting them along with other technological application sectors. We see that sectors more closely related to AI, such as “Control” or “Digital Communication,” absorb AI faster. But more distant technological applications, such as “Medical Technologies,” “Telecommunications,” or “Transport overall” (i.e. non-road transport and other aspects of transport innovation, such as automated driving), also absorb it faster than our two focal energy sectors.

Our analysis is a first step in understanding the impact of AI on the transition, and further

24. The within-firm result rules out the alternative explanation that differences in firm capabilities or location are correlated both with working on clean technologies and having more access to knowledge on the GPT.



Note: The figure illustrates variation in the propensity to absorb AI across sectors in the economy. The x-axis shows the total number of families in sector j ; the y-axis shows the percentage of families in sector j that cite AI. We identify sector-specific families using the technical field classification in PATSTAT's Table 230. We also add on the graph dots representing the electricity and transport sectors as defined in this paper, that is, based on the sets of CPC and IPC codes detailed in Online Appendix Table SI1 and SI2. These are different because the PATSTAT technical field are typically broader. For transport, for example, it includes any family related to non-road transport and other aspects of transport innovation, such as automated driving. For readability, we have excluded computer technology from the graph because it is a strong outlier in both the x and y-axes. The sample used consists of families filed between 1980 and 2018, in any country. For the y-axis we only consider citations made within three years of an AI family being filed to control for potentially differing levels of maturity across sectors.

FIGURE 6
AI Absorption Across the Economy

research is needed to address some limitations. First, our analysis focused on comparing broad categories such as clean and dirty, but further work could develop more granular measures to examine the absorptive capacity of specific energy technologies (e.g., solar or wind). Analyzing the heterogeneous impact on different technologies is important to better understand the extent to which the trends are driven by intrinsic technological factors or endogenous processes that are more amenable to policy intervention. Furthermore, we treat all AI patents the same. However, some AI patents probably have a greater potential to be applied broadly (to be a GPT), while others are likely more narrow. Furthermore, AI algorithms are often not patented. Additional work could include citations to scientific publications and distinguish between broad and narrow AI patents.

Finally, our analysis only looks at knowledge spillovers through citations and does not examine the extent to which these spillovers impact the rate of progress of clean technologies, or indeed if and how fast new technologies drawing on AI actually make it to market. Does the integration of AI make technologies more productive and does it accelerate the rate of subsequent innovation? Are there barriers to deployment which mean that many of these technologies are not actually being used? To answer these questions, future research could look at the impact of AI-based energy innovation on firm productivity, sales and subsequent rate of innovation, as well as the degree to which such innovation leads to technology adoption in markets in which it could be most useful.

Despite its limitations, this paper provides the first empirical analysis of innovation spillovers from AI and ICT to clean and dirty technologies on a global scale. Although policymakers often recognize the potential importance of AI and ICT for clean energy, there has been little research on the topic. Our results, therefore, can help inform energy innovation policy. First, our empirical analysis shows that firms are an essential locus for knowledge spillovers between the GPT and energy applications. This suggests that it is worthwhile to increase the joint development of firm-level capabilities in digital and low-carbon technologies.

Our results also suggest that there is a case to support innovations that specifically draw on AI to advance clean technologies. Indeed, those can help spur a positive feedback loop between more AI absorption in clean and more clean innovations in general. Further research, however, is needed to understand the mechanisms better, particularly the role of different actors in catalyzing spillovers (universities, startups, large firms, regional clusters).

SUPPLEMENTARY MATERIAL

The Online Appendix can be found here: <https://www.lse.ac.uk/granthaminstitute/wp-content/uploads/2022/09/AI-green-patents-supplementary-material.pdf>.

A THEORETICAL DERIVATIONS

A.1 Model derivation

The equilibrium must satisfy the following equations:

1. Competitive equilibrium for the two inputs used in producing the final good. Since the final good is produced competitively, the inputs' relative price must satisfy:

$$\frac{p_{ct}}{p_{dt}} = \left(\frac{Y_{ct}}{Y_{dt}} \right)^{-1/\varepsilon} \quad (17)$$

In addition, we normalize the final good's price to 1:

$$\left(p_{ct}^{1-\varepsilon} + p_{dt}^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}} = 1 \quad (18)$$

2. Profit maximization for input j . This determines labor demand L_{jt} and the inverse demand curve of machine x_{jit} . Specifically, labor demand in each sector must satisfy:

$$(1 - \alpha) p_{jt} L_{jt}^{-\alpha} \int_0^1 A_{jit}^{1-\alpha} x_{jit}^{\alpha} di = w_t \quad (19)$$

And the inverse demand for x_{jit} must satisfy:

$$x_{jit} = \frac{\alpha p_{jt}^{\frac{1}{1-\alpha}}}{p_{jit}} A_{jit} L_{jt} \quad (20)$$

3. Profit maximization for the machine producer. The machine producer is a monopolist maximizing $\pi_{jit} = (p_{jit} - \psi)x_{jit}$ where x_{jit} is given by Equation 20. This gives $p_{jit} = \psi/\alpha$. We follow the original model in normalizing $\psi = \alpha^2$, which yields the following relations:

$$p_{jit} = \alpha \quad (21)$$

$$x_{jit} = p_{jt}^{\frac{1}{1-\alpha}} L_{jt} A_{jit} \quad (22)$$

$$\pi_{jit} = \alpha(1 - \alpha) p_{jt}^{1/(1-\alpha)} L_{jt} A_{jit} \quad (23)$$

4. Profit maximization for research scientists who decide which sector to work in.

Using Equation 22, we obtain the following equilibrium production level of input j :

$$\begin{aligned} Y_{jt} &= L_{jt}^{1-\alpha} \int_0^1 A_{jit}^{1-\alpha} (p_{jt}^{\frac{1}{1-\alpha}} L_{jt} A_{jit})^\alpha di \\ &= (p_{jt})^{\frac{\alpha}{1-\alpha}} L_{jt} A_{jt} \end{aligned} \quad (24)$$

Using Equations 22 and 19, we derive an expression for the relative price of clean and dirty inputs as a function of the relative productivities of the two sectors:

$$\frac{p_{ct}}{p_{dt}} = \left(\frac{A_{ct}}{A_{dt}} \right)^{-(1-\alpha)} \quad (25)$$

Using Equations 24, 17 and 25, we obtain an equation for relative employment in each sector:

$$\frac{L_{ct}}{L_{dt}} = \left(\frac{A_{ct}}{A_{dt}} \right)^{-\phi} \quad (26)$$

where $\phi \equiv (1 - \alpha)(1 - \varepsilon)$.

The expected profit Π_{jt} for a scientist doing research in sector j is the expected profit from becoming a monopolist producer of a machine with productivity $A_{jit} = (1 + \gamma)A_{ji,t-1}$, which is (see Eq 23):

$$\Pi_{jt} = \eta_j(1 + \gamma + \beta_j GPT_t) \alpha (1 - \alpha) p_{jt}^{1/(1-\alpha)} L_{jt} A_{jt-1} \quad (27)$$

Using Equation 27 with Equation 25 and 26, we get the ratio of expected profit from doing research in the clean versus dirty sector given by Equation 9.

Next, we obtain a system of equation to solve to obtain the equilibrium by combining Equations 25, 18 and 26 with market clearing $L_{ct} + L_{dt} = 1$, and the expressions for the advancement of the technology frontier in each sector.

A.2 Proof of Result 1

We defined \bar{A}_{ct-1} as the value of A_{ct-1} for which $f(1,0) = 1$. We want to show that $\frac{d\bar{A}_{ct-1}}{dGPT_t} < 0$ and $\frac{d\bar{A}_{dt-1}}{dGPT_t} < 0$.

$$\begin{aligned} f(1,0) &= \frac{\eta_c}{\eta_d} \frac{1 + \gamma + \beta_c GPT_t}{1 + \gamma + \beta_d GPT_t} \left(1 + (\gamma + \beta_c GPT_t) \eta_c \right)^{-\phi-1} \left(\frac{\bar{A}_{ct-1}}{A_{dt-1}} \right)^{-\phi} = 1 \\ \Rightarrow \bar{A}_{ct-1} &= A_{dt-1} \left(\frac{\eta_c}{\eta_d} \frac{1 + \gamma + \beta_c GPT_t}{1 + \gamma + \beta_d GPT_t} \right)^{1/\phi} \left(1 + (\gamma + \beta_c GPT_t) \eta_c \right)^{-\frac{\phi+1}{\phi}} \\ \frac{d\bar{A}_{ct-1}}{dGPT_t} &= \frac{1}{\phi} \bar{A}_{ct-1} \left(\underbrace{\frac{\eta_c}{\eta_d} \frac{(\beta_c - \beta_d)(1 + \gamma)}{(1 + \gamma + \beta_c GPT_t)(1 + \gamma + \beta_d GPT_t)}}_{\sim 0} - (\phi + 1) \underbrace{\frac{\eta_c \beta_c}{1 + (\gamma + \beta_c GPT_t) \eta_c}}_{< 0} \right) \end{aligned}$$

The first term goes as GPT_t^{-2} whereas the second one goes as GPT_t^{-1} . Thus, the sign of the derivative is dominated by the second term, which is negative if and only if $\phi < 1$. The converse derivation works for \bar{A}_{dt-1} .

We now want to show that $\frac{d\bar{A}_{ct-1}}{d\beta_c} < 0$ and $\frac{d\bar{A}_{dt-1}}{d\beta_c} > 0$.

$$\frac{d\bar{A}_{ct-1}}{d\beta_c} = -\bar{A}_{ct-1} GPT_t \underbrace{\frac{\eta_c(1 + \phi(1 + \gamma + \beta_c GPT_t)) - 1}{\phi(1 + \gamma + \beta_c GPT_t)(1 + \eta_c(\gamma + \beta_c GPT_t))}}_{>0} < 0$$

The term in bracket is positive because under the assumption that $\phi < 1$, both the numerator and denominator are negative. Finally:

$$\begin{aligned} f(0, 1) &= \frac{\eta_c}{\eta_d} \frac{1 + \gamma + \beta_c GPT_t}{1 + \gamma + \beta_d GPT_t} \left(1 + (\gamma + \beta_d GPT_t) \eta_d \right)^{\phi+1} \left(\frac{\bar{A}_{dt-1}}{A_{ct-1}} \right)^\phi = 1 \\ \Rightarrow \bar{A}_{dt-1} &= A_{ct-1} \left(\frac{\eta_d}{\eta_c} \frac{1 + \gamma + \beta_d GPT_t}{1 + \gamma + \beta_c GPT_t} \right)^{1/\phi} \left(1 + (\gamma + \beta_d GPT_t) \eta_d \right)^{-\frac{\phi+1}{\phi}} \\ \frac{d\bar{A}_{dt-1}}{d\beta_c} &= -\frac{\bar{A}_{dt-1} GPT_t}{\phi(1 + \gamma + \beta_c GPT_t)} > 0 \end{aligned}$$

A.3 Proof of Result 2

We start with studying the behavior of B_j^* with respect to GPT_t and b_j .

$$B_j^* = \eta_j b_j GPT_t \frac{\alpha(1 - \alpha)}{2\psi} p_{jt}^{1/(1-\alpha)} L_{jt} A_{jt-1}$$

GPT_t and b_j occupy symmetric positions in the equation, so the proof is the same for both variables. We thus proceed studying the behavior with respect to GPT_t .

$$\frac{dB_j^*}{dGPT_t} = (\eta_j b_j \frac{\alpha(1 - \alpha)}{2\psi} p_{jt}^{1/(1-\alpha)} L_{jt} A_{jt-1}) \left(1 + \frac{1}{1 - \alpha} \frac{GPT_t}{p_{jt}} \frac{dp_{jt}}{dGPT_t} + \frac{GPT_t}{L_{jt}} \frac{dL_{jt}}{dGPT_t} \right) \quad (28)$$

WLOG, we describe what happens in the clean equilibrium $s_c = 1$ (the dirty equilibrium can then be analyzed symmetrically). Using Equation 25 together with $A_{jt} = (1 + (\gamma + b_j B_j^* GPT_t) \eta_j s_j) A_{jt-1}$, we see that:

$$\frac{p_{ct}}{p_{dt}} \equiv r_p = \left(\frac{A_{ct}}{A_{dt-1}} \right)^{-(1-\alpha)}$$

This tells us that in the clean equilibrium, $\frac{dp_{jt}}{dGPT_t} = \frac{dp_{jt}}{dA_{ct}} \frac{dA_{ct}}{dGPT_t}$. In the clean equilibrium, we already know that $\frac{dA_{ct}}{dGPT_t} \geq 0$. Since A_{dt-1} is fixed and $-(1 - \alpha) < 0$, the relative price ratio $r_p \rightarrow$

0 as A_{ct} increases. To understand how this affects $\frac{dp_{jt}}{dA_{ct}}$, take Equation 18 (the normalization of the price of the final good) and rewrite it as:

$$\begin{aligned} (p_d^{1-\varepsilon}(r_p^{1-\varepsilon} + 1))^{\frac{1}{1-\varepsilon}} &= 1 \\ p_d &= \frac{1}{(r_p^{1-\varepsilon} + 1)^{1/(1-\varepsilon)}} \\ \lim_{r \rightarrow 0} p_d &= \frac{1}{r} \rightarrow \infty \\ \lim_{r \rightarrow 0} p_c &= r p_d = 1 \end{aligned}$$

These limits imply that $\frac{dp_c}{dGPT_t}$ is negative but goes to 0 (since p_c asymptotes to 1), and $\frac{dp_d}{dGPT_t} > 0$.

We follow a similar reasoning to examine the behavior of equilibrium labor allocations. From Equation 26, we have:

$$\frac{L_{ct}}{L_{dt}} \equiv r_L = \left(\frac{A_{ct}}{A_{dt-1}} \right)^{-\phi}$$

This tells us that in the clean equilibrium, $\frac{dL_{jt}}{dGPT_t} = \frac{dL_{jt}}{dA_{ct}} \frac{dA_{ct}}{dGPT_t}$. With A_{dt-1} fixed and $-\phi > 0$, the relative labor ratio r_L increases as A_{ct} increases. Given the market clearing condition $L_{ct} + L_{dt} = 1$, this implies that $L_{ct} \rightarrow 1$ and $L_{dt} \rightarrow 0$ as GPT_t , and therefore, A_{ct} increases. Hence, $\frac{dL_{jt}}{dGPT_t} \rightarrow 0$.

Thus, Equation 28 now gives us:

$$\frac{dB_c^*}{dGPT_t} \rightarrow \eta_c b_c \frac{\alpha(1-\alpha)}{2\psi} A_{ct-1}$$

Hence, in the clean equilibrium, investments in absorptive capacity by the clean sector increase with the GPT stock, and this is even more so if b_c (the intrinsic absorptive capacity) and A_{ct-1} (the prior stock) are higher.

For the dirty sector, i.e., the sector which is not favored by the equilibrium, investments in absorptive capacity also have a positive relationship to the GPT . This is because $\frac{dp_{dt}}{dGPT_t} > 0$ and does not asymptote, unlike $\frac{dL_{dt}}{dGPT_t}$. However, the derivative remains small because $L_{dt} \rightarrow 0$.

A.4 Proof of Result 3

We now consider the role of the energy stock in investments towards absorptive capacity, allowing for an aging factor that reduces the intrinsic absorptive capacity of more mature technologies.

$$B_j^* = \eta_j \frac{b_j}{A_{jt-1}^\delta} GPT_t \frac{\alpha(1-\alpha)}{2\psi} p_{jt}^{1/(1-\alpha)} L_{jt} A_{jt-1}$$

with $\delta > 0$, the aging paramter.

$$\frac{dB_j^*}{dA_{jt-1}} = (\eta_j b_j GPT_t (1-\delta) \frac{\alpha(1-\alpha)}{2\psi} p_{jt}^{1/(1-\alpha)} L_{jt} A_{jt-1}^{-\delta}) \left(1 + \frac{1}{1-\alpha} \frac{A_{jt-1}}{p_{jt}} \frac{dp_{jt}}{dA_{jt-1}} + \frac{A_{jt-1}}{L_{jt}} \frac{dL_{jt}}{dA_{jt-1}} \right) \quad (29)$$

The reasoning we developed in proof A.3 regarding the derivatives of prices and labor with respect to GPT_t and their limiting behavior carries over to the behavior of these derivatives and limits with respect to A_{jt-1} . Hence, in the clean equilibrium, we have:

$$\frac{dB_c^*}{dA_{ct-1}} \rightarrow (\eta_c b_c GPT_t (1-\delta) \frac{\alpha(1-\alpha)}{2\psi} A_{jt-1}^{-\delta})$$

Clearly, if $\delta = 0$, this derivative is positive. However, if $\delta > 1$, then the aging effect - impeding absorption of the new GPT - is larger than the “building upon the shoulders of giants” effect (innovation opportunities arising from having a larger stock of past knowledge). In this case, the derivative is negative, indicating that effort in absorbing the GPT will decrease with the maturity of the technology.

REFERENCES

- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. “The Environment and Directed Technical Change.” *American Economic Review* 102 (1): 131–66.
- Acemoglu, Daron, Ufuk Akcigit, and William R Kerr. 2016. “Innovation Network.” *Proceedings of the National Academy of Sciences*, 201613559.
- Aghion, Philippe, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen. 2016. “Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry.” *Journal of Political Economy* 124 (1): 1–51.
- Aghion, Philippe, Benjamin F Jones, and Charles I Jones. 2018. “Artificial Intelligence and Economic Growth.” In *The Economics of Artificial Intelligence: An Agenda*, 237–282. University of Chicago Press.
- Agrawal, Ajay, John Mchale, and Alexander Oettl. 2018. “Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth.” In *The Economics of Artificial Intelligence: An Agenda*, 149–174. University of Chicago Press.
- Bose, Bimal K. 2017. “Artificial Intelligence Techniques in Smart Grid and Renewable Energy Systems – Some Example Applications.” *Proceedings of the Ieee* 105 (11): 2262–2273.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. “The Productivity J-curve: How Intangibles Complement General Purpose Technologies.” *American Economic Journal: Macroeconomics* 13 (1): 333–72.
- Cockburn, Iain M, Rebecca Henderson, and Scott Stern. 2018. “The Impact of Artificial Intelligence on Innovation.” *NBER Working Paper* 24449.
- Cohen, Wesley M, and Daniel A Levinthal. 1990. “Absorptive Capacity: A New Perspective on Learning and Innovation.” *Administrative Science Quarterly*, 128–152.
- Crafts, Nicholas. 2021. “Artificial Intelligence As a General-purpose Technology: An Historical Perspective.” *Oxford Review of Economic Policy* 37 (3): 521–536.
- Dechezleprêtre, Antoine, Ralf Martin, and Myra Mohnen. 2017. “Knowledge Spillovers from Clean and Dirty Technologies.” *Grantham Research Institute on Climate Change and the Environment Working Paper No. 135*.
- Dosi, Giovanni. 1982. “Technological Paradigms and Technological Trajectories: A Suggested Interpretation of the Determinants and Directions of Technical Change.” *Research Policy* 11 (3): 147–162.
- Fouquet, Roger. 2008. *Heat, Power and Light: Revolutions in Energy Services* [in En]. Edward Elgar Publishing, January.
- . 2017. “Make Low-carbon Energy an Integral Part of the Knowledge Economy.” *Nature* 551 (7682): S141.
- Grübler, Arnulf, Nebojša Nakićenović, and David G Victor. 1999. “Dynamics of Energy Technologies and Global Change.” *Energy Policy* 27, no. 5 (May): 247–280.
- Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg. 2005. “Market Value and Patent Citations.” *Rand Journal of Economics*, 16–38.

- Helpman, Elhanan, and Manuel Trajtenberg. 1994. "A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies." *NBER Working Paper* 4854.
- . 1996. "Diffusion of General Purpose Technologies." *NBER Working Paper* 5773.
- Inaba, Takashi, and Mariagrazia Squicciarini. 2017. "Ict: A New Taxonomy Based on the International Patent Classification."
- Jaffe, Adam B, and Gaétan De Rassenfosse. 2017. "Patent Citation Data in Social Science Research: Overview and Best Practices." *Journal of the Association for Information Science and Technology* 68 (6): 1360–1374.
- Jaffe, Adam B, Manuel Trajtenberg, and Michael S Fogarty. 2000. "Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors." *American Economic Review* 90 (2): 215–218.
- Jha, Sunil Kr, Jasmin Bilalovic, Anju Jha, Nilesh Patel, and Han Zhang. 2017. "Renewable Energy: Present Research and Future Scope of Artificial Intelligence." *Renewable and Sustainable Energy Reviews* 77:297–317.
- Johnstone, Nick, Ivan Haščič, and David Popp. 2010. "Renewable Energy Policies and Technological Innovation: Evidence Based on Patent Counts." *Environmental and Resource Economics* 45 (1): 133–155.
- Kalemli-Ozcan, Sebnem, Bent Sorensen, Carolina Villegas-sanchez, Vadym Volosovych, and Sevcen Yesiltas. 2015. "How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts and Aggregate Implications." *NBER Working Paper* 21558.
- Koroteev, Dmitry, and Zeljko Tekic. 2021. "Artificial Intelligence in Oil and Gas Upstream: Trends, Challenges, and Scenarios for the Future." *Energy and AI* 3:100041.
- Lanzi, Elisa, Elena Verdolini, and Ivan Haščič. 2011. "Efficiency-improving Fossil Fuel Technologies for Electricity Generation: Data Selection and Trends." *Energy Policy* 39 (11): 7000–7014.
- Lee, Mekyung. 2020. "An Analysis of the Effects of Artificial Intelligence on Electric Vehicle Technology Innovation Using Patent Data." *World Patent Information* 63:102002.
- Lipsey, Richard G, Kenneth I Carlaw, and Clifford T Beker. 2005. *Economic Transformations: General Purpose Technologies and Long-term Economic Growth*. Oup Oxford.
- Lyu, Wenjing, and Jin Liu. 2021. "Artificial Intelligence and Emerging Digital Technologies in the Energy Sector." *Applied Energy* 303:117615.
- Napolitano, Lorenzo, Evangelos Evangelou, Emanuele Pugliese, Paolo Zeppini, and Graham Room. 2018. "Technology Networks: The Autocatalytic Origins of Innovation" [in En]. *Royal Society Open Science* 5, no. 6 (June): 172445.
- OECD. 2016. "Patent Search Strategies for the Identification of Selected Environment-related Technologies." *OECD Environment Directorate*.
- Perez, Carlota. 2009. "Technological Revolutions and Techno-economic Paradigms" [in En]. *Cambridge Journal of Economics* 34, no. 1 (September): 185–202.

- Pichler, Anton, François Lafond, and J Doyne Farmer. 2020. “Technological Interdependencies Predict Innovation Dynamics.” *Technological Interdependencies Predict Innovation Dynamics* (March 2, 2020).
- Popp, David, Jacquelyn Pless, Ivan Haščič, and Nick Johnstone. 2020. “Innovation and Entrepreneurship in the Energy Sector.” *NBER Working Paper* 27145.
- Raza, Muhammad Qamar, and Abbas Khosravi. 2015. “A Review on Artificial Intelligence Based Load Demand Forecasting Techniques for Smart Grid and Buildings.” *Renewable and Sustainable Energy Reviews* 50:1352–1372.
- Rolnick, David, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-dupont, Natasha Jaques, Anna Waldman-brown, et al. 2019. “Tackling Climate Change with Machine Learning.” *Arxiv Preprint Arxiv:1906.05433*.
- Rosenberg, Nathan. 1979. “Technological Interdependence in the American Economy.” *Technology and Culture* 20 (1): 25–50.
- Rosenberg, Nathan, and Manuel Trajtenberg. 2010. “A General-purpose Technology at Work: The Corliss Steam Engine in the Late-nineteenth-century United States.” In *Studies on Science and the Innovation Process: Selected Works of Nathan Rosenberg*, 97–135. World Scientific.
- Trajtenberg, Manuel. 2018. “AI As the Next GPT: A Political-economy Perspective.” *NBER Working Paper* 24245.
- World Intellectual Property Organization. 2019. “Artificial Intelligence.” *WIPO Technology Trends* 2019.